Center for Public Enterprise

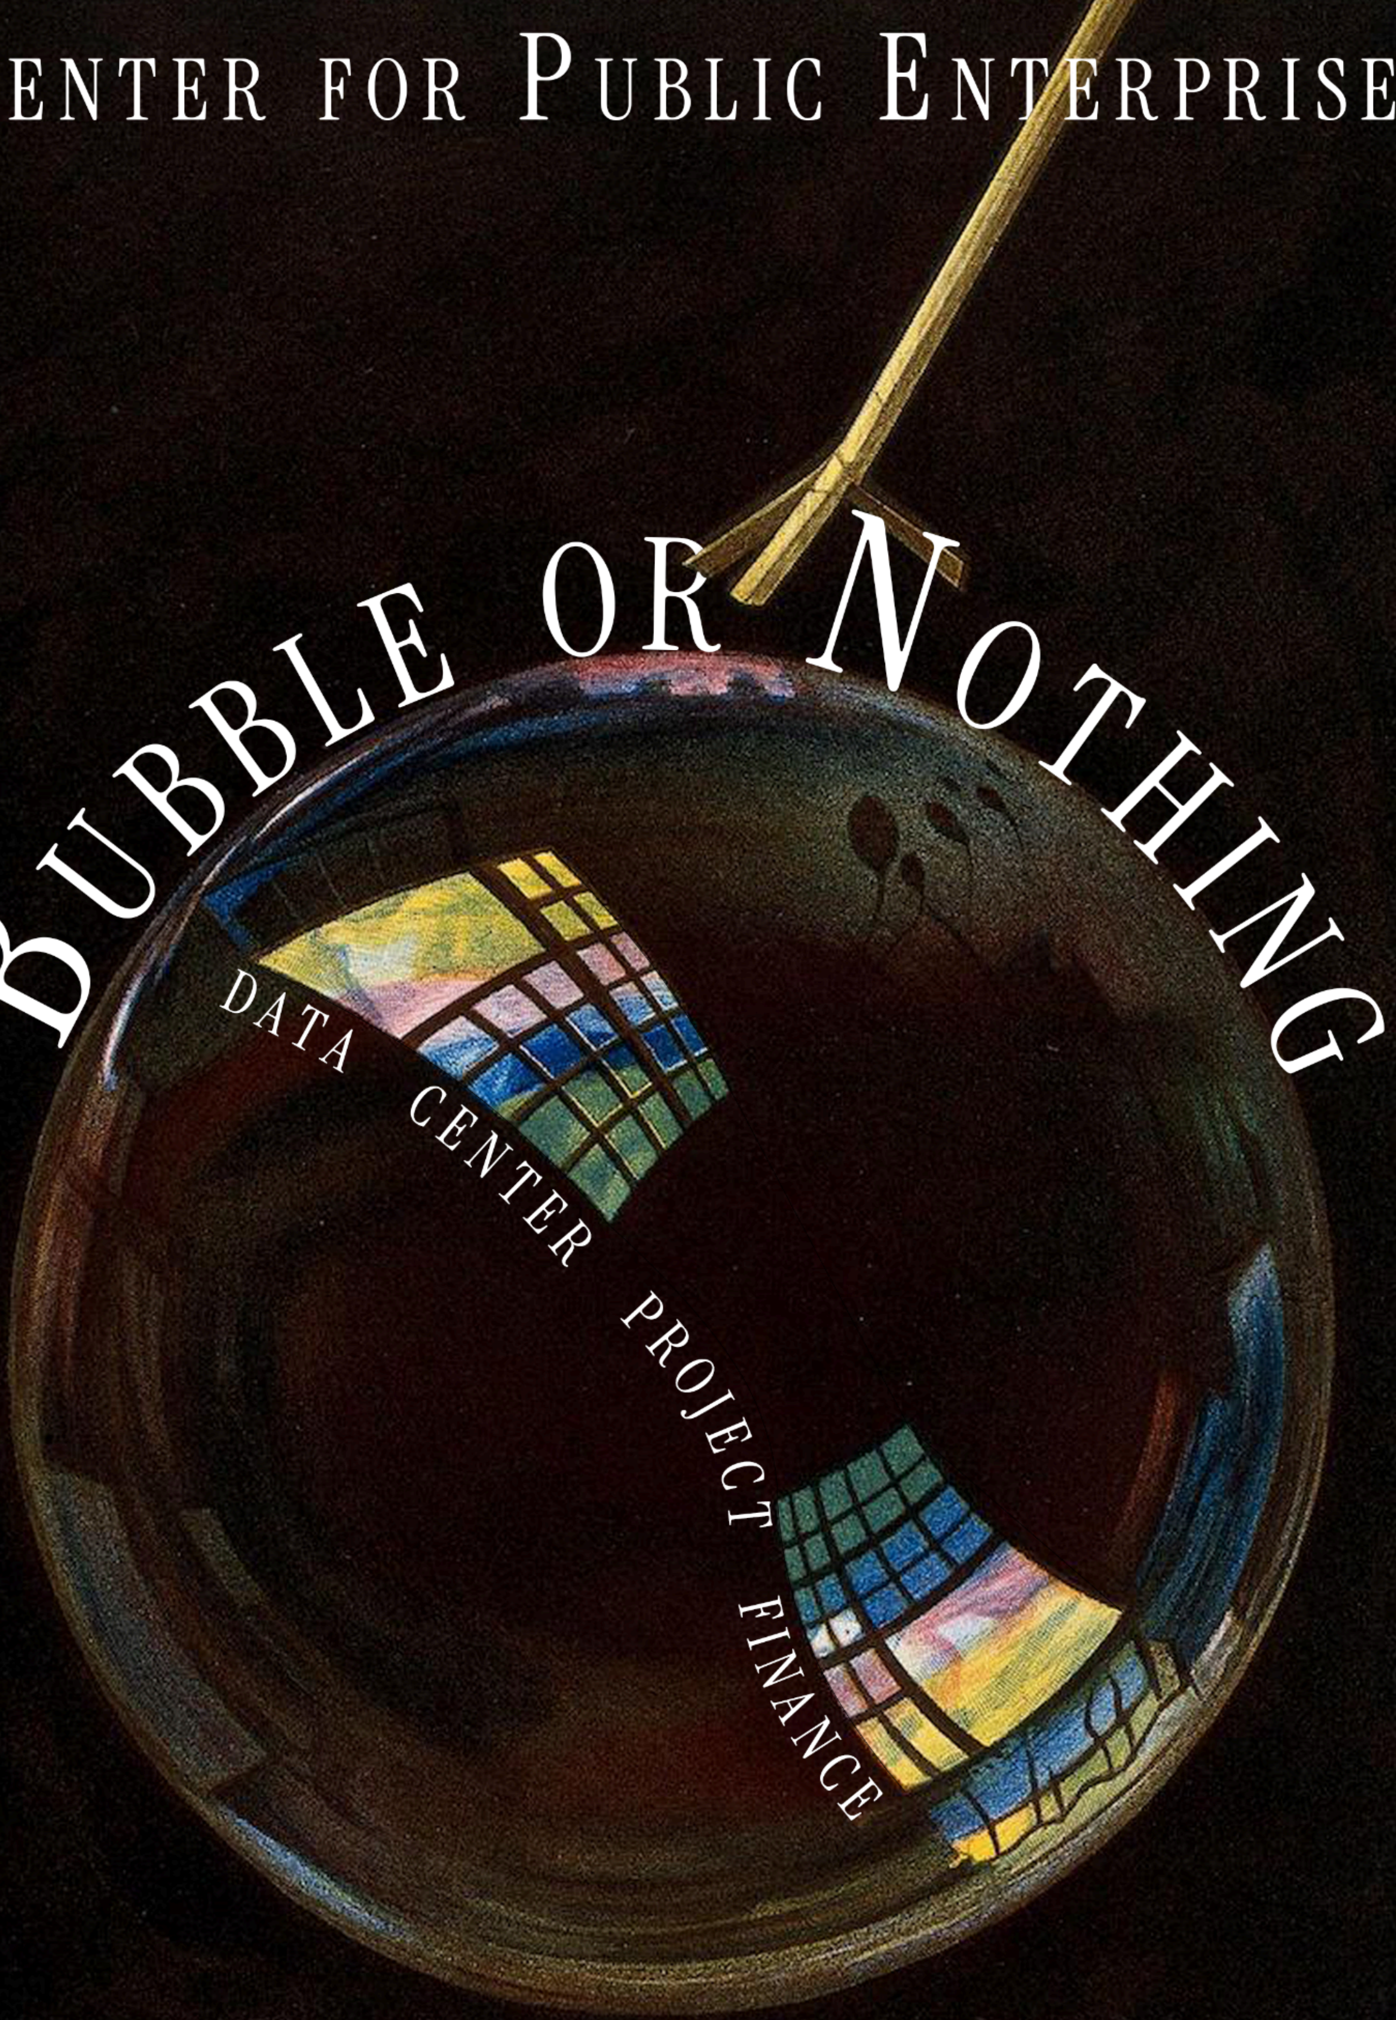

Bubble or Nothing

Data Center Project Finance

Advait Arun

Executive Summary

Should economic conditions in the tech sector sour, the burgeoning artificial intelligence (AI) boom may evaporate—and, with it, the economic activity associated with the boom in data center development.

Policymakers concerned about the deployment of clean energy and compute-focused infrastructure over the long term need a framework for managing the uncertainty in this sector's investment landscape—and for understanding the local and regional impacts of a market correction that strands data centers and their energy projects. This framework requires understanding how a potential downward market correction in the tech sector might occur and, if so, how to sustain investment in critical energy infrastructure assets during potentially recessionary conditions.

Data centers themselves are an asset with the characteristics of both real estate and infrastructure: Data centers have tenants, chiefly large tech companies, that are undertaking expensive long-term capital investment plays with fast-depreciating assets and minimal cash flow to show for them. A careful review of these characteristics suggest that the sector faces the following salient risks:

- **Cash flow uncertainty persists as the cost of providing AI inference services continues to rise.** Leading AI inference service providers are not particularly differentiated from one another; this competitive market structure suppresses market participants' pricing power and prevents them from recovering rising costs.
- **The collateral value of a graphical processing unit (GPU), the sector's keystone asset, looks poised to fall in the near-term.** The value of chips fluctuates depending on uncertain user demand as well as the supply dynamics and technical specifications of new GPUs, now released yearly. The cash flow that GPU collateral can demand is suppressed due to the sector's competitive market structure and the uncertain depreciation schedule of existing GPUs.
- **Data center tenants will undertake multiple cycles of intense and increasingly expensive capital expenditure within a single lease term, posing considerable tenant churn risks to data center developers.** This asset-liability mismatch between data center developers and their tenants will strain developers' creditworthiness without guarantees from market-leading tech companies.
- **Circular financing, or "roundabouting," among so-called hyperscaler tenants—the leading tech companies and AI service providers—create an interlocking liability structure across the sector.** These tenants comprise an incredibly large share of the market and are financing each others' expansion, creating concentration risks for lenders and shareholders.
- **Debt is playing an increasingly large role in the financing of data centers.** While debt is a quotidian aspect of project finance, and while it seems like hyperscaler tech companies can self-finance their growth through equity and cash, the lack of transparency in some recent debt-financed transactions and the interlocked liability structure of the sector are cause for concern.

The first half of the report describes how those four trends interact across the industry: In the short-term, cash flows in the AI sector are wholly insufficient to service liabilities, and, as the AI sector grows, it does not look like this condition will change in the near-term. Given hyperscalers' consistent ability to self-finance new investments out of equity, and (for the most part) their low debt-to-equity ratios, it is not yet clear that the data center boom is yet in such a dire situation. To be sure, these leverage ratios do not account for off-balance sheet vehicles—and debt is playing an increasing role in backstopping some of the latest deals, particularly for non-investment grade issuers. **Each subsection of the first half of the report describes these trends in detail and, aided by T-Chart visualizations, showcases how assets and liabilities are distributed across the relevant market participants.**

The second half of the report proceeds to explain how a market correction might occur and how it would cascade through the sector, making use of a consolidated T-Chart and drawing on the work of pioneering financial economist Hyman Minsky.

The report concludes with a strategic framework for how policymakers ought to evaluate the consumer, regional and energy infrastructure-related impacts of a market correction. Policymakers should be wary of extending tax incentives that do not pay off or yield any benefit, as well as of hitching local budgets to this one growth industry. **More importantly, though, policymakers should prepare an investment strategy centered on acquiring distressed energy infrastructure assets and repurposing them to serve future demand.**

Table of Contents

Table of Figures

I. Introduction

Should economic conditions in the tech sector sour, the burgeoning AI boom may evaporate—and, with it, lots of data center demand. Not only does a prospective market correction create potential stranded assets of data centers, which will lose value as their lease revenue disappears, but it threatens to strand all the energy projects and efficiency innovations that data center demand might have called forth. Estimates of data center load growth vary wildly, but expectations of data center consumption by 2030 range from between 325 and 580 terawatt-hours (TWh), or between 7 to 12 percent of all U.S. electricity consumption.[1] This range suggests substantial uncertainty for the future of U.S. power demand and data center construction.

The volatility of the data center boom has immediate consequences for the public sector due to its outsized effects on the health of state and municipal finances. While state and local policymakers may see the data center boom as a source of local growth, not all localities are adequately reaping the potential tax benefits. At least 10 states already lose over $100 million annually due to longstanding tax exemptions for data centers, and some state tax exemptions preempt localities' ability to collect tax revenues.[2] Data center development is concentrated in certain pockets of the U.S. and, therefore, will not equally support all communities.

In the wider context of U.S. industrial policy, policymakers concerned about clean energy deployment and compute-focused infrastructure deployment over the long term need a framework for managing this uncertainty and understanding its local and regional impacts—especially if a market correction strands data centers and their energy projects. This framework requires understanding how a potential downward market correction in the tech sector might occur and, if so, how to sustain investment in critical infrastructure assets.

The tech sector's ballooning growth belies its fragility; a downward correction within the next few years is plausible. Insofar as energy costs are a small percentage of data centers' total development costs, this report zooms out to understand data centers holistically as an asset class.

An analysis of the key project finance pathways that drive capital flows in the sector suggest that policymakers should be more concerned about the financial fundamentals of data centers and their tenants—the tech companies that are buoying the economy.

Data centers themselves are an asset with the characteristics of both real estate and infrastructure, and thus are often financed off of developers' balance sheets via project finance.[3] Data center developers secure land and energy in order to build facilities for tenants—usually tech companies. Those tech companies are split between cloud services providers and artificial intelligence (AI) companies. While tech companies themselves expect cash inflows to meet their own obligations in the aggregate, their AI-related services often have challenging or uncertain cash flow profiles. Smaller tech companies, which have less internal cash flows to draw from, can rely on their fast-depreciating graphics processing unit (GPU) assets for collateral, to improve their cash flow profiles against which they can borrow.

[1] Goldsmith and Byrum (2025). The authors acknowledge that this range is extremely wide and provide methodology for evaluating this uncertainty.
[2] Tarczynska (2025).
[3] Arun (2024).

With this in mind, a major portion of the booming demand for high-end chips looks more like a scramble for collateral and a strategy to maintain market position rather than a reflection of healthy growth. Therefore, well-capitalized tech companies with wider access to data and alternative revenue streams—chiefly the market leaders, including Google (Alphabet), Meta, Microsoft, and Amazon—seem better-positioned to continue leading the market than their smaller competitors. These companies, as well as their AI-exclusive peers—OpenAI and CoreWeave in particular—are also referred to as "hyperscalers" because they all build the largest data centers. But widespread "roundabouting" or "circular financing" between these hyperscalers, as demonstrated by the circularity of their payments to and equity investments in one another, should also be cause for policymaker worry.

These concerns can fly under the radar because data centers are having their moment in the sun: There has been a marked appreciation in tech stocks over the past three years—not to mention a concomitant reweighting of stock market indices that has left the market incredibly concentrated on the tech sector. This appreciation buoys the sector. But the sector's uncertain revenue profile and the resultant (and growing) asset-liability mismatches between data center developers and their tenants suggest that data centers' continued bankability on capital markets is uncertain in the event of a market correction in the tech sector.

Consumers of AI services are the ultimate revenue source against which the boom is being financed. If those revenues are real and materialize as predicted over the near-term, then these asset-liability mismatches become easier to manage. However, as this report will explain, these revenue projections are incredibly uncertain.

Section II of this report will explain the unique aspects of data center project finance, focusing on data center developers and their tenants in considerable detail before diving into the accounting and financial treatment of GPUs, the complex role of hyperscaler tech companies and "neocloud" companies in providing services, and the rising valence of debt finance across the sector. Each subsection will provide readers with a financial accounting framework that explains how assets and liabilities are distributed between the key market players.

Section III will walk through the drivers of a potential market correction and will suggest how it might play out. It will connect these drivers to a broader understanding of financial bubbles and bailouts, drawing from the work of financial economist Hyman Minsky.

Section IV will offer policymakers a framework for understanding how the risks of a downward market correction will affect consumers, local and regional state budgets, and the energy industry that the data center boom depends on, and, for the latter, will suggest policy recommendations for preserving investments in critical infrastructure. Smart investment planning at the state and local level will be necessary to avoid getting caught unawares by a potential market correction.

Our research is drawn from news sources, academic research, policy work, and some off-the-record conversations. Many reviewers also checked this report's claims. This report will focus less directly on data centers that support cloud functions and more so on AI-focused data centers that house GPU-based infrastructure. While many of these sources stress the scale of the data center investment boom, this report focuses more on the financial dynamics worth policymakers' consideration. The truth is that it is still too early to know much about the scale of the boom and its constituent transactions. But understanding how capital flows in this investment boom helps outline its direction of travel.

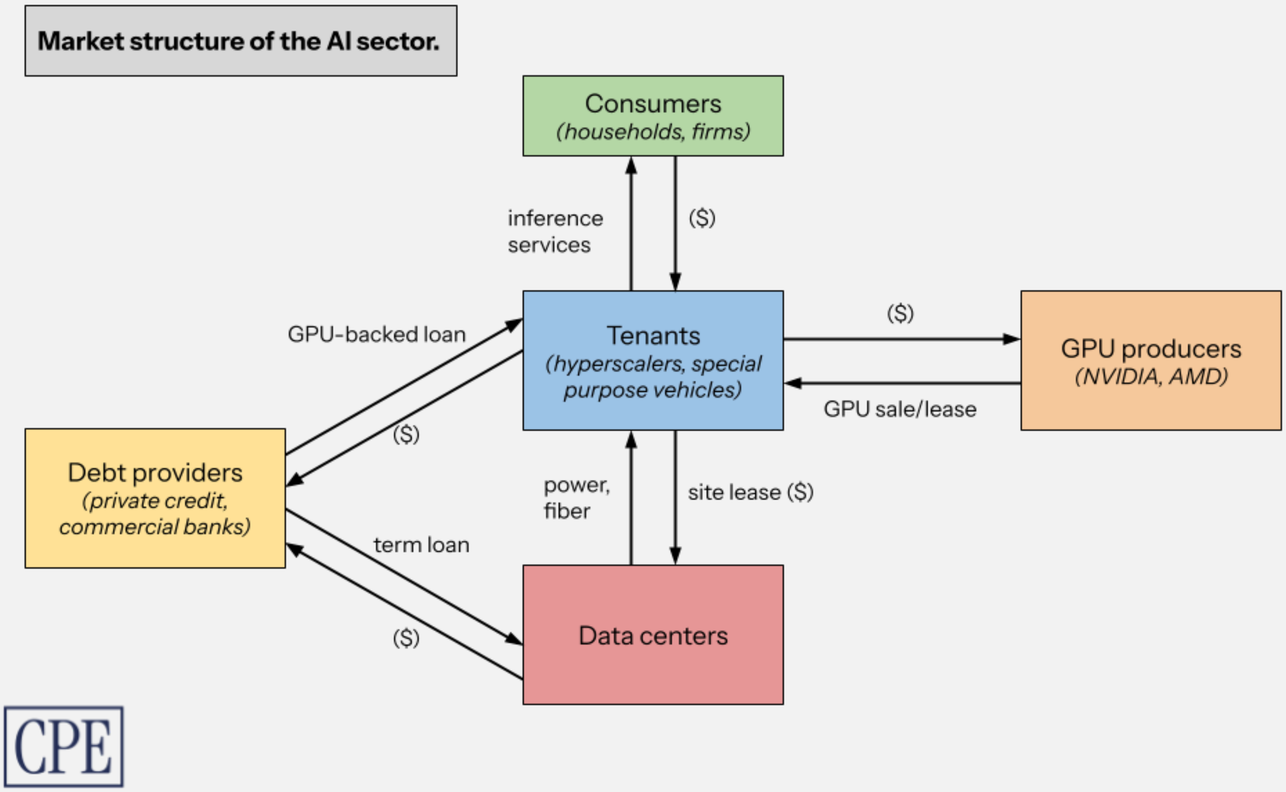

Figure 1.1. *Market structure of the AI sector.*

Table 1.1. *Putting it all together.*

	Data centers		Tenants (Hyperscalers)		Debt providers		GPU producers		Consumers	
	Assets	**Liabilities**	**A**	**L**	**A**	**L**	**A**	**L**	**A**	**L**
stock	Data center facility (20–30)	Data center term loan (20–30)			Data center term loan (20–30)					
stock			GPUs (3–6)	GPU loan (3–5)	GPU loan (3–5)		GPUs (3–6)			
flow				Depreciation (3–6)				Depreciation (3–6)		
flow	Tenant lease income (~10)	Opex + power costs		Tenant lease payments (~10)			GPU sale revenue			
flow			Revenue from inference services / take-or-pay contracts (3–5)	GPU rental costs (4–5)			GPU rental income (4–5)		income	Spending on inference services

II. Data Center Project Finance

The tech sector is the engine of the American economy. As of this year, the technology sector's breakneck investment into data centers is the only tailwind to U.S. economic growth.[4] The "Magnificent Seven" (Mag7) tech stocks—Apple, Alphabet (Google), Amazon, Meta (Facebook), Microsoft, NVIDIA, and Tesla—are growing far faster than the rest of the S&P stock market index.[5] The Mag7's second-quarter earnings this year rose 26 percent while the rest of the S&P 500 rose just 1 percent.[6] And stock market indices themselves have grown incredibly concentrated: the Mag7 now makes up over 50 percent of the market-cap weight of the S&P 500 index.[7] Tech is half the stock market and, despite occasional hiccups, the sector is projected to continue growing at a breakneck pace.

This incredible run for the tech sector is thanks to the emergence of "artificial intelligence" (AI) technologies, particularly "generative" AI technologies such as large language models (LLMs).[8] OpenAI's release of ChatGPT-3 in late 2022, a groundbreaking LLM, sparked a race among the rest of the tech sector to build their own LLMs. Incumbent tech firms and startups alike anticipated upside for themselves chiefly through growth in market share that can be turned into increased services revenue. This growth model, prioritizing market share, is emblematic of Silicon Valley: Tech firms saw what they believed to be a revolutionary new technology and seized up on it to claim the future market for that technology for themselves.

For the publicly traded tech companies of the Mag7, this investment race into LLMs is funded by their already existing and quite lucrative business lines in cloud computing, web services, advertisements, and software development. But many startups receiving venture capital and early-stage equity funding also jumped into the LLM development and AI sector.[9] All of a sudden, "AI" became the next big thing, and the interest has not waned since. Not only do many investors treat it as a world-changing technology that promises both a boom in worker productivity while threatening existing modes of working and living, but many policymakers see dominance in the design and uptake of AI products as a key pillar of the geopolitical competition with China.

Data center capital expenditure is currently the chief motor of growth in the U.S. economy: AI investment accounts for more than 40 percent of U.S. GDP growth this year, and one economist stated that AI capital expenditure contributed more to U.S. growth than did consumer spending over the first half of 2025.[10] The result of this shift in expectations is the breakneck capital expenditures on data centers that seems to be acquiring a momentum of its own—a classic example of "animal spirits" breaking down firms' reluctance to invest. To be sure, there are some

---

[4] Sharma (2025) and Mims (2025b).
[5] Adinolfi (2025b).
[6] Shalett (2025).
[7] Amarnath (2025b) and Amarnath (2025a).
[8] Large language models deploy statistical and reinforcement learning techniques on large volumes of natural language-based data to provide structured and statistically generated text or speech outputs to users.
[9] Initially, many hoped to design their own proprietary models, but the landscape has since changed. Hyperscalers now dominate the market for proprietary "foundational" models, and smaller companies and startups have since shifted to fine-tuning inference and "retrieval augmented generation" for existing models. The report covers this in some detail in the subsections on inference services and on hyperscalers.
[10] Sharma (2025) and Mims (2025b).

indications that demand projections for the sector's future growth are variable[11]—if not overstated. But capital markets are diving into the data center sector nonetheless.

This embrace of large-scale fixed capital investment is a fundamental shift for the sector, which previously prided itself on being capital-light. Google (Alphabet), Meta, Microsoft, and Amazon are expected to spend over $350 billion in AI infrastructure this year and over $400 billion in 2026.[12]

Data centers are a longstanding form of technology infrastructure: Tech companies lease them out as centralized spaces for web services processing, cloud computing, and cloud storage, mostly through the construction of CPU (central processing unit) and server racks. But AI technologies require different kinds of data centers—ones that rely on GPUs (graphical processing units) to train and deploy the kinds of predictive statistical modeling softwares that undergird LLMs.

Data center development is usually executed through a specialized off-balance-sheet vehicle (a special purpose vehicle, or SPV) which separates the data center's owner from its tenant—even if the data center is commissioned and purpose-built by a single tech company to serve its own needs. Tech companies will decide their data center tenancy requirements depending on the kinds of services they seek to provide, whether cloud services, AI model training services, or AI inference workload services. Thus, the industry classifies data center projects as "colocation" and "hyperscaler" projects depending on their intended tenants.[13] Colocation data centers, which have multiple tenants, come in two forms: a wholesale colocation data center is one that was built to serve large tenants which supply much of their own CPU and GPU equipment, while a retail colocation data center is one that was built to serve smaller tenants and offers those tenants limited options to customize their hardware configuration within their leased space. Hyperscaler data centers are large projects built chiefly for the tech giants. While the Mag7 tech giants are almost always the sole tenants of hyperscalers, they are also often anchor tenants for wholesale colocation data center projects.

Data centers are, in other words, a kind of real estate asset for technology companies, many of which are tenants and only some of which—the hyperscalers—are large and equity-rich enough to be their own landlords or the sole tenants of a data center. Data center developers that are not also hyperscalers, such as Skanska, AECOM, Equinix, STACK Infrastructure, and Digital Realty, are separate from the cloud services and LLM services firms that rent from them.[14]

[11] Goldsmith and Byrum (2025).
[12] Elder (2025b) and Acton and Morris (2025).
[13] Equinix (2024).
[14] STACK Infrastructure is owned by private credit firm Blue Owl Capital—a useful acquisition for the private credit firm, which has recently piled into a few high-profile data center transactions.

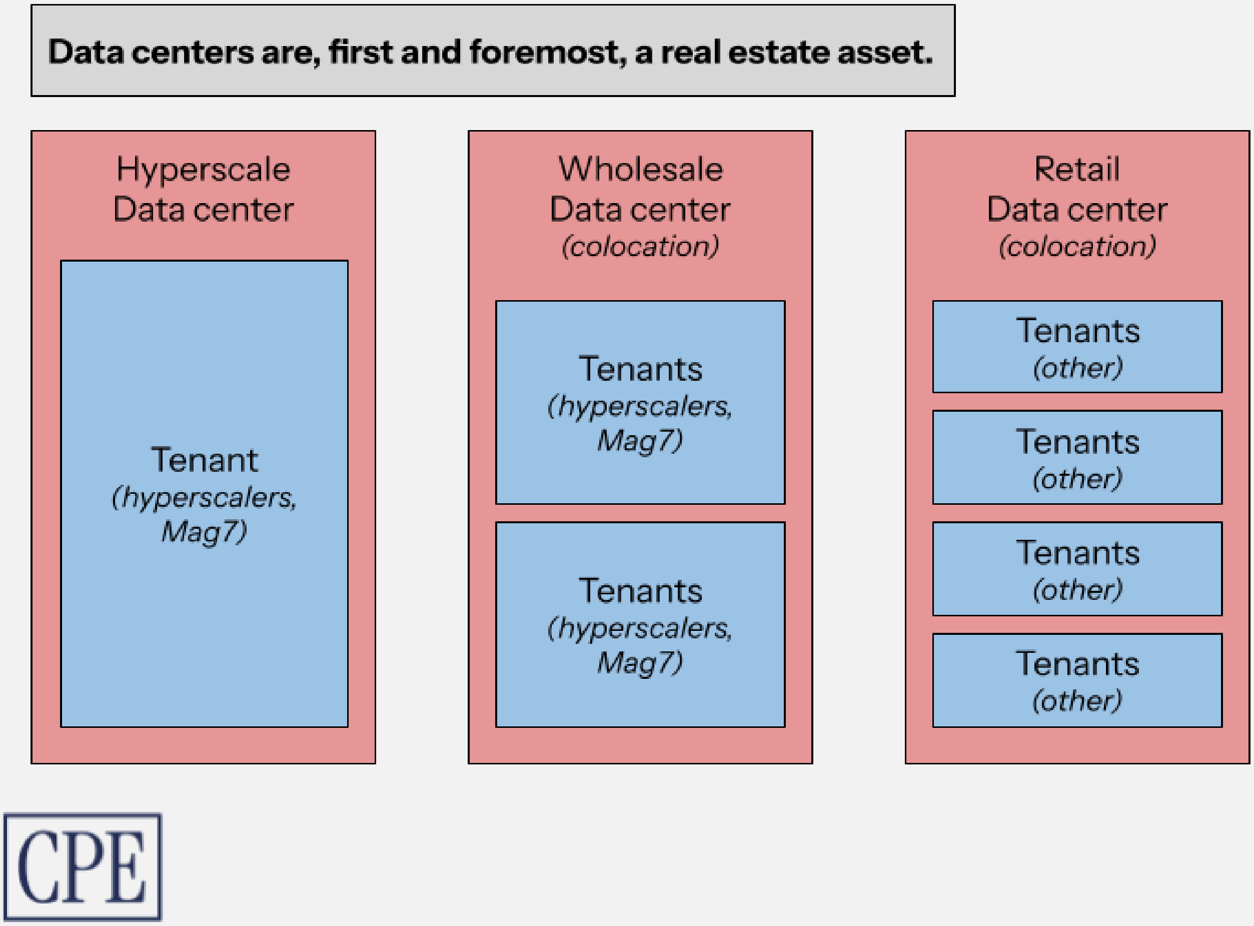

Figure 2.1. *Data centers are, first and foremost, a real estate asset.*

This first half of this section will cover salient aspects of data center developers' financing structures, tenants cash flow and collateral, and some standout examples of transactions. Introducing the dynamics of the sector from the perspective of data center developers' financial and transactional relationships to all their tenants helps illuminate what risks lie where. The second half of this section will clarify the divergence between hyperscaler tech companies and their smaller competitors, the growing role of debt in the sector, and insights from credit rating agencies. It will conclude with a visualization of all the identified assets and liabilities across the sector.

A. Data centers and their tenants

Treating data centers as a real estate asset class allows us to peer into the data center development business model via its peculiar project finance characteristics. Data center developers usually rely on "mini-perm" loans to finance construction: loans with terms of two to five years that support a project's construction and its first few years of operation as it reaches full cash flow.[15] When the mini-perm loan matures, developers can refinance it into a term loan with a

[15] Rivera and Brozynski (2025).

10–20 year maturity on the asset-backed security (ABS) market for project finance loans and on the commercial mortgage-backed security (CMBS) or "commercial real estate" market. This wider market access at refinancing provides developers greater flexibility structuring their longer-term financing.

Access to the ABS market allows data center developers to securitize their lease income and recycle the proceeds of securitization into further investment.[16] And access to the CMBS market in particular allows commercial and regional banks to play a role in providing liquidity to data center operations.[17] Data center-related CMBS issuance hit a high of $4.5 billion earlier this year.[18] And data centers already account for 61 percent of the $79 billion market for digital infrastructure securitizations across both ABS and CMBS asset classes.[19] Banks face notable concentration risks from this arrangement: 50 percent of the market is concentrated among the top 20 developers and operators.[20]

Data center developers face unique refinancing risks, too. As their mini-perm construction loans reach maturity, they need to have certainty to their long-term cashflow projections to access ABS and CMBS markets. That means having tenants with stable cash flows. However, project finance lawyers have mentioned that many data center project finance loans are backed not just by the value of the real estate but by tenants' cash flows on "booked-but-not-billing" terms[21]—meaning that the promised cash flow need not have materialized. Lenders will need to decide the risk premia they charge to developers facing the risk of contracted but unpaid cash flows.

Developers also face an asset-liability mismatch between their tenants' lease terms and their loan terms. Data center tenants have varying lease terms depending on the tenant; retail tenants have the shortest lease terms, anywhere between one to three years, wholesale tenants fall in the middle with five or more years terms, and hyperscalers have the longest leases.[22] Tenants will commit at different points in the development process, but larger data center developers may require hyperscaler tenants to pre-sign leases well in advance of project completion, ensuring they can secure low-cost construction finance, whereas data center developers building retail and smaller wholesale facilities may have more runway to find tenants and tighter borrowing conditions. Most developers require tenants to put up rent deposits of six months to a year, to insure against the cash flow uncertainty that might arise from an unexpected lease cancellation.[23] (This is a drain on tenants' equity.) But tenants, especially larger ones, have an incentive not to cancel—they are often bringing in proprietary equipment and their own technology into the data center, and incurring high sunk costs to do so.[24] They also may prefer to continue undertaking capital expenditures so as not to fall behind their competitors in their race for market share.

Either way, if the majority of developers use mini-perm loan structures, it is possible that their first refinancing does not match up with their tenants' lease and cash flow schedules. Cash flows

[16] Dakin (2025) and Elder (2025b).
[17] Rivera and Brozynski (2025).
[18] CBRE (2025).
[19] Carpenter (2025).
[20] Rivera and Brozynski (2025). Just as a real estate downturn would harm commercial and regional banks that undertake lots of local lending functions, so too would a data center downturn. Section III of this report will discuss these concentration risks in more detail.
[21] Rivera and Brozynski (2025).
[22] Equinix (2024) and Adam et al. (2025).
[23] Papaspanos et al. (2025).
[24] Papaspanos et al. (2025).

at the time of refinancing will set the terms of the ABS or CMBS term loan. Data center developers may find it hard to refinance or, in the event of a foreclosure, sell their projects when tenants' proprietary equipment has limited resale or residual value, complicating its use as collateral.

The good news, for now, is that vacancies seem to be low. Real estate broker Jones Lang LaSalle (JLL) finds that, as of summer 2025, colocation vacancies have neared zero percent and that the data center construction pipeline is 73 percent pre-leased: most developers already have tenants.[25] "Any meaningful loosening of market conditions remains a few years away at minimum," according to JLL.

Assets and Liabilities

Data center developers face refinancing risk on their mini-perm loans covering construction and early asset operations. Their liability, the loan, needs to match assets, their tenants' cash flows—which hyperscalers have on hand but their smaller competitors may not. More important than the mini-perm loan is the term loan, which covers the operational life of the data center and could last between 10 and 20 years.[26] The duration of term loans may not match the duration of tenants' leases: While hyperscalers may be able to sign longer-duration leases, most smaller tenants may not be able to sign leases with terms longer than five to seven years—meaning that tenants are refinancing their leases multiple times within the span of one term loan. Data center developers' ability to access term loans at reasonable interest rates is contingent on lenders' belief that developers can manage and minimize tenant churn in that timeframe. When data center securitizations reach maturity, tenant churn and vacancy rates will determine developers' ability to refinance them.

Table 2.1. *Data centers and their tenants.*

	Data centers		Tenants		Debt providers	
	Assets	**Liabilities**	**A**	**L**	**A**	**L**
stock	"Powered shell" facility + land (20–30)	Data center term loan (20–30)			Data center term loan (20–30)	
stock	"Powered shell" facility + land (2-5)	Data center mini-perm loan (2-5)			Data center mini-perm loan (2-5)	
flow	Tenant lease income (~10)	accounts payable		Tenant lease payments (~10)		

[25] Batson (2025).
[26] Barth et al. (2025).

B. Cash flow uncertainty

It's clear that data centers need tenants with ample and certain cash flow, or at least with enough assets on their balance sheet during their lease terms. This subsection focuses on the problem with the first requirement: AI is nowhere near profitable. (The subsequent subsections will cover costs and assets.) Inadequate consumer spending on AI services is the risk that could unwind investment across the sector.

The sector's revenue growth and revenue projections don't match up; **the sector lacks certainty that sufficient cash flow will materialize to meet the obligations it is incurring for new investment.** A leading partner at Sequoia Capital stated last year that "there was a $500 [billion] gap between the revenue expectations implied by technology companies' AI infrastructure buildout, and actual revenue growth in the AI ecosystem."[27] A recent Bain report stated that the scaling trend of the AI sector required $2 trillion in new revenue—more than five times the size of the existing software subscription market—and, "even with AI-related savings, investors are still $800 billion short in annual revenue required to profitably fund the data centers of 2030."[28]

It is worth stating that there are many unknowns about how, precisely, leading tech and AI companies earn their revenue—or the breakdown of which products and services (e.g., inference service subscriptions, data harvesting, ads, enterprise-level plans) generate what share of their AI-related revenues. The leading tech companies do not formally separate their AI-related income from their broader income statements. The inability to paint a clear picture of revenue flows contributes to this report's judgments of uncertainty in the sector.

Hyperscalers (which the report will discuss in more detail in a later subsection) have ample cash flow due to their status as tech giants. They can draw on their profitable cloud-related and web services business lines to support their breakneck investments into AI models, inference, LLM development, and data center development.[29] So, in the short-term, it does not matter as much to them whether their AI-related investments are profitable as a business unit. Yet they continue to lose money in their AI business lines—their revenue falls so far short of their capital expenditures in that sector: Back-of-the-envelope math using hyperscalers' public statements and publicly available news sources suggests that hyperscalers have invested over $560 billion into AI technology and data centers between 2024 and 2025 and have reported revenues of just $35 billion—in other words, they have made no profit.[30] xAI, a newer hyperscaler, is reportedly spending $1 billion *per month* and is only expected to earn $500 million this year.[31] Recent forensic accounting suggests that OpenAI lost *$11.5 billion* between July and September 2025, or $3.8 billion per month.[32]

[27] Kinder (2024).
[28] Crawford and Hoecker (2025) and Brown and Whelan (2025).
[29] Acton and Morris (2025).
[30] Zitron (2025a).
[31] Nguyen (2025).
[32] Rosoff (2025).

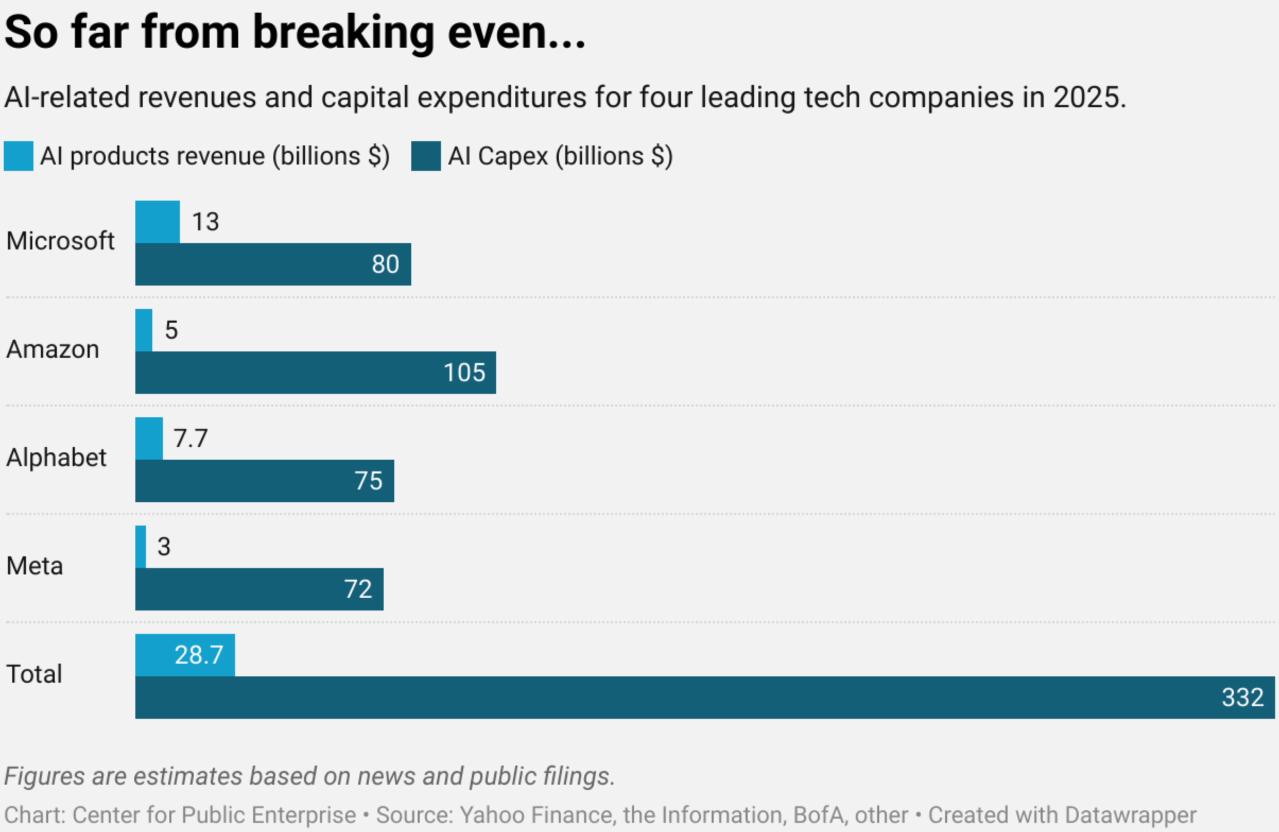

Figure 2.2. *So far from breaking even...*

It is hard to untangle a coherent story about the cash flow growth in the AI sector. Measurements of revenues—usually, annualized revenues—relative to annual capital expenditures are the best indication of the sector's performance in the present. The revenue gap figures stated here, however, give no indication as to the sector's future growth. (Other prominent tech companies, such as ride-sharing app Uber, went over a decade without being profitable.) Optimistic projections of the sector's growth potential abound, but, with varying time scales and varying levels of certainty, these projections are hard to compare. As this report will detail, the commoditization of AI inference services and the concomitant pressure of price competition may prevent demand growth from transforming into revenue growth. We take for granted only that cash flows are inadequate in the nearterm and argue that they appear incredibly uncertain over the medium-term—and leverages the certainty of uncertainty as its starting point for judging the health of the sector.[33] The biggest players in the sector are investing either in

[33] A recent analysis of companies deploying AI found that 95 percent made no returns on integrating AI into their business lines. This data point should not be used to buttress the broader argument about uncertain cash flows. While the analysis is interesting, it does not cover a very large number of firms, and it stands to reason that "integrating AI" is an open-ended task with no clear relationship to productivity or profit. What's more interesting is how AI tools fit into firms' cost structures. Whether firms are are integrating other companies' AI technologies into their workflows (likely via a "wrapper" or an integrated development environment) for productivity improvements or designing and deploying their own models themselves for internal and customer-facing use, they rely in some way on data centers to train the models they use and to answer inference queries. See: Challapally (2025) and Mims (2025a).

anticipation that revenue projects are real or that they will be well-positioned to capture profits through a rise in markups in the event of sectoral consolidation. As analyst David Cahn puts it, "the end customer is not funding the supply chain. Rather, the supply chain is funding its own buildout, in the hope that more paying customers will eventually arrive … [The tech giants are] taking on as much demand risk as they possibly can, and driving the supply chain toward greater and greater [capital expenditure] escalation."[34]

As such, the sheer lack of profit across the AI economy is, in general, a threat to data center tenants' ability to make lease payments should capital investment flow into the sector halt. Especially when data center developers can secure long-term financing via "credit tenant lease financing" structures—where the collateral for their debt comes from the creditworthiness of their tenants—the lack of cash flow threatens a reassessment by lenders, thereby threatening both tenant and landlord.[35]

Profits are a function of revenues and costs. Revenues are uncertain and inadequate—and, as the next subsection will detail, costs are rising. Notable hedge fund manager David Einhorn is wary of the uncertainty in the sector: "The numbers that are being thrown around are so extreme that it's really, really hard to understand them. I'm sure it's not zero, but there's a reasonable chance that a tremendous amount of capital destruction is going to come through this cycle."[36] And, as the CEO of leading infrastructure investor I Squared explained his reluctance to invest into data center infrastructure, "when you start unpeeling the onion, it's much more volatile and uncertain."[37]

Assets and Liabilities

The threat of inadequate consumer spending on inference services is the keystone risk that threatens to trigger all the other refinancing risks discussed in this report. If households and firms do not increase their spending on AI services in line with tenants' growing capital expenditures—covered in the next subsections—then tenants will not have adequate income with which to amortize their capital expenditures and their lease payments to the data centers they are renting from.[38]

Table 2.2. *Cash flow uncertainty between tenants and customers.*

	Tenants		Consumers	
	A	**L**	**A**	**L**
flow	<u>Revenue from AI services</u>	Cost of AI services	income	<u>Spending on AI services</u>

[34] Cahn (2025) and Cahn (2024).
[35] Norton Rose Fulbright (2024).
[36] Fiegerman and Reinicke (2025). CPE has previously cited Einhorn before in our work covering the shale oil and gas boom and bust cycle and its lessons for policymakers interested in deploying geothermal energy. Einhorn's firm, Greenlight Capital, successfully shorted Lehman Brothers in 2008 and made bold predictions about Pioneer Energy, a shale oil extraction firm, during the height of the shale boom. See: Arun (2025b).
[37] Hoffman (2025).
[38] Tenants such as hyperscalers may be better off insofar as they have existing income from other business verticals to draw from.

C. The rising costs of AI training and inference

AI-related business models can be separated into two broad categories, training and inference, the latter of which is more important and commands the lion's share of AI-related business.

AI training is the process of developing and training models using large datasets such that those models can be deployed for and used by customers. Training often requires the highest-end GPUs, high volumes of power and consistent grid access for volatile power load fluctuations,[39] and advanced cooling infrastructure—all of which is to say that it asks a lot of data centers while it earns no revenue. New models are more and more expensive to train; some analysts have judged that each successive iteration of ChatGPT was three to five times more expensive to train than its predecessor.[40] Power demand for training "frontier" models is rising *exponentially*.[41] The rise of large open-source and proprietary "foundational" generative AI models may lead to the decline of training for new models and a shift toward fine-tuning inference processes that use foundational models; the report will discuss the preponderance of "foundational" models in the forthcoming subsection on hyperscalers.

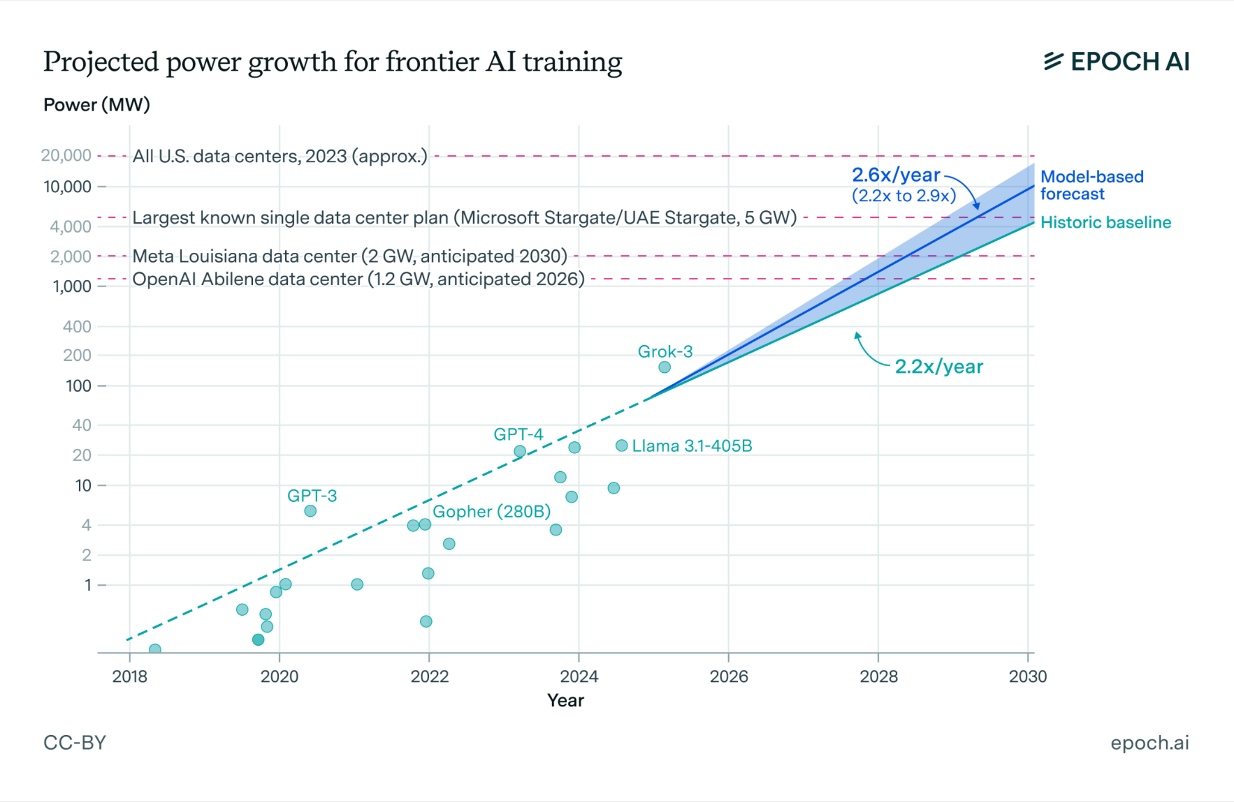

Figure 2.3. *Projected power growth for frontier AI training. See: You and Owen (2025).*

[39] Walton (2025a).
[40] Brown and Whelan (2025).
[41] You and Owen (2025).

The conventional revenue stream in AI technology comes from **inference**, the process of deploying a trained model for widespread customer use, allowing the model to "inference" accurate responses to user prompts, thereby providing value to customers.

Inference-focused data center tenants do not necessarily require the highest-end hardware. But, like cloud-focused tenants, they must site themselves close to the population centers they serve to minimize latency. They must also rent from data centers that can promise year-round uptime and reliable power supply—specifically 99.999 percent uptime (referred to as "five nines" uptime), which translates to about 5 minutes of downtime per year. And, most importantly, as the cost of training rises, inference services must earn more and more income to pay back that upfront cost.[42]

Some analysts have claimed that inference costs will fall rapidly as demand increases and as GPU technology and data center efficiency improve, providing value to customers and generating further uptake from users. These claims are misleading, however. While the cost of inference per "token"—a measurement of a unit of data contained in a prompt—is falling, the total cost of inference continues to rise as the amount of tokens per prompt rises (the effect of both more complex prompts and better model training), as inference models are now programmed to double- or triple-check their outputs, and as rising inference demand forces inference-focused data center tenants to build "scale-out" GPU clusters to manage this rising workload.[43] In other words, efficiency gains on the margin from rising inference compute demand are currently offset by further capital expenditure on GPUs.

To summarize, inference loses money. The CEO of productivity company Notion mentioned that his profit margins dropped 10 percent due to payments for inference costs.[44] And coding firms such as Cursor and Replit that use "wrappers" of OpenAI and Anthropic inference models—thus pegging their end-user prices to the cost of purchasing OpenAI and Anthropic inference models—have had to hike their prices to keep up.[45]

Another possible reason that prices for model inference are not falling is that that many inference-focused firms, motivated to compete for market share in the inference subsector, continue to buy the highest-end GPUs to prove their relevance to their users, who seem willing to pay for accuracy; the price of high-quality inference is thus linked to the price of the latest GPU.[46]

One analyst suggested that there's a prisoner's dilemma of sorts in inference pricing: If every inference application prices its services based on quality and charges by usage, the market might remain stable, but—because market share is more important than margins for the equity investors and venture capital investors supporting these inference firms—every firm has a greater incentive to offer flat-rate pricing with unlimited use, triggering a race to the bottom. "So everyone defects. Everyone subsidizes power users. Everyone posts hockey stick growth charts. Everyone eventually posts 'important pricing updates.'"[47]

---

[42] Brown and Whelan (2025).
[43] Mims (2025a).
[44] Mims (2025a).
[45] Journalist Christopher Mims suggested that this dynamic "raises a thorny question about the state of the AI boom: How long can it last if the giants are competing with their own customers?" Mims (2025a).
[46] Szyszka (2025) and Ding (2025).
[47] Ding (2025).

Expectations that companies' profits would rise on the back of falling inference cost projections have not borne out due to the growing capital expenditure needs of inference—both on the user side, through more complex prompt and the demand for more accurate responses, as well as on the market structure side, through the competition for market leadership vis-a-vis the use of high-end GPUs and consumer demand for state-of-the-art inference products.[48]

Assets and liabilities

Even as households and firms spend on inference services, their spend needs to cover the rising costs of model training and inference.

Table 2.3. *The rising costs of AI training and inference.*

	Tenants		Consumers	
	A	**L**	**A**	**L**
flow	Revenue from AI services	Cost of AI training and inference	income	Spending on AI services

D. Cost compression in the "neocloud" sector

There is also a growing **"compute as a service"** or **"neocloud"** subsector which aims to rent out GPU capacity for training and inference (also known as "GPUs as a service," or "GPUaaS"). These "neocloud" companies, including bigger names such as CoreWeave, Nebius, and Lambda, provide GPUs for rent to market actors looking to access AI training and inference capabilities at competitive pricing relative to the hyperscalers (which can charge higher prices for proprietary services running on higher-end GPUs). These companies buy GPUs from NVIDIA and AMD, the two major GPU producers, and often rent their services via virtual machine through the cloud to market actors. While neoclouds first entered the market to compete with hyperscalers, but it turned out that even hyperscalers find value in neoclouds: OpenAI, for example, recently announced a deal with CoreWeave where it would make use of CoreWeave's GPU compute power. The problem faced by neocloud companies is that they offer essentially identical services. They, too, are competing for market share and, by virtue of their similarities, must compete on price and on the quality of GPUs, which they all source from the same handful of companies.

The next subsection will detail how GPUs lose pricing power over time. This presents a risk for neocloud companies in particular: Both neocloud companies and hyperscalers alike are in the market to purchase new GPUs to service their demand, but hyperscalers can use their market share and pricing power to lay claim to the best inventory first. Neocloud companies' business models are thus structurally dependent on hyperscalers' decision-making.

Untangling the interplay between user demand for inference, the cost of inference, and rise of neocloud inference providers is challenging. Regardless of how "efficient" inference services get, there is a palpable tension between the persistence of higher costs for high-end inference services and, akin to the prisoner's dilemma described in the previous subsection, the threat of

[48] Mims (2025a), Szyska (2025), and Ding (2025).

falling rental revenues for inference providers in the absence of being able to undertake more capital expenditure on more or newer GPUs.

Assets and Liabilities

GPU producers sell to both neocloud and hyperscaler tenants, both of which offer similar inference services. While hyperscalers also purchase inference services from neocloud companies, they are also the first movers in the market to purchase new GPUs.

Table 2.4. *Cost compression in the "neocloud" sector.*

	Tenants (Neocloud)		Tenants (Hyperscalers)		GPU producers	
	Assets	**Liabilities**	**A**	**L**	**A**	**L**
flow	Hyperscaler service rental payments (accounts receivable)			Neocloud service rental payments (accounts payable)		
flow		GPU purchase costs (accounts payable)		GPU purchase costs (accounts payable)	GPU sale revenue (accounts receivable)	
other		*market follower*	*market leader*			

E. GPU depreciation and collateral risk management

GPU market dynamics define the sector's pricing structure. The cost of "renting compute," or renting GPU processing power, from neocloud companies, particularly for the use of capacity from an NVIDIA H100 chip plummeted between 2023 and 2024, particularly as H100 order books cleared out and GPU capacity came online, flipping a seller's market environment of constrained GPU supply into a buyer's market glut—at the same time as NVIDIA released higher-end chips such as the H200 and Blackwell, shifting inference demand away from H100s. An H100 which used to cost $30,000 and commanded a rental price as high as $8 per hour now costs as little as $1 per hour to rent.[49] As the previous subsection made clear, neocloud companies in particular are disproportionately exposed to the risk of falling asset values—but this risk is a threat to the whole sector.

This loss of GPU pricing power is where the **depreciation** of GPU assets starts to bite the sector. This report defines depreciation to mean the change in *marketable* value of existing GPUs.[50]

[49] Elder (2025a).
[50] This value is theoretical and remains subject to considerable uncertainty. A measurement of marketable value should capture how much an asset would sell for in the event that a particular firm was liquidating its

Firms' internal accounting tries to capture this value for tax and liability management purposes, as well as for proper equity valuation.[51] Physical deterioration or obsolescence of GPUs is something that accounting practices, in theory, are intended to capture in recorded book values.

Project development lawyers and engineers alike state that GPUs depreciate over three to four years[52]—although it could take longer if its utilization rate were low. But the internal accounting treatment of GPU depreciation varies. Different companies depreciate GPUs at different rates—CoreWeave, for example, depreciates its GPUs over six years while Nebius depreciates over four years.[53] CoreWeave's six-year accounting treatment is designed to reduce the reported value of annual expenses and thereby reduce the annual revenue that individual GPUs need to return, artificially forcing the "hurdle rate" of investing in a GPU down and thereby smoothing operating profits.[54] Of course, this accounting change might be ambitiously optimistic if high utilization of new GPUs erodes their capacities over a shorter time frame (e.g., three years).

It matters a lot for how these companies report their expenses: Analysts at Cerno Capital estimated that data center depreciation expenses for Microsoft, Google (Alphabet), Meta, and Amazon would fall 54 percent, from $51 billion to $28 billion, if they extended the useful life of all their data center assets to six years.[55] Extending depreciation schedules is a canny risk management strategy.[56]

The relationship between internal depreciation accounting and actual technological deterioration, however, is less important than the impact of the arrival of new higher-end GPUs on the market.[57] GPU market-leader NVIDIA has historically released new GPU models every two years—but now plans to release new GPUs on a one-year cycle.[58] Pushing new GPUs onto the market more frequently will considerably decrease the marketable values of all existing deployed GPUs absent the kind of demand growth that makes these investments realize cash flow, making it harder to treat existing GPUs as quality collateral.[59]

The value of depreciation is linked to the marketable (or residual) value of the GPUs, in the event that GPUs can be sold to pay rent to a data center in the event that cash flows come up short. Faster depreciation means cash flows and asset values fall in tandem: If a firm must depreciate its assets faster than expected, that means its core earnings growth is no longer as good as expected—suggesting that the firm borrowed to finance those assets with the wrong growth projections in mind.

---

holdings of that particular asset. Of course, that fire-sale price would change sharply for the worse if enough firms were all selling those assets simultaneously.

[51] The sale price of an asset has nothing to do with depreciation—to sell or monetize an asset at a price higher than its depreciated book value is akin to capital gains; the asset is "in the money."

[52] Rivera and Brozynski (2025), Tunguz (2025), and Shilov (2024).

[53] It's not just GPU depreciation accounting that's changing. The hyperscalers have also raised the depreciation schedules of server equipment and networking gear on the grounds that internal efficiency improvements improved their useful life. See: Mackereth (2025) and Yao (2025).

[54] Mackereth (2025).

[55] Mackereth (2025)

[56] There is no evidence of Enron-like accounting fraud in this sector.

[57] Mackereth (2025).

[58] Yao (2025). New NVIDIA chips increase the actual rate of depreciation of all existing GPUs because the sale value of those GPUs will fall. Lower initial sale value crimps those GPUs' value as collateral, pushing firms to depreciate those GPUs faster on an accounting basis.

[59] Cerno Capital analysts are quite pessimistic: "Unlike a factory asset with predictable utility, the value of GPU clusters may violently appreciate or depreciate, based on breakthroughs or breakdowns in scaling laws and AI paradigms. This binary risk profile suggests that the true useful life of these assets is unknowable ex-ante, rendering current accounting assumptions vulnerable to swift obsolescence or write-down." See: Mackereth (2025).

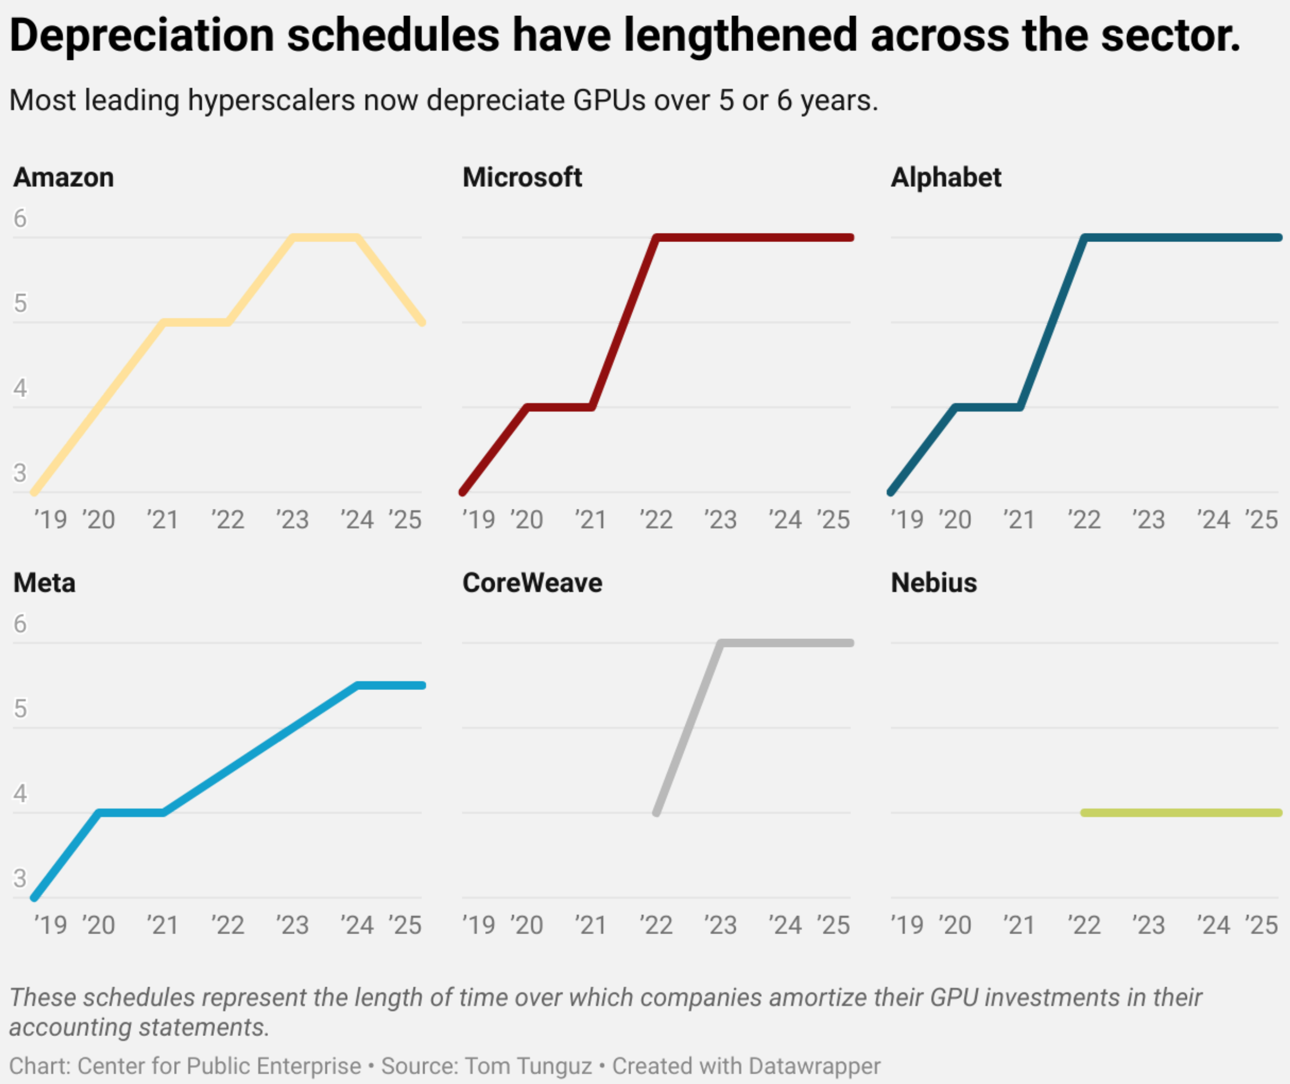

Figure 2.4. *Depreciation schedules have lengthened across the sector.*

Last year, Barclays analysts saw GPU depreciation as risky enough to merit marking down the earnings estimates of Google (Alphabet), Microsoft, and Meta as much as 10 percent, arguing that consensus modeling was severely underestimating the earnings write-offs required and therefore overpricing these companies' share values between 5 and 25 percent.[60] Some analysts, however, think consensus models overestimate the danger of depreciation, insofar as older GPUs can still be repurposed for cost-sensitive "bulk" inference tasks, potentially extending their useful life as long as 7 years; put another way, if the rental market grows substantially relative to the present rate of GPU depreciation, market actors will have ready use for their older or less efficient chips because consistent demand will prop up their marketable values.[61] This

[60] Fox (2024).
[61] LG (2025a).

disagreement centers around, first, the degree to which high utilization rates force faster deterioration and, second, whether demand will nonetheless prop up asset values.

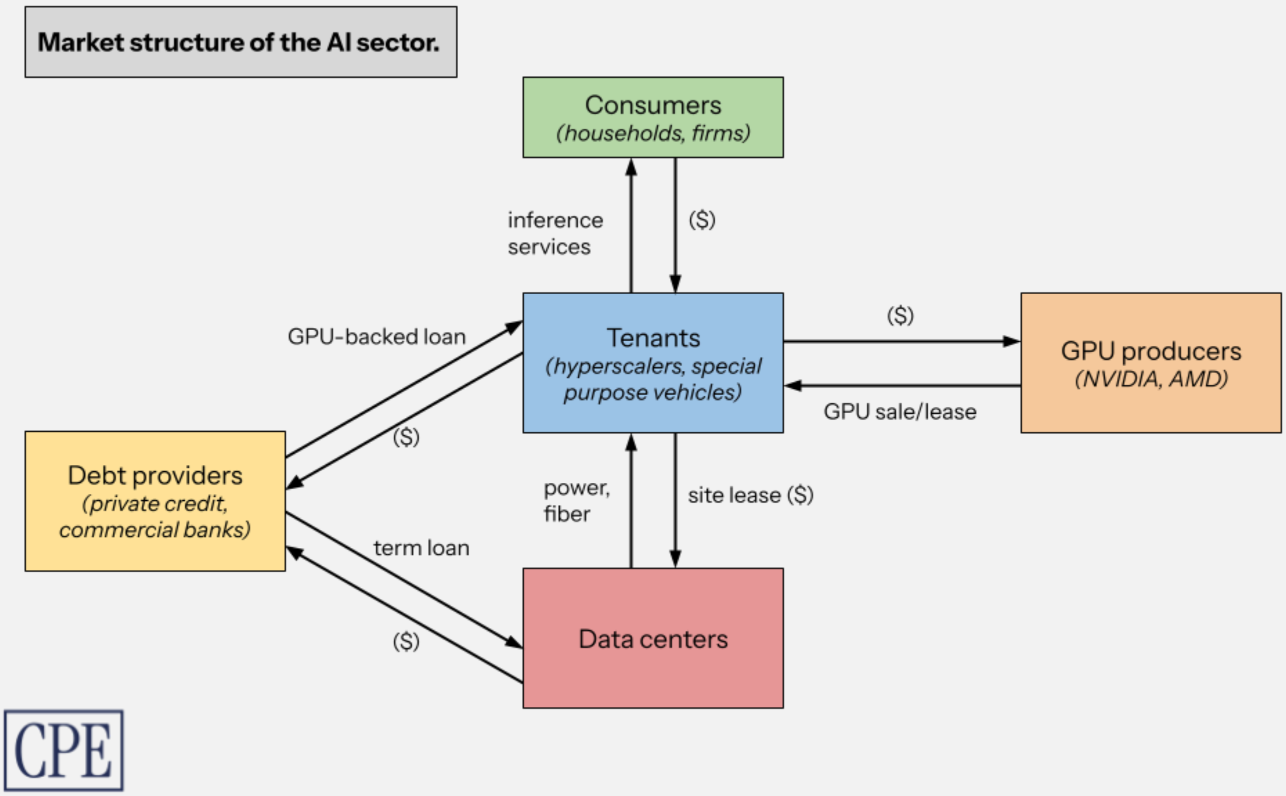

Figure 2.5. *Market structure of the AI sector.*

To be sure, depreciation allowances are paper savings that are only relevant if the assets actually need to be marked to market in the event of an asset sale. But, for inference providers looking to provide high-quality service, the capital investments they will make in new GPUs to keep up with their competitors as well as the deterioration in their existing capital stock's marketable value will, together, cut against the savings they claim to generate via their internal depreciation accounting if demand projections do not hold up.[62] In the event that low demand forces a service provider to go bankrupt, the marketable value of GPUs falls both due to a glut in spare capacity and the revaluation of their sale price due to lower demand.

Companies have different ways to mitigate the risk of holding a quickly depreciating asset on their balance sheets. CoreWeave is shifting from a "pay-as-you-go" GPU rental model to

[62] Innovations in GPU technology, such as photonics-based GPU architectures and quantum computing processors, could also significantly transform the market: New technologies would require hyperscalers and data center developers to alter the architecture of their data centers. If they were sufficiently cheaper and/or more powerful and had uptake from hyperscalers, they would seriously cut into the value of the conventional GPU architectures developed by companies like NVIDIA and AMD. While NVIDIA and AMD could even be the developers of these technologies, they would need to ensure the backward compatibility of their GPU designs with their customer base's infrastructure. According to investor and market strategist Alex Turnbull, it is more likely that new competitors or new sources of GPU supply eat into NVIDIA's and AMD's market share, especially if demand for inference services plateaus. See: Turnbull (2025).

“take-or-pay” rental contracts, where customers pay upfront to rent the compute power of a GPU over a period of months or years, billed monthly.[63] By locking in these contracts, CoreWeave thus reduces its exposure to customer churn. But the weighted average life of these take-or-pay contracts is four years—nearly as long as the asset’s depreciation timeline. It’s not clear if customers may not want to lock themselves into renting compute from a fast-depreciating asset. (CoreWeave’s two main customers, Microsoft and NVIDIA, represent over 60 percent of the company’s revenue. It may offer them slightly different contracts than what it advertises publicly.) It is unclear if demand for this kind of contract is sustainable if competitors can undercut CoreWeave with cheaper pay-as-you-go rental structures. Amazon Web Services, for example, recently slashed the price of its pay-as-you-go and take-or-pay rental structures between 30 to 45 percent.[64]

The other strategy is simply to avoid buying GPUs. OpenAI is now leasing NVIDIA’s GPUs through a special purpose vehicle (SPV) rather than buying them outright—keeping GPU depreciation off of OpenAI’s balance sheet and repaying NVIDIA through lease payments.[65] This strategy, visualized in Figure 2.6, is capital-efficient, essentially converting an upfront capital expenditure into longer-term operational expenditure. OpenAI is replicating it in its arrangement with AMD, where OpenAI is allowed either to buy GPUs directly or rent them via a third-party SPV.[66] Elon Musk’s xAI is using the latter strategy: the SPV raises equity and debt to purchase GPUs and leases the GPUs back to xAI, payments from which help amortize the purchase cost.[67] (Both xAI and NVIDIA own equity in the SPV, which is raising debt to purchase GPUs from NVIDIA.)

When these leasing contracts expire in the coming years, there will be a glut of chips available on the market.[68] Not just low-end chips: if ambitious startups using higher-quality GPUs fail in the coming years, they will help flood the market, too.[69] In other words, there is persistent if temporarily suppressed longer-term downward pressure on chip rental costs.

The result of the push to keep GPUs off of tenants’ balance sheets is that NVIDIA, rather than its customers, now has to eat the depreciation on its products. Perhaps that is part of why NVIDIA is shifting to a yearly product release cycle—to pay for its own assets’ depreciation by tapping into the consistent (for now) market for high-end inference. (The xAI SPV faces a similar depreciation risk—but bankruptcy would chiefly hurt the private lenders to the vehicle rather than the equity owners.)

One financial analyst framed these trends aptly:[70]

> “The major question driving deals today is "Do I want to be long expected GPU useful life and own chips now (because shortages will persist) or short it and rent (because they will be obsolete due to NVDA's [NVIDIA’s] next gen)” ... because of the uncertainty both sides have positive optionality even if in the aggregate the combination of the two will certainly lose money. This is also why NVDA is a core financier for all these deals: they are already massively long "can Nvidia make a good next-gen AI chip": by committing to

---

[63] Elder (2025a).
[64] Yun (2025).
[65] Park (2025).
[66] Acton and Morris (2025).
[67] Park (2025).
[68] Kinder (2024).
[69] Corrigan (2025).
[70] Quantian (2025).

> buyback eg unused CRWV [CoreWeave] chips they have effectively sold calls against their upside-case revenue but not downside-case, which is a lovely hedge."

Not only does this framing capture the potential value of renting, but, in referencing NVIDIA's promise to buy back up to $6.3 billion in unused CoreWeave capacity through 2032,[71] it illuminates NVIDIA's role in the ecosystem. The GPU producer is making a bet on its own growth—it makes the chips that everyone wants, after all—and is financing it by selling the equivalent of a call option. Companies that buy NVIDIA chips will grow with NVIDIA and support NVIDIA's growth—and, in the event demand does not materialize and the call option is "out of the money," NVIDIA's downside is capped.

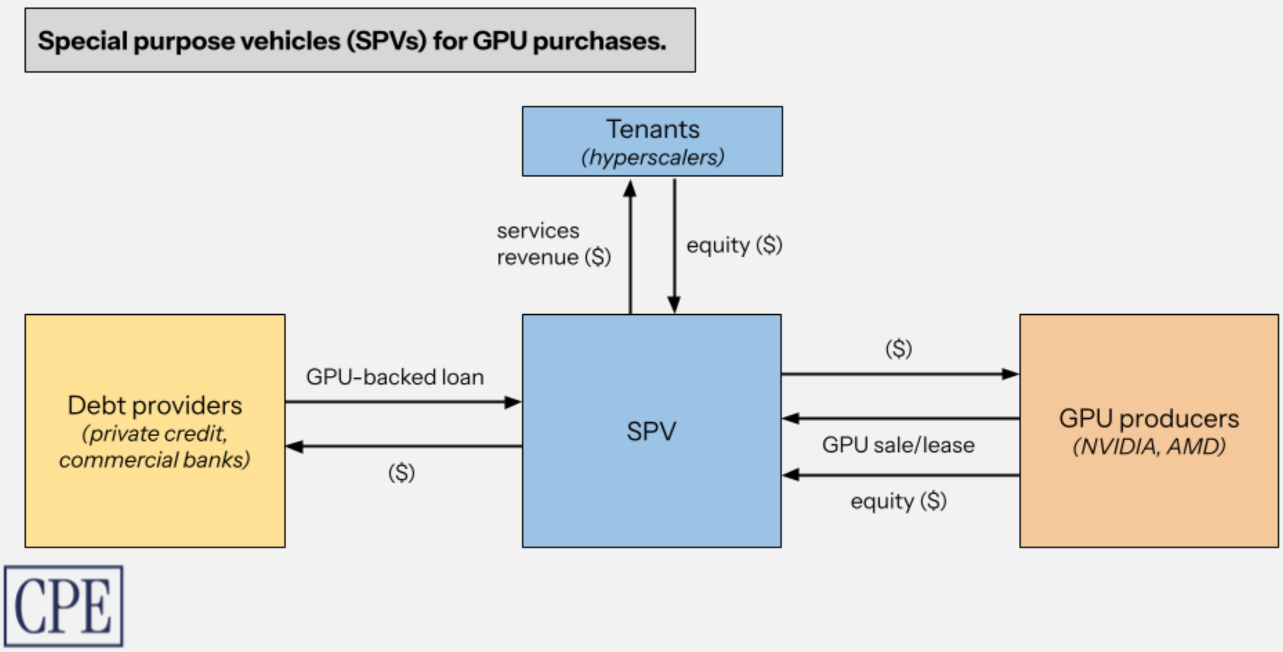

Figure 2.6. *Special Purpose Vehicles (SPVs) for GPU purchases.*

The asset value of a GPU is ultimately constrained by the short-term product cycles of NVIDIA and AMD, the GPU's technological specifications, and the GPU's utilization rate; the fact that tenants have demonstrated little capability to control product cycles is mismatch enough, but, being stuck in a buyer's market, they cannot control demand, either.

Assets and Liabilities

Some tenants, neocloud companies in particular, offer take-or-pay contracts to lock in the earning potential of their GPU assets upfront. But their GPUs' depreciation schedules also do not necessarily match the duration of their rental contracts. CoreWeave's GPU rental structure exemplifies the mismatch: Even if Coreweave claims its GPUs depreciate over six years (for whatever internal accounting reason) offering customers a four-year take-or-pay GPU rental contract locks customers into renting compute from an asset that still loses much of its value a

[71] Wile (2025a) and Forgash and Ghosh (2025).

year after new NVIDIA and AMD chips arrive on the neocloud GPU rental market.[72] This is even more true for companies like Lambda and Nebius, which have more conventional four-year GPU depreciation schedules.

This contract-driven strategy is an adequate way for neocloud companies to derisk their exposure to customer churn, but, if customers have no flexibility under this lease structure, they may switch to a different, cheaper GPU rental provider if an alternative is available. Not only are alternatives available, given the lack of inference product differentiation between many neocloud companies, but the potential for successive waves of GPU gluts will make it hard for tenants to fix their price offerings. Neocloud companies are, ultimately, price-takers—they are vulnerable.

Table 2.5. *GPU depreciation and collateral risk management, with contracts.*

	Tenants		GPU producers		Consumers	
	Assets	**Liabilities**	**A**	**L**	**A**	**L**
stock	GPUs (3–6)		GPUs (3–6)			
flow		Depreciation (3–6)		Depreciation (3–6)		
flow	Revenue from inference services / take-or-pay contracts (3–5)				income	Spending on inference services / take-or-pay contracts (3-5)

The trend to rent GPUs from NVIDIA through long-term leases rather than to buy them, to keep the cost of GPU depreciation off a neocloud company's balance sheet, suggests that nobody actually wants to own GPUs—its quality as a buy-and-hold asset deteriorates over successive product cycles and rapid wear-and-tear and is contingent only on consistent demand growth.[73] When these longer-term GPU leases all start expiring in the next five or six years, not only will the price of rentals collapse as GPU capacity is freed up on the market, but GPU producers such as NVIDIA and AMD may not be able to count on lease renewals if the companies purchasing those long-term leases—OpenAI and other hyperscalers—no longer have the ability to pay for them. After all, they continue to have nothing else on the asset side of their balance sheet save the promise of future cash flows—here, not even GPUs!

[72] Elder (2025a).
[73] Perhaps this dynamic suggests the commodification of GPUs. Keynes, in Chapter 17 of *The General Theory*, sets out a framework for understanding how the liquidity premium and "carrying cost" of a commodity constrain its ability to serve as a store of value. New GPUs may currently command an incredible liquidity premium—they are quite valuable to have around—but they also have immense carrying costs given their market structure. GPUs may become a commodity, but they may never be a good store of value.

Table 2.6. *GPU depreciation and collateral risk management, with GPU rentals.*

	Tenants		GPU producers	
	Assets	**Liabilities**	**A**	**L**
stock			GPUs (3–6)	
stock	*GPU lease residual value (4–5)*			*GPU lease residual value (4–5)*
flow				Depreciation (3–6)
flow		GPU rental costs (4–5)	GPU rental income (4–5)	
flow	Revenue from inference services / take-or-pay contracts (3–5)			

F. GPU-backed loans

Some neocloud companies such as CoreWeave and Lambda are using GPU-backed loans to purchase new GPUs[74]—and this kind of loan should raise red flags. Not only do neoclouds face serious long-term cash flow uncertainties, but, as described above, analysts worry that the collateral value of their GPUs could collapse within three years due both to depreciation and to NVIDIA's GPU release schedule. Using a GPU-backed loan to purchase new GPUs—where the physical GPU and its asset value are secured as collateral for the loan—seems akin to using one's credit card to pay off the debt on another credit card—all you have done is move the debt around and push off default until the promised revenue stream arrives, or doesn't.[75]

Market actors understand their competitiveness, in part, based on their access to new GPUs, the acquisition of which is, in part, financed by their old GPUs. Thus, the acquisition of new GPUs is driven by a strategy of stockpiling more of the latest models than their competitors can in the hope that developing a dominant market position will eventually allow them to achieve sufficient revenues to pay off their earlier investments. Accessing improved GPU performance may be less important than denying competitors access to a limited resource. As more and more market actors deploy this risk management strategy concurrently, they both temporarily raise the prices of new GPUs and increase their collective exposure to a collapse in GPU value.

[74] Reuters (2024).
[75] As one writer put it, "The math makes AI spending seem less like a flywheel than a hamster wheel." See: Beam (2025).

GPU-backed loan volume.

Notable venture debt deals collateralized by GPUs, sorted by date.

Company	Debt (millions $)	Date
Crusoe	225	2025, May
Fluidstack	10,000	2025, Apr
CoreWeave	7,500	2024, May
Lambda	500	2024, Apr
Crusoe	200	2023, Oct
CoreWeave	2,300	2023, Aug

as of August 2025

Table: Center for Public Enterprise (2025) • Source: Pitchbook, the Information • Created with Datawrapper

Figure 2.7. *GPU-backed loan volume.*

Data center developers are not usually exposed to their tenants' GPU acquisition and finance decisions unless they are working with a hyperscaler as their anchor tenant, where the hyperscaler's choices constrain the data center developer's creditworthiness. Still, data center developers are exposed to their tenants' cash flow and collateral value problems by virtue of the fact that their tenants' assets have a weak revenue profile while rapidly losing value. The exposure may be indirect, but it remains material.

As mentioned in the previous subsection, some companies, including market leader OpenAI as well as Elon Musk-led xAI, are trying to sidestep this by renting GPUs from NVIDIA. While this tactic keeps GPU balances and infrastructure spending off tenants' balance sheets, it still exposes them to the threat of lost value from the residual value of the lease itself being repriced.

The growing value of GPU-backed loans despite these well-publicized concerns about collateral value suggest a signaling issue: the marketable value of older GPUs might be suppressed relative to their economic value—whether due to delayed demand, NVIDIA product release cycles, and other previously discussed factors—whereas the value of newer GPUs might be inflated due to hyperscaler demand for the latest NVIDIA and AMD GPUs. Neocloud companies may be losing value when they respond to short-term market dynamics by getting rid of old compute earlier, bringing in new compute faster than needed, and signing leases to keep depreciation off their balance sheets. Then again, they may not be able to take duration risk with these assets. NVIDIA again stands to benefit: Leasing out GPUs to neoclouds lets them own assets that, while perhaps over-depreciated in the short-term, they can resell for a premium over the longer term if demand picks up.

The success of AI service providers seems to depend both on both demand growth and firms' ability to transition from price-taker to price-setter—which they may not all be able to do given their similar AI services. Price-setting power remains limited by market structure.

Assets and Liabilities

Here, tenants are caught between their GPU suppliers and their creditors: The terms on their GPU-backed loans may not match the rate at which their GPUs depreciate, which is itself a function of GPU producers' product cycles and broader demand for inference services. Creditors are now exposed to their tenants' asset-liability mismatch.

Table 2.7. *GPU-backed loans to tenants.*

	Tenants		GPU producers		Debt providers	
	Assets	**Liabilities**	**A**	**L**	**A**	**L**
stock		GPU-backed loan (3–5)	GPUs (3–6)		GPU-backed loan (3-5)	
flow		Depreciation (3–6)		Depreciation (3–6)		
flow	Revenue from inference services / take-or-pay contracts (3–5)		GPU sale revenue			

At this point, we can bring data centers themselves back into the picture—after all, data centers are housing these tenants. Tenant churn is a salient risk: Because tenants' GPU depreciation schedules do not necessarily match the duration of their leases to data centers,[76] it is likely that tenants are undertaking multiple cycles of costly capital expenditures on high-end GPUs within the duration of one lease term, cutting into the cash flows they need to make lease payments; if they are already prepaying some of their lease upfront for their landlord's security, they have already cut into their own equity. This report has already detailed the challenges that data center tenants face in preserving the value of their GPU collateral. If tenants using GPU-backed loans cannot preserve their collateral value, to say nothing of their already-unstable cash flows, they will be unable to renew their leases.

[76] There is also a depreciation mismatch between data center infrastructure assets, including the energy resources, all of which have longer depreciation schedules between 10 and 30 years, and the key assets inside of data centers—the GPUs—which depreciate on schedules between two to four years. To be sure, these sets of assets are often financed separately: Data center developers finance the long-duration assets while their tenants usually finance the short-duration assets.

Table 2.8. *Loans across the sector.*

	Data centers		Tenants		Debt providers	
	Assets	**Liabilities**	**A**	**L**	**A**	**L**
stock	"Powered shell" facility + land (20–30)	Data center term loan (20–30)			Data center term loan (20–30)	
stock	"Powered shell" facility + land (2-5)	Data center mini-perm loan (2-5)			Data center mini-perm loan (2-5)	
stock			GPUs (3–6)	GPU-backed loan (3–5)	GPU-backed loan (3-5)	
flow				Depreciation (3–6)		
flow	Tenant lease income (~10)	accounts payable		Tenant lease payments (~10)		

G. Standout examples

It's worth providing some examples of companies that are demonstrating the above dynamics: Oracle, CoreWeave, and Fermi.

Oracle, a hyperscaler with longstanding market leadership in cloud computing, is quickly moving into the neocloud business, recently inking a massive $300 billion deal to secure NVIDIA GPUs to lease to OpenAI on a data center it is leasing from Crusoe—visualized in Figure 2.8 and Table 2.9.[77] Recent news suggests Oracle sits at the intersection of all the aforementioned headwinds facing neocloud companies.[78] Its gross margins on its AI capital expenditure have fallen to between 10 and 20 percent—far less than the gross margins of other retail and software businesses. This low figure reflects the sheer capital expenditure required to purchase NVIDIA's high-end Blackwell chips and build data centers, the company's internal depreciation accounting, low pricing for its GPU rental services in order to stay competitive against other neoclouds, and perhaps also a lower-than-expected data center utilization rate. None of these factors are per se dangerous for Oracle in the near-term: AI requires expensive hardware that most cloud and software businesses do not require, and utilization rates may rise over time with demand. But the fact that Oracle is a price-taker, not a price-setter—it must try to undercut its direct competitors in a sector where it does not provide a substantially unique service—exemplifies the broader struggle neocloud companies face.

[77] Wile (2025a).
[78] Gardizy et al. (2025).

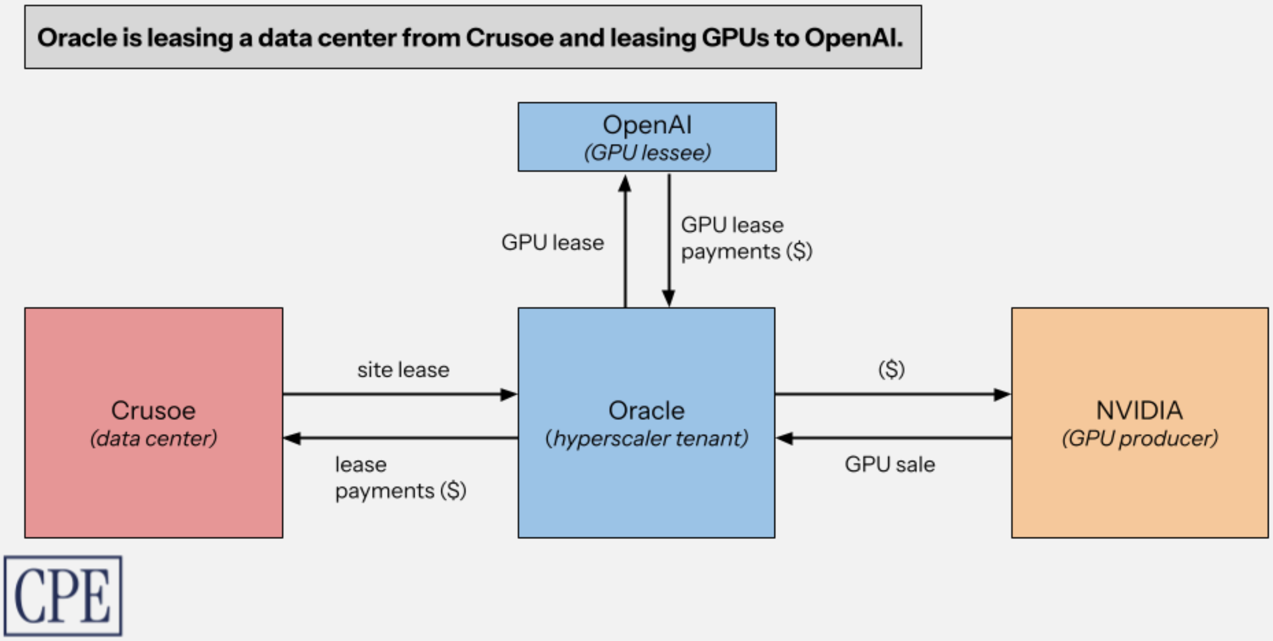

Figure 2.8. *Oracle is leasing a data center from Crusoe and leasing GPUs to OpenAI.*

Table 2.9. *Oracle is leasing a data center from Crusoe and leasing GPUs to OpenAI.*

	Crusoe		Oracle		OpenAI		NVIDIA	
	A	**L**	**A**	**L**	**A**	**L**	**A**	**L**
stock	Data center		GPUs				GPUs	
stock				*GPU lease residual value*	*GPU lease residual value*			
flow				Depreciation				Depreciation
flow			GPU rental income	GPU purchase cost	Revenue from inference services	GPU rental cost	GPU sales	
flow	Tenant lease income	accounts payable		Tenant lease payment				

CoreWeave, Oracle's competitor and another hyperscaler, secured over $12 billion of loans to purchase data center capacity. It raised much of that debt at interest rates between 10 to 14 percent, but one of its loans cost an eye-watering 17 percent.[79] These high rates—triple the cost of investment-grade corporate debt—reflect the risk of CoreWeave's GPU investments and force CoreWeave into targeting growth at or above those loan rates. (CoreWeave does not have an investment-grade credit rating.[80]) CoreWeave's debt is on the balance sheet of an SPV, limiting CoreWeave's direct risks, and the transaction is fairly transparent. But the SPV's debt is collateralized by CoreWeave's existing GPU holdings and its service contracts. CoreWeave's growth trajectory, however, may not be sustainable if its two key revenue sources—Microsoft and NVIDIA, over 70 percent of CoreWeave's revenues—do not meet their payment commitments. CoreWeave is on the hook for over $56 billion in data center lease payments, which will last around 10 years.[81] CoreWeave's recent acquisition of one of its own primary data center providers, Core Scientific, should allow it to eliminate up to $10 billion in lease costs as well as tap into the project finance loan market.[82]

These companies' growth is certainly possible if stocks continue climbing—equity appreciation can effectively collateralize this sizable leverage—and might be attractive to the likes of retail and private equity investors, which care more about shorter-term equity value growth than about long-term asset solvency, but this bet is less safe from the perspective of long-term debt investors in infrastructure, who invest in stable cash flows rather than in equity appreciation. Loans at high interest rates, like CoreWeave's, also demonstrate these companies' relative insensitivity to interest-rate fluctuations—a Federal Reserve rate change would not substantially change its capital planning with respect to this loan.

On the data center side, some data center developers are securing finance well in advance of securing anchor tenants. Fermi, a data center developer incorporated as a real estate investment trust (REIT) is the latest and largest example: it recently launched an IPO (initial public offering) on the stock market to raise capital from retail investors.[83] It plans to use the proceeds of its IPO as equity collateral for the construction of an 11 GW data center megacampus in Amarillo, Texas, with its own nuclear power plants to boot. The only issue is that Fermi is raising finance well in advance of establishing its ability to remain a going concern. Fermi has no confirmed or committed anchor tenants for its data center, no constructed on-site infrastructure, no grid connection, no financing for its nuclear power plants, and no historical revenue. It has its ground lease and land use rights in Texas, about $450 million in cash, equity, preferred equity, and convertible debt, and a loan of $100 million to secure gas turbines. Yet its recent IPO still valued the company at $15 billion. Fermi is creating its own leverage—but is exposing its retail investors to all the risks of its massive plans.

H. Focusing on the hyperscalers

The majority of the AI sector's capital investments come from hyperscalers, which comprise the leading tech companies and the largest AI-focused startups (e.g., OpenAI and CoreWeave). The tech companies—the Magnificent Seven—are long-established, well-capitalized companies working across multiple different internet-related sectors, with no issues raising equity or corporate debt financing on the market. Unlike many newer technology firms and AI companies

[79] Elder (2025a).
[80] Fitch (2025a).
[81] Brown and Whelan (2025).
[82] Ainvest (2025).
[83] Giacobone (2025).

in particular, they are publicly traded companies with widely used cloud services products, the use of which continues to grow. Hyperscalers already use more than half of all the sector's computing power but, by the end of the decade, could use up to 70 percent of it.[84] The data center sector is nearly coterminous with the hyperscalers and their investment decisions simply due to the sheer volume of hyperscaler capital expenditure.[85]

Hyperscalers have more "longitudinal" access to all kinds of proprietary and collected data with which they can conduct model training and demonstrate more accurate inference capacities.[86] This longitudinal data access allows them to build proprietary "foundational" models (such as llama3, ChatGPT, and Grok) that outclass the capabilities of their competitors, which must either pay companies like OpenAI to bootstrap their existing foundational models or build their own models entirely from scratch, incurring significant data purchase and structuring expenses. This is the case regardless of whether these foundational models are proprietary (e.g., OpenAI's ChatGPT) or open-source (e.g., Meta's Llama), the latter of which are still prohibitively expensive to build. Some analysts judge that the sheer cost of even training new foundational models is now too high for non-hyperscalers to enter the market, suggesting that startups can no longer compete with Meta.[87]

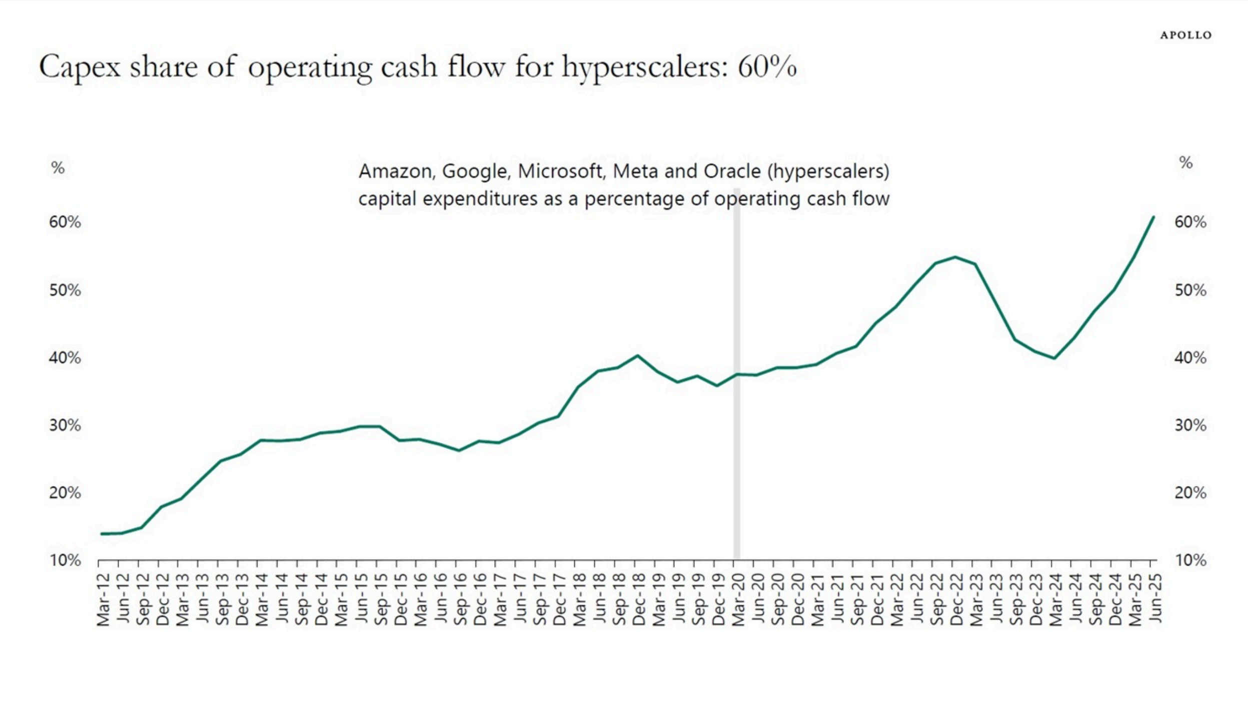

Figure 2.9. *Capex share of operating cash flow for hyperscalers. See: Sløk (2025).*

Their sheer size and capitalization allows them to access the best GPUs first, from producers like NVIDIA and AMD. In fact, OpenAI worked closely with AMD to determine the specifications of

[84] Barth et al. (2025).
[85] Elder (2025b) and Acton and Morris (2025).
[86] Alloway and Weisenthal (2025).
[87] Cheah (2024).

AMD's next chips.[88] And because hyperscalers have control over internet-spanning platforms, they can fold other services into their inference businesses and charge slightly higher rates.

Analysts and policymakers are not particularly concerned that hyperscalers dominate the data center market, but they are growing more concerned that the hyperscalers are "roundabouting" one another. Roundabouting is a term from the dot-com bubble of 2001, when the startups were all investing into and cross-selling products with each other.[89] Other terms for roundabouting include "circularity," "circular financing," and "vendor financing." Whatever the term, the data center boom—really, the AI boom—is undoubtedly that.

Billions of dollars, back and forth.

A stylized list of recent deals between leading AI companies.

Relationship	Description
Nvidia and OpenAI	Nvidia agrees to invest up to $100 billion in OpenAI.
OpenAI and Oracle	OpenAI inks a $300 billion cloud deal with Oracle.
Nvidia and CoreWeave	Nvidia buys $6.3 billion of cloud services from CoreWeave.
OpenAI and CoreWeave	OpenAI to pay CoreWeave as much as $22.4 billion.
OpenAI and AMD	OpenAI agrees to deploy billions of dollars worth of AMD chips.
Oracle and Nvidia	Oracle is spending $40 billion on Nvidia chips.
Nvidia and Intel	Nvidia invests $5 billion in Intel and plans to co-develop chips.

These deals represent a selection of the relationships between some of the most prominent AI players.

Table: Center for Public Enterprise • Source: Bloomberg, author's additions • Created with Datawrapper

Figure 2.10. *Billions of dollars, back and forth. (As of October 2025.)*

OpenAI's deals with NVIDIA, AMD, and Oracle exceed $1 trillion in value.[90] Oracle is spending $40 billion on purchasing NVIDIA's high-end GPUs to power a data center it has leased for 15 years to

[88] Acton and Morris (2025).
[89] Forgash and Ghosh (2025).
[90] Forgash and Ghosh (2025).

support OpenAI, for which OpenAI is paying Oracle $300 billion over the next five years.[91] NVIDIA has also invested $100 billion in OpenAI, to be repaid over time through OpenAI's lease of NVIDIA's GPUs.[92] SoftBank is financing both NVIDIA and OpenAI.[93]

But CoreWeave is also a major player. OpenAI is paying CoreWeave over the next five years to rent its NVIDIA GPUs; the contract is valued at $11.9 billion and OpenAI has committed to spending at least $4 billion through April 2029.[94] OpenAI already has a $350 million equity stake in CoreWeave.[95] NVIDIA has committed to buying any of CoreWeave's unsold cloud computing capacity by 2032 in a $6.3 billion order, after NVIDIA already took a 7 percent stake in CoreWeave's IPO.[96] And Microsoft represented over 60 percent of CoreWeave's revenues as of early 2025—an extraordinarily high share.[97]

NVIDIA owns much of this market. One analyst judged that half of NVIDIA's revenue comes from just these above companies.[98] Another put NVIDIA's share of revenue from data centers at 88 percent.[99] Microsoft, for its part, spends nearly half of its annual capital expenditure on NVIDIA chips.[100] xAI is buying GPUs from NVIDIA through an SPV that NVIDIA is co-sponsoring.

Describing CoreWeave's business model, Financial Times columnist Bryce Elder distilled the absurdity of this roundabouting:[101]

> "Imagine a caravan maker. It sells caravans to a caravan park that only buys one type of caravan. The caravan park leases much of its land from another caravan park. The first caravan park has two big customers. One of the big customers is the caravan maker. The other big customer is the caravan maker's biggest customer. The biggest customer of the second caravan park is the first caravan park. Sorry, not caravans. GPUs."

This industry-wide roundabouting practice reveals something about the nature of equity financing in this sector: **These agreements are effectively risk-sharing structures where the collateral justifying the risk-taking is another company's future growth.** When NVIDIA, for instance, takes an equity stake in a company and that company spends on NVIDIA in some capacity, they can grow together and can quite literally split the dividends. The sector has cross-collateralized itself.[102] Analysts' concerns that these companies' growth figures are a mirage may be validated if there is a sudden stop in cash flows—or when many of OpenAI's and NVIDIA's purchase commitments with each other and with other hyperscalers come due within the next five years.

Not only is the investment growth from this highly leveraged sector the primary driver of GDP growth in the U.S. economy, but its constituent market leaders are each others' primary sources of demand growth. This would not necessarily be a cause for concern except that, as discussed

[91] Wile (2025a), Hammond and Kinder (2025a), and Kinder (2024).
[92] Forgash and Ghosh (2025).
[93] Wile (2025a) and Wile (2025b).
[94] Reuters (2025a) and Wile (2025a)
[95] Forgash and Ghosh (2025).
[96] Wile (2025a) and Forgash and Ghosh (2025).
[97] Fiegerman and Reinicke (2025) and Usman (2025).
[98] Bratton (2025b).
[99] Martin (2025).
[100] Bratton (2025b).
[101] Elder (2025a).
[102] This dynamic is not unheard of in business history; it amounts to an agreement between specific firms to attempt to survive a market correction together, presuming that the result of a market correction is not the emergence of a monopolist.

above, these hyperscalers are consistently losing money on their AI business verticals. While they can of course cross-subsidize their losses here from their other revenue streams, the lack of cash flow and the rising capital expenditure needs of the inference business combined with the scramble for the newest GPUs despite collateral problems suggests instability at the core of their strategy.

Assets and liabilities

The persistent roundabouting in the sector between hyperscalers *and* their suppliers appears at first like the opposite of an asset-liability mismatch: It produces a financial accounting identity of sorts, as equity investments between hyperscalers and GPU producers count as assets' on both balance sheets. However, these investments also count as liabilities—hyperscalers each have the return expectations of the others! Because large tenants and GPU producers have inter-correlated capital structures, the value of these crossed equity investments will rise and fall together. This widespread risk-sharing could be described as a kind of "risk mutualism" which, at least for now, ends up supporting greater capital expenditures.

Table 2.10. *Roundabouting.*

	Tenants (Hyperscalers)		GPU producers	
	A	L	A	L
stock	GPUs (3–6)	GPU loan (3–5)	GPUs (3–6)	
stock	**Equity investment in GPU producer**	**Equity investment from GPU producer**	**Equity investment in tenant**	**Equity investment from tenant**
flow		Depreciation (3–6)		Depreciation (3–6)
flow	Revenue from inference services / take-or-pay contracts (3–5)		GPU sale revenue	

I. Debt's role in the sector

It cannot be overstated how much debt financing is now pouring into the sector. AI-related debt is now the largest segment of the market for investment-grade debt—at a whopping $1.2 trillion—and AI-focused companies comprise a higher share of the investment-grade issuance market than U.S. banks.[103] If data center investment demand projections are to be believed, the potential for over $150 billion in ABS and CMBS issuances for data center financing would triple the size of those markets.[104]

[103] Mutua (2025).
[104] Elder (2025b).

While hyperscalers—especially those with access to public markets—can generally finance more of their investments in training and inference out of their equity rather than by raising debt, to keep their leverage low and their balance sheets flexible, their smaller competitors and the hyperscalers that are not publicly traded are now raising record volumes of debt for capital expenditures. Private equity, private credit and asset management institutions such as Blackstone, Pimco, Magnetar Capital, Apollo, BlueOwl Capital, Carlyle, and Blackrock are all getting into the business of issuing debt to data center developers.[105] As of early 2025, private debt funds—non-bank financial institutions providing structured credit products—had loaned around $450 billion to the technology sector, an increase of $100 billion from early 2024.[106]

The conventional narrative is that hyperscalers can finance costly site leases and GPU purchases out of equity—whereas their non-hyperscaler competitors rely much more heavily on less transparent private credit markets for financing.[107] Hyperscalers' debt accounts for only two percent of the investment-grade bond market.[108] They seem to take on less leverage relative to their competitors, minimizing their exposure to creditors.

However, these cheery measurements of hyperscalers' lack of leverage ignores that they are issuing debt—just not always on their balance sheets.

Hyperscalers can use sale-leaseback transactions to commission and build single-tenant data centers off their balance sheets. This is the structure of the $29 billion deal Meta just inked with Blue Owl Capital and PIMCO: Meta will spend on preparing a site for its data center through an SPV it co-owns with Blue Owl Capital (Meta owns 20 percent, Blue Owl 80 percent) and the SPV will secure debt from PIMCO—which beat out KKR and Apollo to become the anchor lender to the deal.[109] When the data center is complete, Meta will lease the data center back from the SPV as its sole tenant for 15 to 20 years and, thanks to PIMCO's involvement, its debt will be more easily tradeable on capital markets.[110] This sale-leaseback is capital efficient for hyperscalers, the return profiles of which are not designed to carry long-lived real estate assets on their balance sheets. Deals like these likely keep topline estimates of hyperscalers' leverage low, even when hyperscalers are very obviously driving debt issuance among their partners. In other words, this transaction allows Meta to raise billions in debt without posting that on its balance sheet.

Some data center developers have begun issuing junk bonds (high-yield, non-investment-grade bonds) to finance their capital expenditures. Morgan Stanley, which underwrote Meta's SPV, is also underwriting data center developer TeraWulf's $3.2 billion junk bond issuance. TeraWulf's data center has one primary tenant, a neocloud company called Fluidstack, the revenues from which are collateral for the junk bonds.[111] (This structure would qualify as credit tenant lease financing.) Crucially, unlike other junk bond issuances such as CoreWeave's earlier this year, the collateral for TeraWulf's junk bonds is guaranteed by Google: Not only has Google committed to repaying creditors if Fluidstack, the tenant, defaults, but it has also secured warrants for up to 14

[105] Kinder (2024) and Robbins and Bodley (2025).
[106] Fishlow (2025).
[107] Elder (2025b).
[108] The Economist (2025).
[109] Arroyo and Benitez (2025a, 2025b). Morgan Stanley underwrote the deal for Meta.
[110] Arroyo and Benitez (2025b).
[111] Weinman and Gurumurthy (2025a, 2025b).

percent of TeraWulf that it can choose to convert into equity at the end of construction.[112] This guarantee is an expression of Google's demand for Fluidstack's neocloud services and a vote of confidence in TeraWulf's future growth prospects.[113] It also demonstrates that hyperscalers like Google, which may claim to avoid leverage themselves, are providing the collateral, guarantees, and demand required to allow other firms—neoclouds in particular—to raise debt themselves.

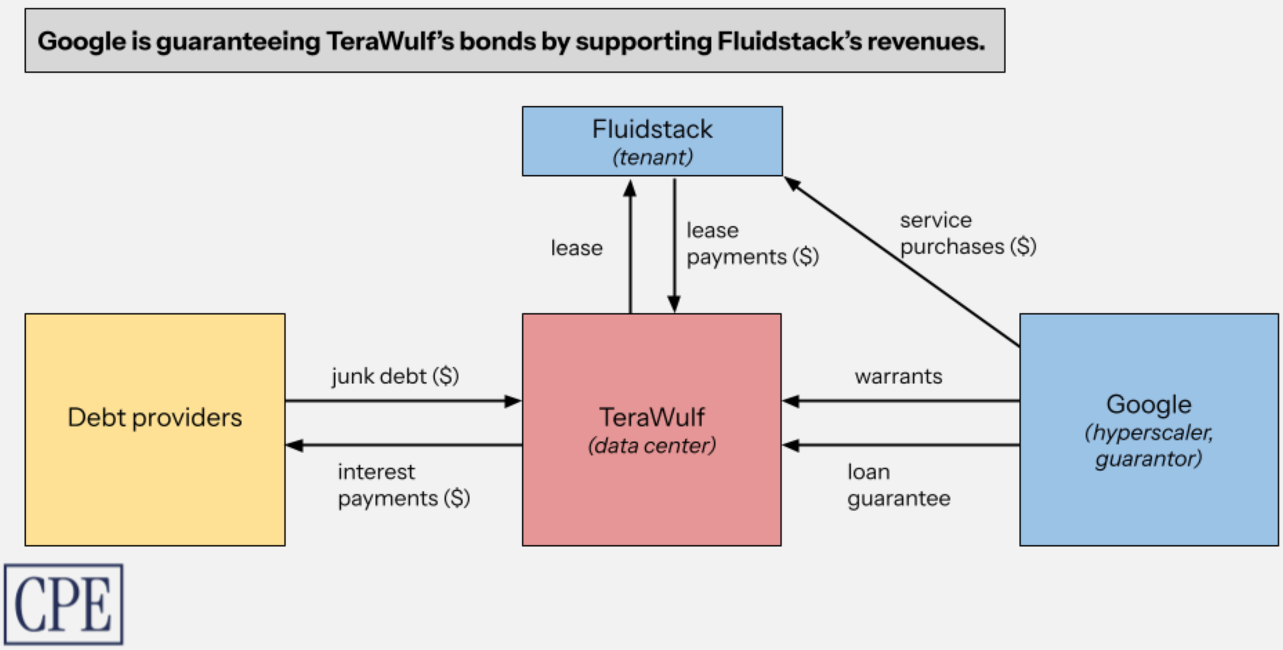

Figure 2.11. *Google is guaranteeing TeraWulf's bonds by supporting Fluidstack's revenues.*

Oracle recently inked an $18 billion debt deal for a four-year loan to construct a data center in New Mexico.[114] The deal does not seem to involve junk bonds; it is a syndicated transaction led by American and Japanese investment banks financing a data center built by STACK Infrastructure (owned by private credit firm Blue Owl Capital) in which Oracle is the sole tenant. This four-year loan, which is likely a mini-perm loan covering construction and early tenancy—it seems to be an interest-only facility until the end of construction—is priced at only 2.5 percentage points (250 basis points) above the risk-free Treasury rate—a remarkably tight spread reflecting Oracle's creditworthiness to the syndicated lenders. Oracle also recently inked a $38 billion deal to build two data centers, one in Texas and one in Wisconsin, with a similar four-year lease term requiring interest-only payments until the end of the projects' construction periods, again priced at 2.5 percentage points above the risk-free Treasury rate.[115] At both data centers, Oracle will lease its GPU capacity to OpenAI, similar to the Crusoe deal mentioned in the previous subsection on standout examples.

---

[112] This deal is worth discussing—but it has some caveats. Fluidstack can call off its lease if construction is delayed by more than 180 days, and, in doing so, it would void Google's backstop, which only applies after Fluidstack's lease begins. (The data center is expected to be completed by the end of 2026.) Google may also be able to pull its guarantee under other conditions. Creditors remain exposed to construction risk. See: Weinman and Gurumurthy (2025a, 2025b).
[113] Weinman and Gurumurthy (2025a) and Thomson and Bergen (2025).
[114] Seligson et al. (2025). The maturity can be extended up to two years.
[115] Amodeo (2025).

There are other examples; the market is growing. xAI's SPV is buying GPUs with debt. Oracle's data center servicing OpenAI has $15 billion in debt and equity from the likes of Crusoe (the site owner and developer), Blue Owl Capital, and JP Morgan—which provided $9.6 billion of the debt.[116] Analysts expect OpenAI to continue to tap debt markets even as it continues to raise more equity.[117] CoreWeave's CEO called debt "the fuel for this company."[118] (It is very likely that most, if not all, of these transactions include at least some recourse to the corporate balance sheets of hyperscaler tenants to protect lenders. Microsoft, Oracle, and NVIDIA, among other players, are ultimately underwriting a large volume of data center debt, even if indirectly.[119])

Morgan Stanley analysts note that most of the private credit investment in the AI space will likely be backing non-hyperscalers—in other words, the likeliest losers in the event of a crash. Spreads on recent data center-related ABS issuances price in a higher likelihood of tenant failure for data centers with fewer investment-grade tenants (which includes those issuing junk bonds).[120]

Some analysts have suggested that the heavy use of off-balance sheet special purpose vehicles (SPVs) signals a desire among data center developers and hyperscalers to shift risky borrowing off their balance sheet and hide them from public scrutiny.[121] This critique has some merit given the incredible fragility of cash flows across the sector and the amount of roundabouting between hyperscalers. That being said, SPVs are a quotidian way of conducting non-recourse project finance for capital-intensive infrastructure and technology investments; debt is a necessary part of any investment cycle and companies in all sectors seek to minimize the extent to which long-term debt weighs down their corporate balance sheets. Additionally, as transactions like TeraWulf's confirm, it is not the case that this sector's debt is entirely non-recourse: Parent guarantees from hyperscalers matter quite a lot for a project's creditworthiness.

The growing share of private credit funds piling into the sector poses transparency concerns for analysts seeking to understand what risks are on whose balance sheets. And private credit is a less liquid asset class, making it more exposed to rapid value deterioration in the event of a market correction. But focusing on debt itself as dangerous is misleading: The real threat to the sector, if any, is the instability of cash flows and collateral values.

J. From the perspective of credit rating agencies

Credit rating agencies evaluate the debt issued to finance data center construction and operation, providing that debt with a rating that testifies to a data center developer's creditworthiness to the long-term investors in ABS and CMBS markets. A study of Moody's and Fitch ratings guidelines for data center infrastructure provides some useful guidelines for understanding the dynamics discussed in this paper.[122]

Data center developers do best with rating agencies and can take on the most debt when their projects can lead the market by outpacing potential competitors in their specific geographies. The rating agencies are clear that, especially for data center developers working with hyperscalers, differentiated and unique AI service offerings—businesses with a "moat"—are much better and more creditworthy than a "commodity offering" of AI technology where services are

[116] Kinder (2024).
[117] Acton and Morris (2025) and Hammond and Acton (2025).
[118] Brown and Whelan (2025).
[119] Cahn (2025).
[120] Elder (2025b).
[121] Mills (2025).
[122] Moody's (2025) and Fitch Ratings (2025b).

not unique. This criterion suggests that rating agencies may not appreciate treating GPUs as commodities for rental as neoclouds are doing. The rating agencies also approve of overbuilding capacity within data centers: Doing so means minimal downtime for customers. This suggests that the leading tech companies will have the upper hand over the rest of the market: Unlike their competitors, many of them can offer whole platforms of services they can integrate into any inference product they release, and they have the cash flow strength to overbuild capacity.

This being said—data centers acting as "carrier hotels" for their tenants are less risky for debt investors than data centers providing bespoke "turnkey" services to tenants. By this line of reasoning, hyperscalers—which need "turnkey" services due to their large and sometimes customized requirements for GPUs and server racks—are more risky to investors than regular colocation data centers are. But hyperscalers' creditworthiness seems to more than make up for the attendant concentration risks their data centers carry.

Fitch's recent rating and analysis of CoreWeave's debt provides additional color to how lenders are viewing this sector.[123] While CoreWeave is not an investment-grade borrower (it carries a BB- rating), Fitch seems mostly positive about the company. Yet the agency expresses concern about (1) CoreWeave's overwhelming dependence on its top two customers (Microsoft and NVIDIA, accounting for 77 percent of its 2024 revenue), (2) notes the fuzziness of CoreWeave's long-term revenue projections farther than five years out, (3) implies that the company needs to build a better "moat" in the face of competition, and (4) stresses that the mismatch between CoreWeave's take-or-pay GPU rental contracts with customers (3–5 years) and its data center site leases (3–15 years) raises the risk that CoreWeave must make site lease commitments without revenue certainty.

In February 2024, Moody's changed its rating criteria for **hybrid bonds**—complex but flexible financial instruments that blur the line between equity and debt such as callable capital instruments, debt-for-equity swaps, perpetual bonds, and combined issuances of subordinate debt and equity warrants—making it far easier for corporate borrowers and investors to treat those hybrid instruments like equity and suppress their leverage. The change allowed many companies, utilities in particular, to raise lots more capital from yield-seeing investors comfortable with subordinate positions in the capital stack without threatening their credit ratings.[124] This rating shift allows utilities to more easily invest in the power infrastructure required for a data center buildout, especially as the electricity sector's overall capital spending could increase up to $1.1 trillion between 2025 and 2029.[125] Telecom companies have also benefited.[126] The spread of hybrid instruments over US Treasuries recently hit a post-2008 low of 145 basis points over the risk-free rate, and hybrid instrument issuance now dwarfs the issuance of preferred equity.[127] Moody's may have made this change due to industry concerns about the tension between placating investors concerned about leverage and meeting booming investment demand.

K. Putting it all together

Under these market dynamics, the booming demand for high-end chips and the meteoric rise in NVIDIA sales look less like a reflection of healthy growth for the tech sector and more like a

[123] Fitch Ratings (2025a).
[124] Love (2024).
[125] Walton (2025b).
[126] Vossos (2025).
[127] Vossos (2025).

scramble for high-value collateral and a strategy to maintain market position among a set of firms with limited product differentiation. If high demand projections for AI technologies materialize, collateral ostensibly depreciates less quickly as older GPUs retain their marketable value over their useful life—but, otherwise, this combination of structurally compressed profits and rapidly depreciating collateral is evidence of a snake eating its own tail.

For now, the dominance of publicly traded hyperscalers and the sheer size of their deals contribute to a working, if birds-eye, understanding of how capital is flowing. We have visualized it below, in Table 2.11. (Here, we assume that GPUs are debt-financed, hence the use of "GPU loan," which is inclusive of GPU-backed loans.)

Table 2.11. *Putting it all together.*

	Data centers		Tenants (Hyperscalers)		Debt providers		GPU producers		Consumers	
	Assets	**Liabilities**	**A**	**L**	**A**	**L**	**A**	**L**	**A**	**L**
stock	Data center facility (20–30)	Data center term loan (20–30)			Data center term loan (20–30)					
stock			GPUs (3–6)	GPU loan (3–5)	GPU loan (3–5)		GPUs (3–6)			
flow				Depreciation (3–6)				Depreciation (3–6)		
flow	Tenant lease income (~10)	Opex + power costs		Tenant lease payments (~10)			GPU sale revenue			
flow			Revenue from inference services / take-or-pay contracts (3–5)	GPU rental costs (4–5)			GPU rental income (4–5)		income	Spending on inference services

III. Market Corrections

Despite concerns that analysts are raising about the "roundabouting" and circular payments among hyperscalers and the threat of GPU depreciation, the industry continues to raise record volumes of capital. Oracle's September deal with OpenAI saw its market value jump $244 billion, and AMD saw its market value increase by $63 billion the day it announced its deal with OpenAI.[128] Fermi's recent IPO put its value at $15 billion.[129] Computer hardware companies like Dell continue to expect strong annual sales growth over the coming years on the back of the AI boom.[130]

Hyperscalers' willingness to use unique deal structures has kept their equity valuations rising, allowing them to fund more investment out of prospective cash flow and take on more leverage as needed. This leverage appears low for now.[131] However, it is not clear if measurements of leverage or debt exposure account for the use of off-balance-sheet financing vehicles. Moveover, leverage is not necessarily a threat to the sector or to the economy unless cash flow and demand in the AI sector collapse. That collapse would trigger a market correction as AI companies, including hyperscalers, and data center developers find themselves in a position where they are effectively "margin called"—in other words, where they can only finance continued expansion until they post enough collateral to pay off their existing liabilities and to prove that they are solvent.

There are not many pathways by which an AI demand shock is passed through to most consumers beyond a sharp decline in the stock market (in the manner of a conventional recession). Still, short of a stock market collapse, the sector remains strewn with asset-liability mismatches that will constrain the ability of capital to flow smoothly through the sector. This section will discuss how a broader market correction might operate, making use of the economist Hyman Minsky's financial instability hypothesis.

A. Refinancing risks

Assessing the various asset-liability mismatches described in the previous section allows us to understand exactly how these mismatches will cascade. This subsection combines the T-Charts from the previous section into four periodized T-Charts, where each period corresponds to the realization of each one of the previously described mismatch. (While these T-Charts exclude the equity investments characteristic of roundabouting, we discuss how circular financing plays into the refinancing waterfall in the next subsection). This series of T-Charts is our attempt to schematize how the investment boom might collapse by identifying the refinancing deadlines that cut through the sector's asset-liability mismatches. The boxes outlined in red highlight the financial source of instability in each period.

[128] Hammond and Kinder (2025a).
[129] Basil et al. (2025).
[130] Bass (2025).
[131] Edwards (2025) and the Economist (2025).

Table 3.0. PERIOD 0: CASH FLOW UNCERTAINTY. *Falling spending on inference services cuts into AI service providers' revenues.*

	Data centers		Tenants (Hyperscalers)		Debt providers		GPU producers		Consumers	
	Assets	**Liabilities**	**A**	**L**	**A**	**L**	**A**	**L**	**A**	**L**
stock	"Powered shell" facility + land (20–30)	Data center term loan (20–30)			Data center term loan (20–30)					
stock			GPUs (3–6)	GPU loan (3–5)	GPU loan (3–5)		GPUs (3–6)			
flow				Depreciation (3–6)				Depreciation (3–6)		
flow	Tenant lease income (~10)	Opex + power costs		Tenant lease payments (~10)			GPU sale revenue			
flow			Revenue from inference services / take-or-pay contracts (3–5)	GPU rental costs (4–5)			GPU rental income (4–5)		income	Spending on inference services

Table 3.1. PERIOD 1: INABILITY TO AMORTIZE GPU INVESTMENTS. *Without income, AI service providers cannot pay the costs of renting or financing the purchase of GPUs. GPU producers suffer revenue losses due to their customers' inability to pay for GPUs. Lenders to AI service providers write down their loans.*

	Data centers		Tenants (Hyperscalers)		Debt providers		GPU producers		Consumers	
	Assets	**Liabilities**	**A**	**L**	**A**	**L**	**A**	**L**	**A**	**L**
stock	"Powered shell" facility + land (20–30)	Data center term loan (20–30)			Data center term loan (20–30)					
stock			GPUs (3–6)	GPU loan (3–5)	GPU loan (3–5)		GPUs (3–6)			
flow				Depreciation (3–6)				Depreciation (3–6)		
flow	Tenant lease income (~10)	Opex + power costs		Tenant lease payments (~10)			GPU sale revenue			
flow				GPU rental costs (4–5)			GPU rental income (4–5)		income	

Table 3.2. PERIOD 2: TENANT CHURN. *AI service providers can no longer make lease payments to data centers, breaking their leases.*

	Data centers		Tenants (Hyperscalers)		Debt providers		GPU producers		Consumers	
	Assets	**Liabilities**	**A**	**L**	**A**	**L**	**A**	**L**	**A**	**L**
stock	"Powered shell" facility + land (20–30)	Data center term loan (20–30)			Data center term loan (20–30)					
stock							GPUs (3–6)			
flow								Depreciation (3–6)		
flow	Tenant lease income (~10)	opex + power costs		Tenant lease payments (~10)						
flow									income	

Table 3.3. PERIOD 3: TERM LOAN REFINANCING RISKS. *Data center owner-operators cannot amortize or refinance their term loans in the absence of tenants, putting them at risk of default.*

	Data centers		Tenants (Hyperscalers)		Debt providers		GPU producers		Consumers	
	Assets	**Liabilities**	**A**	**L**	**A**	**L**	**A**	**L**	**A**	**L**
stock	"Powered shell" facility + land (20–30)	Data center term loan (20–30)			Data center term loan (20–30)					
stock							GPUs (3–6)			
flow								Depreciation (3–6)		
flow		opex + power costs								
flow									income	

Table 3.4. PERIOD 4: STRANDED ASSETS. *GPU producers have no customers to sell to and data center landlords have facilities with no tenants.*										
	Data centers		Tenants (Hyperscalers)		Debt providers		GPU producers		Consumers	
	Assets	**Liabilities**	**A**	**L**	**A**	**L**	**A**	**L**	**A**	**L**
stock	"Powered shell" facility + land (20–30)									
stock							GPUs (3–6)			
flow								Depreciation (3–6)		
flow		opex + power costs								
flow									income	

B. The Financial Instability Hypothesis

Economist Hyman Minsky, a theorist of financial markets, coined the financial instability hypothesis to spell out how speculative debt issuance during times of economic growth can metastasize into a dangerous liability for the entire economy. The threat lies in what Minsky aptly calls "Ponzi finance," when market actors must borrow to pay interest:[132]

> "For Ponzi units, the cash flows from operations are not sufficient to fulfill either the repayment of [principal] or the interest due on outstanding debts by their cash flows from operations. Such units can sell assets or borrow. Borrowing to pay interest or selling assets to pay interest (and even dividends) on common stock lowers the equity of a unit, even as it increases liabilities and the prior commitment of future incomes. A unit that Ponzi finances lowers the margin of safety that it offers the holders of its debts."

The frequent comparisons between the data center boom and previous technology-driven bubbles, such as the railroad and dotcom bubbles, compels us to consider the relevance of the financial instability hypothesis.[133] Indeed, commentators have already picked up on the potential applicability of this hypothesis—particularly the notion of a "Minsky moment," when a cash flow shortfall among Ponzi finance units forces industry-wide, or even economy-wide, de-leveraging.[134] While this risk is real, this report aims to clarify how exactly it might manifest.

In the short-term, cash flows in the AI sector are wholly insufficient to service liabilities, and, as the AI sector grows, it does not look like this condition will change in the near-term. Given hyperscalers' consistent ability to raise new equity and the rising value of that equity, their ability to self-finance new investments, and (for the most part) their low debt-to-equity ratios, it is not yet clear that the data center boom is yet in such a dire situation. To be sure, these leverage ratios do not account for off-balance sheet vehicles—and debt is playing an increasing role in backstopping some of the latest deals, particularly for non-investment grade issuers such as CoreWeave and TeraWulf.

But debt is also playing an increasingly important role in the stock market as equity investors seek to magnify their returns from investments in the tech sector. The number of leveraged stock market exchange-traded funds (ETFs) is at a record high, suggesting the increasing use of debt finance to magnify stock market returns.[135] In October, three leveraged fund issuers requested permission from the Securities and Exchange Commission to set up new ETFs to capture magnified returns from companies including NVIDIA and CoreWeave.[136] A market correction would disproportionately decimate investors in these funds.

Minsky warned that a bubble is often preceded by the increasing use of leverage to purchase equity stakes of multiple forms in one sector. The fact that the market has reached this point already through the rise of tech-focused leveraged ETFs is especially dangerous because any rush to service debt, caused by the failure of even one weak firm, will also hurt the equity valuations of stronger firms as investors sell shares of the latter—including leveraged shares—to cover their losses from the former. The result is a self-sustaining downward spiral.[137]

[132] Minsky (1992).
[133] Mims (2025b) and the Economist (2025).
[134] Kedrosky (2025).
[135] Roberts (2025).
[136] Adinolfi (2025a).
[137] Minsky, Hyman (2008). Page 119.

It remains the case that Ponzi finance units can continue to stay liquid, if not solvent, so long as they can continue borrowing. In the meantime, they lower the "margin of safety" that they are offering their creditors, and the leverage of the entire economy increases. As Minsky explains, "the fragility of the financial system is related to the ratio of debt payments to operations income for the various sectors and the extent to which units are dependent upon refinancing their positions in long assets in smoothly functioning short-term financial markets."[138] This ratio is climbing as the contract and debt obligations of the AI sector balloon out of proportion to the uncertain prospective cash flow returns that collateralize them[139]—and refinancing deadlines loom.

Thus, despite the unknowns, Minsky's hypothesis still provides a useful framework for understanding how a market correction might play out. If there is a sudden stop in new lending to data centers, perhaps caused by volatile capital costs or due to a high-profile default, Ponzi finance units "with cash flow shortfalls will be forced to try to make position by selling out position"—in other words, to force a fire sale—which is "likely to lead to a collapse of asset values."[140]

Box 3.1. How do we know if the AI boom is really a bubble?

A 2019 paper written by senior managing directors at State Street Associates, an asset manager, and the CEO of Windham Capital Management offers an interesting criteria for defining an asset bubble: Whether or not investment in a boom sector is a "crowded trade." The authors derive a measurement of (1) asset centrality—"the extent to which a group of similar assets drives the variability of returns across a broader universe of which it is a member"—as a proxy for locating sectors subject to crowded trading, and they combine this measurement with an indicator of a sector's (2) relative value to "identify when a bubble transitions from its inflationary phase to its deflationary phase" due to market participants' sudden reevaluation of its real value. The authors combine and apply these tools to stock trading data to identify the creation of bubbles, particularly the dotcom bubble and the real estate bubble, and to suggest that they can be predicted before they occur.

The paper's methodology is intricate and compelling, but it may be less applicable to the current AI boom due to the inability to directly measure returns on private markets, through which increasing amounts of debt and equity investment are flowing. But the heuristics involved remain useful, especially given the concentration of key stock market indices around the tech sector and the rapid rise in tech companies' valuations. Further research into whether the AI boom meets this paper's criteria would be welcome.

See: Kinlaw et al. (2019). "Crowded Trades: Implications for Sector Rotation and Factor Timing." IPR Journals.

A sudden stop to the "roundabouting" might also be dangerous, particularly because it would engender a run on the equity valuations of many of the publicly traded hyperscalers at

[138] Minsky (2008). Page 160. This book is one of his major publications on the relationship between financial market instability and firms' capital planning decisions.
[139] These cash flows are, in Minsky's terms, the "quasi-rents" of the capital assets these firms invest in. As detailed previously, these quasi-rents are structurally suppressed in the AI sector due to competition among inference providers.
[140] Minsky (1992).

once—insofar as their equity investments are all correlated. These companies' equity valuations are their ticket to financing their capital expenditures, either through their own cash flow or through debt. Without a booming equity market, their short-term corporate bond spreads will rise—and the spreads on debt for their long-term investments as well as their SPVs will certainly start climbing.[141]

Minsky's hypothesis and the history of recent liquidity crises in the banking sector suggest that hyperscalers would, in a crash, sell off their best assets to meet their liabilities—those assets being their most valuable high-end GPUs—and otherwise slow their purchases of new assets. Whether this portfolio unwinding is successful or not for the struggling seller depends on the marketable value of those GPUs contingent on their depreciation schedule and on persistent cash flow-generating demand for AI-related services. A GPU fire sale may not be as successful in a potential market correction driven by cash flow and/or demand uncertainty, especially if there are no buyers for older GPUs.

The report previously discussed how GPU producers like NVIDIA may gain some advantage from leasing GPUs rather than selling them if they believe that GPU lessees are mispricing the residual value of their lease in the event that there is a robust market for used GPUs. If there is a GPU fire sale, however, GPU producers/lessors would also be caught on the wrong side of their own trade as the value of used GPUs falls even further due to the supply glut of high-end GPUs.

GPU owners that have taken out loans to finance the GPU might choose not to sell their assets at all and produce their way out of insolvency through continuing to offer GPU rental services—but, like an agricultural debt crisis, this decision threatens to further depress the price of GPU rentals if the whole sector chooses to respond this way, making it harder for any one neocloud services provider to amortize their debts.

AI-focused data centers, which host GPUs, are not necessarily stranded assets if a tenant defaults: So long as the data center has the adequate energy and fiber optic connections to provide reliability to tenants, new tenants can simply replace or buy out the GPUs of their predecessor. Data centers only begin to look like stranded assets in the event of a demand-based market correction where prospective tenants vanish. When data center ABS financings mature and need to be rolled over, data centers without a high percentage of investment-grade tenants will be in danger.[142]

It is not yet clear what creditors would do if a data center itself defaults. Converting an AI-focused data center into a cloud-focused one, which has more stable revenue streams, is expensive and complex: Doing so would require replacing infrastructure that supports GPUs with infrastructure that supports CPUs.

Considering the available evidence, the weak point in the system appears to be the private credit lenders and the "venture debt" groups writing expensive loans to every actor in the system—to data center developers; to hyperscaler and non-hyperscaler tenants, neocloud or otherwise; and to the SPVs set up to insulate hyperscalers from their own capital investment plans. These lenders are exposed to almost every part of the "circular" AI economy.[143] Due to their status as debt financing institutions, these lenders have limited upside and virtually unlimited downside.

---

[141] Mutua (2025).
[142] Elder (2025b).
[143] Available evidence suggests that private credit does not seem to be very exposed to the GPU makers themselves. If GPU makers like NVIDIA and AMD are raising debt, it is mostly short-term commercial paper; NVIDIA's long-term debt as of July 2025 only represented a fifth of its total liabilities. See: NVIDIA (2025).

For them to stop lending into the AI sector—precipitating the feared "sudden stop" in new lending—would be to prevent themselves from refinancing each others' debt and, in rendering themselves illiquid, to bring about their own collapse.

The final aspect of Minsky's framework is what happens *after* an investment bubble bursts. The likely outcome for the sector that faces a sharp market correction is consolidation. Any widespread defaults among AI-related businesses and SPVs will leave revenue-producing capital assets and technologies stranded, losing value in the absence of demand—the perfect targets for a rollup by hyperscalers. While the hyperscalers will lose considerable value in the event of a market correction, it stands to reason that their dominance of cloud and web services businesses, not to mention internet advertising, will enable most of them to continue leading the economy as the world's preeminent internet and software companies. Insofar as certain hyperscalers like Microsoft already exercise special control over AI-specific firms like OpenAI, it is plausible that a collapse in OpenAI's value would see it folded into Microsoft. The firms that survive a market correction secure what is essentially monopolistic control over the remaining consumer demand in their sector; this easier access to more certain cash flows eases their leverage constraints over the longer term as the economy recovers.

Minsky's framework carries broader public policy implications. As political economist Mike Beggs has pointed out, the true "Minsky moment" is not the bust, but the bailout—the moment when political and financial sector leaders have to decide, under public scrutiny, how to distribute losses across the economy and what sectors or firms to stabilize with targeted injections of public financial support.[144] The next section of this report offers policymakers a framework for considering this question.

[144] Beggs (2018) and Feygin (2023).

IV. Risk Management for Policymakers

The release of the federal government's AI Action Plan has prompted tech companies and the federal government alike to publicly consider how federal policymaking can support the sector's growth. While the federal government has denied plans to provide loan guarantees to the sector, federal financial support for the AI sector remains plausible. This section will first discuss the changing federal policy landscape and the impact that financial assistance through loan guarantees and tax credits could have.

The remaining subsections will focus on the policy risks of a market correction. First, there's a hit to consumers due to economic contraction and a passthrough of tighter lending conditions. Second, there's a hit to state budgets. Third, there's a hit to energy investments and infrastructure across the energy and utility sector. Finally, there are additional headwinds impacting the sector's health worth highlighting. This section will flesh out these impacts in order to give policymakers a window into the types of risk mitigation frameworks they should be developing.

A. The possibility of federal support for the AI sector

In early November, the CFO of OpenAI suggested that the company would be open to a federal "guarantee" for GPU purchases, especially in light of the uncertainty in their depreciation rates and the challenge of keeping them on tech companies' balance sheets.[145] Although vague, the remarks elicited considerable media coverage and forced both the federal government and OpenAI CEO Sam Altman to stress that there are no plans for federal financial backstops of AI-related capital expenditure.[146] That's not to say the industry might not desire it; OpenAI has previously petitioned the federal government for loan guarantees to help finance the construction of semiconductor manufacturing facilities and the purchases of U.S.-made GPUs.[147] A week before the CFO's comments concerning a backstop, OpenAI stated its desire that the federal government expand the CHIPS Act tax credits for semiconductor manufacturing to the AI sector, to cover the costs of building AI-focused data centers, AI-related servers, and electrical grid components.[148]

How could loan guarantees and tax credits affect the sector's finances and redistribute risk across market actors' balance sheets? A loan guarantee for debt taken out in the course of data center construction would attenuate the cash flow risks that data center developers face from less creditworthy tenants and from tenant lease churn. This guarantee would help insulate the broader ABS and CMBS markets backstopping data center term loans by preventing defaults on individual term loans.[149]

---

[145] Jin (2025).
[146] Reinicke and Semenova (2025).
[147] Krouse (2025) and Eastland and Ghaffary (2025).
[148] Eastland and Ghaffary (2025).
[149] However, this would not be the same as the kinds of temporary liquidity guarantees and interest payment coverage that Fannie Mae and Freddie Mac provide to purchasers of residential mortgage-backed securities. This is because Fannie and Freddie are dealers in that market: They buy residential mortgages, securitize them, and sell them—meaning that they are in the right position in the market to paper over mismatches between any losses on individual mortgages and losses to purchasers of their securities. Where data centers are concerned, this guarantee proposal more likely covers individual term loans in the primary market for lending to data centers—derisking those loans before they are securitized into ABS and CMBS structures for secondary market investors.

Alternatively, a guarantee on any shorter-term debt that tenants—AI inference providers, including "neoclouds"—use to purchase GPUs would put the federal government on the hook for any deterioration in their collateral value. This guarantee attenuates the risk that tenants earn insufficient cash flow to amortize their GPU purchases and supports the GPU-backed loan market irrespective of the market value of GPU collateral. Depreciation schedules across the sector may not matter as much.

However—and much more interestingly—a guarantee on GPU financing would place the federal balance sheet in the hands of NVIDIA and AMD: If every new GPU design that they put on the market contributes to the value destruction of old GPU models, a GPU financing guarantee would allow NVIDIA and AMD to ignore the market consequences of their own production and to accelerate technical innovation at the most rapid pace possible—all thanks to the federal government, which would be shielding all of NVIDIA's and AMD's customers from asset value deterioration and allowing them, too, to purchase GPUs without thought to risk management.

On one hand, a GPU guarantee might support the policy objective of maximizing technical advancement in GPUs. On the other hand, when NVIDIA already releases new technology at a yearly clip, when costs for data center construction and GPU purchases are ballooning into the tens of billions of dollars, and when it takes years to interconnect data centers to the energy grid,[150] it is likely that the costs of this policy would be far, far higher than any returns it provides to the federal balance sheet or the taxpayer. It would amount to the federal government "going for broke," as it were, to realize the revenues that are ostensibly on the other side of this investment bubble.

Finally, a tax credit for data center construction, server production, and grid components could significantly lower the breakeven costs of construction and installation across the industry. Tax equity monetization and tax credit transfer markets may still be required, to smooth over mismatches between when developers incur costs and when they earn revenues—and to insure against the possibility that developer or producer tax liabilities are less than the value of a tax credit. Tax credits would also serve as collateral for investors (e.g., through a tax credit-backed bridge loan or mini-perm loan). These are useful aspects of tax credits that would support capital expenditure across the AI sector. However, the sector's unclear cash flow profiles and the potential for unexpected losses may make it challenging for developers to monetize the full value of their tax credits: Valuing expected depreciation and taxable losses will be extremely complicated, and tax equity investors may not want to be on the hook for those risks.

And, unlike the tax credits for clean energy resources, which directly shift some of the burden of electricity costs from ratepayers (facing the brunt of rate hikes) onto the federal budget, it is not clear that tax credits for the AI sector would significantly lower any salient costs for consumers: AI inference services are an infinitesimal share of consumer spending compared to, say, housing and energy costs, which an AI tax credit would not mechanically decrease.

None of these avenues for federal financial support are prima facie implausible, especially given the close attention political leaders are paying to the tech sector and the future of AI-related technologies. The federal government has already taken an equity stake in semiconductor company Intel and has expressed willingness to develop financing packages for companies across the semiconductors and critical minerals sectors. The federal government also published an AI Action Plan in July that directs federal agencies to support AI innovation, rapid energy access for data centers, and streamlined data center construction—including through the

[150] Clavenna (2025a).

provision of federal land for data center construction.[151] And the recently passed budget reconciliation package, the One Big Beautiful Bill Act, reinstates 100 percent upfront bonus depreciation expensing for most kinds of capital investments.[152] While the sector continues to grow at a breakneck pace, political leaders may be inclined to find more ways to fuel that growth through the budgetary and investment tools they have available.

B. Consumer demand

It is not immediately clear how exposed most households and firms are to an AI-led contraction in the market, but there are quite a few warnings that exposure might be quite high. AI-focused companies comprise over 80 percent of the gains in the U.S. stock market this past year.[153] The growing concentration of the leading AI-focused companies on the stock market and across its most popular index funds, however, indicates that our money is now being put where their mouths are, too. Stock ownership represents around 30 percent of the net worth of American households[154]—we are all incredibly exposed to the data center boom. The rise of private credit ETFs, pension fund investments in private equity and private credit, and the deepening relationships between banks and private credit all suggest that household assets are mechanically tied up in this boom. A market correction would pose a sharp negative "wealth effect" for most households. As Beggs put it, "A financial crisis reveals how much everybody's livelihoods depend on the confidence of private investors."[155]

Households and firms are more directly exposed to a market correction via their exposure to the regulated banking system. While regulated banks participate in the data center market, they transact mainly in the longer-term CMBS and ABS financings, downstream of the private debt ecosystem, which does most of the debt origination and structuring. Still, any sudden stop in the AI sector severely threatens liquidity across the rest of the country's consumer economy, particularly if local or regional banks are investors in local data center projects due to the outsized geographic concentration of data center-related spending (discussed further in the next subsection). Policymakers should be prepared to support buttressing these banks' liquidity in the event of a downturn.

Analysis from Employ America stressed that there is no countervailing economic dynamism to cushion the impact of a potential market correction in the near-term aftermath: "The losses in a [tech sector] bust will simply be too large and swift to be neatly offset by an imminent and symmetric boom elsewhere. Even as housing and consumer durables ultimately did well following the bust of the 90s tech boom, there was a 1-2 year lag, as it took time for long-term rates to fall and investors to shift their focus."[156]

C. State and local finances

Local policymakers who are counting on tax revenue from data centers to support their budgets should reconsider data center developers' reliability. Such jurisdictions should not only be aware of the solvency of data centers' tenants and the ability of data center owners themselves to

[151] Anderson et al. (2025).
[152] Arun and Feygin (2025).
[153] Sharma (2025).
[154] The Economist (2025).
[155] Beggs (2018).
[156] Amarnath (2025b). Employ America does not view an "AI bust" as imminent due to otherwise positive capital expenditure trends in the sector.

service their debt over the long-term, but also of whether a slowdown would also indirectly impact taxpayer incomes due to job losses.[157]

Tax exemptions for data center development already cause states considerable budgetary losses: At least 10 states already lose over $100 million *annually* through uncapped, blanket, and decades-long exemptions on sales taxes.[158] Some states exempt data centers from property taxes and state income taxes, too, and state-level tax exemptions can sometimes preempt local policymakers' taxation decisions.[159] The rapid growth of data centers across the country will frustrate state and local policymakers' ability to accurately project tax and public budget impacts from data center development; even now, as the boom continues, policymakers should be wary of making too many tax concessions. Federal policymakers weighing the merits of further supporting data center development through federal tax credits and loan guarantees should recognize the sheer scale of state- and local-level subsidies that data center developers already receive.

In New Mexico, the commissioners of Doña Ana County recently approved a sweeping set of tax breaks for a $165 billion data center project built by STACK Infrastructure, comprising four data centers.[160] In exchange for tax relief, STACK Infrastructure will pay the county $300 million—less than one percent of the total value of their projected investments. The tax breaks work through a local industrial revenue bond: The county owns the data center site and leases it to STACK Infrastructure (similar to a sale-leaseback transaction) until the bond matures in thirty years, after which STACK Infrastructure takes ownership. The county's ownership exempts the $165 billion investment from property taxes for the term of the bond, and built into the bond is also a sales tax exemption.[161]

Data centers are regionally concentrated across a few markets.[162] 80 percent of data center capacity is concentrated across just 15 states. Many data centers already operate across the Northeast corridor, the Midwest, California, and Texas—but proposed data center construction is now concentrated in Northern Virginia (including Washington, DC), California, Texas, and Phoenix, AZ. A market correction in the AI sector therefore threatens to affect some geographies much more severely than others—suggesting that there will be regional budgetary impacts through loss of data center property and sales tax revenues, income and sales tax revenues from lost jobs and lower spending, and higher financing costs for future bond issuances. These impacts will likely be incurred at the state level and, as explained below, at the utility level, and therefore they may be less immediately visible on municipal or county-level balance sheets. There may be some lag in the passthrough between state budgetary shocks and municipal balance sheet stress.

[157] Just because data center developers and hyperscalers see returns on equity markets does not mean their investments in physical infrastructure are yielding cash flows or tax revenues.
[158] Tarczynska (2025).
[159] Tarczynska (2025).
[160] Adler (2025a).
[161] Adler (2025b).
[162] Campbell and Beckler (2025) and Roberts (2025).

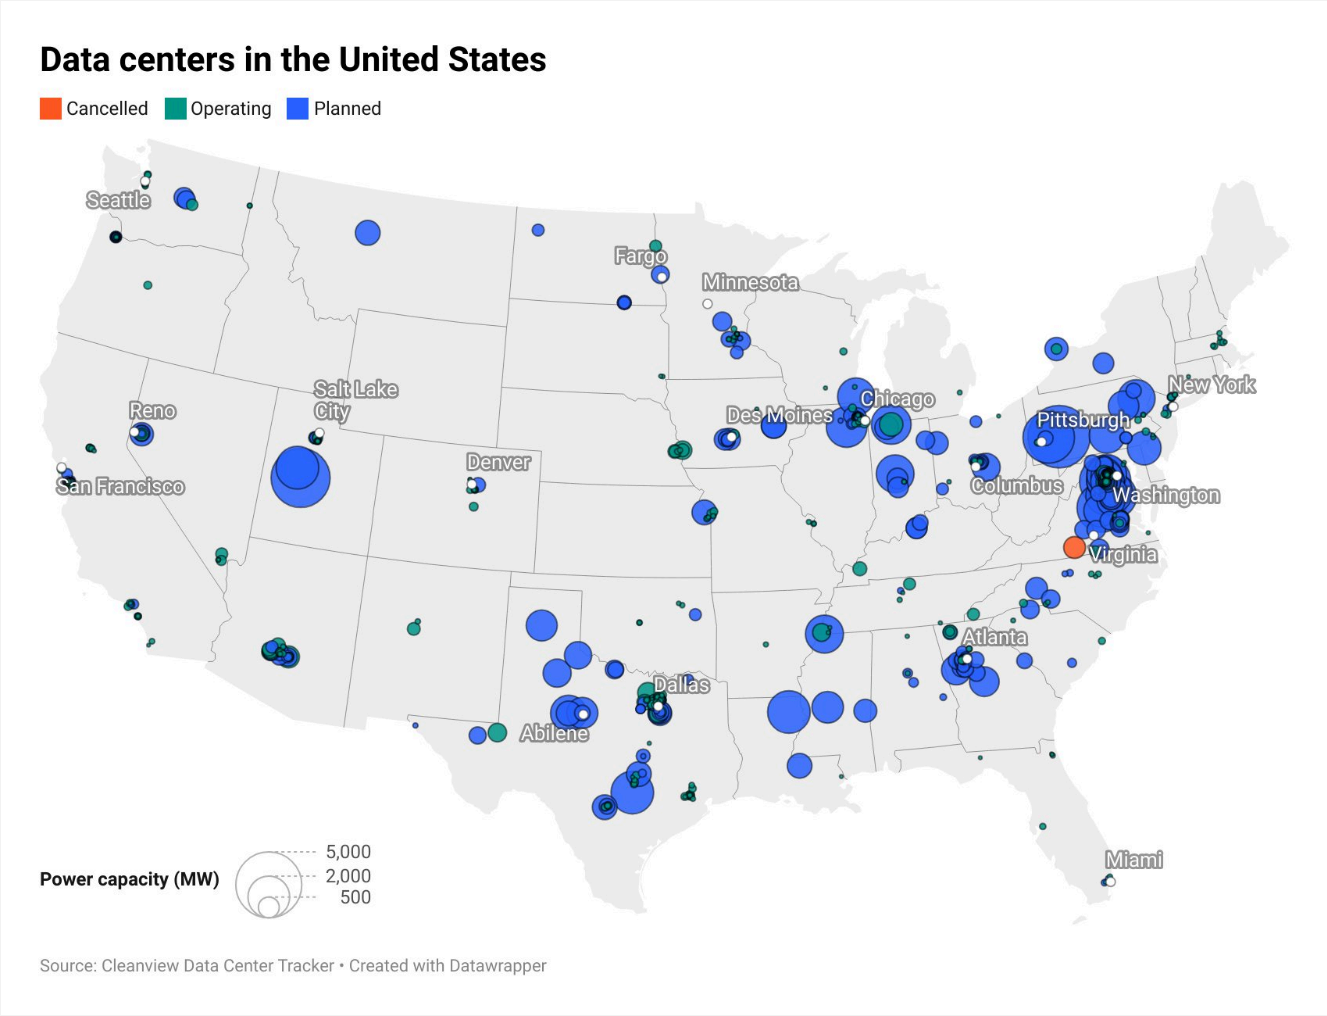

Figure 12. *Data centers in the United States. See: Thomas (2025).*

California's state budget is a glaring example of the threat of a market correction. California's budget, which already depends overwhelmingly on the tax revenues from the tech sector, has a dangerously "inverted capital structure"—a capital structure in which movements in the value of assets and liabilities are strongly correlated. A market correction in the tech sector will quickly deplete the state's resources, while at the same time increasing its obligations due to a particularly severe local economic downturn.[163]

D. Power infrastructure

Data centers, as "powered shells" for intensive computation infrastructure, may have longer-term value to policymakers independent of their short-term value to the market and of their contributions to tax revenues. Expected demand from data centers is bringing into existence a whole host of innovations in supercapacitors, uninterruptible power systems, cooling systems, batteries (including flow batteries)[164], grid-enhancing technologies, and demand-response

[163] Sumagaysay (2024) and Pettis (2015).
[164] Penrod (2025).

technologies. Data centers also have valuable land and grid interconnection rights for the energy generation projects they are supporting. It would be dangerous for long-term technological capacity and for energy technology development if a severe downturn in data center investment also constrained the uptake of these innovative technologies.

In addition to constraining technological innovation and deployment, the prospect of vacant data centers represents a severe budgetary threat to utilities. Regulated utilities around the country have based their power procurement decisions and rate-setting decisions on data center demand growth; other power purchasers may not be able to pay utilities and power developers for power at the high rate that data center developers promised. The loss of data centers' electricity offtake would place these sunk infrastructure development costs onto the bills of other ratepayers and otherwise contribute to financial distress across the utility sector.

Given the importance of power infrastructure for the United States' longer-term energy capacity, affordability, and emissions goals, policymakers must think seriously about how to cushion the impact of a data center market correction for both ratepayers and power developers. Policymakers should develop investment platforms and power offtake warehouses that can plug the hole in missing energy demand at a slightly lower cost—all while securing other sources of power demand. For example, there are all kinds of advanced manufacturing and minerals or metals refining processes that rely on cheaper and reliable power: creative policymakers can use a market correction to connect power-hungry but otherwise lower-margin enterprises like these to power assets that would otherwise be stranded—and, in doing so, the public sector can capture value from facilitating the reinvestment process.

More broadly, policymakers should identify the assets that have long-term cash flow-generating potential and, in the event of a market crash, should take steps to invest in and even purchase them outright when no other market actors want to hold onto them.[165] This principle is generalizeable, and policymakers have already used it to great effect: The post-2008 Troubled Asset Relief Program (TARP) used federal liquidity to stabilize the banking and auto sectors of the U.S. economy by purchasing struggling assets until their long-term cash flows stabilized. When markets and asset values stabilize, particularly when cash flows pick back up, public agencies can make a return from having saved struggling assets. Indeed, TARP returned value to the U.S. taxpayer.

The above recommendation amounts to a public backstop for critical assets. If data centers' technology, services, or licenses have critical long-term value, the public sector should preserve what investments it can and capture that value by finding ways to refinance those assets at lower price points and longer durations. (This recommendation could even encompass older GPUs, in the case that those assets find steady market demand in the future.[166])

It's important that the data center boom, if it is indeed a bubble, becomes a "productive bubble" that leaves the U.S. economy with the ability to take advantage of a more advanced capital stock.[167] Doing so requires policymakers to manage the boom-bust cycle in a way that preserves crucial investments in recessionary conditions—and manages the downside distributional consequences safely, with a view to maintaining the U.S.'s future growth potential and the

[165] Policymakers working to support the federal government's AI Action Plan may be interested in investment planning strategies like these.
[166] CoreWeave itself grew quickly by acquiring the GPUs and processing assets of insolvent cryptocurrency miners after that sector saw a sharp market correction in 2018. It is not clear that a GPU recycling market or ecosystem exists at scale. See: Roundy (2024).
[167] Janeway (2021).

financial stability of the consumer economy. The stakes of this recommendation are high: Political leaders and policymakers will have to negotiate which asset-holders and market actors get liquidated—and which get subsidized, their losses underwritten by public support.[168]

E. Headwinds facing the sector

There are external factors that set macroeconomic constraints on the scale of the data center boom and its broader economic relevance. The current lack of energy supply to match data centers' power needs could stall projects that are not backed by developers with enough capital to procure all the necessary power components well in advance of construction—imposing a physical constraint on any data center investment boom.[169] Gas turbine manufacturers, backlogged as they are, have expressed reticence to significantly expand production to avoid overinvesting into a potential change in demand trajectories.[170] Utilities are complaining about "ghost" interconnection queue submissions from data center developers.[171] The inability to rapidly build and interconnect new energy sources has almost certainly throttled the speed of data center investment.[172] If hyperscalers purchase leading-edge GPUs but cannot wire them into data centers on schedule, those GPUs may lose value even if they do not deteriorate from lack of use due to price pressures from new GPUs coming to market.[173] The fact that data center deployment is now primarily constrained by energy access—there's a reason data center sizes are measured in gigawatts—suggests that the AI sector's growth is now dependent on energy supply growth.

The scale of the AI boom relative to GDP might look higher than it otherwise could have been due to the negative impact tariffs and federal policy uncertainty are having on growth across the rest of the economy. Indeed, the AI sector is exempt entirely from tariffs, and the data center investment boom may never have materialized had the sector faced the same tariff rates as other sectors do.[174] Continued trade policy uncertainty remains a serious threat to the sector.

These are all indications that data centers' investment demand projections are constrained by economic reality. Some financial analysts have noted that data center capital expenditure may *fall* in the coming years.[175]

---

[168] Beggs (2018) and Feygin (2023).
[169] Elder (2025b) and Clavenna (2025a).
[170] Zeitlin (2025) and Arun (2025a).
[171] Martucci (2025).
[172] Clavenna (2025a, 2025b).
[173] Gupta (2025). As described in Section II, the value of a GPU asset is caught between the price pressure of new GPUs coming to market and technological deterioration from utilization.
[174] Politano (2025b). Politano later stated that "the AI industry is already basically getting bailed out (because they're exempt from the tariffs the rest of us have to pay)." See: Politano (2025a).
[175] Quintanilla (2025).

V. Conclusion

The proponents of AI technology are convinced that it will reshape the world—and have put their money where their mouths are. By some accounts, the data center boom is now larger than previous technology booms—the railroad and dotcom bubbles are most frequently referenced[176]—and, as a share of GDP, it is comparable to the housing bubble, too.[177] This boom might involve a unique technology—but data centers are still an asset class akin to commercial real estate, with its attendant cash flow, collateral, and demand risks. The dominance of hyperscalers initially limited the amount of debt across the sector, but stratospherically rising capital expenditure needs for GPU purchases and data center construction have made debt a crucial if invisible foundation for the industry's continued expansion.

The messy details of data center project finance should not be the sole purview of Wall Street private credit firms—after all, our exposure to the sector only grows more concentrated by the day. Mapping out how precisely capital is flowing through the sector should help financial policymakers and industrial policy thinkers understand the promises and pitfalls of the data center boom for the U.S. economy and for its innovation landscape. Political leaders should be prepared to tackle the downside distributional questions raised by the instability of this data center boom.

It should be a cause for concern that the AI sector is not only the main source of growth in an otherwise sclerotic economy but is also a concentrated set of hyperscalers engaging in "circular" transactions with shaky long-term cash flow-generating potential.[178] This sparkling sector is no replacement for industrial policy and macroeconomic investment conditions that create broad-based sources of demand growth and prosperity—and this report has not even covered the risks to the broader energy, manufacturing, and construction sectors that data centers rely on in the first place. Public agencies should tread carefully. In lieu of a major rebalancing of U.S. economic growth and a diversification of its engines, policymakers should prepare an investment strategy that makes use of potentially recessionary periods to acquire assets and repurpose them to serve future demand.

The fact that the data center boom is threatened by, at its core, a lack of consumer demand and the resulting unstable investment pathways, is itself an ironic miniature of the U.S. economy as a whole. Just as stable investment demand is the linchpin of sectoral planning, stable aggregate demand is the keystone in national economic planning. Without it, capital investment crumbles.

---

[176] Mims (2025b) and the Economist (2025).
[177] Amarnath (2025a).
[178] As Keynes put it, "when the capital development of a country becomes a by-product of the activities of a casino, the job is likely to be ill-done." See: Keynes (1936).

VI. References

Aakash Gupta [@aakashg0]. 2025. “Satya just told you the entire AI trade thesis is wrong and nobody is repricing anything. Microsoft has racks of H100s collecting dust because they literally cannot plug them in. Not ‘won't,’ cannot. The power infrastructure does not exist. Which means every analyst model that's.” Tweet. Twitter, November 3. https://x.com/aakashg0/status/1985176339712970900.

Acton, Michael, and Stephen Morris. 2025. “Investors Glimpse Pay-off for Big Tech's Mammoth Spending on AI Arms Race.” *The Financial Times*, August 1.

Adinolfi, Joseph. 2025a. “The Race to Launch Ever-Riskier Leveraged ETFs in the U.S. Is Heating Up.” Morningstar, Inc., October 23. https://www.morningstar.com/news/marketwatch/20251023222/the-race-to-launch-ever-riskier-leveraged-etfs-in-the-us-is-heating-up.

Adinolfi, Joseph. 2025b. “These 3 Charts Show How the Rise of Technology Stocks Has Transformed the U.S. Equity Market.” MarketWatch, September 17. https://www.marketwatch.com/story/these-3-charts-show-how-the-rise-of-technology-stocks-has-transformed-the-u-s-equity-market-6c1c0678.

Adler, Maxwell. 2025a. “AI Data Center Gets Approval for $165 Billion Phantom Debt Deal.” Bloomberg News, September 19. https://www.bloomberg.com/news/articles/2025-09-19/ai-data-center-gets-approval-for-165-billion-phantom-debt-deal.

Adler, Maxwell. 2025b. “New Mexico AI Data Centers Nears Tax Break With $165 Billion of Phantom Debt.” Bloomberg News, September 4. https://www.bloomberg.com/news/articles/2025-09-04/ai-data-centers-near-tax-break-with-165-billion-of-phantom-debt.

“AI Inference vs. Training – What Hyperscalers Need to Know.” 2025. *EdgeCore Digital Infrastructure*, June 23. https://edgecore.com/ai-inference-vs-training/.

Ainvest. 2025. “CoreWeave's GPU-Backed Debt Strategy Inspires AI Startups to Borrow Big.” July 8. https://www.ainvest.com/news/coreweave-gpu-backed-debt-strategy-inspires-ai-startups-borrow-big-2507/.

Alex Turnbull [@alexbhturnbull]. 2025. “Low key cannot believe we are doing the fracking maths all over again and everyone is falling into the same camps...” Tweet. Twitter, November 3. https://x.com/alexbhturnbull/status/1985251734005137736.

Alloway, Tracy, and Joe Weisenthal. 2025. *Josh Wolfe on AI and the Breaking of Silicon Valley's Social Contract*. Bloomberg Odd Lots. September 8. https://www.bloomberg.com/news/articles/2025-09-08/josh-wolfe-on-ai-and-the-breaking-of-silicon-valley-s-social-contract.

Amarnath, Skanda. 2025a. “Financial Conditions Monitor: A Cohesive Fed and a Tech Boom That Needs To Be Respected.” Employ America, September 22. https://www.employamerica.org/financial-conditions-monitor/financial-conditions-monitor-september-22-2025/.

Amarnath, Skanda. 2025b. “Financial Conditions Monitor: Tech + Mag 7 > 50% of S&P 500. AI Will Drive The The Business Cycle. Powell Gives The Hawks A Voice.” Employ America, November 3. https://www.employamerica.org/financial-conditions-monitor/financial-conditions-monitor-november-3-2025/.

Amodeo, Jeannine. 2025. “Record $38 Billion Debt Sale Nears for Oracle-Tied Data Centers.” Bloomberg News, October 23.

https://www.bloomberg.com/news/articles/2025-10-23/record-38-billion-debt-sale-nears-for-oracle-tied-data-centers.

Anderson, Hope, Henry Y. Huang, and John Oltean. 2025. "White House Unveils Comprehensive AI Strategy: 'Winning the Race: America's AI Action Plan.'" White & Case LLP, August 14. https://www.whitecase.com/insight-alert/white-house-unveils-comprehensive-ai-strategy-winning-race-americas-ai-action-plan.

Arroyo, Carmen, and Laura Benitez. 2025a. "Blue Owl Seals Largest Private Capital Deal for Meta's AI Growth." *Bloomberg News*, October 16. https://www.bloomberg.com/news/articles/2025-10-16/blue-owl-seals-largest-private-capital-deal-for-meta-s-ai-growth.

Arroyo, Carmen, and Laura Benitez. 2025b. "How Pimco Outmaneuvered Apollo and KKR to Win $29 Billion Meta Deal." *Bloomberg News*, August 19. https://www.bloomberg.com/news/articles/2025-08-19/how-pimco-outmaneuvered-apollo-kkr-to-win-29-billion-meta-deal.

Arun, Advait. 2024. "Project Finance." *Center for Public Enterprise*, September 20. https://publicenterprise.org/report/project-finance/.

Arun, Advait. 2025a. "Committing to the Drill Bit." *Center for Public Enterprise*, April 29. https://publicenterprise.org/report/committing-to-the-drill-bit/.

Arun, Advait. 2025b. "The Natural Gas Turbine Crisis." Heatmap News, February 26. https://heatmap.news/ideas/natural-gas-turbine-crisis.

Arun, Advait, and Yakov Feygin. 2025. "Leveling the Playing Field." Commentary. *Center for Public Enterprise*, August 11. https://publicenterprise.org/leveling-the-playing-field/.

Ashare, Matt. n.d. "Cloud Data Centers Get Bigger, Denser amid AI Building Boom." Utility Dive. Accessed October 8, 2025. https://www.ciodive.com/news/cloud-ai-data-center-aws-microsof-google-oracle/743218/.

Barth, Adam, Chhavi Arora, Gayatri Shenai, Jesse Noffsinger, and Pankaj Sachdeva. 2025. "The Future of US Hyperscale Data Centers." McKinsey, August 8. https://www.mckinsey.com/industries/public-sector/our-insights/the-data-center-balance-how-us-states-can-navigate-the-opportunities-and-challenges.

Basil, Arasu Kannagi, Echo Wang, and Timothy Gardner. 2025. "Rick Perry's Fermi Cashes in on AI Frenzy, Valued at $15 Billion in Nasdaq Debut." Reuters, October 1. https://www.reuters.com/business/rick-perrys-data-center-reit-set-nasdaq-debut-after-683-million-ipo-2025-10-01/.

Bass, Dina. 2025. "Dell Hikes Estimates for Next Four Years on Strong AI Demand." Bloomberg News, October 7. https://www.bloomberg.com/news/articles/2025-10-07/dell-raises-estimates-for-next-four-years-on-booming-ai-demand.

Batson, Andrew. 2025. "North America Data Center Report Midyear 2025." JLL. https://www.jll.com/en-us/insights/market-dynamics/north-america-data-centers.html.

Beam, Christopher. 2025. "'I Believe It's a Bubble': What Some Smart People Are Saying About AI." *Bloomberg News*, October 9. https://www.bloomberg.com/news/articles/2025-10-09/why-experts-are-warning-the-ai-boom-could-be-a-bubble.

Beggs, Mike. 2018. "The Minsky Millennium." Jacobin, November 27. https://jacobin.com/2018/11/the-minsky-millennium.

Bratton, Laura. 2025a. "Amazon Stock Falls as Raymond James Downgrades Shares, Citing Tariff Headwinds and 'limited' AI Monetization." Yahoo Finance, April 21.

https://finance.yahoo.com/news/amazon-stock-falls-as-raymond-james-downgrades-shares-citing-tariff-headwinds-and-limited-ai-monetization-144747355.html.

Bratton, Laura. 2025b. “Big Tech’s Spending Drove Nvidia’s Rise.” Yahoo Finance, May 24. https://finance.yahoo.com/news/big-techs-spending-drove-nvidias-rise-154027146.html.

Brown, Eliot, and Robbie Whelan. 2025. “Spending on AI Is at Epic Levels. Will It Ever Pay Off?” The Wall Street Journal, September 26. https://www.wsj.com/tech/ai/ai-bubble-building-spree-55ee6128.

Cahn, David. 2024. “The Game Theory of AI CapEx.” *Sequoia Capital*, July 16. https://sequoiacap.com/article/ai-optimism-vs-ai-arms-race/.

Cahn, David. 2025. “Circular Deals & Supply Chain Dynamics.” Substack. *David Cahn’s Substack*, October 13. https://dcahn.substack.com/p/circular-deals-and-supply-chain-dynamics.

Campbell, Dakin. 2025. “Why the Biggest Risk in AI Might Not Be the Technology, but the Trillion-Dollar Race to Build It.” Business Insider, October 7. https://www.businessinsider.com/big-tech-ai-capex-infrastructure-data-center-wars-2025-10.

Campbell, Dakin, and Hannah Beckler. 2025. “Are You Living near a Data Center? Our Interactive Map Shows Where Construction Is Skyrocketing.” Business Insider, September 29. https://www.businessinsider.com/data-center-locations-us-map-ai-boom-2025-9.

Carpenter, Scott. 2025. “Data Centers to Propel Infra Securitizations Past $110 Billion by 2026 End.” *Bloomberg News*, August 25. https://www.bloomberg.com/news/articles/2025-08-25/data-centers-to-propel-infra-securitizations-past-110-billion-by-2026-end.

CBRE. 2025. “2025 Global Data Center Investor Intentions Survey.” August 6. https://www.cbre.com/insights/reports/2025-global-data-center-investor-intentions-survey.

Challapally, Aditya, Chris Pease, Ramesh Raskar, and Pradyumna Chari. 2025. “The GenAI Divide: State of AI in Business 2025.” *MIT NANDA*, August 18. https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf.

Cheah, Eugene. 2024. “$2 H100s: How the GPU Bubble Burst.” Substack. Latent Space, October 10. https://www.latent.space/p/gpu-bubble.

Clavenna, Scott. 2025a. “Are We in an AI-Energy Bubble?” *Latitude Media*, March 12. https://www.latitudemedia.com/news/are-we-in-an-ai-energy-bubble/.

Clavenna, Scott. 2025b. “The AI Bubble Checklist.” *Latitude Media*, October 8. https://www.latitudemedia.com/news/the-ai-bubble-checklist/.

Close, Dan, and Margot A. Kleinman. 2025. “Powering Growth: Data Centers Impact Muni Issuers’ Risk and Reward.” Nuveen, April 9. https://www.nuveen.com/en-us/insights/municipal-bond-investing/data-centers-impact-muni-issuers-risk-and-reward.

Corrigan, Hope. 2025. “Startups Are Using Nvidia’s AI GPUs as Collateral to Secure Loans of up to $10 Billion from Financial Institutions.” *PC Gamer*, July 10. https://www.pcgamer.com/hardware/startups-are-using-nvidias-ai-gpus-as-collateral-to-secure-loans-of-up-to-usd10-billion-from-financial-institutions/.

“Covenant Lite #29: Meta’s $29 Billion Bet with Apollo on AI Data Centers.” 2025. Substack. *Covenant Lite*, July 20. https://covenantlite.substack.com/p/covenant-lite-29-metas-29-billion.

Craske, Ben. 2025. “Top 10: Construction Companies Building Data Centres.” Data Centre Magazine, October 1. https://datacentremagazine.com/top10/top-10-construction-companies-building-data-centres.

Crawford, David, and Anne Hoecker. 2025. “$2 Trillion in New Revenue Needed to Fund AI’s Scaling Trend - Bain & Company’s 6th Annual Global Technology Report.” Bain, September 23. https://www.bain.com/about/media-center/press-releases/20252/$2-trillion-in-new-revenue-needed-to-fund-ais-scaling-trend---bain--companys-6th-annual-global-technology-report/.
Ding, Ethan. 2025. “Tokens Are Getting More Expensive.” Substack. *Mandates*, July 31. https://ethanding.substack.com/p/ai-subscriptions-get-short-squeezed.
Eastland, Maggie, and Shirin Ghaffary. 2025. “OpenAI Asks US to Expand Chips Act Tax Credit to AI Data Centers.” *Bloomberg News*, November 7. https://www.bloomberg.com/news/articles/2025-11-07/openai-asks-us-to-expand-chips-act-tax-credit-to-ai-data-centers.
Edwards, Jim. 2025. “AI Isn’t in a Bubble—the Cash (and the Hype) Are Real, These Analysts Say.” Fortune, October 9. https://fortune.com/2025/10/09/ai-isnt-in-a-bubble-analysts-say/.
Elder, Bryce. 2025a. “Eight Odd Things in CoreWeave’s IPO Prospectus.” *The Financial Times*, March 27.
Elder, Bryce. 2025b. “What’ll Happen If We Spend Nearly $3tn on Data Centres No One Needs?” *The Financial Times*, July 30.
Equinix. 2024. “Hyperscale vs. Colocation.” June 26. https://blog.equinix.com/blog/2020/08/27/hyperscale-vs-colocation/.
Feygin, Yakov. 2023. “SVB, Bailouts, and How to Regulate Capitalism.” Building a Ruin, March 13. https://building-a-ruin.ghost.io/svb-regulation-bailouts-capitalism/.
Fiegerman, Seth, and Carmen Reinicke. 2025. “Why Fears of a Trillion-Dollar AI Bubble Are Growing.” Yahoo Finance, October 4. https://finance.yahoo.com/news/why-fears-trillion-dollar-ai-130008034.html.
Fishlow, Olivia. 2025. “Private Credit-Powered AI Boom at Risk of Overheating, UBS Says.” *Bloomberg News*, August 18. https://www.bloomberg.com/news/articles/2025-08-18/private-credit-powered-ai-boom-at-risk-of-overheating-ubs-says.
Fitch, Asa. 2025. “How Nvidia Adapted Its Chips to Stay Ahead of an AI Industry Shift.” Tech. *Wall Street Journal*, February 27. https://www.wsj.com/tech/ai/how-nvidia-adapted-its-chips-to-stay-ahead-of-an-ai-industry-shift-044edaad.
Fitch Ratings. 2025a. “Fitch Rates CoreWeave’s Proposed New Notes ‘BB-’/’RR4.’” July 21. https://www.fitchratings.com/research/corporate-finance/fitch-rates-coreweave-proposed-new-notes-bb-rr4-21-07-2025.
Fitch Ratings. 2025b. “Project Finance Digital Infrastructure Rating Criteria.” August 8. https://www.fitchratings.com/research/infrastructure-project-finance/project-finance-digital-infrastructure-rating-criteria-08-08-2025.
Forgash, Emily, and Agnee Ghosh. 2025. “OpenAI, Nvidia Fuel $1 Trillion AI Market With Web of Circular Deals.” *Bloomberg News*, October 7. https://www.bloomberg.com/news/features/2025-10-07/openai-s-nvidia-amd-deals-boost-1-trillion-ai-boom-with-circular-deals.
Fox, Matthew. 2024. “There’s a Hidden Risk Lurking for AI Stocks in 2025.” Markets Insider, November 15. https://markets.businessinsider.com/news/stocks/ai-stocks-risk-nvidia-gpus-blackwell-tech-outlook-depreciation-accounting-2024-8.
Gardizy, Anissa, Anita Ramaswamy, and Cory Weinberg. 2025. “Internal Oracle Data Show Financial Challenge of Renting Out Nvidia Chips.” The Information, October 7.

https://www.theinformation.com/articles/internal-oracle-data-show-financial-challenge-renting-nvidia-chips.
Giacobone, Bianca. 2025. “Is Fermi America’s IPO All Hype?” *Latitude Media*, September 10. https://www.latitudemedia.com/news/is-fermi-americas-ipo-all-hype/.
Goldsmith, Ian, and Zach Byrum. 2025. “Powering the US Data Center Boom: Why Forecasting Can Be So Tricky.” *World Resources Institute*, September 17. https://www.wri.org/insights/us-data-centers-electricity-demand.
Hammond, George, and Michael Acton. 2025. “OpenAI Targets 10% AMD Stake via Multibillion-Dollar Chip Deal.” Artificial Intelligence. *The Financial Times*, October 6. https://www.ft.com/content/bfafd06e-0a92-4add-9ae5-622e3c2c8f29.
Hammond, George, and Tabby Kinder. 2025a. “OpenAI’s Computing Deals Top $1tn.” Artificial Intelligence. *The Financial Times*, October 7. https://www.ft.com/content/5f6f78af-aed9-43a5-8e31-2df7851ceb67.
Hammond, George, and Tabby Kinder. 2025b. “Oracle to Buy $40bn of Nvidia Chips for OpenAI’s New US Data Centre.” Oracle Corp. *The Financial Times*, May 23. https://www.ft.com/content/a9cd130f-f6bf-4750-98cc-19d87394e657.
Hoffman, Liz. 2025. “I Squared CEO on the AI Boom: ‘I Politely Decline.’” Semafor, October 21. https://www.semafor.com/article/10/20/2025/i-squared-ceo-on-the-ai-boom-i-politely-decline.
Janeway, William. 2021. “Productive Bubbles.” *Noema*, July 27. https://www.noemamag.com/productive-bubbles.
Jin, Berber. 2025. “Exclusive | OpenAI Isn’t Yet Working Toward an IPO, CFO Says.” Wall Street Journal, November 5. https://www.wsj.com/tech/ai/openai-isnt-yet-working-toward-an-ipo-cfo-says-58037472.
Kabir, Usman. 2025. “Microsoft (MSFT) Hits $13B AI Revenue Run Rate, Plans $80B Investment.” Yahoo Finance, April 13. https://finance.yahoo.com/news/microsoft-msft-hits-13b-ai-180154434.html.
Kakashii [@kakashiii111]. 2025. “CoreWeave and the Never-Ending GPU Depreciation — A Masterclass in Accounting Elasticity CoreWeave and Nebius, two public companies riding the AI infrastructure wave, both operate nearly identical business models: providing high-performance GPUs — such as NVIDIA’s H100 and H200 https://t.co/pmRGsSnuEx.” Tweet. Twitter, July 6. https://x.com/kakashiii111/status/1941922959792239011.
Kedrosky, Paul. 2025. “Minsky Moments and AI CapEx.” Paul Kedrosky, October 17. https://paulkedrosky.com/minsky-moments-and-ai-capex/.
Keynes, John Maynard. 1936. *The General Theory of Employment, Interest, and Money*.
Kinder, Tabby. 2024. “Wall Street Frenzy Creates $11bn Debt Market for AI Groups Buying Nvidia Chips.” *The Financial Times*, November 4. https://www.ft.com/content/41bfacb8-4d1e-4f25-bc60-75bf557f1f21.
Kinlaw, William, Mark Kritzman, and David Turkington. 2019. “Crowded Trades: Implications for Sector Rotation and Factor Timing.” *The Journal of Portfolio Management* 45 (5): 46–57. https://doi.org/10.3905/jpm.2019.45.5.046.
Krouse, Sarah. 2025. “Trump Adviser David Sacks Says There Won’t Be a Federal Bailout for AI Companies.” Wall Street Journal, November 6. https://www.wsj.com/tech/ai/trump-adviser-sacks-says-there-wont-be-a-federal-bailout-for-ai-391f2024.
LG. 2025a. “How Long Do GPUs Last Anyway? A Look Into Hyperscalers’ Depreciation Policies, GPUaaS Unit Economics.” Substack. *Applied Conjectures*, May 18. https://appliedconjectures.substack.com/p/how-long-do-gpus-last-anyway-a-look.

LG. 2025b. "TSMC: And Then There Was One." Substack. *Applied Conjectures*, June 29. https://appliedconjectures.substack.com/p/tsmc-and-then-there-was-one.

Love, David. 2024. "Moody's Policy Shift Fuels Hybrid Bond Growth across Key U.S. Sectors." Nasdaq, December 19. https://www.nasdaq.com/articles/moodys-policy-shift-fuels-hybrid-bond-growth-across-key-us-sectors.

Mackereth, Oscar. 2025. "Accounting for AI: Financial Accounting Issues and Capital Deployment in the Hyperscaler Landscape." *Cerno Capital - Investment Management*, April 10. https://cernocapital.com/accounting-for-ai-financial-accounting-issues-and-capital-deployment-in-the-hyperscaler-landscape.

Martin, Nik. 2025. "AI Bubble about to Pop as Returns on Investment Fall Short?" DW, October 11. https://www.dw.com/en/ai-bubble-about-to-pop-as-returns-on-investment-fall-short-chatgpt-microsoft-nvidia-grok-elon-musk/a-74636881.

Martucci, Brian. 2025. "A Fraction of Proposed Data Centers Will Get Built. Utilities Are Wising Up." Utility Dive, May 15. https://www.utilitydive.com/news/a-fraction-of-proposed-data-centers-will-get-built-utilities-are-wising-up/748214/.

Mills. 2025. "Axios Markets." Axios, October 3. https://www.axios.com/newsletters/axios-markets-c9496d70-9fcd-11f0-b981-993edcc2a8f0.html?utm_source=newsletter&utm_medium=email&utm_campaign=newsletter_axiosmarkets&stream=business.

Mims, Christopher. 2025a. "Cutting-Edge AI Was Supposed to Get Cheaper. It's More Expensive Than Ever." Tech. *Wall Street Journal*, August 30. https://www.wsj.com/tech/ai/ai-costs-expensive-startups-4c214f59.

Mims, Christopher. 2025b. "Silicon Valley's New Strategy: Move Slow and Build Things." The Wall Street Journal, August 1. https://www.wsj.com/tech/ai/silicon-valley-ai-infrastructure-capex-cffe0431.

Minsky, Hyman P. 1992. *The Financial Instability Hypothesis*. no. Working Paper No. 74 (May). https://www.levyinstitute.org/pubs/wp74.pdf.

Minsky, Hyman P. 2008. *John Maynard Keynes*. McGraw Hill.

Moody's. 2025. "Data Centers." https://www.moodys.com/web/en/us/insights/credit-risk/data-centers.html.

Mutua, Caleb. 2025. "JPMorgan Says $1.2 Trillion Debt Tied to AI Tops Bank High Grade." Bloomberg News, October 7. https://www.bloomberg.com/news/articles/2025-10-07/jpmorgan-says-1-2-trillion-debt-tied-to-ai-tops-bank-high-grade.

Naprys, Ernestas. 2025. "Why Are $30,000 AI Chips Now Renting for Just $1 per Hour?" Cybernews, February 24. https://cybernews.com/tech/gpu-rental-prices-are-crashing/.

Nguyen, Khac Phu. 2025. "Musk's $13B AI Gamble: xAI Burns $1B/Month with Just $500M in Revenue." Yahoo Finance, June 18. https://finance.yahoo.com/news/musks-13b-ai-gamble-xai-164425784.html.

Norton Rose Fulbright. 2024. "Credit Tenant Lease (CTL) Financing." https://www.nortonrosefulbright.com/en/knowledge/publications/3ff191c8/credit-tenant-lease-financing.

NVIDIA. 2025. "NVIDIA Corporation - Financial Reports." July 25. https://investor.nvidia.com/financial-info/financial-reports/.

Papaspanos, John, Matt Gurch, Clint Rancher, Spencer Withers, and Ben Zeter. 2025. "Financing the Future: Trends and Challenges in Data Center Infrastructure for AI." Baker Botts, March 7.

https://www.bakerbotts.com/thought-leadership/publications/2025/march/financing-the-future-trends-and-challenges-in-data-center-infrastructure-for-ai.

Park, J. 2025. "OpenAI Leases NVIDIA GPUs to Cut Costs, Avoid Obsolescence." The Chosun Daily, September 29. https://www.chosun.com/english/industry-en/2025/09/29/NND7FSJ4CRBCVCNR26NSGUO3FA/.

Penrod, Emma. 2025. "Data Centers Could Bring Alternative Battery Types into the Mainstream, Developers Say." Utility Dive, June 3. https://www.utilitydive.com/news/data-center-flow-zinc-battery-xl-eos-prometheus/751144/.

Pettis, Michael. 2015. "Inverted Balance Sheets and Doubling the Financial Bet." Carnegie Endowment for International Peace, January 22. https://carnegieendowment.org/china-financial-markets/2015/01/inverted-balance-sheets-and-doubling-the-financial-bet?lang=en.

Pfeffer, Brandon. 2025. "A Guide to Project Finance for Data Centers." LinkedIn, August 4. https://www.linkedin.com/pulse/guide-project-finance-data-centers-brandon-pfeffer-cma-jq0kf/.

Politano, Joseph. 2025a. "Real ones also know that the AI industry is already basically getting bailed out (because they're exempt from the tariffs the rest of us…." Post. Bluesky, November 6. https://bsky.app/profile/josephpolitano.bsky.social/post/3m4xynp6xg22o.

Politano, Joseph. 2025b. "The Tariff Exemption Behind the AI Boom." Apricitas, June 15. https://www.apricitas.io/p/the-tariff-exemption-behind-the-ai.

Quantian. 2025. "This is also why NVDA is a core financier for all these deals: they are already massively long 'can Nvidia make a good next-gen AI chip': by…." Post. Bluesky, October 15. https://bsky.app/profile/quantian.bsky.social/post/3m3auul3ms22l.

Quintanilla, Carl. 2025. "Goldman sees hyperscaler capex growth slowing down dramatically, beginning in Q4." Post. Bluesky, September 8. https://bsky.app/profile/carlquintanilla.bsky.social/post/3lyedxvckbs2c.

Rattner, Nate, and Jason Dean. 2025. "Tech Giants Double Down on Their Massive AI Spending." Tech. *Wall Street Journal*, February 7. https://www.wsj.com/tech/ai/tech-giants-double-down-on-their-massive-ai-spending-b3040b33.

Reinicke, Carmen, and Alexandra Semenova. 2025. "OpenAI Funding Questions Put Stress on Stock Market's AI Rally." *Bloomberg News*, November 6. https://www.bloomberg.com/news/articles/2025-11-06/openai-funding-questions-put-stress-on-stock-market-s-ai-rally.

Reuters. 2024. "Lambda Secures $500 Mln Loan with Nvidia Chips as Collateral." April 4. https://www.reuters.com/technology/lambda-secures-500-mln-loan-with-nvidia-chips-collateral-2024-04-04/.

Reuters. 2025a. "CoreWeave, Nvidia Sign $6.3 Billion Cloud Computing Capacity Order." Business. September 15. https://www.reuters.com/business/coreweave-nvidia-sign-63-billion-cloud-computing-capacity-order-2025-09-15/.

Reuters. 2025b. "OpenAI Hits $12 Billion in Annualized Revenue, The Information Reports." July 30. https://www.reuters.com/business/openai-hits-12-billion-annualized-revenue-information-reports-2025-07-31/.

Rivera, Christy, and Christine Brozynski. 2025. "Data Center Financing Structures." Norton Rose Fulbright: Project Finance, June 11.

https://www.projectfinance.law/publications/2025/june/data-center-financing-structures/.
Robbins, Jacob, and Michael Bodley. 2025. "As Venture Debt Gambles on GPUs, Not All Are Sold on Silicon-Backed Loans." PitchBook, August 27. https://pitchbook.com/news/articles/ai-venture-debt-gpu-chip-backed-loans.
Roberts, Billy. 2025. "Data Center Infrastructure in the United States, 2025 (Map)." https://research-hub.nrel.gov/en/publications/data-center-infrastructure-in-the-united-states-2025-map/.
Roberts, Lance. 2025. "Leveraged ETFs: Yet Another Sign of Rampant Speculation." Investing.Com, October 22. https://www.investing.com/analysis/leveraged-etfs-yet-another-sign-of-rampant-speculation-200668904.
Rosoff, Matt. 2025. "Microsoft Earnings Suggest $11.5B OpenAI Quarterly Loss." The Register, October 29. https://www.theregister.com/2025/10/29/microsoft_earnings_q1_26_openai_loss/.
Roundy, Jacob. 2024. "Chip Recycling: Addressing e-Waste in the AI Hardware Industry." TechTarget, December 18. https://www.techtarget.com/searchdatacenter/tip/Chip-recycling-Addressing-e-waste-in-the-AI-hardware-industry.
Seligson, Paula, Jeannine Amodeo, and Laura Benitez. 2025. "Banks Lend $18 Billion for Oracle-Tied Data Center in New Mexico." *Bloomberg News*, November 7. https://www.bloomberg.com/news/articles/2025-11-07/banks-lend-18-billion-for-oracle-tied-data-center-in-new-mexico.
Shalett, Lisa. 2025. "Hidden Risks Beneath the Bullish Narrative." Morgan Stanley, August 20. https://www.morganstanley.com/insights/articles/tech-stock-rally-investment-risks-2025.
Sharma, Ruchir. 2025. "America Is Now One Big Bet on AI." US Economy. *The Financial Times*, October 6. https://www.ft.com/content/6cc87bd9-cb2f-4f82-99c5-c38748986a2e.
Shilov. 2024. "Datacenter GPU Service Life Can Be Surprisingly Short — Only One to Three Years Is Expected According to Unnamed Google Architect." Tom's Hardware, October 24. https://www.tomshardware.com/pc-components/gpus/datacenter-gpu-service-life-can-be-surprisingly-short-only-one-to-three-years-is-expected-according-to-unnamed-google-architect.
Sløk, Torsten. 2025. "Hyperscaler Capex Spending." *Apollo Academy*, October 27. https://www.apolloacademy.com/hyperscaler-capex-spending/.
Snodin, Ben, David Owen, and Luke Emberson. 2025. "The Combined Revenues of Leading AI Companies Grew by over 9x in 2023-2024." Epoch AI, April 3. https://epoch.ai/data-insights/ai-companies-revenue.
Stylianou, Nassos, Sam Learner, Tim Bradshaw, et al. 2025. "Inside the Relentless Race for AI Capacity." *The Financial Times*, July 31. https://ig.ft.com/ai-data-centres/.
Sumagaysay, Levi. 2024. "Tech Is Growing as a Source of California Tax Revenue — That's the Good News and the Bad News." Economy. *CalMatters*, January 11. https://calmatters.org/economy/2024/01/ca-tech-tax-withholding/.
Syafiq, Muhd. 2025. "The Rise of Neocloud in the AI Landscape." JLL. https://www.jll.com/en-us/insights/the-rise-of-neocloud-in-the-ai-landscape.html.
Szyszka, Ewa. 2025. "Future AI Bills of $100k/Yr per Dev." Substack. Kilo Code Blog, August 26. https://blog.kilocode.ai/p/future-ai-spend-100k-per-dev.
Tarczynska, Kasia. 2025. "Cloudy with a Loss of Spending Control: How Data Centers Are Endangering State Budgets." Publications. *Good Jobs First*, April 24.

https://goodjobsfirst.org/cloudy-with-a-loss-of-spending-control-how-data-centers-are-endangering-state-budgets/.

TeraWulf Inc. 2025. "TeraWulf Inc. Announces Proposed Offering of $3.2 Billion of Senior Secured Notes." October 14. https://investors.terawulf.com/news-events/press-releases/detail/118/terawulf-inc-announces-proposed-offering-of-3-2-billion?utm_source=chatgpt.com.

The Economist. 2025. "What If the AI Stockmarket Blows Up?" September 7. https://www.economist.com/finance-and-economics/2025/09/07/what-if-the-ai-stockmarket-blows-up.

Thomas, Michael. 2023. "These Data Centers Are Getting Really, Really Big." Substack. Distilled, March 8. https://www.distilled.earth/p/these-data-centers-are-getting-really.

Thomas, Michael. 2025. "These Data Centers Are Getting Really, Really Big." Substack. Distilled, October 16. https://www.distilled.earth/p/these-data-centers-are-getting-really.

Thomson, Amy, and Mark Bergen. 2025. "Google Eyes Another Crypto Mining Stake In AI Data Center Rush." *Bloomberg News*, September 25. https://www.bloomberg.com/news/articles/2025-09-25/google-eyes-another-crypto-mining-stake-in-ai-data-center-rush.

Tunguz, Tomasz. 2025. "Circular Financing: Does Nvidia's $110B Bet Echo the Telecom Bubble?" Tom Tunguz, October 3. https://tomtunguz.com/nvidia_nortel_vendor_financing_comparison/.

Vossos, Tasos. 2025. "Boom in Risky Hybrid Debt Is Spreading From US Around the World." *Bloomberg News*, October 9. https://www.bloomberg.com/news/articles/2025-10-09/boom-in-risky-hybrid-debt-is-spreading-from-us-around-the-world.

Walton, Robert. 2025a. "Data Center Activity 'Exploded' in Texas, Spiking Reliability Risks: Monitor." Utility Dive, July 14. https://www.utilitydive.com/news/data-center-activity-has-exploded-in-ercot-spiking-grid-reliability-risk/752780/.

Walton, Robert. 2025b. "Investor-Owned Utilities Could Spend $1.1T between 2025 and 2029: EEI." Utility Dive, October 9. https://www.utilitydive.com/news/investor-owned-utilities-spending-more-than-ever-eei/802315/.

Weinman, Aaron, and Gowri Gurumurthy. 2025a. "'Google Backstop' Adds New Twist to Data Center Financing Frenzy." *Bloomberg News*, October 16. https://www.bloomberg.com/news/articles/2025-10-16/-google-backstop-adds-new-twist-to-data-center-financing-frenzy.

Weinman, Aaron, and Gowri Gurumurthy. 2025b. "TeraWulf Floats $3.2 Billion of Junk Bonds for New York Data Center." *Bloomberg News*, October 14. https://www.bloomberg.com/news/articles/2025-10-14/terawulf-floats-3-2-billion-of-junk-bonds-for-ny-data-center.

Wile, Rob. 2025a. "Big AI's Reliance on Circular Deals Is Raising Fears of a Bubble." NBC News, October 6. https://www.nbcnews.com/business/economy/openai-nvidia-amd-deals-risks-rcna234806.

Wile, Rob. 2025b. "Chipmaker Nvidia to Invest up to $100 Billion in OpenAI." NBC News, September 22. https://www.nbcnews.com/business/business-news/chipmaker-nvidia-invest-100-billion-openai-rcna232974.

Yao, Allen. 2025. “The Great GPU Trade.” Substack. Residual Thoughts, July 16. https://getresidualthoughts.substack.com/p/the-great-gpu-trade.

You, Josh, and David Owen. 2025. “How Much Power Will Frontier AI Training Demand in 2030?” Epoch AI, August 11. https://epoch.ai/blog/power-demands-of-frontier-ai-training.

Yun, Channy. 2025. “Announcing up to 45% Price Reduction for Amazon EC2 NVIDIA GPU-Accelerated Instances | AWS News Blog.” Amazon Web Services, June 5. https://aws.amazon.com/blogs/aws/announcing-up-to-45-price-reduction-for-amazon-ec2-nvidia-gpu-accelerated-instances/.

Zeitlin, Matthew. 2025. “The Gas Turbine Crisis May Be Ending.” Heatmap News, September 2. https://heatmap.news/energy/mitsubishi-heavy-industry-gas-turbines.

Zitron, Ed. 2025a. “CoreWeave Is A Time Bomb.” Where’s Your Ed At, March 17. https://www.wheresyoured.at/core-incompetency/.

Zitron, Ed. 2025b. “The Hater’s Guide To The AI Bubble.” Where’s Your Ed At, July 21. https://www.wheresyoured.at/the-haters-gui/.

Zitron, Ed. 2025c. “Why Everybody Is Losing Money On AI.” Ed Zitron’s Where’s Your Ed At, September 5. https://www.wheresyoured.at/why-everybody-is-losing-money-on-ai/.

Acknowledgments
The author would like to thank Fatima Ahmad, Ray Cai, Chris Schenker, Jared Jageler, and Tom Cui for their feedback and insights throughout the editing of this paper.

Author Contact
Advait Arun, Senior Associate for Capital Markets
advait.arun@publicenterprise.org

Executive Director
Paul Williams
paul.williams@publicenterprise.org

Cover Photo
"*Soap bubble*," Plate VI in "Le Monde Physique." 1882. Mezzotint by M. Rapine, after an image by Alexandre-Blaise Desgoffe. Wellcome Collection, London.

About the Center for Public Enterprise
Center for Public Enterprise is a 501(c)(3) nonprofit organization focused on expanding public sector capacity to deliver broad economic development. Our work focuses on various sectors of the economy, including housing, energy, and finance. For more information, visit our website at publicenterprise.org.