AI Data Centers Are Semiconductor Fabs on Steroids. Why Would the Boom-Bust Cycle Be Different?

Mathew M. Hubbard | October 29, 2026

I've spent enough time in the semiconductor industry to recognize what happens when everyone sees the same demand forecast and decides to build capacity.

Looking at the enormous investments being made in AI data centers, I can't help thinking we've seen this before.

For all the attention AI is receiving as a revolutionary technology, the economics of the infrastructure supporting it look remarkably familiar.

The economics of AI compute infrastructure increasingly resemble those of semiconductor manufacturing, particularly for operators selling compute capacity. Enormous capital requirements, high fixed costs, utilization-driven economics, and equipment that can become economically obsolete long before it stops working.

In some respects, the economics may be even more challenging. The semiconductor industry has spent more than 50 years demonstrating that extraordinary long-term growth doesn't prevent repeated periods of overcapacity.

Why would AI infrastructure be any different?

The Economics Look Familiar

Building a semiconductor fab requires billions in capital, much of it committed years before the resulting capacity generates revenue.

Once operational, utilization becomes critical. The difference between running at 95% capacity and 65% can be the difference between attractive returns and substantial losses. The fixed costs don't disappear when demand slows. Neither does depreciation.

AI data centers face much the same problem. Billions are being committed to facilities, power infrastructure, cooling, networking, and accelerators based on expectations of future demand.

Once operational, those assets must generate enough economic value to recover their investment before the next generation of technology changes the competitive equation. And just like a fab, an AI data center doesn't stop depreciating when utilization falls.

The Equipment Problem May Be Worse

Semiconductor fabs have one significant advantage. They can remain productive for decades. A fab that can no longer compete at the leading edge may continue producing analog devices, sensors, power management ICs, and other products that don't require the latest process technology.

Older, fully depreciated equipment can become a competitive advantage. AI accelerators may not enjoy the same economic life.

© 2026 Mathew M. Hubbard. All rights reserved. • mathewhubbard.com

Each new generation promises better performance and lower cost per unit of computation. Older accelerators can still perform useful work, particularly for inference, but that doesn't guarantee they'll generate the returns originally expected.

Consider a simplified example. A $20 billion investment in accelerators with a five-year economic life represents $4 billion in annual straight-line depreciation, assuming no residual value. Shorten that economic life to three years, and annual depreciation rises to roughly $6.7 billion.

That's a 67% increase before considering financing, power, maintenance, or other operating costs. Real-world economics are more complicated, but the underlying problem remains.

The question isn't how long the equipment will operate. It's how long it will remain economically competitive.

Fifty Years of History

The semiconductor industry has been through this cycle repeatedly. Demand accelerates. Capacity tightens. Prices improve. Manufacturers respond by investing in new capacity.

The problem is that competitors are generally looking at the same market conditions and reaching similar conclusions. Because capacity takes time to build, investment decisions are made long before anyone knows what demand will look like when the equipment finally comes online.

Eventually, the new capacity arrives. Sometimes demand absorbs it. Sometimes it doesn't. Utilization falls. Pricing weakens. Margins compress. Capital spending gets cut. Then the cycle begins again.

I've watched versions of this play out at least six times during my own career.

It doesn't necessarily mean anyone made a bad decision. Each investment can look perfectly reasonable when evaluated independently. The problem is what happens when everyone makes that reasonable decision at approximately the same time. That's what makes these cycles so difficult to avoid.

Exponential Growth May Actually Drive the Cycle

One argument I frequently hear is that AI demand will continue growing exponentially. I suspect it will. But that doesn't eliminate the risk of overcapacity. It may actually increase it.

Semiconductor demand has compounded dramatically over the past 50 years. It didn't prevent the boom-bust cycle. Some of the industry's most aggressive periods of expansion have also been accompanied by sharp investment and correction cycles. That doesn't establish that faster growth always produces shorter cycles, but it demonstrates that rapid growth is no protection against overcapacity.

When demand grows rapidly, everyone wants capacity in place before their competitors. Forecasts become more aggressive, investment accelerates, and companies begin planning around growth rates that may not be sustainable. Because everyone responds to similar market signals, capacity tends to arrive in waves.

Rapid growth can absorb enormous amounts of new capacity. But it can also encourage the industry to build even faster.

© 2026 Mathew M. Hubbard. All rights reserved. • mathewhubbard.com

Exponential growth doesn't necessarily protect a capital-intensive industry from boom-bust cycles. It may be one of the forces driving them.

AI infrastructure has another complication. Operators aren't just forecasting how much AI customers will consume. They're forecasting how much computing capacity will be required to deliver it. Those are very different forecasts.

If improvements in algorithms, models, and hardware allow ten times as much AI activity using only three times the computing capacity, that's tremendous progress. But it could be a serious problem for anyone who invested assuming the relationship would be closer to one-to-one.

The Hyperscalers May Make the Cycle Worse

Microsoft, Google, Amazon, and Meta have a significant advantage over traditional infrastructure providers. They don't necessarily need an attractive standalone return on every data center. The infrastructure supports broader businesses, and these companies have the financial resources to continue investing when the direct economics might not justify it.

That could make the cycle different. But I'm not convinced it makes it less severe. It could delay the correction and encourage even more capacity to be built before the economics force a change. For independent operators whose business depends on selling computing capacity, the consequences could be particularly painful.

AI Can Succeed While Infrastructure Investors Lose Money

I don't question that AI will become an increasingly important part of the global economy. What I question is the assumption that extraordinary growth guarantees attractive returns for everyone financing the infrastructure.

The semiconductor industry has already demonstrated that those are two entirely different things. An industry can grow dramatically for decades while repeatedly experiencing overinvestment, falling prices, poor utilization, and destroyed capital.

AI data centers share many of those economic characteristics, potentially amplified by faster equipment obsolescence and unprecedented investment levels.

If anything, the extraordinary growth expectations surrounding AI may be making the next capacity correction more likely, not less.

Maybe AI demand will grow fast enough to absorb all the capacity being planned. Maybe this time really is different. But after 50 years of semiconductor boom-bust cycles, I'd want a pretty convincing explanation of why.

Because exponential demand growth has never been enough to repeal the economics of overcapacity.

Companion article

AI Data Centers and Airlines: Extraordinary Growth Does Not Guarantee Extraordinary Returns — mathewhubbard.com (insert permanent URL after publication).

© 2026 Mathew M. Hubbard. All rights reserved. • mathewhubbard.com

Sources and further reading

SIA 2025 Factbook

SIA: Seven Downturns and Fab Utilization

Goldman Sachs: Tracking Trillions

Online publication and PDF edition

Read the article online: https://mathewhubbard.com/articles/ai-data-centers-semiconductor-fabs.html

PDF edition: https://mathewhubbard.com/articles/pdf/ai-data-centers-semiconductor-fabs.pdf

© 2026 Mathew M. Hubbard. All rights reserved. • mathewhubbard.com