

Here's How the AI Crash Happens

Matteo Wong, Charlie Warzel

The AI boom is visible from orbit. Satellite photos of New Carlisle, Indiana, show greenish splotches of farmland transformed into unmistakable industrial parks in less than a year's time. There are [seven](#) rectangular data centers there, with 23 more on the way.

Inside each of these buildings, endless rows of fridge-size containers of computer chips wheeze and grunt as they perform mathematical operations at an unfathomable scale. The buildings belong to Amazon and are being used by Anthropic, a leading AI firm, to train and run its models. According to one estimate, this data-center campus, far from complete, already demands more than 500 megawatts of electricity to power these calculations—as much as hundreds of thousands of American homes. When all the data centers in New Carlisle are built, they will demand more power than two Atlantas.

The amount of energy and money being poured into AI is breathtaking. Global spending on the technology is projected to hit \$375 billion by the end of the year and half a trillion dollars in 2026. [Three-quarters of gains](#) in the S&P 500 since the launch of ChatGPT came from AI-related stocks; the value of every publicly traded company has, in a sense, been buoyed by an AI-driven bull market. To cement the point, Nvidia, a maker of the advanced computer chips underlying the AI boom, yesterday became the first company in history to be worth \$5 trillion.

Here's another way of thinking about the transformation under way: Multiplying Ford's current market cap 94 times over wouldn't quite get you to Nvidia's. Yet 20 years ago, Ford was worth nearly triple what Nvidia was. Much like how Saudi Arabia is a petrostate, the U.S. is a burgeoning AI state—and, in particular, an Nvidia-state. The number keeps going up, which has a buoying effect on markets that is, in the short term, good. But every good earnings report further entrenches Nvidia as a precariously placed, load-bearing piece of the global economy.

America appears to be, at the moment, in a sort of benevolent hostage situation. AI-related spending now contributes more to the nation's GDP growth than all consumer spending combined, and by another [calculation](#), those AI expenditures accounted for 92 percent of GDP growth during the first half of 2025. Since the launch of ChatGPT, in late 2022, the tech industry has gone from making up 22 percent of the value in the S&P 500 to roughly one-third. Just yesterday, Meta, Microsoft, and Alphabet all reported substantial quarterly-revenue growth, and *Reuters* [reported](#) that OpenAI is planning to go public perhaps as soon as next year at a value of up to \$1 trillion—which would be one of the largest IPOs in history. (An OpenAI spokesperson told *Reuters*, “An IPO is not our focus, so we could not possibly have set a date”; OpenAI and *The Atlantic* have a corporate partnership.)

Many people believe that growth will only continue. “We're gonna need stadiums full of electricians, heavy equipment operators, ironworkers, HVAC technicians,” Dwarkesh Patel and Romeo Dean, AI-industry analysts, [wrote](#) recently. Large-scale data-center build-outs may already be reshaping America's energy systems. OpenAI has announced that it intends to build at least 30 gigawatts' worth of data centers—more power than all of New England requires on even the hottest day—and CEO Sam Altman has said he'd eventually like to build a gigawatt of AI infrastructure every week. Other major tech firms have similar ambitions.

Listen to the AI crowd talk enough, and you'll get a sense that we may be on the cusp of an infrastructure boom. And yet, something strange is happening to the economy. Even as tech stocks have skyrocketed since 2022, the companies' share of net profits from S&P 500 companies has hardly budged. Job openings have fallen despite a roaring stock market, 22 states are in or near a recession, and despite data centers propping up the construction industry, U.S. manufacturing is in decline.

It's clear that AI is both drowning out and obscuring other stories about the wobbling American economy. That's a concern. But even worse: What if AI's promise for American business proves to be a mirage? What happens then?

The yawning gap between data-center expenditures and the rest of the economy has caused whispers of *bubble* to rise to a chorus. A growing number of financial and industry analysts have pointed out the enormous divergence between the historic investments in AI and the tech's relatively modest revenues. For instance, according to *The Information*, OpenAI likely made [\\$4 billion](#) last year but lost \$5 billion (making the idea of a \$1 trillion IPO valuation that much more staggering). From July through September, Microsoft's investments in OpenAI resulted in [losses](#) totaling more than \$3 billion. For that same time period, Meta reported rapidly growing costs due to its AI investments, spooking investors and sending its stock down 9 percent.

Much is in flux. Chatbots and AI chips are getting more efficient almost by the day, while the business case for deploying generative-AI tools remains shaky. A recent [report](#) from McKinsey found that nearly 80 percent of companies using AI discovered that the technology had no significant impact on their bottom line. Meanwhile, nobody can say, beyond a few years, just how many more data centers Silicon Valley will need. There are researchers who believe there may already be enough electricity and computing power to meet generative AI's requirements for years to come.

[Read: The AI industry is radicalizing](#)

The economic nightmare scenario is that the unprecedented spending on AI doesn't yield a profit anytime soon, if ever, and data centers sit at the center of those fears. Such a collapse has come for infrastructure booms past: Rapid construction of canals, railroads, and the fiber-optic cables laid during the dot-com bubble all created frenzies of hype, investment, and financial speculation that crashed markets. Of course, all of these build-outs did transform the world; generative AI, bubble or not, may do the same.

This is why OpenAI, Google, Microsoft, Amazon, and Meta are willing to spend as much as possible, as rapidly as possible, to eke out the tiniest advantage. Even if a bubble pops, there will be winners—each company would like to be the first to build a superintelligent machine. For now, many of these tech companies have cash to burn from their other ventures: Alphabet and Microsoft both made more than \$100 billion in profit over the previous fiscal year, while Meta and Amazon both made more than \$50 billion. But at some point in the near future, data-center spending will likely outpace even these enormous cash flows, reducing Big Tech's liquidity and worrying investors. And so, as the AI arms race continues to escalate, the companies are beginning to [raise](#) outside money—in other words, take on debt.

Here is where the bubble dynamics get complicated. Tech firms don't want to formally take on debt—that is, directly ask investors for loans—because debt looks bad on their balance sheets and could reduce shareholder returns. To get around this, some are partnering with private-equity titans to do some sophisticated financial engineering, Paul Kedrosky, an investor and a financial consultant, told us. These private-equity firms put up or raise the money to build a data center, which a tech company will repay through rent. Data-center leases from, say, Meta can then be repackaged into a financial instrument that people can buy and sell—a bond, in essence. Meta recently did just this: Blue Owl Capital raised money for a massive Meta data center in Louisiana by, in essence, issuing bonds backed by Meta's rent. And multiple data-center leases can be combined into a security and sorted into what are called “tranches” based on their risk of default. Data centers represent an [\\$800 billion market](#) for private-equity firms through 2028 alone. (Meta has [said](#) of its arrangement with Blue Owl that the “innovative partnership was designed to support the speed and flexibility required for Meta's data center projects.”)

In this way, the data-center financing ends up being a real-estate deal as much as an AI deal. If this sounds complicated, it's supposed to: The complexity, investment structure, and repackaging make exactly what is going on hard to parse. And if the dynamics also sound familiar, it's because not two decades ago, the Great Recession was precipitated by banks packaging risky mortgages into tranches of securities that were falsely marketed as high-quality. By 2008, the house of cards had collapsed.

Data-center build-outs aren't the same as subprime mortgages. Still, there is plenty of precarity baked into these investments. Data centers deteriorate rapidly, unlike the more durable infrastructure of canals, railroads, or even fiber-optic cables. Many of the chips inside these buildings become obsolete within a few years, when Nvidia and its competitors release the next wave of bleeding-edge AI hardware. Meanwhile, the returns on scaling up chatbots are, at present, diminishing. The improvements made by each new AI model are becoming smaller and smaller, making the idea that Silicon Valley can spend its way to superintelligence more tenuous by the day.

The people who are paying attention to this cycle are getting anxious. On a scale from one to 10, the AI-bubble concern is: *people posting [memes of Christian Bale's character from The Big Short, squinting in disbelief at his computer monitor](#)*. If tech stocks fall because of AI companies failing to deliver on their promises, the [highly leveraged](#) hedge funds that are invested in these companies could be forced into fire sales. This could create a vicious cycle, causing the financial damage to spread to pension funds, mutual funds, insurance companies, and everyday investors. As capital flees the market, non-tech stocks will also plummet: bad news for anyone who thought to play it safe and invest in, for instance, real estate. If the damage were to knock down private-equity firms (which are invested in these data centers) themselves—which manage [trillions and trillions of dollars](#) in assets and constitute what is basically a global shadow-banking system—that could produce another major crash.

For now, money is still pouring into the AI industry. But there's also something circular about these investments. To wit: OpenAI has agreed to [pay](#) \$300 billion to Oracle for new computing capacity, Oracle is paying Nvidia tens of billions of dollars for [chips](#) to install in one of OpenAI's data centers, and Nvidia has agreed to [invest](#) up to \$100 billion in OpenAI as it deploys Nvidia chips. Attempts to illustrate these circular investments have produced a series of byzantine charts that one software engineer referred to on X as "[the technocapital hyperobject at the end of time](#)." The consensus seems to be that although this is legal, it likely cannot go on forever.

Maybe it will all work out. Three years ago, the generative-AI industry made functionally no revenue; today, it produces tens of billions of dollars annually, a rate of growth that, eventually, could catch up with all of this spending. Generative-AI tools are currently used by hundreds of millions of people, and it's hard to imagine that simply ceasing overnight. Perhaps OpenAI or Anthropic will pull off superintelligence, allowing them to, in the words of the *Bloomberg* [columnist](#) Matt Levine, "create God and then ask it for money."

Data centers take time to approve and build; power plants and transmission lines take perhaps even more. Labor is limited, supply chains hit snags, investment waxes and wanes—meaning that even if these data centers are built at the tremendous scale desired by Altman and his competitors, construction and energy constraints may keep the boom from growing too irresponsibly.

In any case, as we approach the end of 2025, data centers have become a peculiar cultural object. Their immense scale is a physical reminder of the economic dominance of Silicon Valley companies and their seemingly unchecked ambition. The uneasiness they inspire economically is rooted in memories of 2008 but also of the tech industry's own financial chicanery, specifically the 2022 crypto crash, which was facilitated by a circular-payment scheme of its own. (FTX, a crypto exchange, and Alameda Research, a hedge fund, both co-founded by Sam Bankman-Fried, were found to be propping each other up: Alameda bought FTX's bespoke cryptocurrency, and FTX lent Alameda money from its customers' accounts.) And so, in some way, the externalities of the data-center boom, be they [environmental](#) or economic, are tied up in fears of what happens not when these tech companies fail, but when they *succeed*.

[Read: AI is taking drinking water from the desert](#)

Boom and bust can feel like two sides of the same coin: Consider also that if AI companies deliver on their massive investments, it would likely mean producing a technology so capable and revolutionary that it wipes out countless jobs and sends an unprecedented shock wave through the global economy before humans have time to adapt. (Perhaps we will be unable to adapt at all.) If they fail, there will likely be

unprecedented financial turmoil as well.

The biggest lesson of the past two decades of Silicon Valley is that Meta, Amazon, and Google—and even the newer AI labs such as OpenAI—have remade our world and have become unfathomably rich for it, all while being mostly oblivious or uninterested in the fallout. They have chased growth and scale at all costs, and largely, they've won. The data-center build-out is the ultimate culmination of that chase: the pursuit of scale for scale itself. In all scenarios, the outcome seems only to be real, painful disruption for the rest of us.