

Bigger than the railroads, on a spreadsheet

The chart is easy to read and easy to misuse. Six circles, one for each American infrastructure boom, sized by average annual spending as a share of GDP. Canals, 0.66 percent. Railroads, 2.24. Electrification, 0.50. Highways, 1.13. Telecom and fiber, 1.10. Then a circle for AI, 2025 to 2032, at 3.63 percent, with a note that the AI figure is estimated.[1][4]

The Wall Street Journal ran that chart under a headline calling the AI build-out the biggest economic bet in U.S. history. The deck said data-center spending is greater than the canals, the railroads, and the grid combined.[3][4] Add the three averages in the underlying table and the sum is 3.40. The AI scenario sits 0.23 points above it. The arithmetic holds. The economics are the part worth arguing about.

The number comes from a conference draft by Stijn Van Nieuwerburgh of Columbia, written for the fall 2026 Brookings Papers on Economic Activity. The Journal reported it. The draft is public. What follows is an attempt to say what the 3.63 percent is, what it is not, and which part of the build-out an investor can actually watch.

A scenario, not a print from the national accounts

Van Nieuwerburgh is not reporting what the United States spent. He is pricing a pipeline.

A representative 200 megawatt AI training campus, in his estimate, costs about \$8.2 billion. Roughly a third of that is the building and the power hookup. Roughly two-thirds is the IT system: GPU racks, networking, storage, integration. Scale that unit cost and a gigawatt campus is about \$41 billion. A continuous 200 megawatt load uses about as much electricity in a year as 170,000 average American households. Racks that used to draw 5 to 10 kilowatts now draw well above 100, and the new ones want liquid cooling straight to the chip. Old cloud halls do not take this load without a rebuild.[1]

He starts from Cleanview's project list as of July 2026: about 57 gigawatts already operating in the United States, and a planned pipeline of about 509 gigawatts. The pipeline is not a forecast. Projects slip, shrink, and die. His central case assumes 182.8 gigawatts of additional capacity becomes operational from 2025 through 2032, another 117.2 gigawatts finishes after 2032, and 226.9 gigawatts is never built. He is throwing out almost half the announced pipeline before he starts counting dollars. The case is sensitive to a handful of huge multi-phase projects and to projects that have no completion date at all.[1]

He then applies \$41 million per megawatt in 2025 dollars, lets that cost rise 4 percent a year, and spreads each project's spend across the completion year and the three years before it. Nominal GDP is set at \$30.762 trillion in 2025 and is also grown at 4 percent a year, so the ratio is not being juiced by a stagnant denominator. Investment that falls inside 2025-2032, including money spent before 2032 on plants that switch on later, comes to \$10.3 trillion. Divided by cumulative GDP over those eight years, that is 3.63 percent. The average of the eight annual ratios is the same 3.63 percent. Both facts are in the draft. Neither fact means the spending arrives in a smooth line.[1]

The path in his Table 5 is lumpy. Investment over GDP is 3.7 percent in 2025, 4.5 in 2026, 4.2 in 2027, 3.2 in 2028, 2.2 in 2029, 2.5 in 2030, 3.6 in 2031, and 5.1 in 2032. Anyone who treats 3.63 as "the"

number for the rest of the decade is averaging a hump, a dip, and a second hump. The second hump is partly pre-spending on plants that are not even finished inside the window.[1]

He is explicit that this would be larger, relative to the economy, than the canal, railroad, electrification, highway, and telecom booms he lines up against it. He also says the full build-out would double the electricity use of the entire U.S. residential sector. That is a physical claim sitting under a financial chart, and it is the reason power, not model quality, is now the binding constraint on where these campuses can be built.[1]

The other number in the same newspaper

The Journal's own body text does not live entirely inside the 3.63 circle. It cites an estimate that U.S. AI investment hits 1.9 percent of GDP in 2026, and says the late-19th-century railroad boom was the last time one new industry's build-out took a larger share. The archive text I read has a hole where the estimator's name should be. It was a link, and the link did not survive extraction. A secondary write-up of the same article names Goldman Sachs. I have not read that note, and I am not going to invent what it counts. The gap with the conference draft is large enough to matter either way.[4][6]

In the draft, 2026 investment is \$1.432 trillion against an assumed GDP of \$31.992 trillion, which is 4.5 percent, not 1.9. The same draft says capital spending at Oracle, Microsoft, Amazon, Meta, and Alphabet rose from about \$97 billion in 2020 to more than \$400 billion in 2025, and is projected to exceed \$800 billion in 2026, the first year that spending would top the five firms' combined operating cash flow. Set \$800 billion against the draft's own 2026 GDP and you get about 2.5 percent. That capex includes spending that is not AI, and it excludes data centers the hyperscalers lease rather than build. It is a third object, sitting between that 1.9 percent and the pipeline scenario's 4.5.[1][4]

A fourth object is what the Commerce Department can already see. Through July 2026, the Journal reports, a seasonally adjusted \$37 billion had been spent on private data-center construction, about \$9 billion more than in the first seven months of 2025. Private construction of everything else (houses, apartments, shopping centers, the rest) was about \$46 billion below the year-earlier level over the same months. Thirty-seven billion dollars of shells is not \$1.15 trillion of campuses-plus-chips, which is the draft's 2025 scenario cell. Construction is the third of the campus that is concrete and switchgear. The chips, which are most of the money, do not show up in that construction series.[1][4]

The chart and the construction print are not in conflict. They are not the same measurement. The 3.63 percent is what the announced pipeline costs if a large share is finished and equipped on the draft's timetable. The 1.9 percent is a nearer-term estimate from a bank, reported by the Journal, of investment that looks more like what companies are guiding. The \$37 billion is buildings in the ground. Collapsing those three into "we are already spending 3.63 percent of GDP on AI" is the mistake the chart invites.

The Journal says the projection is tricky and that total spending might end up substantially lower. That sentence should sit next to the headline, not under it.[4]

There is a second version of the author's own number, and it is still live. The abstract on SSRN, dated as written on 23 September 2026 and revised on 2 October, still says a 188 gigawatt build-out finished by 2032 would cost nearly \$9 trillion, or 3.2 percent of GDP a year. The conference draft that the chart uses says 183 gigawatts operational by 2032, \$10.3 trillion in the measurement window, 3.63 percent. The unit cost is the same order (\$41 billion a gigawatt). The dollar gap is mostly the window: the \$10.3 trillion includes spending before 2032 on plants that finish after, which a "completed by 2032" total leaves out. Treat 3.63 as the chart's figure. Do not treat it as the only figure with the author's name on it. [1][5]

How the old booms were counted

The historical bars are not a single official series copied six times. Appendix C of the draft says the sources differ in scope and accounting, and that the estimates are approximate measures of scale, not harmonized national accounts. That caveat is doing more work than the chart design allows.[1]

Canals use H. Jerome Cranmer's construction-expenditure estimates. Over 1836-1841, annual investment rose from \$4.4 million to a peak of \$14.3 million in 1840, then fell to \$11.7 million in 1841. Against nominal GDP from Johnston and Williamson, the average is 0.66 percent, and the peak is 0.90 percent in 1840. Six years, and then the boom was over. A lot of the later canal mileage was a bad investment. The average does not say that.[1]

Railroads use Robert Gallman's revised construction series, in constant 1860 dollars, as reported by Paul Rhode, set against Gallman's national-product estimates. Over 1870-1890 the average is 2.24 percent. The annual share exceeded 4 percent in 1870 and again in 1881. The series is construction. It leaves out at least some spending on locomotives and rolling stock. The railroad bar is therefore a floor on the railroad boom, not a ceiling. If you want a fair fight with an AI figure that is two-thirds chips, the railroad number is the one that was trimmed.[1]

Electrification uses Melville Ulmer's gross capital spending for electric light and power, private and municipal plant, from 1905 to 1925. It excludes land, and it excludes the electrical machinery bought by the people who used the power. Average 0.50 percent. Peak 0.96 percent in 1924. Calling this "the grid" in a newspaper deck is fair as a label and unfair as a total. The complementary machines in factories and houses were part of the economic bet and are not in the circle.[1]

Highways use the Federal Highway Administration's capital outlays by federal, state, and local governments from 1956 to 1973: land, engineering, construction, reconstruction, rehabilitation. Average 1.13 percent. Peak 1.32 percent in 1958. This one is closer to a complete public-works series than the others. It is also the boom people remember as nation-building, and it is a third the size of the AI scenario on this metric. Memory and the ratio are not the same thing.[1]

Telecom and fiber use BEA private fixed investment in communications equipment and communications structures, 1996-2003. That catches fiber and the switches and routers bought across the economy. Average 1.10 percent. Peak 1.39 percent in 2000, which is also the year the overbuild became obvious. The fiber is still in the ground. A large share of the companies that laid it are not. The draft cites Couper, Hejkal, and Wolman on that rise and collapse. The chart has no way to show a boom whose average was modest and whose peak still ruined the people who financed the peak.[1]

Add all five historical averages and you get 5.63, which is more than 3.63. The deck does not claim AI beats the sum of every prior boom. It claims AI beats canals, railroads, and the grid. That is a choice of comparators. It is a legal choice given the table. It is also the choice that makes the new circle look lonely at the top.

Gross dollars are not the same bet

The draft's own warning on the comparison is the most important paragraph in the paper, and it does not fit in a bubble.

Historical comparisons of gross investment need caution because the assets have different economic lives. Canals, railroads, and fiber often provided service for decades. GPUs may be economically obsolete in three to six years. Per dollar of gross investment, the AI build-out therefore creates less net

capital than many of the earlier booms. Buildings, power connections, and the rest of the non-IT plant are long-lived. They are the minority of the cost. In the draft's cost split, IT equipment is 68 percent and is given a six-year life in the central case. The other 32 percent is given twenty years.[1]

That cut runs in one direction. A second cut runs the other way, and the draft states it just as plainly. The \$10.3 trillion counts the initial build and the initial fit-out. It does not count the refresh. Capacity commissioned in 2025 and 2026 may need a substantial GPU replacement around 2030, while the building and the power connection are still usable. Refresh cycles would raise cumulative gross investment. The estimates include neither component-specific depreciation nor future replacement spending.[1]

The chart is too kind to AI if you care about capital that is still there when the debt matures, and too kind to the historical booms if you care about gross dollars that will have to be spent again. Both cuts are real. The first is the one a picture of six circles cannot show. A railroad spike is not a depreciating accelerator. Treating them as the same kind of "infrastructure spending" because both can be divided by GDP is how a spectacular ratio gets mistaken for a spectacular stock of wealth.

The fiber boom is the closest cousin, and even that cousin is kinder than this one. Fiber from the late 1990s still carries traffic. A GPU from that era is scrap. The useful life is the whole argument. If the economic life of the IT share is three years rather than six, the draft's own sensitivity says the annual revenue required to earn a 10 percent unlevered return, at a 50 percent cash-flow margin, rises from about \$3.7 trillion to about \$6.0 trillion. Same buildings. Worse chips. A 60 percent higher hurdle.[1]

What has to come back

The revenue exercise is not a forecast of AI sales. The draft says so, in those words. It asks what mature annual revenue the 2025-2032 completion cohorts would have to produce to recover the investment, replace depreciating capital, and pay investors a given unlevered return.[1]

He takes the capacity finished between 2025 and 2032, leaves out capacity that was already running before 2025, and assigns some undated projects into the late years of the window. That slice is 182.7 gigawatts. Initial cost about \$8.54 trillion. Because money goes in before revenue comes out, he grosses the capital up for the construction period. At commissioning the capital base is about \$9.61 trillion.[1]

At a 10 percent unlevered return, a six-year life on the IT share, and a twenty-year life on the rest, the annual capital charge is 19.37 percent of that base. The cohorts together need \$1.86 trillion a year of operating cash flow once they are all mature. At a 50 percent operating cash-flow margin, that is \$3.725 trillion of mature annual revenue by 2032. Against his assumed 2032 GDP of \$40.481 trillion, that is 9.2 percent of the economy, every year, just to give this capital an ordinary unlevered return. Not a windfall. A 10 percent unlevered return.[1]

The range around that sentence is wide, and he prints it. At a 5 percent required return and a 67 percent margin, required revenue is \$2.29 trillion. At a 15 percent return and a 33 percent margin, it is \$6.72 trillion. The central case is a point in a table, not a law.[1]

He then puts the central case in market units, which is the useful part. Spread over 182.7 gigawatts, \$3.725 trillion is about \$20.4 million per megawatt-year of facility power, or about \$5.50 per installed GPU-hour at full utilization, on his representative rack. He says current on-demand prices for high-end Nvidia capacity commonly run from roughly \$6 to more than \$10 per GPU-hour, with lower prices on long reservations and spot. At 80 percent utilization the required billed price is about \$6.90. At 70

percent it is about \$7.90. On today's price list, the central case is not absurd. That is his conclusion, and it is narrower than the headline. The open question is whether demand will absorb something like 180 gigawatts of new capacity at high utilization, and whether providers keep pricing power as hyperscalers, clouds, and model labs compete the rate down.[1]

The growth rate that gets you there from here is the part that should make a person sit still. The draft puts current combined annual revenue at OpenAI and Anthropic around \$100 billion. Reaching \$3.725 trillion by 2032 from that base is roughly 80 percent growth a year. Compounded over six years the arithmetic is about 83 percent. Call it 80, as he does. Eighty percent a year, for six years, from an already large base, while the product gets cheaper and more people are allowed to sell it. Possible. Not the base case a bondholder should underwrite because a chart of past canals looked smaller.[1]

One more cut. The \$3.725 trillion is the revenue required from these completion cohorts, not a claim that OpenAI and Anthropic personally have to book all of it. Hyperscalers, other clouds, and whoever rents the racks can be the collectors. The cash still has to appear somewhere in the chain, at high utilization, at a price that has not collapsed. Moving the revenue from the lab to the landlord does not shrink it.

Who is actually on the hook

The interesting risk is not "AI works" versus "AI does not." It is that the debt is being written against a stack whose useful life is shorter than the bond, inside vehicles designed so the obligation does not show up on the tenant's balance sheet until the tenant has to pay it.

The five hyperscalers can no longer fund the guided pace out of cash generation. The draft's Figure 1 is the transition: internal cash is still large, and in 2026 it is no longer enough. Once that happens, the marginal dollar comes from leases, joint ventures, project debt, private credit, and special-purpose vehicles. Those structures expand the amount of capital that can be raised. They also raise leverage on the asset even when the hyperscaler's reported leverage stays low, and they hide contingent obligations. The draft's abstract is blunt about the exposures: tenant concentration, technological obsolescence, execution delays, uncertain residual values.[1]

He reports a Morgan Stanley estimate, which I have not checked against the original note, that more than half of the roughly \$2.9 trillion needed for hyperscalers' incremental compute from 2025 to 2028 comes from outside capital, with something like a 60/40 equity-debt split across the full investment. Inside the debt, private credit is the largest slice, about \$800 billion, then corporate debt of roughly \$200 billion, then structured finance of about \$150 billion. The 60/40 split can understate leverage at the asset. Hyperscalers point a lot of their own cash at the chips. Shells, power, and related real estate get project-level debt. Reported corporate leverage and system leverage are different numbers. Watching only the first one is how this stays quiet.[1]

Meta's Hyperion financing is the worked example, and it is worth reading slowly because it is a template, not an anecdote.

The project is about 2.0 gigawatts and about \$30 billion. That is roughly \$15 million per megawatt for the campus, in the neighborhood of the draft's \$11 million per megawatt for facility plus power, and it excludes the IT gear, which Meta finances separately as tenant. Meta owned it, then sold an 80 percent equity stake to Blue Owl for about \$2.5 billion. The joint venture, Beignet, raised \$27 billion of debt in October 2025. S&P rated it A+, one notch below Meta. The draft calls it the largest individual investment-grade corporate debt issue in U.S. history.[1]

Debt against assets is about 90 percent. The debt-service coverage ratio is 1.12 times. Meta's own year-end 2025 leverage was about 25 percent on book equity and about 4 percent on market equity. Its operating cash flow of \$130 billion was at least 18 times its 2026 debt service. Same economic neighborhood. Completely different capital structure, because the debt does not sit on Meta.[1]

The leases are five terms of four years, starting in 2029 and running to the bond's 2049 maturity. At each renewal Meta can walk away from some or all of the campus. If it walks, the assets are offered for sale, and Meta covers any shortfall against a contractual minimum value. That residual-value guarantee is what lets Meta keep an exit while still giving lenders something that resembles a long lease. Under the GAAP treatment the draft describes, Meta does not recognize the future lease or the guarantee on its balance sheet before the relevant commitment becomes effective, even though one of the two will, in the end, have to be met. Reported leverage understates the economic obligation. The draft says this is not a Hyperion quirk. Moody's, as cited there, puts hyperscaler lease commitments at roughly \$970 billion, of which about \$660 billion is future leases not yet on the balance sheet.[1]

The structure is not cheaper money. Beignet's debt came at 6.58 percent, at least 100 basis points above what the draft thinks Meta would have paid on comparable unsecured debt at the time. Over the life of the financing that spread is more than \$5 billion of extra interest, and it comes back to Meta through the lease. Meta is paying up to keep an asset-light profile: less reported capital intensity, less reported leverage, a multiple that still looks like a software company's. For a firm valued partly on high growth and low capital intensity, that can be a rational purchase. It is not a sign that the project is lightly financed. It is a sign that the leverage was moved to where the multiple does not have to look at it.[1]

A footnote in the draft cites a separate Journal tally, which I have not read on its own: for Microsoft, Alphabet, Amazon, and Meta, about \$2.4 trillion of off-balance-sheet liabilities (\$904 billion of future lease commitments and \$1.52 trillion of future purchase commitments, mostly chips) against \$604 billion recognized on the balance sheet (\$356 billion of long-term debt and \$248 billion of current lease obligations). On that tally, about one-fifth of \$3 trillion in data-center-related obligations is currently recognized. If the tally is even roughly right, the balance sheet is the minor document.[1]

The circularity is not subtle. A footnote describes a \$500 billion credit facility stood up in August 2026 by Apollo, BlackRock, Blackstone, Brookfield, Goldman Sachs, and KKR to finance Nvidia chips, with Nvidia providing a 25 percent residual-value guarantee, exposure of as much as \$125 billion. The supplier of the short-lived asset is backstopping the lenders who finance the customers who buy the asset. Google and Broadcom, the draft says, entered a similar arrangement around TPU purchases. When the residual value of the chip is the thing everyone is unsure about, a guarantee from the company whose next chip destroys that residual value is not outside capital in any ordinary sense. It is the industry lending to itself and calling the lender a fund.[1]

Public market prices have already stopped treating data-center landlords as dull infrastructure. The draft's rolling betas on listed data-center REITs rise from about 0.5 against the S&P 500 in 2021-2022 to about 1.0 by 2026. Their beta to the non-AI part of the index rises too. These stocks no longer look like a hedge against an AI disappointment. They look like part of the same bet, with a REIT wrapper.[1]

Van Nieuwerburgh's own policy line is narrower than the commentary around him. It would be premature, he writes, to conclude that AI infrastructure already poses systemic risk comparable to earlier credit booms. Off-balance-sheet structures matter because they can make correlated exposures hard to see before a downturn, not because they automatically raise the total amount of risk. Strong growth, high utilization, and better models could support the cash flows. The downside he names is specific: uncertain demand, fast technical change, execution bottlenecks, and high leverage, together, if expectations get revised. Historical booms, railroads included, became financially consequential when

demand fell short, financing tightened, or the debt matured on a different clock from the technology. The question he wants asked is not "is this systemic yet." It is under what conditions project-level losses become correlated and move through firms and lenders at once.[1]

I think that is the right question, and I think the Hyperion coverage ratio is a hint about the answer. A debt-service coverage of 1.12 times does not need a depression. It needs a delay, a lower renewal rate, or a residual value that misses the contractual floor by enough to matter. The guarantee then pulls the loss back to the tenant who had arranged not to show it. The loss was never gone. It was waiting for the option date.

What the spending is already doing

The scenario can be wrong by a trillion dollars and the build-out can still be large enough to shove other parts of the economy around. That part is no longer a model. It is reporting, and it should be labeled as reporting.

The construction split is the cleanest fact. Data-center shells up about \$9 billion year over year through July 2026. Everything else in private construction down about \$46 billion. The Richmond Fed, the Journal says, has reported that data-center construction is straining labor supply in its district. Hyperscalers are tying up scarce electricians and scarce power, and other projects feel it in price and in delay.[4]

One industrial story is specific enough to keep, and too thinly sourced to treat as a national pattern. The Journal says Mississippi was in the running for an aluminum smelter and an estimated 1,000 permanent jobs. A data center was announced near one proposed site, near Vicksburg, and tied up the electricity the smelter needed. The smelter went to Oklahoma. The sourcing is "a person familiar with the operator's decision." One plant, one unnamed person. It is still the right kind of example, because the constraint was power, not a lack of demand for aluminum. A site-selection consultant who works with heavy industry, Didi Caldwell, told the Journal that data centers are pushing up land costs and crowding out manufacturing. That is a practitioner's claim, not a census. It matches the mechanism.[4]

Jobs created by the build-out are real and oddly distributed. LinkedIn, as cited by the Journal, credits AI with more than 750,000 U.S. jobs from 2023 through 2026 to date. The median salary on AI-related listings was around \$180,000, against \$80,000 for all listings. Employers added 117,000 jobs at data centers since the start of 2024, and that count leaves out construction, much of which is not permanent. In the Washington area, IBEW Local 26's political coordinator told the Journal that union electricians had risen from 9,000 to 17,500. The Journal's sketches (an apprentice at about \$30 an hour who left an entry-level IT job, a more experienced electrician at \$62 an hour plus overtime on Amazon-related work) are color. The wage gap between the listing median and the trades is the point. A lot of the "AI jobs" number is white-collar listings. A lot of the physical build is a temporary construction boom that pays well and then leaves a building with a small permanent staff.[4]

The wealth effect is the other side of the same trade. The Journal, citing the Federal Reserve, says U.S. household holdings of corporate equities and mutual funds were \$63 trillion in the second quarter, nearly double the level at the end of 2022. That has held up consumer spending while inflation-adjusted wage growth has faltered, and it has accrued mostly to households that already owned stock. In Silicon Valley the same money is in houses. A local agent told the Journal the luxury market was the best he had seen since 2000, and described a five-bedroom listing at \$9.9 million that went into contract above \$13 million in two weeks, to an AI entrepreneur. Nationwide, the Journal notes, home sales have been

in a four-year slump. The boom is not a housing boom. It is a boom in the zip codes where the equity sits.[4]

Inflation is showing up in the equipment. Import prices for computers, peripherals, and semiconductors were 20 percent higher in August than a year earlier, the Journal reports. That feeds into consumer hardware. The same article says data-center investment is pushing up wages in related fields, and quotes Fed chairman Kevin Warsh naming hyperscaler borrowing as one reason long-term rates are higher, which feeds mortgage costs. Electricity bills have risen in some places with a lot of data centers. I am taking those lines from the Journal. I have not read the speeches.[4]

Put the pieces next to each other and the distribution is not mysterious. Construction workers and chip sellers get a boom. Households with stock get a wealth effect. Households trying to buy a house, or to run a factory that needs the same substation, get the crowding out. The chart of percent-of-GDP does not have a bar for that split. The split is most of what the boom feels like if you are not the one issuing the bond.

What I would watch

I would not watch the 3.63. It will not be revised on a calendar, and arguing with it is a way to miss the prints that can falsify it.

The first print is energized gigawatts against announced gigawatts. The draft already discards 227 gigawatts of the pipeline. If cancellations and slips run hotter than that, the scenario falls toward the bank estimates, and the "bigger than the railroads" sentence becomes a description of a press release. If energized capacity tracks the central case, the revenue hurdle stops being theoretical. Cleanview-style project lists, interconnection queues, and utility filings are duller than a bubble chart and much harder to fake.[1]

The second print is the price per GPU-hour as the 2026 and 2027 cohorts come online, and the utilization under that price. The draft's own arithmetic says today's rental prices can cover the central return only if utilization stays high and the rate does not break. A CEO saying demand is "insatiable" is not the print. The print is the rate card, the reservation discount, and whether spot is clearing. If the billed price has to rise because utilization fell, the capital is not earning what the model assumed, even if the headline investment number was right.[1]

The third print is the guarantees, not the reported debt. Future leases that are not recognized until they start. Residual-value guarantees that are not recognized until a tenant walks or a chip is sold below the floor. Purchase commitments for chips that will be obsolete on a schedule the buyer does not control. The draft's policy ask is measurement and disclosure while the capital structure is still being invented: project-level debt, joint-venture equity, future leases, residual-value guarantees, tenant concentration, renewal options, assumed recovery values. That is not a call for a new regulator to bless the boom. It is a call to stop letting the interesting liability live in a footnote until the option date.[1]

A fourth print, smaller and nearer, is whether hyperscaler capex stays above operating cash flow after 2026, and whether the gap is funded with bonds a holder can read or with another Beignet. The crossover is the draft's Figure 1. If it reverses, the off-balance-sheet machinery matters less, because the cash is there. If it widens, every new campus is a negotiation about who shows the debt.[1]

I do not know whether the revenue shows up. Neither does the author of the chart. He priced a pipeline, compared gross investment to gross investment, and then spent the rest of the paper on the ways that comparison fails: asset life, hidden leverage, tenant concentration, a refresh cycle the \$10.3 trillion

leaves out. The Journal was right that nothing in the American record matches this share of GDP on his scenario. It was also right, in the smaller type, that the spending might come in lower, and that a lot of it is debt.

The railroads were a bet on a network that was still there when the bonds came due. A large share of this one is a bet on a chip that will not be. That can still be a sensible use of capital if the cash from the chip, while it lasts, pays for the building and the next chip. It is a different bet than the chart implies. It is being financed as if the difference were an accounting detail.

Sources

[1] https://www.brookings.edu/wp-content/uploads/2026/09/4c_Van-Nieuwerburgh.pdf — Van Nieuwerburgh, Financing the AI Buildout, BPEA conference draft, September 2026 [3] <https://www.wsj.com/economy/the-ai-build-out-is-becoming-the-biggest-economic-bet-in-u-s-history-c60716dd> — Putzier and Lahart, WSJ, 23 September 2026 [4] <https://archive.ph/gSScH> — archive.ph snapshot of the WSJ article [5] https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6444363 — Van Nieuwerburgh, SSRN abstract 6444363, revised 2 October 2026 [6] <https://www.chaincatcher.com/en/article/2291984> — ChainCatcher summary of the WSJ article, 24 September 2026