



THE BROOKINGS PODCAST ON ECONOMIC ACTIVITY

“The \$10 trillion question: Financing the AI buildout”

**Washington, D.C.
Thursday, October 8, 2026**

Participants:

Aaron Klein
Miriam K. Carliner Chair, Economic Studies
Senior Fellow, Center on Regulation and Markets
The Brookings Institution

Stijn Van Nieuwerburgh
Earle W. Kazis and Benjamin Schore Professor of Real Estate, Professor of Finance
Graduate School of Business
Columbia University

Janice Eberly
James R. and Helen D. Russell Professor of Finance
Kellogg School of Management, Northwestern University
Co-editor, *Brookings Papers on Economic Activity*

Jón Steinsson
Marek Professor of Public Policy and Economics
University of California, Berkeley
Co-editor, *Brookings Papers on Economic Activity*

Episode Summary:

AI is often described as a software story, but the boom is also a massive physical investment in data centers, power systems, and specialized chips. A new BPEA paper projects that this buildout will total \$10.3 trillion from 2025 to 2032, an average of 3.63% of U.S. GDP per year and larger relative to the economy than past booms in railroads, highways, and telecommunications. On this episode of the Brookings Podcast on Economic Activity, Aaron Klein talks to the paper's author, Stijn Van Nieuwerburgh about what this means for transparency and risk.

[music]

EBERLY: Welcome to season nine of the Brookings Podcast on Economic Activity. I'm Jan Eberly, the James R. and Helen D. Russell Professor of Finance at Northwestern University. And this season, we'll be listening to discussions about papers from the fall 2026 *Brookings Papers on Economic Activity* conference, hosted by Brookings in September.

STEINSSON: And I'm Jón Steinsson, Marek Professor of Public Policy and Economics at the University of California, Berkeley. We are the co-editors of the *Brookings Papers on Economic Activity*, a semi-annual conference and journal that pairs rigorous research with real-time policy analysis to address the most urgent economic challenges of the day.

EBERLY: Once again, we have a great lineup of timely topics, energy markets, inflation, migration's effects on labor and housing markets, financing AI, and more. Over the course of this season, we'll hear a lot more about the conference papers and importantly, how decision-makers should be thinking about the topics they explore.

STEINSSON: This episode is about the AI boom, but it's about an aspect of the AI boom that is not always at the center of attention, how we're actually paying for all the AI investments that are going on at the moment. Aaron Klein, senior fellow in Economic Studies at Brookings, speaks with author Stijn Van Nieuwerburgh of Columbia Business School about his new paper, "Financing the AI Build-Out."

EBERLY: The numbers here are staggering. Van Nieuwerburgh estimates the U.S. data center build-out will require something like 3.6% of GDP in annual investment through 2032, bigger relative to the economy than the railroad boom, the interstate highway system, or the build-out of the electrical grid.

STEINSSON: Amazingly, companies like Meta and Microsoft are now spending more on AI investments than they're bringing in through operating cash flow. So increasingly, they're turning to leases, joint ventures, and private credit to make up the difference. This should actually sound familiar to those that followed previous investment booms like the internet or fiber optic booms in the early 2000s.

EBERLY: Van Nieuwerburgh walks through one deal in particular, Meta's Hyperion Data Center project, to show an example of how these structures actually work and why they can make leverage that is debt harder to see from the outside. He doesn't label this a bubble, which is always tricky, but he does flag exactly what regulators and investors should be watching.

STEINSSON: Let's turn it over to Aaron and Stijn.

[2:46]

KLEIN: Thank you, Jon and Jan. It's a pleasure to be working for you again. And hello, Stijn. Thanks for joining.

VAN NIEUWERBURGH: Hi, Aaron.

KLEIN: Where are you based out of?

VAN NIEUWERBURGH: I'm based out of Columbia University in New York City.

KLEIN: Oh, well, capital of the world. And it may be that the United States is trying to become the next capital of the economic world as it relates to AI.

[3:08]

Your paper is fascinating, and I want to just dive right in. Because you estimate the AI data center build-out will require 3.6% of GDP in annual investment between now and 2032. That's a massive number. How does that compare to past infrastructure booms like canals, railroads, or the interstate highway system?

[3:28]

VAN NIEUWERBURGH: So indeed, like my central estimate is 3.6% of GDP per year from 2025 until 2032. And on that same approximate basis, the railroad boom averaged about 2.2% of GDP during the peak, about 20 years between 1870 and 1890.

Highways were about 1.1% of GDP, in the 1950s. The telecom and fiber boom, or the internet boom, if you want, was about 1.1% of GDP. Canals about 0.66, and electrification about 0.5% of GDP. So that AI estimate is about 50% larger than the largest of these prior capex booms, the railroad one, and more than three times the internet boom.

So, you know, clearly a huge number. Now, I would add two qualifications to that. The first one is that historical estimates are not always perfectly harmonized across these episodes. Some of these periods were a long time ago. Um, for example, in the railroad number, we did have some years where, uh, peak investment was over 4% of GDP. The railroad number does not include the rolling stock, the locomotives, so maybe it's a little bit underestimated.

But the point is here, it's not about sort of false precision, it's that the technological boom that we're going through is really massive and of impact to the macroeconomy.

[4:50]

KLEIN: So let me push you a little bit on that, because when I dug through the paper, I saw that you kind of had this estimate a little bit over \$4 billion for, for 100 megawatt AI training campus.

And that kind of seemed like a constant estimate throughout this time period, and I'm always reminded of Moore's law, that computing power is always getting cheaper over time. And so how did you estimate thinking, I know it's only six years down the road, but every few years computing power seems to get quite a bit cheaper?

[5:22]

VAN NIEUWERBURGH: Yeah. So first I want to clarify the assumption. So you're right, the paper uses about \$4.1 billion for 100 megawatt facility in 2025, and then it

assumes that that dollar cost rises by 4% per year. We're not assuming a constant nominal price. And just to be clear, about two-thirds of that cost of that 41 million per megawatt comes from the IT equipment, including all the fiber and all the peripherals, and one-third comes from the the real estate, the data center box itself and all the equipment, as well as the power infrastructure.

Now, I think you're right that computing historically, its efficiency is evolving very rapidly. Right? And the cost per unit of compute is typically falling. But the paper really sizes these data centers in megawatts, not in FLOPS, in floating point operations. So the way I think about this is that maybe better chips can deliver more FLOPS per watt, whereas, you know, these these frontier developers, they're gonna be training larger and larger models as these costs per FLOP are falling.

We're gonna have more and more inference where you and I are pinging our apps and our ChatGPT window and asking more and more questions. And so basically, the cost of the leading-edge technology rack can keep rising even as the performance per dollar keeps coming down. And so they're basically these two opposite effects.

The efficiency reduces the capacity that we need for a fixed workload, but the lower computing costs then induce a much higher computing demand. Right? So this is sort of, uh, the rebound effect, or people sometimes refer to this as Jevons Paradox, where when something becomes cheaper, sometimes we want to do more of it.

You know, the way the—a data center developer put this to me once is, you know, he said the the use case line is 12 miles long. And so there is a lot of things we haven't computed yet that once the cost of compute comes down, we'll do more of it.

[7:14]

KLEIN: Yeah. I I come from the banking world, and I always remember this anecdote that a banker told me when they set up online banking so that you could check your bank balance, they assumed they'd save all this money from reduced call centers because before people were picking up the phone and dialing to find out how much was in their bank account.

And it turned out it was actually more expensive to find it online because even though it was cheaper every time somebody looked, people were looking so much more often because it was so much easier for them to look as opposed to picking up the phone and calling and waiting online. So that, that type of function is is pretty wise to think, and then you kind of just assume the two balance each other out.

Let me push you on another assumption, which is this 4% nominal growth rate, which is in there for GDP and and you mentioned in your last question. Is that a belief of 2% inflation and 2% real growth, or is there something else? Say, inflation's at 3%, which is what we've kind of had for the last few years, and we only have 1% real growth.

And kind of correlated to that, nobody ever really predicts a recession in their models, but you go seven years down the road, the odds of having zero recessions

over longer time periods falls to eventually zero. How is your model sensitive to inflation, real economic growth, and the probability of a recession?

[8:34]

VAN NIEUWERBURGH: Okay, good question. So that 4% path, just to be transparent, it's not a finely-tuned macro forecast. It's just a convenient normalization. It's historically what nominal GDP growth has been, about 2% inflation, about 2% real economic growth.

Now, if we want to assume 3% inflation or 3% nominal growth, we can grow the GDP from 2025 to 2032 at a slightly higher rate. What that will do mechanically is it basically will increase my denominator in my investment-to-GDP ratio, and it will lower my headline number of 3.6% of GDP. So for example, if you assume 5% nominal growth, maybe because you want to bake in 3% inflation instead of 2%, then that number becomes 3.5% instead of 3.6% of GDP. So it's not overly sensitive.

And there's of course a question if we think 3% inflation affects the entire economy, shouldn't it also affect the cost of the AI build-out? So again, if you apply the same cost escalation in the numerator as you do to GDP in the denominator, the ratio will not change. And that's that's why I picked that 4% number, because remember, I was increasing the cost of AI by 4%. I also wanted to increase GDP by 4% so that mechanical inflation effects wouldn't distort this ratio.

Now, when it comes to recessions, you know, it's a good question. You could imagine recession would reduce AI demand. It could tighten financing constraints, it could cancel projects. So it could shrink that numerator potentially, you know, substantially.

So again, I'm not trying to forecast the business cycle three to eight years from now. I do think there is a possibility, like you said, of a recession. I think that would lead to both lower GDP and lower AI. I would imagine that the numerator would fall more than the denominator in such a scenario.

[10:30]

KLEIN: Yeah. Historically, after these big economic expansions in investment, there's usually a bit of overinvestment and some sort of recession. The dot-com boom and bust in 2001 is the most recent of the list that you've had. But you go back, you can find the railroads and canals had a little bit of a boom-bust cycle as well.

Let's pivot for a second from the macro, because it was fascinating to hear that your underlying estimate isn't that sensitive to a 1% change in nominal GDP. That's, that's a big change in the economy or the inflation rate. And your model's still saying we're getting three times the boom here in AI that we're gonna get from highways, the interstate boom, which was truly transformative boom.

[11:14]

VAN NIEUWERBURGH: Yeah. I think one more point about that is that I think there's people out there that believe that we're about to have a growth miracle, and

so maybe GDP growth will not be growing at, you know, 2% real, but at 4% real, so 6% nominal growth instead of 4 or even more. You can get my headline number of 3.6% investment-to-GDP all the way down to 3% if you assume very, very strong growth.

But it's gonna be hard to get it below that. So again, it's not that sensitive, even if you're willing to assume that AI will drive major GDP growth in the U.S. economy for the next eight years.

[11:47]

KLEIN: Let's get down to a specific deal, because I thought your paper did a great job of simultaneously discussing the big picture and then walking us through a transaction.

And you walked through a very specific deal, which was Meta's Hyperion Data Centers as a case study. And what does the deal reveal about how these projects are actually financed?

[12:06]

VAN NIEUWERBURGH: Yeah. So the reason I chose Hyperion is not only because it's a massive deal, it's also, I think, a template for how a lot of these things are financed. And it really reveals how you can have a credit-worthy technology company like Meta obtain capacity, new compute capacity, without reporting the full project debt on its own balance sheet. Right?

So the story is interesting. Meta initially was gonna build this facility on balance sheet, fully owned, the way hyperscalers have traditionally built their own data centers. But in the last few years, what has been happening is a-- I call it as the great risk shift away from build towards rent. Right? Where now you have third-party landlords, infrastructure funds, building, and and owning these facilities. So Meta actually sold a 20% stake in that Hyperion data center to this private credit firm called Blue Owl for \$2.5 billion, and it retained 20% ownership in the equity.

And then the two of them combined formed this joint venture called Beignet. It's a special purpose vehicle; the only asset it owns is that data center in Richland Parish, Louisiana, called Hyperion. It's a huge data center. It's about two gigawatts in the initial phase, which is the phase that I'm focused on, with potential expansion all the way to five gigawatt.

And this Beignet entity then decided to issue a large piece of debt. It issued \$27 billion of debt against this Hyperion asset valued at roughly \$30 billion. Right? So think about that. That's 90% debt-to-asset ratio at the project level and, you know, with debt service coverage of around, you know, 1.15, let's say.

So Meta, you know, itself in the process remained very lightly levered because it only had that 20% joint venture stake. None of that Beignet debt that was issued had to be reported on Meta's own balance sheet. And and so Meta could continue to be a very low leverage company, all the while taking on a massive amount of 90%

leverage on this data center Now, nobody would build this data center without a long-term commitment from Meta that it would rent that facility for the long term.

[14:18]

So sometimes you would find hyperscaler tenants signing 10-year leases, 15-year leases, 20-year leases. Meta did something a little different, which I found fascinating, which is it signed five 4-year leases. So you would say, "Well, five 4-year leases, isn't that also a 20-year lease?"

Not quite, because every four years Meta has an opportunity to walk away from each and every one of these nine buildings that are on that campus, and it can decide not to renew these leases. So Meta has a lot of flexibility in this contract. It can decide if, for whatever reason, 4 or 8 or 12 years from now it doesn't need as much capacity anymore, it can just walk away.

Now wouldn't the bondholders then be very worried that they're gonna be left in the in the cold? Well, Meta also had to promise to sign what's called a residual value guarantee. And this guarantee basically says if Meta decides not to renew a lease, let's say eight years from now, it has to essentially pay a fine or a lease non-renewal fee if you like, which is this residual value guarantee. It's an amount of dollars that would allow the debt on that particular building to be paid back.

And so if you think about it from the perspective of the bondholders who are financing this \$27 billion debt, they are, you know, essentially, , certainly the people selling these bonds would like you to think of them as, you know, a very safe investment because either Meta is gonna pay the lease and then that rental payment will be more than enough to pay the interest and the principal of the debt, or Meta will not renew the lease and it'll pay a chunk of cash that would allow the debt to be repaid as well.

And so as long as Meta remains solvent, it doesn't go into bankruptcy, this debt will be repaid.

[16:02]

KLEIN: I come from a world of financial regulation. I'm a veteran of the subprime mortgage disaster and the great financial crisis. And one part that's often lost, the narrative focuses on Bear Stearns and AIG and Fannie Mae, and the big companies, Lehman Brothers, that went bankrupt. But the very first canary in the coal mine were these things called structured investment vehicles, or SIVs. And essentially they allowed the very large banks to create massive exposure to subprime mortgages through these off-balance sheet investments.

The Federal Reserve and other bank regulators blessed this accounting system as protecting the banks because the risk was off-balance sheet. That's what you said. Meta's creating all this off-balance sheet risk, so it doesn't look like it has to report these four-year option leases or these residual value payments, and Meta looks like it's relatively under-leveraged.

And then the Project Deal looks like it's not 90% debt levered because there's Meta standing in the shadows behind it. What we found out in the banking crisis was the bank's off-balance sheet vehicles got taken on-balance sheet because they had a little bit of reputational risk here. They couldn't just stiff their various investors that they'd promised all these things to, and that was the very first sign there were problems.

Then-Treasury Secretary Paulson had a government proposal for a super SIV to save the SIVs, and it kind of all went nowhere. But it helped explain why one minute the regulators said the banks had limited exposure to subprime, and the next minute, Oh, the subprime system is no longer contained. It's threatening the global financial system and we need to bail out the banks or else the world's gonna implode.

You can see a world in which if all the AI firms went down, Google, Meta, Amazon, I mean, these companies in are in a very different way too big to fail, and now they're hiding their leverage. Do you see parallels here between how Meta and the other hyperscalers are using this off-balance sheet accounting and the way the banks did it in the 2000s, financing another boom of mortgages?

[18:13]

VAN NIEUWERBURGH: Yeah, I see a real parallel here. And and just like you, I was a student of the subprime mortgage crisis, and for years after that crisis was teaching a class on securitization. And just like you said, the idea of securitization is to disperse the risks away from the levered banking system onto, ideally, unlevered investors, pension funds, sovereign wealth funds, mutual funds.

And the promise of securitization in the case of the subprime mortgage crisis was never fully realized because, you know, the banks were holding on to often, you know, first loss tranches of these securitizations or through guarantees, took that exposure back on their balance sheet, maybe for reputational reasons, as you pointed out.

The parallel here is that legal separation, like Meta is accomplishing with this Hyperion example we just talked about, can understate through economic recourse. Right? And so maybe there is a state of the world where the hyperscalers could decide that they need to support a strategically important Hyperion project rather than enforcing every legal option and their sort of bankruptcy remoteness that they achieve through this structure.

So, you know, Hyperion, it's a specific operating asset. It has an investment-grade tenant. It has long-dated project debt. So I think that is a key distinction with the subprime mortgage crisis. Right?

[19:28]

I think where in the subprime mortgage crisis we got into trouble was that we had a lot of maturity mismatch. We had a lot of these SIVs invest in long-term mortgages financed with short-term debt. In particular, I'm thinking of the asset-backed commercial paper, a huge market at the time, hundreds of billions all the way up to over a trillion dollars back in 2006 and '07 and '08.

And so the big difference here, as at least for now, is that I do not see a lot of short-term liability exposure. Typically, like in my example of Hyperion, the SPV would finance a long-term asset, this data center, with long-term debt, and so you don't have that same maturity mismatch. Furthermore, Meta is, at least for now, an investment-grade credit, you know, with a public credit rating. It's currently has a very strong balance sheet, is a very profitable company.

So, you know, things would have to go really haywire for Meta to go to go bankrupt. Again, it's not impossible, but it seems like a remote possibility, at least from today's vantage point.

The only place where you see a little more short-term debt is when people start to finance GPUs. Right? So in addition to financing data centers, the real estate, people are increasingly financing the GPU. And, you know, GPUs are chips, and chips depreciate relatively quickly. Right? And so, you know, that's maybe a five-year vehicle. It's not a 20-year investment. It's maybe a four or five or six-year investment. And so that's the place where you see a bit more short-term, uh, exposure and, and of course, the technological risk that these chips may be worth less than they currently are.

So there are parallels and there are differences, I guess, is what I'm trying to say here.

[21:05]

KLEIN: Let's push on that a little bit because I remember you mentioned the asset-backed security. That was a massive bailout the Federal Reserve did to corporate America, companies like GE and and others.

But some of that mismatch in the securitization structure was that where the banks were peddling these long-term assets as being short-term and rolloverable. And on the private credit side, I saw a familiar name to me in this Hyperion deal of Blue Owl, which is one of the first private credit funds to have massive problems as investors try to withdraw.

So as you point out, the debt and the asset mismatch is there, but the private capital funders, the people who are in Blue Owl, have a right to sell their shares of Blue Owl, so to speak, their credit exposure and, and they, Blue Owl put a pause on that.

I saw that Blue Owl was the fund in your example, and they've had some problems. Is this is this deal related to Blue Owl's broader problems, or do you see anything from the Blue Owl experience that might raise this mismatch between investors in private credit and now the assets that they're holding in private credit?

[22:14]

VAN NIEUWERBURGH: Yeah, it's a good question. I mean, I would separate the manager from the project here. Right? So Blue Owl, I think, did indeed face genuine pressure in its credit platform with all these redemption requests. And typically these funds have gates, right, so they can restrict what fraction of investors can redeem in

a given quarter. That gives them, in some sense, time to resolve the underlying problems.

And and these redemptions are usually, you know, relatively minor sort of each quarter. Of course, if you have, you know, 2% of your investors who want their money back every quarter for 20 quarters in a row, that's that's still a big chunk of your investment.

You know, the other point is this is an equity investment that Blue Owl is making in Hyperion. I do not think it's the same private credit fund that faced the redemption requests that is investing in Hyperion. But I think you're making a broader point, which is, you know, yes, these are long-term investments, but if investors in these private credit funds have the right to to redeem some of their investments in the shorter run, that also creates a mismatch.

And I think, you know, structurally, it's very important to align the timing here, right, and to basically, you know, tell the investors, "Look, this is a... You're in, you're in this asset for the long run. You can't get out for the next 10 or, or 15 years." And, you know, otherwise we indeed sort of recreate this, this run risk through through the back door.

At the same time, I think I want to point out that as long as Meta doesn't go bankrupt, this specific asset, this Hyperion asset, will generate enough cash flows to pay back the debt. And so that's, like, a relatively, I would say, a relatively safe, well-funded investment-grade credit tenant that stands behind these these investments.

[23:50]

KLEIN: Right. They're, they're just not so liquid., You made an accounting point, and my Senator Paul Sarbanes was my first boss and and I served as his chief economist in Sarbanes-Oxley. The accounting reform legislation is near and dear to my heart. Those may recall that Enron and WorldCom and other major companies faced an accounting crisis where we found that the accountants weren't playing it straight. So a regulatory agency, the Public Company Accounting Oversight Board, PCAOB, was set up to oversee accounting for the Big Four accounting firms.

Do you think the PCAOB is aware of some of these accounting gimmicks? Is there something they can do here? And do you see a conflict of interest with the accounting firms that are supposed to be doing the accounting for these projects, which are separate, as you say, the Beignet that may have gotten away from Meta's balance sheet. But if I'm an accounting firm, Meta's a client, you know, I'm I'm gonna want to keep Meta happy here. And now I'm doing a separate accounting for this separate entity known as as Beignet.

Do you see any concerns in the, in how the accounting is being done for these? Walk through that gimmick a little bit that you mentioned and and whether the accounting regulator you think has flagged this at all.

[25:04]

VAN NIEUWERBURGH: Yeah. So let me first explain what I call the probability vacuum. So remember this Hyperion deal, Meta signed five consecutive 4-year leases. Only the first lease, the current lease, needs to be reported on its balance sheet as a lease. The next 16 years' worth of leases, right, the next four 4-year leases are future leases. And under GAAP accounting rules, if a liability is not at least 70% likely, it need not be reported on balance sheet.

So there is essentially very large future lease commitments that Meta has in this deal, which are not currently reflected on its balance sheet. They will be reflected if and when Meta embarks on that next lease.

Now remember, at the same time, if Meta does not renew the lease, it owes this residual value guarantee to the structure. That too is not at least 70% likely to materialize, and so that too does not need to be reported on Meta's balance sheet.

And so Meta essentially has for sure one of two obligations. Four years from now and eight years and 12 years from now, either it's gonna renew the lease or it's gonna pay the residual value guarantee. One of these two things is guaranteed to happen, but neither of them is reflected on its balance sheet because neither of them is 70% likely.

[26:23]

So there's some probability math that is literally not adding up here. Right? And again, there's nothing wrong per se, it's just these are what the GAAP accounting rules are prescribing. Meta is following these rules. There is no fraud. There is no allegation of impropriety. It's it's FASB in this case that decides these recognition and consolidation rules.

So FASB could change those rules. Maybe it should. The SEC could require public company disclosure and and sort of enforce reporting obligations about this so that at least investors are clear about what future liabilities may be coming their way. The companies themselves are reporting some of these future obligations and footnotes in their 10-Ks and 10-Qs, so you know, the, the diligent investor can find that evidence. But maybe we should make all of this a little bit more transparent. Right?

In addition, you make a, you know, an important point, which is this client-paid model that remains, I think, a structural concern. You know, Sarbanes-Oxley, I think, improved that independence by prohibiting some non-audit services and some requiring some audit committee pre-approvals and oversight and rotating some key partners and empowering the PCAOB.

And I think those safeguards sort of reduce the conflict of interest, but they did not eliminate sort of the commercial fact that the audit firm is paid by the issuer. And, you know, these leases and these guarantees and these consolidation decisions, they are judgment-heavy. Right? And and they deserve, I think, close audit committee and PCAOB attention.

[27:54]

KLEIN: Let's talk about another potential scenario here, which is if AI demand doesn't grow as fast as, as expected. There've been plenty of booms that were said to materialize another world. The demand doesn't grow as fast as expected, or a new technology comes up and sweeps beyond it.

The canal boom that you cite didn't quite last as long as some people thought because this railroad technology came to be, and all of a sudden shipping things on canals wasn't quite the deal it was thanks to the locomotive.

I try to think of risk a little bit like the second law of thermodynamics. Right? It can't be destroyed, it can simply be moved around. Who bears this risk if there's something materially different and this boom build-out happens and it doesn't click? Does it come back to the banking system? Is it stuck in the tech world of Google and Meta? Is it Blue Owl and the private capital investors? Where, as you dove into the nuance of these sector, did you find the risk actually residing?

[28:58]

VAN NIEUWERBURGH: Yeah. This is the key question, and I think we need precise answer to this question, and that is precisely what I'm working on right now. That's what I spent most of my summer doing, is to try to map what are all these financial liabilities that these companies have vis-à-vis one another, and who is ultimately bearing the risk associated with this AI build-out?

And it's a difficult question, in part because of the opacity of these systems that we were describing before. And we don't really have a reporting system that is set up to answer that question, and maybe we should build such a system.

Now, if we think about this migration from on to off balance sheet financing for data centers, that does suggest that a bunch of the risk has shifted from these tech firms towards the rest of the world, all the other investors. Right? And so the first loss usually falls on, on the equity investors in these projects. So if you're a landlord and your tenant does not renew your data center lease and nobody else needs that GPU facility anymore because maybe now there is quantum computing or there is edge computing that where everybody's computing on their phone instead of in your data center, then that loss will fall onto you.

You know, similarly, if the construction is delayed and you issue debt, like you're Hyperion, you've already issued the debt, but now, the community doesn't want to give you permits to actually keep building. You know, you have this mismatch of when the debt is due and when your revenues will be coming in. Again, that's-- the first loss of that falls on the equity investor, and if the loss is large enough, it starts to fall on the debt.

[30:31]

And remember, if you have 90% leverage, then your cash flows only have to fall by 10% before the debt is also impaired. Right? 10% cash flow drop wipes out the equity. More than that starts to impair the debt at 90% leverage.

So the hyperscalers, even in this off-balance sheet financing, for now they are still bearing a lot of risk, and the reason they're still bearing a lot of risk is because they're typically owning the IT. Right? So when we're saying they're renting their data center, what we mean by that is they're renting the box and the power that comes with the box, but they typically still own the the servers themselves. That's typically how it's done.

Increasingly, as we discussed earlier, the the servers are also being rented. Right? And somebody else owns them. But for now, most of the IT is still being financed by the hyperscalers. There's a reason why the hyperscalers are having \$800 billion of capex in 2026. It's because a lot of that is this IT that they still own themselves.

[31:28]

Now to your question of the banks, I think that is really a key question. Right? So by some tally that I've seen, about \$1.4 trillion of debt has been issued to date on these structures. Of that \$1.4 trillion, about half, \$700 to \$800 billion, is with the banks. The banking industry has a leverage ratio of 90%, roughly, whereas the private credit industry as a whole has a leverage ratio of about 50%.

So it's clearly less risky, less levered than the traditional banking system, but it does not have zero leverage. It's not like mutual funds or pension funds. It does have leverage, and there are connections between the private credit, the, the non-bank system, and the banking, the traditional banking system as well, lines of credit and so forth.

So this is really complex, and it's really the key question because, ultimately, if the music stops playing, the question is, are we gonna get an internet-style drawdown where the equity markets lose a lot of money, but it doesn't spill over to the banking system? Which is, you know, a mild recession like we had in 2001, and we have a bunch of rich people losing a bunch of money, and that's fine, we can deal with that.

Or are we gonna have a 2008, 2009-style financial crisis where the banking system is impaired? And now that spills over to every other sector in the economy that needs credit. Right? And and that ultimately is the key question, and I don't think we currently have the measurement framework to answer that question. I am trying to build it. We should all try to build it. The Federal Reserve and the and and the government and the Treasury Department should all worry about this question.

[33:06]

KLEIN: So we tried to set up a system to do that called the Office of Financial Research, which was supposed to create the data underpinning infrastructure to be able to track all of these assets so we would know the answer to that question. Unfortunately, OFR has not really lived up to what the folks that created it thought it would do. Europe has done some more work on this with something called GLEIF, which allows people to track these assets to figure out who owns and and and operates them, so that we could have access to this data so that we wouldn't be, quote-unquote, "surprised" as to where the risk migrated.

One thing that's different about this AI boom than the past is that the past booms that you refer to have a public infrastructure element. The Eisenhower Interstate System had a lot of federal investment. The canal system and Erie Canal had a lot of state support. There was a public usage and there was a taxpayer element to the infrastructure.

That seems lacking in the AI boom. Is the taxpayer putting up any of this capital, if not through the front door, then through the back? And if there is not much role for taxpayer in financing this, then what role should the government have in regulating this, since in the past, right, the government's rule was, "Well, we set the gas tax, we collect it, we get to set the rules of the road for how you drive on the interstate," whether it's, you know, safety rules or rules about how close exits can be. There's a lot of federal rules there.

What's the system here, and state and local, I might add? What's the government's exposure financially in funding this, and what's then their power level in regulating it?

[34:44]

VAN NIEUWERBURGH: So most of the direct data center chip and project financing indeed is private capital. I think that is an important difference, certainly from the interstate highway system part of electrification.

But the public sector is not absent here. The support shows up through tax expenditures like depreciation rules, and state and local sales tax, and property tax incentives. Right? So there's a lot of subsidy at the local level that data center developers have received, whereby states are competing to attract these projects, and they give especially exemption from use taxes, like on electricity or on the purchase of equipment.

And these can amount to billions of dollars, and they are meaningful, and research has shown that that sometimes makes the difference between whether that state attracts that investment or a different state attracts that investment.

And then property taxes are another, you know, a really meaningful thing. I think data center developers will point out that they are often one of the largest sources of contribution to the property tax base, especially in some of these more rural counties where some of these really large data center campuses are going up. That can be very meaningful support to and, and basically reduce the need to tax, you know, households or other businesses.

[35:59]

Now, I think that this debate has sort of reached a tipping point. The resistance to data centers has grown tremendously in the last nine months, and states are beginning to remove some of these subsidies that they gave. So that's sort of an interesting development.

So what role should the government play in all of this? Well, I've already tried to emphasize measurement and disclosure. Right? I think a lot more can be done in

that realm. I think the other point that hopefully has come through is this idea of prudential monitoring, right, where especially the levered system, the banks, the insurance companies, they should map who holds the debt on these data centers, who supplies leverage to these private credit funds. And so the whole idea here would be to see these correlated exposures before the stress materializes.

And then there's, of course, sort of physical infrastructure regulation for the electricity grid.

You know, the last point I will make is that, you know, we have direct exposure, through productivity, all of us, every taxpayer. Right? So if AI is indeed a strong productivity boom, then that will be good news for the U.S. economy. That will be good news for GDP growth. That will be good news for our fiscal situation. Right? Through progressive taxation, we'll be generating more tax revenue.

And, and so, you know, all of us, the taxpayer, we have AI exposure and, you know, that's, you know, we should be probably in the aggregate, we should probably be rooting for the success of this technology because ultimately it's gonna, it's gonna increase the the the size of the pie.

KLEIN: So Sten, thank you. This has been a fascinating conversation. We're just about at the end. One final question. So one thing you would want policymakers-- I'm in Washington, you're in New York, you're in the finance capital of the world, I'm in the political capital of the world --what's the one thing you'd want policymakers listening to this podcast in Washington to take away from this paper?

[37:49]

VAN NIEUWERBURGH: Yeah. I want them to take away that AI is not just about software and about ChatGPT, it's about a massive physical capital boom that is taking place, and it could be on the order of \$10 trillion over the next eight years. It's really macroeconomically consequential.

And the financing of that massive investment is changing. Right? The hyperscalers began this cycle with a ton of cash and very low leverage, but the spending, the capital spending, has been rising, and more and more of these projects are now being placed in joint ventures and special purpose vehicles and financed through leases and private credit and securitizations.

And these structures can be very useful economically. They bring more capital to the table, but they can also make the leverage harder to see. And a public company may report little debt because the debt is all hiding off balance sheet.

[38:39]

And so my policy recommendation is straightforward. You know, we need comparable disclosure on the full economic capital stack. Who owns each major asset? Who owes the debt? Who signed the lease? Who guaranteed the construction or the residual value? Who bears the loss if ultimately the demand disappoints or the revenues disappoint. Right? And so I think regulators should map

which banks, which insurers, which pension funds, which bond funds, which private credit vehicles hold that exposure.

And the goal here is not to stop the investment. The goal is to have transparency that will help with the pricing. It will help the regulators to see the concentration. And I think it reduces the chance that everyone discovers the same hidden obligations somewhere in a downturn.

And I think the moment to build these systems is now while we're still shaping the contracts that people are signing and while this whole architecture is still taking shape.

KLEIN: Well, thank you so much. This is a fascinating dive into the world behind the AI boom that's the world that may change the economy as we know it. Stijn, thank you very much. This has been a wonderful conversation.

VAN NIEUWERBURGH: Thank you, Aaron.

KLEIN: Back to you, Jan and Jon.

[music]

STEINSSON: Once again, I'm Jon Steinsson.

EBERLY: And I'm Jan Eberly.

STEINSSON: And this has been the *Brookings Podcast on Economic Activity*. Thanks to our guests for this great conversation.

EBERLY: As always, you can find the full papers and the Brookings Papers on Economic Activity conference proceedings on our website at Brookings dot edu slash BPEA.

STEINSSON: And thank you to the team that makes this podcast possible: Ike Blake, supervising producer; Fred Dews, senior producer; Ranga Krishnamurthy, co-producer; Teddy Wansink, video editor; Chris Chang, audio editor. Show art was designed by Katie Merris, and promotional support comes from our colleagues in Brookings Communications.