AI Bubble Burst

Geomatrix (H.K.) Limited · AKAMAI Master Fund · Special report, October 2026

Summary and Recommendation

- **Thesis:** AI's obligations are fixed and front-loaded; its revenue is variable and back-loaded. Like the 2004-6 teaser mortgages, the frontier labs' compute commitments commence later, in a reset wall of $412bn in 2027 and $440bn in 2028, matching the 2008 mortgage reset peak of $410bn.
- **Trigger and timing:** cash-flow worries always trigger the pop, and we believe the market repricing has started, here in mid-October 2026. The solvency stress comes later, as those commitments commence in 2027-28.
- **What we own:** the users of AI (insurers and banks on single-digit to low-teens P/Es), biotech, and non-US markets the bubble has not stretched, led by Japan, our biggest long.

Fancy a conversation, info on AKAMAI? Contact Robert Howe at howe@geomatrix-investment.com.

Steve Connell: I would like to see you put out a report that has everything you know about AI both good and bad, then rearrange the data into an argument. Putting forward a thesis that the bubble is bursting causes you to exaggerate negative facts and not emphasize the positive developments. I far prefer data to theses. There are always opportunities in the stock market.

Market experts and now the public compare the 2025-26 AI bubble to the late-1990s internet and telecom bubble. They know it will pop; the only questions are when and what will be the fallout. Cash-flow worries always trigger the market pop, well ahead of the financing wall: the Minsky Moment[1], "emperor has no clothes," the coyote looking down from off the cliff. Stock prices crack, and momentum traders, now including AI algo agents, sell, sell, sell. We believe that time is NOW, here in mid-October 2026. Like 2000, AI/tech will bounce, not a straight line down, but the very fall in stock prices creates what George Soros calls "reflexivity,"[2] feedback loops from market caps to fund-raising, capex, revenues.

Steve Connell: Negative cash flow is a narrative not supported by data. Mag 7 operating cash flow is $1.1 trillion.

The stock charts are almost identical, and valuations are near the March 2000 highs (varies by metric), when the internet bubble popped. Every bubble goes hyperbolic at the end, as FOMO (Fear of Missing Out) pulls in the doubters alongside retail momentum players and now AI algo agents (see our previous study, "AI buys AI"). Institutions fire under-performing value managers and switch to growth.

[1] **Minsky moment.** *Definition:* the point when a long, debt-fuelled boom tips into collapse because borrowing can no longer be repaid from cash flow, forcing asset sales. Named for economist Hyman Minsky's financial instability hypothesis (hedge, speculative and Ponzi finance); the phrase was coined by PIMCO's Paul McCulley in 1998, during the Russian debt crisis. *Sources:* Hyman P. Minsky, 'The Financial Instability Hypothesis', Levy Economics Institute Working Paper No. 74, May 1992, https://www.levyinstitute.org/pubs/wp74.pdf ; Richard W. Fisher, 'Minsky Moments and Financial Regulatory Reform', Federal Reserve Bank of Dallas, 14 Apr 2010, https://www.dallasfed.org/news/speeches/fisher/2010/fs100414 ; Alessandro Vercelli, 'A Perspective on Minsky Moments', Levy Economics Institute Working Paper No. 579, 2009, https://www.levyinstitute.org/pubs/wp_579.pdf .

[2] **Reflexivity.** *Definition:* George Soros's theory that market participants' biased views move prices, and those prices in turn change the fundamentals they are meant to reflect, creating self-reinforcing boom-bust cycles. *Sources:* George Soros, The Alchemy of Finance: Reading the Mind of the Market (New York: Simon and Schuster, 1987), https://openlibrary.org/books/OL2376425M/The_alchemy_of_finance ; George Soros, 'Fallibility, Reflexivity, and the Human Uncertainty Principle', Journal of Economic Methodology 20(4), 2013, pp. 309-329, https://doi.org/10.1080/1350178X.2013.859415 .

CHART 1

S&P 500 Shiller CAPE: Dot-com (1995-2002) vs AI (2021-2026)

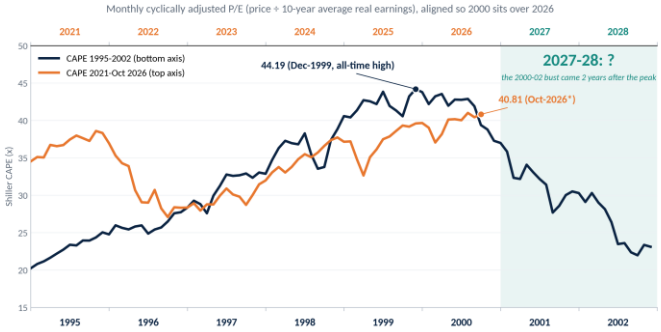

The AI-era CAPE of 40.81 (Oct 2026) is within sight of the dot-com high of 44.19 (Dec 1999). Source: multpl.com; Robert Shiller; Geomatrix (H.K.) Limited.[3,4,5]

Steve Connell: CAPE and market cap vs GDP are discredited ratios. They are paraded by every bear who can't find anything more convincing.

Steve Connell: Cyclically adjusted PE only works for cyclical companies. Applied to growth companies, it gets progressively worse every year. Do you really want to claim that Google META AAPL and AMZN are cyclical companies?

[3] **Shiller CAPE.** *Definition:* CAPE = Shiller cyclically adjusted P/E (real S&P 500 price divided by the 10-year average of real reported earnings), monthly. *Note:* bottom axis 1995-2002 (navy), top axis 2021-2028 (orange); slot 6 aligns 2000 with 2026, and the 2021-26 line ends at Oct 2026, leaving 2027-28 empty. The shaded 2027-28 panel marks that the 2000-02 bust came 2 years after the peak.

[4] **CAPE values.** Dot-com all-time high 44.19 (Dec 1999); Jan 2000 43.77. Latest 40.81 (Oct 2026), a partial month (month to date, multpl table as of 16 Oct 2026).

[5] **Sources.** multpl.com, Shiller PE Ratio by Month, https://www.multpl.com/shiller-pe/table/by-month ; cross-checked against Robert Shiller, ie_data.xls, https://shillerdata.com/ . Data: 'Teaser Wall Chart 4 - SP500 PE overlay data.xlsx'. Geomatrix (H.K.) Limited.

CHART 2

S&P 500 Price-to-Book: Dot-com (1995-2002) vs AI (2021-2026)

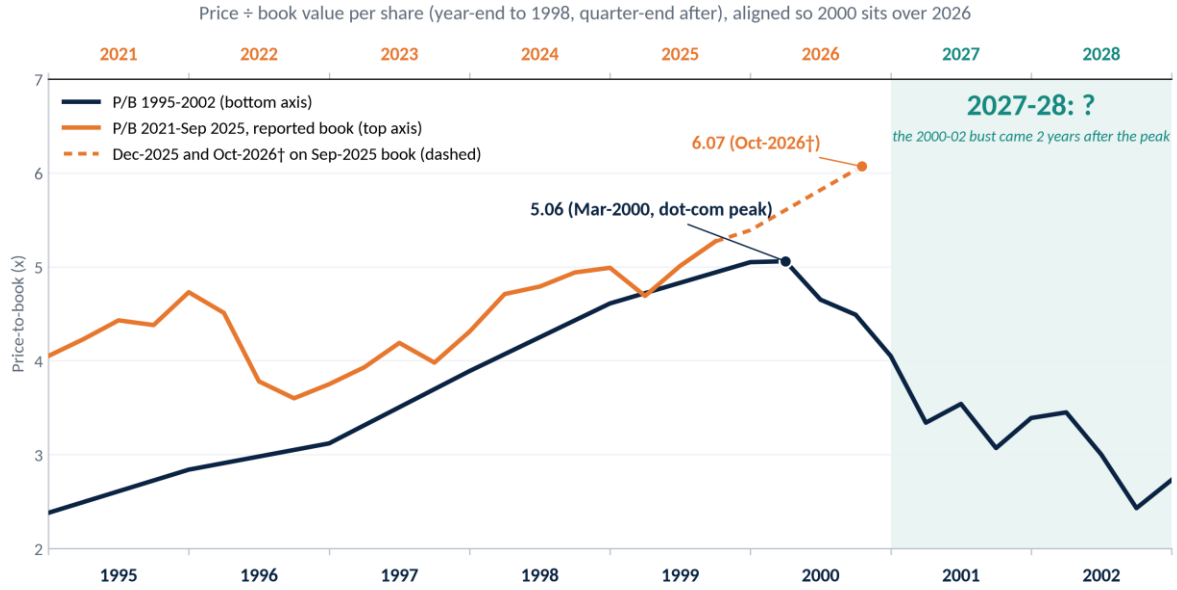

We are 20% above the dot-com peak on price-to-book: 6.07x now vs 5.06x in March 2000. Source: multpl.com (Standard & Poor's); S&P Dow Jones Indices; Geomatrix (H.K.) Limited.[6,7,8,9,10]

[6] **Definition.** *Definition:* price-to-book (P/B) = S&P 500 index level ÷ S&P 500 book value (shareholders' equity) per share, as published by S&P Dow Jones Indices. *Window:* bottom axis 1995-2002 (navy), top axis 2021-2028 (orange); slot 6 aligns 2000 with 2026 as in Chart 1, and 2027-28 is left empty. Each line starts at the prior year-end (Dec 1994, Dec 2020).

[7] **Data and frequency.** *Data:* 1994-1998 are S&P year-end P/B ratios only (no public quarterly S&P 500 book value before Dec 1999), so the navy line is annual to Dec 1998 and quarterly from Dec 1999. 2021-2025 are quarter-end values on reported book through Sep 2025. *Note:* †Dec 2025 (5.39) and Oct 16, 2026 (6.07) are multpl estimates: market price over the latest S&P-reported book value per share (1,270.32, Sep 30, 2025). No 2026 book value is published yet, so Mar, Jun and Sep 2026 quarter-ends are not shown; the dashed segment joins the published points.

[8] **Peak and latest.** *Note:* dot-com peak 5.06x at Mar 31, 2000 (Dec 31, 1999: 5.05x). Latest 6.07x at Oct 16, 2026†, 20% above the dot-com peak. Latest on reported book 5.27x at Sep 30, 2025, 4% above. *Cross-check:* the PCRA year-end S&P P/B for 1999-2002 (5.05, 4.05, 3.39, 2.73) matches multpl exactly.

[9] **Sources.** *Sources:* multpl.com, S&P 500 Price to Book Value by Quarter, https://www.multpl.com/s-p-500-price-to-book/table/by-quarter (source: Standard & Poor's); multpl.com, S&P 500 Book Value Per Share by Quarter, https://www.multpl.com/s-p-500-book-value/table/by-quarter; S&P Dow Jones Indices year-end data 1967-2007 (SP500from1967to2007) in the PCRA R package, https://cran.r-project.org/package=PCRA ; S&P source file https://www.spglobal.com/spdji/en/documents/additional-material/sp-500-eps-est.xlsx . Retrieved Oct 17, 2026 (HST).

[10] **Caveats.** *Caveats:* book value understates today's asset base more than in 2000 (intangibles such as software and R&D are expensed, and buybacks shrink equity), so P/B overstates the gap versus 2000 somewhat. The Oct 2026 value rises with price on a fixed Sep 2025 book and will fall back when newer book values are reported. PCRA notes S&P data may not be redistributed; we use four year-end ratios for internal analysis.

CHART 3

US Market Cap to GDP: Dot-com (1995-2002) vs AI (2021-2026)

Market value of US corporate equities as % of nominal GDP (Buffett indicator), quarterly, aligned so 2000 sits over 2026

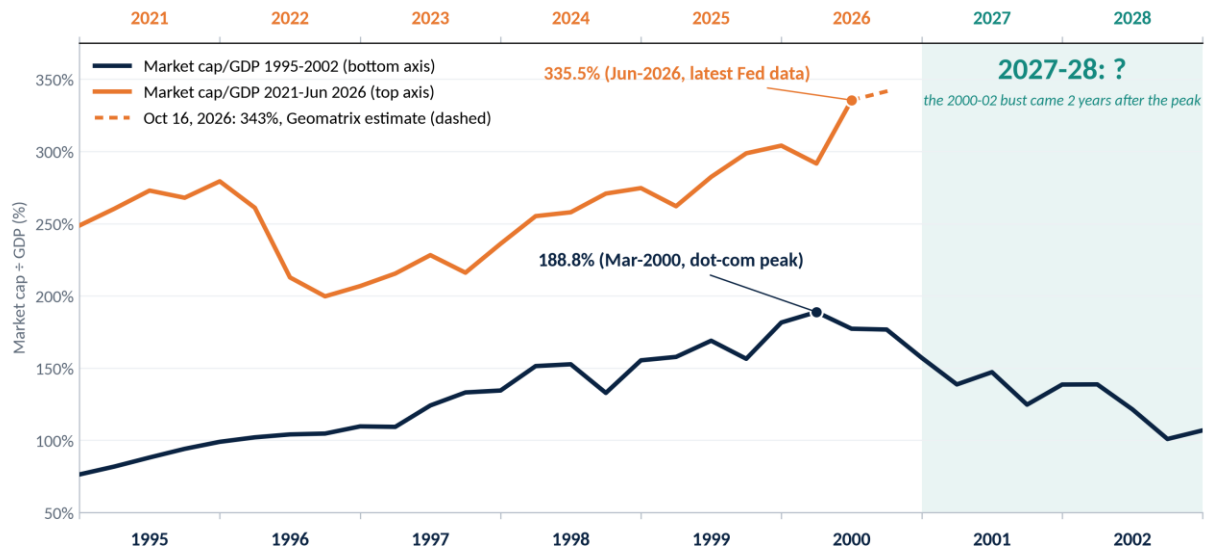

Our stock market is worth 335% of GDP, 1.8x the dot-com peak of 189% in March 2000. Source: Federal Reserve Z.1 and GDP via FRED; Geomatrix (H.K.) Limited.[11,12,13,14,15]

Caveat: Market cap to GDP does not account for enormous increases in operating margins for American companies. Nor does it account for rising overseas operations since Warren Buffett first started using this measurement. As of 2025, overseas sales for the S&P500 had risen to X%, profits Y%. (fill these in from the answer coming from Soros).

Every tech bubble forms around a transforming technology. The biggest ever was railways: new routes, speed and capacity, enormous gains for society. But the financing frenzy inflated stocks, and when cash flows could not service the debt, bankruptcies followed, preceded by stock price crashes.

[11] **Definition.** *Definition:* market cap ÷ GDP = (market value of corporate equities issued by US nonfinancial corporate business, including closely held firms, plus equities issued by US domestic financial sectors, end of quarter) ÷ nominal GDP at a seasonally adjusted annual rate for the same quarter, in %. One definition is used for both eras. ***Window:*** bottom axis 1995-2002 (navy), top axis 2021-2028 (orange); slot 6 aligns 2000 with 2026 as in Chart 1, and 2027-28 is left empty. Each line starts at the prior year-end (Q4 1994, Q4 2020).

[12] **Data and frequency.** *Data:* quarterly, Q4 1994 to Q4 2002 and Q4 2020 to Q2 2026. Q2 2026 is the latest Z.1 and GDP quarter on FRED (Z.1 Q2 2026 release; Q3 2026 GDP is not yet published). ***Method:*** Geomatrix estimate for Oct 16, 2026 = Q2 2026 ratio × S&P 500 close Oct 16, 2026 ÷ close Jun 30, 2026 (7,659.44 ÷ 7,499.36) = 342.6%, with GDP held at Q2 2026 (nominal GDP growth since then would lower it slightly). The S&P 500 stands in for the total market because the Wilshire 5000 series (WILL5000PR, WILL5000IND) are no longer available on FRED.

[13] **Peak and latest.** *Note:* dot-com peak 188.8% at end-Q1 2000 (Mar 31, 2000). Latest 335.5% at end-Q2 2026 (Jun 30, 2026): 147 points and 78% above the dot-com peak (1.78x). Oct 16, 2026 estimate 342.6%, 81% above. ***Cross-check:*** using all corporate equity issues held in the US (FRED BOGZ1FL893064105Q, which adds foreign equities held by US residents) gives 209.7% (Q1 2000) and 379.8% (Q2 2026); the gap versus the dot-com peak is similar (81%).

[14] **Sources.** *Sources:* Board of Governors of the Federal Reserve System, Financial Accounts of the United States (Z.1), via FRED: nonfinancial corporate business; corporate equities; liability, level, https://fred.stlouisfed.org/series/NCBEILQ027S ; domestic financial sectors; corporate equities; liability, level, https://fred.stlouisfed.org/series/FBCELLQ027S ; all sectors; corporate equities; liability, level, https://fred.stlouisfed.org/series/BOGZ1FL893064105Q ; Z.1 release https://www.federalreserve.gov/releases/z1/ . US Bureau of Economic Analysis, gross domestic product, via FRED, https://fred.stlouisfed.org/series/GDP . S&P Dow Jones Indices, S&P 500, via FRED, https://fred.stlouisfed.org/series/SP500 . Downloaded Oct 17, 2026 (HST).

[15] **Caveats.** *Caveats:* the Z.1 nonfinancial series includes closely held companies valued from listed-company multiples, so levels run above listed-only measures (such as Wilshire 5000 ÷ GDP); the comparison between eras is like for like. Z.1 and GDP are revised; the Oct 2026 point is an estimate that will be replaced when Q3 and Q4 2026 data are published. US companies' large foreign earnings make GDP an imperfect denominator in both eras.

CHART 4

Every Great Tech Bubble Gave Back Two-Thirds or More

Stock price index, peak month = 100 (log scale), months before and after each bubble's peak

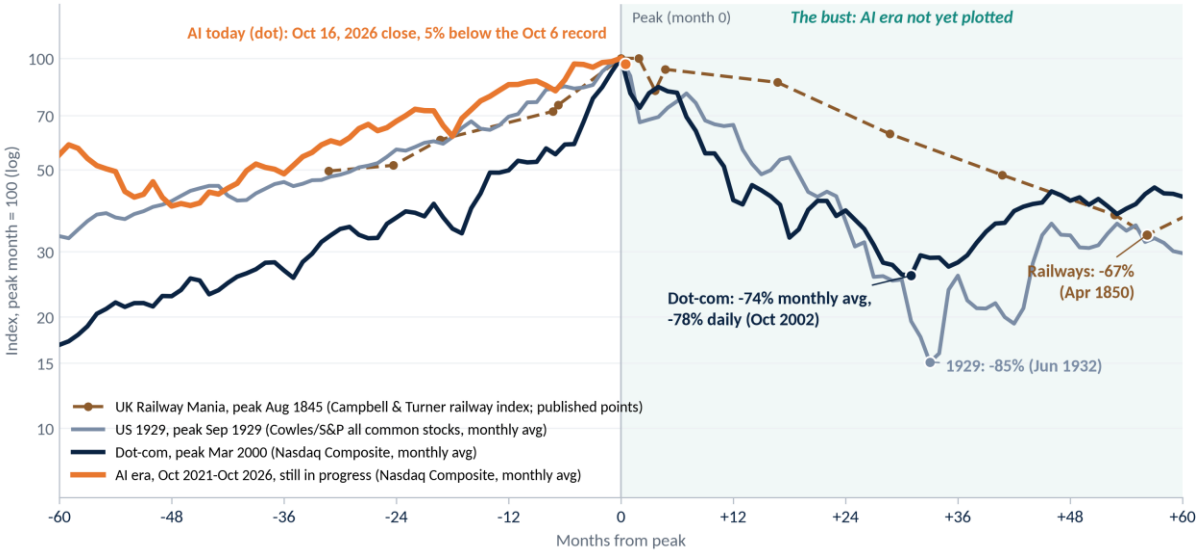

Each great tech bubble gave back two-thirds or more from its peak: railways -67%, 1929 -85%, dot-com -78%. Our AI-era Nasdaq is at month 0. Sources: Campbell and Turner (railways); NBER Macrohistory, Cowles/S&P (1929); Nasdaq via Yahoo Finance; Geomatrix (H.K.) Limited.[16,17,18,19,20]

One reason: over-investment competes away supplier pricing, benefiting users but bankrupting suppliers.

[16] **How to read.** *Method:* each line is a price index rebased so that its peak month = 100, plotted against months from that peak (-60 to +60) on a log scale, so equal vertical distances are equal percentage moves. Peaks: railways 9 Aug 1845 (daily index); 1929, Sep 1929 (monthly average); dot-com, Mar 2000 (monthly average). The AI era is aligned on its highest month to date, Oct 2026 (average of 1-16 Oct closes); that is not a confirmed peak, and the dot shows the Oct 16, 2026 close. *Note:* run-up into the peak: railways +102% (Jan 1843 to Aug 1845, 31 months); 1929 +203%, dot-com +495% and AI era +82% over the 60 months to each peak month.

[17] **Railway Mania.** *Series:* Campbell and Turner all-railway price index (Jan 1, 1843 = 1000), built from daily prices of all listed railway shares in the Railway Times, 1843-50. Only the 15 values printed in the paper (text and Table 1) are plotted, joined by straight lines: peak 2,017.4 on 9 Aug 1845, trough 672.1 on 16 Apr 1850, a fall of 66.7%. No index exists before Jan 1843, so the line starts at month -31. *Sources:* G. Campbell and J. D. Turner, 'The Greatest Bubble in History': Stock Prices during the British Railway Mania, MPRA Paper 21820 (2010), https://mpra.ub.uni-muenchen.de/21820/1/MPRA_paper_21820.pdf . *Caveats:* the older Gayer-Rostow-Schwartz monthly railway index covers only 10 established lines, so we do not use it; year-end points hide intra-year swings.

[18] **US 1929.** *Series:* Cowles Commission / Standard and Poor's index of all common stock prices, monthly averages, 1935-39 = 100 (NBER Macrohistory series m11025a). Peak Sep 1929, 237.8; trough Jun 1932, 35.9 (-84.9%). On daily closes the Dow Jones Industrial Average fell 89% (381.17 on 3 Sep 1929 to 41.22 on 8 Jul 1932). *Sources:* NBER Macrohistory Database, https://data.nber.org/databases/macrohistory/rectdata/11/m11025a.dat ; documentation https://data.nber.org/databases/macrohistory/rectdata/11/docs/m11025a.txt ; Dow closes: Federal Reserve History, Stock Market Crash of 1929, https://www.federalreservehistory.org/essays/stock-market-crash-of-1929 .

[19] **Nasdaq, dot-com and AI era.** *Series:* Nasdaq Composite index, monthly average of daily closes (Nasdaq, Inc., via Yahoo Finance, retrieved Oct 17, 2026 HST). Dot-com: record close 5,048.62 on 10 Mar 2000, low 1,114.11 on 9 Oct 2002 (-77.9%); monthly averages Mar 2000 to Oct 2002, -74.1%. AI era: record close 27,599.79 on 6 Oct 2026; 26,140.03 on 16 Oct 2026 (-5.3%). *Sources:* https://finance.yahoo.com/quote/%5EIXIC/history/ ; cross-check: FRED series NASDAQCOM (Nasdaq OMX), https://fred.stlouisfed.org/series/NASDAQCOM (FRED did not respond at build time).

[20] **Caveats.** *Caveats:* price indices only (no dividends), nominal, and built differently (a full railway-share index, a broad US index, a tech-heavy index). We left out the 2006 housing bubble (Case-Shiller), because FRED did not respond at build time. We also left out US railroads in the 1870s: NBER m11005a shows a modest boom (+37%, Jan 1869 to Jan 1873) rather than a bubble shape. Geomatrix (H.K.) Limited, Oct 17, 2026.

CHART 5

Capex Booms as % of US GDP: Peak Years

Investment in each build-out in its peak year, as a share of US nominal GDP

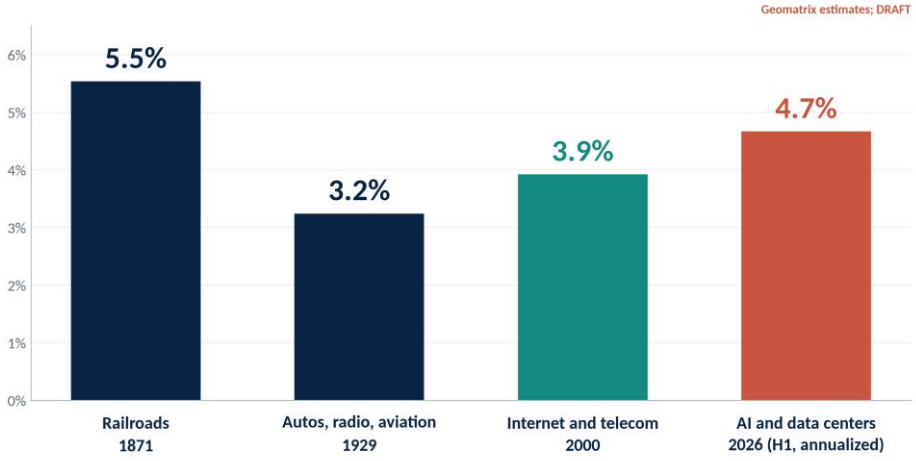

AI and data-center investment, at 4.7% of US GDP (H1 2026, annualized), already tops the 2000 internet and telecom peak of 3.9%. Bars: Railroads 1871, 5.5%[21]; Autos, radio and aviation 1929, 3.2%[22]; Internet and telecom 2000, 3.9%[23]; AI and data centers 2026 (H1, annualized), 4.7%[24]. Geomatrix estimates[25].

[21] **Railroads 1871.** *Definition:* gross capital expenditures of US steam railroads, including land, in current dollars, as a share of US nominal GDP. 1871 is the highest ratio in Ulmer's annual series for 1870-1895, the railroad era's peak. *Components:* railroad capex $429M (Ulmer, Table C-10, col. 1) ÷ nominal GDP $7,751M (MeasuringWorth) = 5.5%. *Sources:* Ulmer (1960), Capital in Transportation, Communications and Public Utilities, NBER, https://www.nber.org/system/files/chapters/c1498/c1498.pdf ; MeasuringWorth, US GDP series, https://www.measuringworth.com/datasets/usgdp/ *Caveats:* Ulmer warns that pre-1910 annual figures may carry large errors; Fishlow (1966) finds Ulmer's cumulative pre-1909 investment short by more than 10% (https://www.nber.org/system/files/chapters/c1578/c1578.pdf); 1871 GDP is a historical reconstruction. *Alternate measures:* 1880s peak 1881 = 3.8% ($450M ÷ $11,902M); 1887 construction peak 2.1%; ~6% for the '1880s' (Kedrosky/Nordvig, https://www.investmentresearchpartners.com/post/chart-of-the-week-8-3-2025); 1850s peak 1854 ~2.6% net (Pereira et al. 2014, https://economics.wm.edu/wp/cwm_wp153.pdf).

[22] **Autos, radio and aviation 1929.** *Definition:* fixed investment in the auto, radio and aviation build-out as a share of US nominal GDP. 1929 is the first year of BEA's national accounts and the 1920s peak for US vehicle output (5.36 million) and radio sales ($842.6M). *Components:* business autos $1,105M; business trucks, buses and trailers $496M; aircraft $41M; air transportation structures $25M; communication equipment including radio and telephone $443M (BEA NIPA Tables 5.4.5 and 5.5.5); highway capital outlay $1,278M (FHWA HF-210); total $3,388M ÷ GDP $104,556M = 3.2%. *Sources:* BEA NIPA, https://apps.bea.gov/national/Release/TXT/NipaDataA.txt ; FHWA Highway Statistics Summary to 1995, Table HF-210, https://www.fhwa.dot.gov/ohim/summary95/hf210.pdf ; peak-year evidence: New York Times, 21 Jan 1930, https://www.nytimes.com/1930/01/21/archives/1929-auto-output-increased-998602-gain-over-1928-production-brought.html ; Gilder Lehrman Institute, https://www.gilderlehrman.org/history-resources/teaching-resource/statistics-american-economy-during-1920s *Caveats:* no published combined estimate exists, so this is our build from official series; BEA does not split radio from telephone equipment; auto and radio factory investment is not separately available and is excluded; the 1930 ratio is similar (3.2%) only because GDP fell. *Alternate measure:* broad version including consumer purchases of motor vehicles and parts ($3,331M) and radios and other audio-visual equipment ($1,002M) = 7.4%.

[23] **Internet and telecom 2000.** *Definition:* US private fixed investment in information and communications technology (ICT) as a share of US nominal GDP. 2000 is the peak year (1999 3.6%, 2001 3.5%). *Components:* computers and peripherals $103.2B, communication equipment $123.3B, software $156.8B, communication structures $18.8B; total $402.1B ÷ GDP $10,251.0B = 3.9%. *Sources:* BEA NIPA Tables 1.1.5, 5.3.5, 5.4.5 and 5.5.5, https://apps.bea.gov/national/Release/TXT/NipaDataA.txt ; https://apps.bea.gov/iTable/?reqid=19&step=2&isuri=1&categories=survey *Caveats:* includes all software and computer investment, not only internet-related; data centers were not reported separately by BEA in 2000 (they sat within commercial structures). *Alternate measures:* telecom only (communication equipment and structures) 1.4%, against published telecom capex estimates of 1.0% (Kedrosky, https://paulkedrosky.com/honey-ai-capex-ate-the-economy/) to 1.18% (Smead Capital, citing FRBSF, https://smeadcap.com/wp-content/uploads/2026/03/Cole-Oasis-Presentation_Manias-in-American-CapEx-Cycles_2026-Oasis.pdf); broad BEA information processing equipment plus software 4.4%.

[24] **AI and data centers 2026 (H1, annualized).** *Definition:* the same ICT definition as the 2000 bar plus BEA's data-center structures, H1 2026 average of Q1 and Q2 seasonally adjusted annual rates, as a share of US nominal GDP on the same basis. This keeps the AI bar like-for-like with 2000: US investment over US GDP. *Components (H1 2026, annual rate):* software $800.0B, computers and peripheral equipment $409.8B, communication equipment $204.6B, data centers (structures) $61.7B, communication structures $27.6B; total $1,503.8B ÷ GDP $32,234.7B = 4.7%. Excluding data centers: 4.5%. *Sources:* BEA NIPA quarterly data (Tables 1.1.5, 5.3.5, 5.4.5, 5.5.5), https://apps.bea.gov/national/Release/TXT/NipaDataQ.txt ; https://apps.bea.gov/iTable/?reqid=19&step=2&isuri=1&categories=survey *Caveats:* covers all US ICT investment, not only AI; H1 2026 annualized, not a full year, and 2026 may not be the peak; quarterly estimates are subject to BEA revision; excludes power and grid investment and investment by US firms abroad. *Alternate measures:* global capex guidance of five hyperscalers (Microsoft, Amazon, Alphabet, Meta, Oracle) ~$825B ÷ US GDP = ~2.5% (range 2.5-2.6%), not like-for-like because it mixes worldwide spend with US GDP (https://www.tmtfinance.com/intel/2026-hyperscaler-capex-tops-us700bn-analysis ; https://www.microsoft.com/en-us/investor/events/fy-2026/earnings-fy-2026-q4 ; https://investor.oracle.com/investor-news/news-details/2026/Oracle-Announces-Q1-Results-Driven-by-Triple-Digit-Growth-in-Cloud-Infrastructure-Revenues/default.aspx); US-located AI investment ~1.8% of GDP (Goldman Sachs, Aug 2026, https://www.goldmansachs.com/insights/articles/global-investment-is-forecast-to-exceed-1-trillion-in-2026); BEA data centers plus computers alone 1.5%.

[25] **Geomatrix estimates.** All four bars are Geomatrix estimates: our arithmetic on the cited published series, with no interpolated or invented inputs. Historical bars use official or academic series; peak years are chosen from each era's annual data. Full data, method and source URLs are in the workbook 'Teaser Wall Chart 3 data.xlsx'. Geomatrix (H.K.) Limited, 17 Oct 2026.

CHART 6

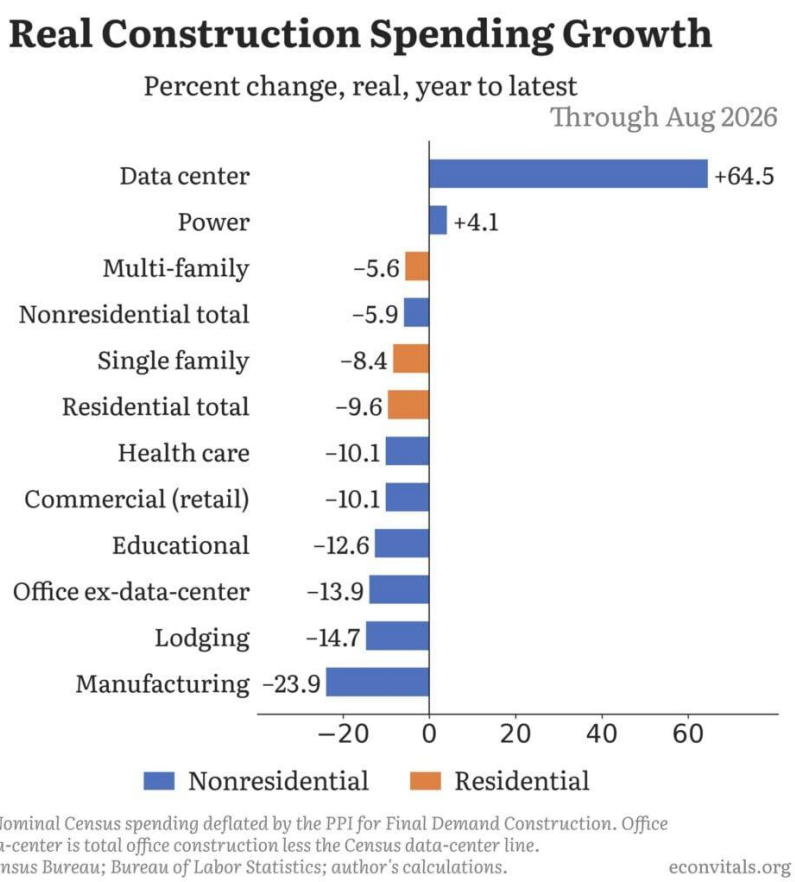

Data-center construction is up 64.5% in real terms while almost every other category shrinks, from offices (-13.9%) to manufacturing (-23.9%). Third-party chart: econvitals.org, "Real Construction Spending Growth" (Census Bureau and BLS data), used with permission.[26]

[26] **Construction spending source.** *Source:* econvitals.org, "Real Construction Spending Growth", percent change, real, year to latest, through Aug 2026. From Robert Howe's file 'AI and the 2-Year Teaser Wall chart 10.docx', saved Oct 17, 2026 (HST). Third-party chart; used with permission. ***Note:*** on-chart note: nominal Census spending deflated by the PPI for Final Demand Construction; office ex-data-center is total office construction less the Census data-center line. U.S. Census Bureau; Bureau of Labor Statistics; author's calculations.

CHART 7

Every Build-out Crashes the Cost to the User

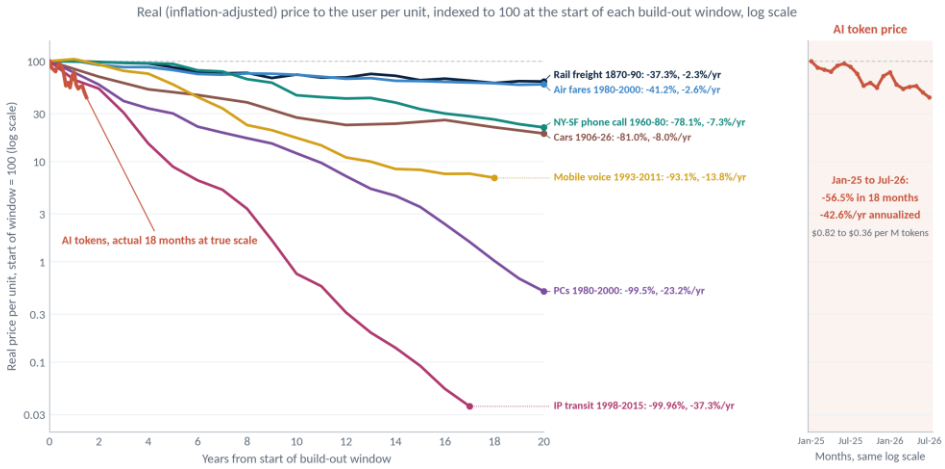

AI token prices paid fell 56.5% in 18 months, a faster annual pace (-42.6%/yr) than any earlier build-out shown. Real decline and real annual rate over each window: Rail freight 1870-90: -37.3%, -2.3%/yr[27]; Cars 1906-26: -81.0%, -8.0%/yr[28]; Air fares 1980-2000: -41.2%, -2.6%/yr[29]; NY-SF phone call 1960-80: -78.1%, -7.3%/yr[30]; PCs 1980-2000: -99.5%, -23.2%/yr[31]; Mobile voice 1993-

[27] **Rail freight 1870-90: -37.3%, -2.3%/yr.** *Series:* Rail: revenue per ton-mile, 13 trunk railroads, 1.80c (1870) to 0.80c (1890); NBER Macrohistory a03030 via FRED, https://fred.stlouisfed.org/data/A0303FUSA259NNBR. All-roads HSUS Df908/Df979: 1.24c (1882) to 0.73c (1900). *Sources:* NBER Macrohistory Database series a03030 (Revenue per freight ton-mile, thirteen railroads), via FRED A0303FUSA259NNBR. *Window:* Most dramatic ~20-yr window available near 1880-1900 in a consistent series: the trunk-line build-out (US track ~53k to ~164k miles). 1880-1900 on HSUS all-roads data is shallower (1882-1902: -32% real); 1865-85 would be -54% real but sits 15 years off the 1880-1900 period. *Cross-check:* all US railroads (HSUS Millennial Df908 Poor's, Df979 ICC): 1.236 c (1882) to 0.927 c (1890) to 0.729 c (1900). https://hsus.cambridge.org/HSUSWeb/toc/treeTablePathIdDf901-910.html ; https://hsus.cambridge.org/HSUSWeb/toc/treeTablePathIdDf965-979.html

[28] **Cars 1906-26: -81.0%, -8.0%/yr.** *Series:* Cars: Raff & Trajtenberg (1996), hedonic quality-adjusted, CPI-deflated new-car price index, Table 2.6, 100 (1906) to 19.0 (1926); odd years interpolated; https://www.nber.org/system/files/chapters/c6065/c6065.pdf. Model T: $850 (1908) to $290 (1925). *Sources:* Raff & Trajtenberg (1996), 'Quality-Adjusted Prices for the American Automobile Industry: 1906-1940', NBER The Economics of New Goods, Table 2.6 (index using constant prices). *Window:* Hedonic (quality-adjusted), CPI-deflated price of a car of constant quality; 1906-1926 is the steepest 20-yr window in the 1906-1940 series (-81%), steeper than 1908-1928 (-76%) and 1920-1940 (-35%). Spans Model T launch (1908) and the moving assembly line (1913). *Cross-check:* Model T touring list price $850 (1908) to $290 (1925), nominal (Model T Ford Club of America, https://www.mtfca.com/encyclopedia/). Raff-Trajtenberg: constant-quality mean price $52,640 (1906) vs $10,002 (1926) in 1993 dollars (Table 2.7). *Flags:* interpolated (geometric mean of adjacent biennial points) (1907, 1909, 1911, 1913, 1915, 1917, 1919, 1921, 1923, 1925).

[29] **Air fares 1980-2000: -41.2%, -2.6%/yr.** *Series:* Air fares: ATA/A4A yield, revenue per passenger-mile, US scheduled airlines, 11.0c (1980) to 13.51c (2000) nominal; ATA Annual Reports 1991 and 2001, https://airlines.org/wp-content/uploads/2014/08/1991.pdf and https://airlines.org/wp-content/uploads/2014/08/2001.pdf. *Sources:* Air Transport Association (now A4A) Annual Reports: 1991 report '1980-1990 Summary' (1980-89, 1 decimal); 2001 report 10-year table (1990-2000). *Window:* Post-deregulation build-out of hub-and-spoke networks and low-cost carriers. Starting 1978 is not steeper in real terms (1979-81 fuel-driven fare rises). *Cross-check:* ATA 2001 report: 'Since deregulation in 1978, in real terms, airline prices have fallen by 38.4 percent.'

[30] **NY-SF phone call 1960-80: -78.1%, -7.3%/yr.** *Series:* Phone: 3-minute daytime station-to-station call, New York-San Francisco, $2.25 (1960) to $1.37 (1980) nominal, step tariff; Historical Statistics of the US, Millennial ed., Dg59, https://hsus.cambridge.org/HSUSWeb/toc/treeTablePathIdDg56-63.html. *Sources:* Historical Statistics of the US, Millennial Edition, Table Dg56-63, series Dg59 (from Bicentennial ed. series R16 to 1970 and FCC/AT&T thereafter). *Window:* Microwave/coaxial long-haul build-out and direct distance dialing. Nominal rate fell $2.25 to $1.37 while CPI rose 2.8x. 1926-1946 would be about -80% real but is far from the 1960-80 build-out. *Cross-check:* Tariff changes: Feb-1960 $2.25, Dec-1965 $2.00, Dec-1967 $1.75, Aug-1968 $1.70, Feb-1970 $1.35, 1973 $1.45, 1975 $1.36, 1976-78 $1.30, 1980 $1.37 (1981: $1.58). *Flags:* rate carried forward (step tariff, no change recorded) (1961, 1962, 1963, 1964, 1965, 1967, 1971, 1979).

[31] **PCs 1980-2000: -99.5%, -23.2%/yr.** *Series:* PCs: BEA PCE price index, personal computers and peripheral equipment (NIPA Table 2.4.4U line 49, DCPPRG, quality-adjusted), 101,536 (1980) to 1,074 (2000), 2017=100; https://apps.bea.gov/national/Release/TXT/NipaDataA.txt. *Sources:* US BEA, NIPA Table 2.4.4U line 49, 'Personal computers/tablets and peripheral equipment' price index (series DCPPRG); also https://apps.bea.gov/iTable/?reqid=19&step=2&isuri=1&categories=underlying. *Window:* BEA's consumer (PCE) PC price index is quality-adjusted (hedonic) and so measures cost per unit of computing to the household. 1982-2002 is marginally steeper (-99.55% vs -99.49% real), not a material difference. *Cross-check:* BEA business investment price index for computers & peripherals (B935RG) 8,402.8 (1980) to 317.0 (2000), -96% nominal. Nordhaus (2007) finds far steeper declines in cost per computation.

2011: -93.1%, -13.8%/yr[32]; IP transit 1998-2015: -99.96%, -37.3%/yr[33]; AI tokens Jan-25 to Jul-26: -56.5%, -42.6%/yr (nominal)[34]. Real prices, deflated by US CPI; sources in the notes; Geomatrix (H.K.) Limited.[35,36]

[32] **Mobile voice 1993-2011: -93.1%, -13.8%/yr.** *Series:* Mobile: voice revenue per minute, CTIA via FCC 15th and 16th Mobile Competition Reports, $0.439 (1993) to $0.047 (2011), last year published; https://docs.fcc.gov/public/attachments/FCC-13-34A1.txt. *Sources:* FCC 16th Mobile Wireless Competition Report (FCC 13-34, 2013), Table 30, from CTIA Wireless Indices (same series as FCC 15th Report, Table 20); also https://docs.fcc.gov/public/attachments/FCC-11-103A1.txt. *Window:* Longest consistent official series (1993-2011, 18 years). The FCC stopped publishing voice RPM after 2011 (bundled unlimited plans), so 1995-2015 cannot be sourced on a consistent basis. Starting at the 1994 peak would give -90% nominal. *Cross-check:* Blended RPM (incl. data) published to 2011 as $0.08; voice RPM is the per-unit price of the service being built out.

[33] **IP transit 1998-2015: -99.96%, -37.3%/yr.** *Series:* Bandwidth: US IP transit, $1,200 (1998) to $0.63 (2015) per Mbps per month; DrPeering.net (W. Norton), survey-based, 2014-15 possibly projected; https://drpeering.net/HTML_IPP/chapters/ch02-Internet-Transit/ch02.1-Internet-Transit-Prices. *Sources:* William B. Norton, DrPeering.net, 'Internet Transit Prices - Historical and Projected' and The Internet Peering Playbook Table 2-2; also https://peering.drpeering.net/white-papers/Internet-Transit-Pricing-Historical-And-Projected.php. *Window:* The only longitudinal public $/Mbps series; starts 1998 (no reliable pre-1998 data), so the window is 17 years. Survey-based market estimates, not audited prices. BLS internet-services CPI (from Dec-1997) is far shallower and not a per-unit price. *Cross-check:* Ending at 2013 (last clearly historical year) would give -99.87% nominal. *Flags:* DrPeering labels the table 'historical and projected'; 2014-15 may be projections (2014, 2015).

[34] **AI tokens Jan-25 to Jul-26: -56.5%, -42.6%/yr.** *Series:* AI: volume-weighted price paid on OpenRouter, $0.82 (Jan-25) to $0.36 (Jul-26) per million tokens, nominal (US CPI rose ~4% over the period); OECD.AI OpenRouter data. A platform proxy for the market, not a list price. *Sources:* OECD.AI, 'Average cost of AI model usage through OpenRouter', https://oecd.ai/en/data?selectedArea=ai-model-usage&selectedVisualization=average-cost-of-ai-model-usage-through-openrouter-across-open-and-closed-weight-models ; data pulled 17 Oct 2026 from https://observatory.oecdai.org/graphql/v1 (table o2__openrouter.average_costs). *Note:* plotted at true time scale (0-1.5 years) in the left panel and zoomed with months on the x-axis in the right panel, on the same log scale; $0.82 to $0.36 per million tokens = -56.5% in 18 months, -42.6%/yr annualized.

[35] **How to read.** *How to read:* left panel = years from each window start (0-20); each line is the real price per unit, start year = 100; end labels give total real decline and real CAGR. AI overlay: the orange stub at 0-1.5 years is the AI token price over its actual 18 months on the same time axis. The right panel repeats it on the identical log scale with months on the axis (time axis stretched about 3x).

[36] **Deflator, windows and method.** Deflator: US CPI, 1800- (Minneapolis Fed / BLS), https://www.minneapolisfed.org/about-us/monetary-policy/inflation-calculator/consumer-price-index-1800-. Windows chosen as the steepest ~20-year span in or near each build-out period with consistent data; wireless (18 yrs) and bandwidth (17 yrs) are limited by data availability. *Method:* no series was dropped. Two windows are shorter than 20 years for data reasons: wireless voice RPM 1993-2011 (FCC stopped publishing after 2011) and IP transit 1998-2015 (DrPeering series starts 1998). *Method:* rail window moved to 1870-1890 (13-railroad NBER series) because all-road national revenue per ton-mile is only available from 1882; 1882-1902 all-roads = -32% real. Autos: Raff-Trajtenberg biennial index; odd years are geometric interpolations (flagged) and are not plotted as separate points. Telephone: step tariff; years without a recorded change carry the prior rate (flagged). *Method:* a frontier-model list-price line is not shown; it would need verified 2026 list prices. Index arithmetic in Python (c5_data.py); full data in 'Teaser Wall Chart 5 - Tech price declines data.xlsx'. Geomatrix (H.K.) Limited, 17 Oct 2026.

CHART 8

Compute (Token) Volume Explodes, but Price Collapsing

Chinese Driving Price to Zero

Infinity x Zero = ?

On OpenRouter, monthly tokens rose ~110x from Jan 2025 to Jul 2026 while the average price paid fell from $0.82 to $0.36 per million tokens, and Chinese-developed models went from 6% to 65% of tokens. OpenRouter is a proxy, not the market; company disclosures are shown as markers.

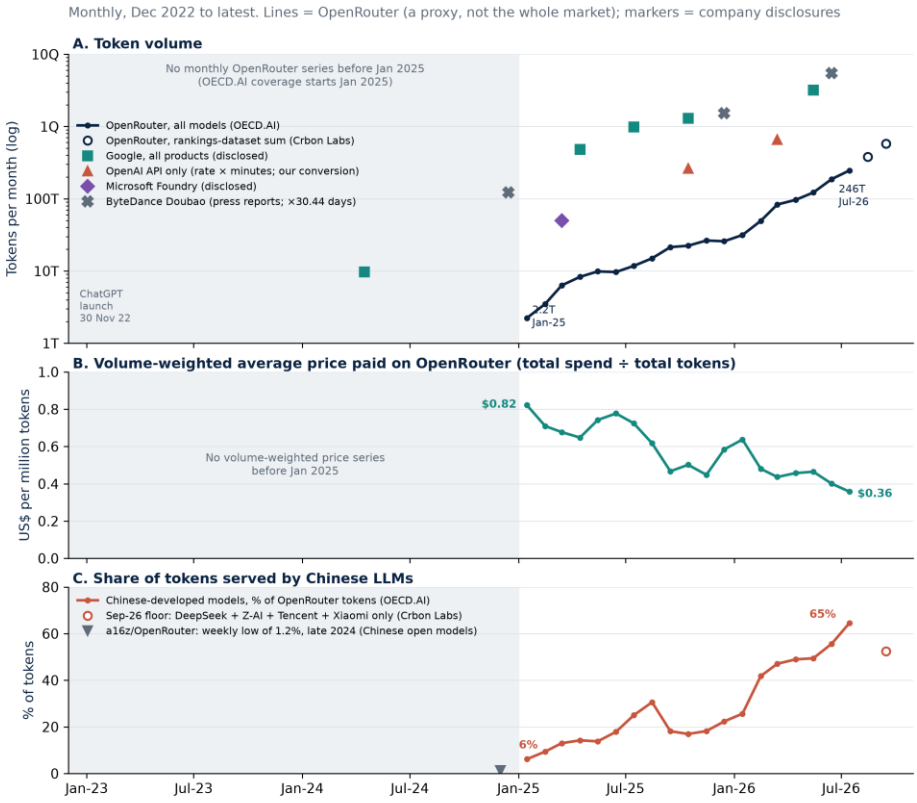

Usage is exploding while the price per token falls, and Chinese-developed models now carry most OpenRouter tokens.

Sources: OECD.AI (OpenRouter data); Crbon Labs; company disclosures; a16z/OpenRouter; Geomatrix (H.K.) Limited.[37,38,39,40,41]

[37] **OpenRouter monthly series.** *Sources:* OECD.AI 'AI model usage' data (OpenRouter), monthly, Jan 2025 to Jul 2026, last update 27 Aug 2026 (tokens, upstream cost, model developer country = CHN), https://oecd.ai/en/data?selectedArea=ai-model-usage&selectedVisualization=average-cost-of-ai-model-usage-through-openrouter-across-open-and-closed-weight-models ; API pull https://observatory.oecdai.org/graphql/v1 (table o2__openrouter.average_costs; chart 34020; pulled 17 Oct 2026). *Note:* tokens 2.23T (Jan-25) to 246T (Jul-26), ~110x; volume-weighted price = total spend ÷ total tokens, $0.82 to $0.36 per million tokens; Chinese-developed models 6% to 65% of tokens. Price is what OpenRouter users paid, incl. cached and reasoning tokens; not list prices. No points are interpolated.

[38] **Aug-Sep 2026 points.** Tokens 379.4T (Aug-26) and 575.9T (Sep-26) are sums of the OpenRouter rankings-daily dataset (top-50 models per day) by Crbon Labs, a different method from the OECD series; the Sep-26 Chinese share of 52.4% is a floor (DeepSeek + Z-AI + Tencent + Xiaomi only). Crbon Labs, 'LLM market share in tokens' (9 Sep 2026, updated 6 Oct 2026), https://crbonfree.com/blog/llm-market-share-tokens

[39] **Company markers.** Google, all products and API: 9.7T tokens/month (Apr-24) and 480T (I/O May 2025), https://blog.google/technology/ai/io-2025-keynote/ ; 980T+ (Q2-25 call), https://blog.google/inside-google/message-ceo/alphabet-earnings-q2-2025/ ; 1.3 quadrillion (Q3-25 call), https://blog.google/inside-google/message-ceo/alphabet-earnings-q3-2025/ ; 3.2 quadrillion+ (I/O 2026), https://blog.google/innovation-and-ai/sundar-pichai-io-2026/ . Microsoft Foundry: record 50T tokens in March 2025 (FY25 Q3 call), https://www.microsoft.com/en-us/investor/events/fy-2025/earnings-fy-2025-q3 . OpenAI API only, excl. ChatGPT: 6B+ tokens/minute (DevDay Oct 2025) and 15B+ (Mar 2026); monthly figures are our conversion (x 43,834 minutes), https://aisafetytracker.org/usage/ . ByteDance Doubao (Volcano Engine, press reports): 4T tokens/day (Dec-24), 50T+ (Dec-25), 180T+ (Jun-26); our x30.44 days conversion, https://aisafetytracker.org/usage/ . *Note:* company token counts use different tokenisers and scopes and are not additive.

[40] **Chinese share context.** a16z/OpenRouter 'State of AI' (Dec 2025): Chinese open-weight models' weekly share 'as low as 1.2%' in late 2024; average 13.0% Nov-24 to Nov-25; 'nearly 30%' some weeks, https://openrouter.ai/state-of-ai

[41] **Coverage and gaps.** OpenRouter is one API gateway skewed to price-sensitive developers; it excludes ChatGPT, the Gemini app and direct lab APIs, so it is a proxy, not the whole market. No market-wide monthly token, price or Chinese-share series exists; nothing monthly for Dec 2022 to Dec 2024; Oracle and Alphabet do not disclose customer splits; Anthropic, DeepSeek, Meta and xAI publish no token volumes; Microsoft stopped token totals after FY25. Data: 'Teaser Wall chart data.xlsx'. Geomatrix (H.K.) Limited.

For now the math is bullish: 110x the tokens at 44% of the price ($0.82 to $0.36) means spend rose ~48x.[42] Demand has been real so far; the risk is the 2027-28 reset wall, not today's usage. And most of the price fall is mix, not like-for-like cuts: Chinese-developed models, at $0.11 per million tokens in Jul 2026, took 65% of tokens, while the average price of all other models barely moved ($0.82 to $0.81).

Jevons Paradox says falling prices make demand surge, but elasticity is the key: a 90% price cut needs a 900% demand surge just to break even.[43]

One well-known bear, Michael Burry of "The Big Short", has called many more downturns than have occurred: at least eight public bear warnings since 2015, against only two actual bear markets (Covid in 2020, and the 2022 "rates riot").[44] He points to accounting gimmicks like stretching AI-chip depreciation from 3-4 years to 5-6.[45] That inflates accounting profits, but bulls point out that old AI chips are still being used productively, the Nvidia H100 introduced 6 years ago still profitable. The issue, however, is not accounting profits but cash flow to service a looming debt/lease wall.

Unlike individual homeowners in 2007-8, and dot.com startups in 2000, the mega-cap hyperscalers earn enormous profts and free cash flow from their core businesses. Bankruptcy will not be an issue, so the contagion to banks from homeowner defaults does not parallel with OpenAI bankrupting the Mag 7, but still a decline in revenue and profits will not favor share prices. OpenAI alone is about 45% of Microsoft's backlog and, we estimate, and of Oracle's[46]. Suppliers finance their own customers (Nvidia invests in OpenAI and backstops CoreWeave),[47] and defaults will result in write-downs.

We are not reliving the internet bust. We are reliving the housing crash, just without the banking sector contagion, so the fallout should be limited.

---

[42] **Implied spend and price mix.** *Data:* OECD.AI OpenRouter series (Chart 8 notes): tokens 2.2301T (Jan-25) to 246.0795T (Jul-26) = 110.3x; price $0.8222 to $0.3578 per million tokens (44% of the start); spend $1.83m to $88.05m a month = 48.0x. *Mix:* Chinese-developed models $0.833 to $0.107 per million tokens (64.5% of Jul-26 tokens); all other models $0.822 to $0.814 (our arithmetic: all models less Chinese-developed). *Caveats:* a platform proxy, not list prices. *Source:* OECD.AI table o2__openrouter.average_costs, pulled 17 Oct 2026; 'Teaser Wall chart data.xlsx'.

[43] **Jevons Paradox.** Named for William Stanley Jevons, The Coal Question (1865), who observed that more efficient use of coal raised, rather than cut, total coal consumption. *Method:* revenue = price x volume. After a 90% price cut, revenue is unchanged only if volume rises 10x (0.1 x 10 = 1), i.e. +900%. Whether demand grows that much depends on the price elasticity of demand; a price-elasticity below 1 means revenue falls.

[44] **Burry bear calls vs actual bear markets.** Count: at least eight public market-wide crash, bubble or bear warnings since 2008: Dec 2015 (crash within months); May 2017 (global financial meltdown); Sep 2019 (index funds a bubble like subprime CDOs); Feb 2021 (speculative bubble; short Tesla); Jun 2021 ("mother of all crashes", tweets later deleted); 31 Jan 2023 ("Sell", retracted 30 Mar 2023: "I was wrong to say sell"); Aug 2023 (13F showed puts on SPY and QQQ, ~$1.6bn notional); Nov 2025 (puts on Nvidia and Palantir, plus the depreciation critique). Not counted: bearish bets in Mar 2020 and the Sep 2022 "bottom not hit" warning, made inside bear markets already under way, and his famous 2005-07 subprime call. Bear markets: S&P 500 closing falls of 20% or more since the March 2009 low: Covid, 19 Feb-23 Mar 2020 (-33.9%), and 3 Jan-12 Oct 2022 (4,796.56 to 3,577.03, -25.4%). Near misses: 2011 (-19.4%) and 2018 (-19.8%). In fairness, his 2021 warnings came 7-11 months before the 2022 bear market began. *Sources:* B. Ritholtz, "Folly of Forecasts, Big Short Edition" (Aug 2023, citing J. Khoo), https://ritholtz.com/2023/08/folly-of-forecasts-big-short-edition/ ; Bloomberg, 18 Jun 2021, https://www.bloomberg.com/news/articles/2021-06-18/crypto-meme-stocks-michael-burry-warns-about-mother-of-all-crashes ; Kitco, 31 Mar 2023, https://www.kitco.com/news/article/2023-03-31/michael-burry-i-was-wrong-say-sell ; Reuters, 14 Aug 2023, https://www.reuters.com/markets/us/burry-famous-big-short-bought-bearish-options-against-sp-nasdaq-100-2023-08-14/ ; CNN, 5 Nov 2025, https://www.cnn.com/2025/11/05/business/nvidia-palantir-michael-burry-stock ; BNY Investments, Historical bouts of volatility, https://www.bny.com/assets/investments/im/documents/manual/perspectives/market-behavior-during-historical-bouts-of-volatility.pdf ; S&P 500 daily closes (S&P Dow Jones Indices).

[45] **Depreciation lives.** Burry (X post, 10 Nov 2025): hyperscalers lengthened the useful lives of compute equipment despite a 2-3 year Nvidia product cycle, which he estimates understates depreciation by ~$176bn over 2026-28 and overstates 2028 earnings by ~27% at Oracle and ~21% at Meta. *Data:* server useful lives per company 10-K filings: Microsoft 3 to 4 years (FY2021), 4 to 6 (FY2023); Alphabet 3 to 4 (2021), 4 to 6 (2023); Amazon 3 to 4 (2020), 4 to 5 (2022), 5 to 6 (2024), a subset back to 5 (2025, citing the pace of AI); Meta 5.5 years from Jan 2025; Oracle 4 to 5 (FY2023), 5 to 6 (FY2025). Nvidia says customers depreciate GPUs over 4-6 years. *Caveats:* the lives cover servers and network equipment as a class, not AI chips alone. *Sources:* CNBC, 11 Nov 2025, https://www.cnbc.com/2025/11/11/big-short-investor-michael-burry-accuses-ai-hyperscalers-of-artificially-boosting-earnings.html ; CNBC, 14 Nov 2025, https://www.cnbc.com/2025/11/14/ai-gpu-depreciation-coreweave-nvidia-michael-burry.html ; CNBC, 25 Nov 2025, https://www.cnbc.com/2025/11/25/nvidia-pushes-back-on-charges-that-ai-investment-is-a-bubble.html ; The Register, 2 Aug 2022 (Microsoft 4 to 6), https://www.theregister.com/off-prem/2022/08/02/microsoft-extends-life-of-cloud-servers-to-six-years/311008 ; Meta 10-K FY2024, https://www.sec.gov/Archives/edgar/data/1326801/000132680125000017/R31.htm

[46] **Backlog concentration.** *Data:* Microsoft says about 45% of its $625bn commercial RPO (Dec 31, 2025) is tied to OpenAI. Oracle's RPO was $664bn (Aug 31, 2026) and it does not disclose customers; 45% is our estimate from the reported ~$300bn OpenAI contract. *Sources:* CNBC, Jan 28, 2026, https://www.cnbc.com/2026/01/28/microsoft-msft-q2-earnings-report-2026.html ; Oracle 10-Q, https://www.sec.gov/Archives/edgar/data/1341439/000119312526389274/orcl-20260831.htm ; Reuters, Sep 10, 2025, https://www.reuters.com/technology/openai-oracle-sign-300-billion-computing-deal-wsj-reports-2025-09-10/ .

[47] **Circular deals.** *Data:* Nvidia planned up to $100bn for OpenAI (Sep 2025), later a $30bn stake (Chart 10 notes), and agreed to buy up to $6.3bn of CoreWeave's unsold capacity through 2032; CoreWeave, Nvidia-backed, rents Nvidia GPUs to OpenAI (contracts up to $22.4bn). *Sources:* Reuters, Sep 22, 2025, https://www.reuters.com/business/nvidia-invest-100-billion-openai-2025-09-22/ ; CNBC, Oct 15, 2025, https://www.cnbc.com/2025/10/15/a-guide-to-1-trillion-worth-of-ai-deals-between-openai-nvidia.html ; Reuters, Sep 25, 2025, https://www.reuters.com/business/coreweave-expands-openai-pact-with-new-65-billion-contract-2025-09-25/ .

In the housing bubble, teaser-rate mortgages let anyone buy a home. Banks ignored credit standards because they quickly sold the loans on as bundled securities (Collateralized Debt Obligations), layered by credit quality, with rates rising as quality fell: pick your poison.

The 2-year teasers written in 2004-6 reset in 2006-8 and, unlike the 2002-4 wave, could not be refinanced on rising home prices.[48] The resets needed the bubble to keep inflating; even slower price gains stopped the merry-go-round.

Today's data-center leases work the same way: no payments until the center turns on, in 2027-28.[49] What if compute (token) demand slows? Corporations are moving data centers in-house, safe from hackers, cutting compute costs to zero for the tuned models that run the business; only new development needs frontier models. And we consumers will soon run everyday AI on our laptops and phones, paying for tokens only when we stretch, which few will. So volume will disappoint.

Timing: the market might justreprice now on cash-flow worries; the solvency stress itself comes when the commitments commence in 2027-28. **What would prove us wrong:** hyperscalers raising capex guidance again with matching revenue growth, or AI revenue scaling fast enough to meet the 2027-28 commitments as they start.

[48] **Teaser resets.** *Definition:* a "2/28" subprime hybrid ARM carries a fixed teaser rate for 2 years, then resets to a floating rate; borrowers typically refinanced before the reset. Fed Chairman Bernanke (House testimony, 20 Sep 2007): subprime ARMs originated in late 2005 and 2006 performed worst; the sharp deceleration in home prices since 2005 left recent borrowers with little equity, and "low home equity has made refinancing--the typical way for many subprime borrowers to avoid large scheduled interest rate resets--difficult or impossible for many." *Sources:* https://www.federalreserve.gov/newsevents/testimony/bernanke20070920a.htm ; Congressional Research Service, RL33775, Alternative Mortgages: Causes and Policy Implications of Troubled Mortgage Resets, https://www.everycrsreport.com/reports/RL33775.html

[49] **Data-center leases.** *Data:* Oracle FY2026 10-K (May 31, 2026): $260 billion of additional lease commitments, substantially all data centers, generally expected to commence between Q1 fiscal 2027 and fiscal 2029, for 15-19 years, not on the balance sheet. Meta 10-Q (Jun 30, 2026): $278.99 billion of leases not yet commenced (2026-2036), plus about $68 billion signed in July 2026 expected to commence in 2027 and 2028. Microsoft: $329.1 billion of leases not yet commenced (Bloomberg Law, Oct 2026). *Method:* under ASC 842 a lease is recorded only at commencement, when the facility is made available; rent generally starts then. Caveats: terms vary by contract; Oracle has also guaranteed up to $3.3 billion of one lessor's borrowing. *Sources:* https://www.hudson-labs.com/research/oracle-data-center-leases-orcl ; https://www.edgar.tools/blog/oracle-leases-not-yet-commenced ; https://www.sec.gov/Archives/edgar/data/1326801/000162828026050705/R17.htm ; https://news.bloomberglaw.com/artificial-intelligence/microsoft-reports-over-130-billion-in-new-data-center-leases

CHART 9

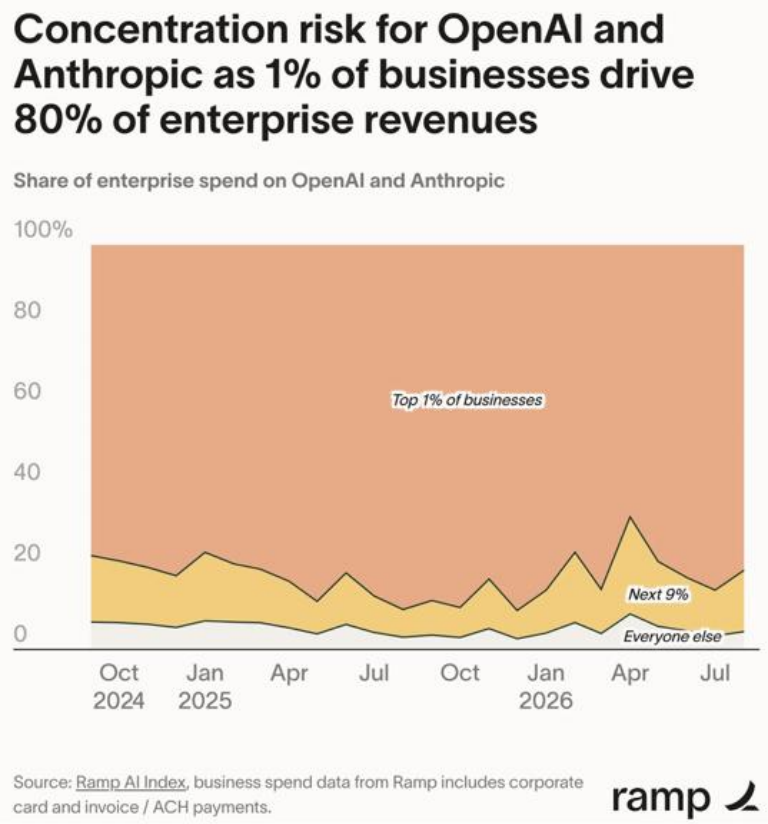

The top 1% of businesses drive about 80% of enterprise spend on OpenAI and Anthropic, so our AI demand rests on a narrow base. Third-party chart: Ramp, Ramp AI Index, used with permission.[50]

Even if data-center volumes DO rise 900%, compute prices are crashing: zero switching costs, no moat, constant technological improvements in chips and software that cut costs, speed response, use less energy. Chip demand is still strong. H100 rental prices are stable or rising. But the compute (token) prices from use of those chips is falling. That is where the rubber meets the road at corporate and consumer real money flowing into the ecosystem. At the bottom are the Chinese competitors are gaining dominant share and driving token prices inexorably to zero.[51].That is the story of every technology since the dawn of man. This time is different? No, this time is orders of magnitude worse in terms of price declines, combined with the S curve of adoption and usage intensity leveling off.

[50] **Ramp AI Index source.** *Source:* Ramp, Ramp AI Index, "Concentration risk for OpenAI and Anthropic as 1% of businesses drive 80% of enterprise revenues"; share of enterprise spend on OpenAI and Anthropic by business size band, 2024-2026 as shown. On-chart note: business spend data from Ramp includes corporate card and invoice / ACH payments. From Robert Howe's file 'AI and the 2-Year Teaser Wall chart 11.docx', saved Oct 17, 2026 (HST). Third-party chart; used with permission. *Caveats:* Ramp customers only, a sample of US businesses, not all enterprise AI spend.

[51] **Chinese models and token prices.** *Data:* on OpenRouter, Chinese-origin models have taken at least 30% of US companies' tokens every week since Feb 8, 2026, peaking at 46%, against 11% averaged over the prior 12 months and 4.5% in H1 2025 (CNBC, Jul 7, 2026). OpenRouter list prices in Oct 2026: DeepSeek V4 Pro $0.18 and Qwen3 235B $0.09 per million input tokens (promotional). See Chart 8 for OpenRouter volume, price and Chinese share. *Sources:* https://www.cnbc.com/2026/07/07/chinese-ai-models-costs-us-openai-anthropic.html ; https://openrouter.ai/collections/discounted-models

CHART 10

AI Finances AI - Payback Wall Hits 2027-28

We estimate ~45% of the big-4 cloud backlog ($1.05T of $2.35T) is tied to OpenAI and Anthropic commitments; 55% of Nvidia's Data Center revenue comes from its Hyperscale customers. The same names sit at both ends of the loop.

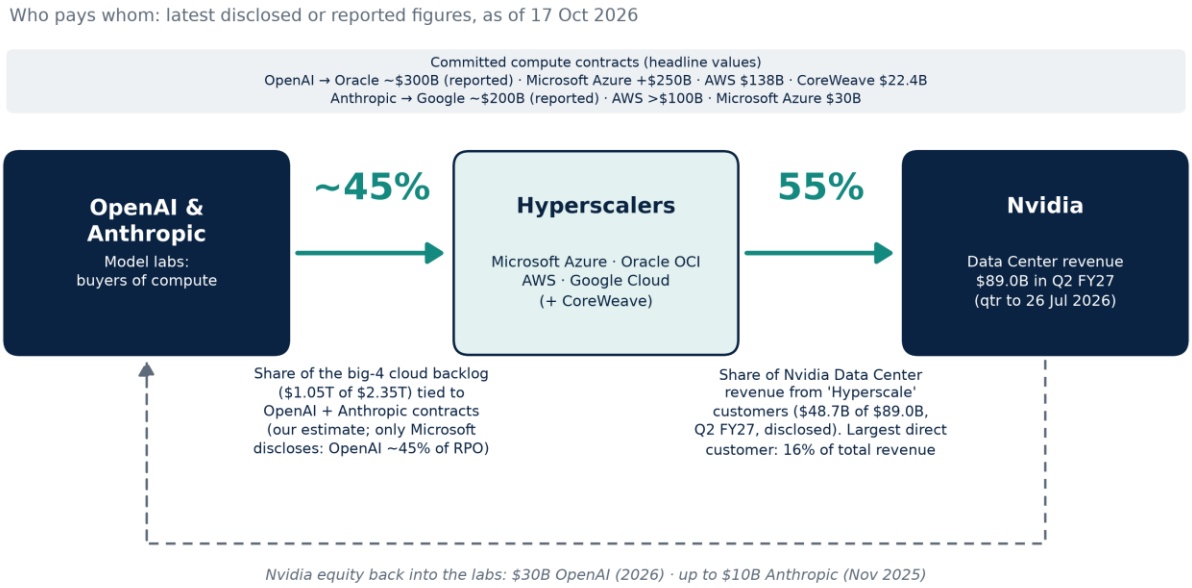

The labs' commitments are the hyperscalers' backlog, and the hyperscalers' spending is Nvidia's revenue. Sources: company filings and press reports; Geomatrix estimates.[52,53,54,55]

[52] **~45% of the big-4 cloud backlog (our estimate).** *Method:* sum of OpenAI and Anthropic headline commitments at Microsoft, Oracle, Google Cloud and AWS ($1,049B) ÷ sum of their latest backlogs ($2,352B) = 44.6%. Geomatrix estimate. Headline values span different dates and terms and are not all confirmed as booked backlog; CoreWeave is excluded from the ratio. Only Microsoft discloses a share (OpenAI ~45% of $625B commercial RPO, Dec-25). *Sources:* backlogs: Microsoft commercial RPO $678B (30 Jun 2026), https://www.microsoft.com/en-us/investor/events/fy-2026/earnings-fy-2026-q4 ; Oracle RPO $664B (31 Aug 2026, 10-Q), https://www.sec.gov/Archives/edgar/data/1341439/000119312526389274/orcl-20260831.htm ; Google Cloud backlog $514B (30 Jun 2026), https://blog.google/company-news/inside-google/message-ceo/alphabet-earnings-q2-2026/ ; Amazon performance obligations, mainly AWS, $496B (30 Jun 2026, 10-Q), https://www.sec.gov/Archives/edgar/data/1018724/000101872426000026/amzn-20260630.htm . *Sources:* commitments: Microsoft-OpenAI $281B (Microsoft CFO: ~45% of $625B commercial RPO at 31 Dec 2025; includes the incremental $250B Azure commitment of 28 Oct 2025; $ figure is our 45% x $625B), https://www.microsoft.com/en-us/investor/events/fy-2026/earnings-fy-2026-q2 ; Microsoft-Anthropic $30B (18 Nov 2025, announced), https://www.reuters.com/technology/anthropic-commits-30-billion-microsoft-azure-compute-2025-11-18/ ; Oracle-OpenAI ~$300B over ~5 years from 2027 (WSJ via Reuters, 10 Sep 2025; reported, not disclosed by Oracle), https://www.reuters.com/technology/openai-oracle-sign-300-billion-computing-deal-wsj-reports-2025-09-10/ ; Google Cloud-Anthropic ~$200B over 5 years (The Information via Reuters, 5 May 2026; reported, not disclosed by Alphabet), https://www.reuters.com/business/anthropic-commits-spending-200-billion-googles-cloud-chips-information-reports-2026-05-05/ ; AWS-OpenAI $38B (Nov 2025) + $100B expansion over 8 years (Amazon Q2-26 10-Q), https://www.sec.gov/Archives/edgar/data/1018724/000101872426000026/amzn-20260630.htm ; AWS-Anthropic more than $100B over 10 years, up to 5 GW (Amazon Q2-26 10-Q; floor), https://www.sec.gov/Archives/edgar/data/1018724/000101872426000026/amzn-20260630.htm ; CoreWeave-OpenAI contracts up to $22.4B (25 Sep 2025; not in the big-4 total), https://investors.coreweave.com/news/news-details/2025/CoreWeave-Expands-Agreement-with-OpenAI-by-up-to-6-5B/default.aspx .

[53] **55% of Nvidia Data Center revenue (disclosed).** Nvidia 'Hyperscale' revenue $48.710B ÷ Data Center revenue $89.023B, Q2 FY27 (quarter to 26 Jul 2026) = 54.7%; Q1 FY27 57.2% ($43.050B ÷ $75.246B). 'Hyperscale' is Nvidia's own customer category. Q2 FY27 total revenue $96.221B; AI Clouds, Industrial & Enterprise $40.313B. Source: Nvidia Q2 FY27 CFO commentary, https://www.sec.gov/Archives/edgar/data/1045810/000104581026000073/q2fy27cfocommentary.htm . *Note:* largest direct customer 16% of total revenue (Q2 FY27); three direct customers 16% + 15% + 13% = 44% (H1 FY27), https://s201.q4cdn.com/141608511/files/doc_financials/2027/NVDA-2027-Q2-10Q-Final-including-exhibits.pdf ; two direct customers 22% + 14% = 36% (FY26 to 25 Jan 2026), https://www.sec.gov/Archives/edgar/data/1045810/000104581026000021/nvda-20260125.htm . Customers are unnamed in the filings.

[54] **Nvidia equity back into the labs.** $30B into OpenAI (2026; CNBC, 4 Mar 2026, reported/CEO statement), https://www.cnbc.com/2026/03/04/nvidia-huang-openai-investment.html ; up to $10B into Anthropic (Nov 2025, announced), https://www.reuters.com/technology/anthropic-commits-30-billion-microsoft-azure-compute-2025-11-18/ .

[55] **Chart 10 data.** Latest disclosed or reported figures as of 17 Oct 2026. Status labels: Disclosed = in company filing or call; Reported = press, not confirmed by the company; Estimate = our arithmetic. Full data and source URLs: 'Teaser Wall chart data.xlsx'. Geomatrix (H.K.) Limited.

CHART 11

The Expanded AI Capital Web: Who Owes Whom (2 Oct 2026)

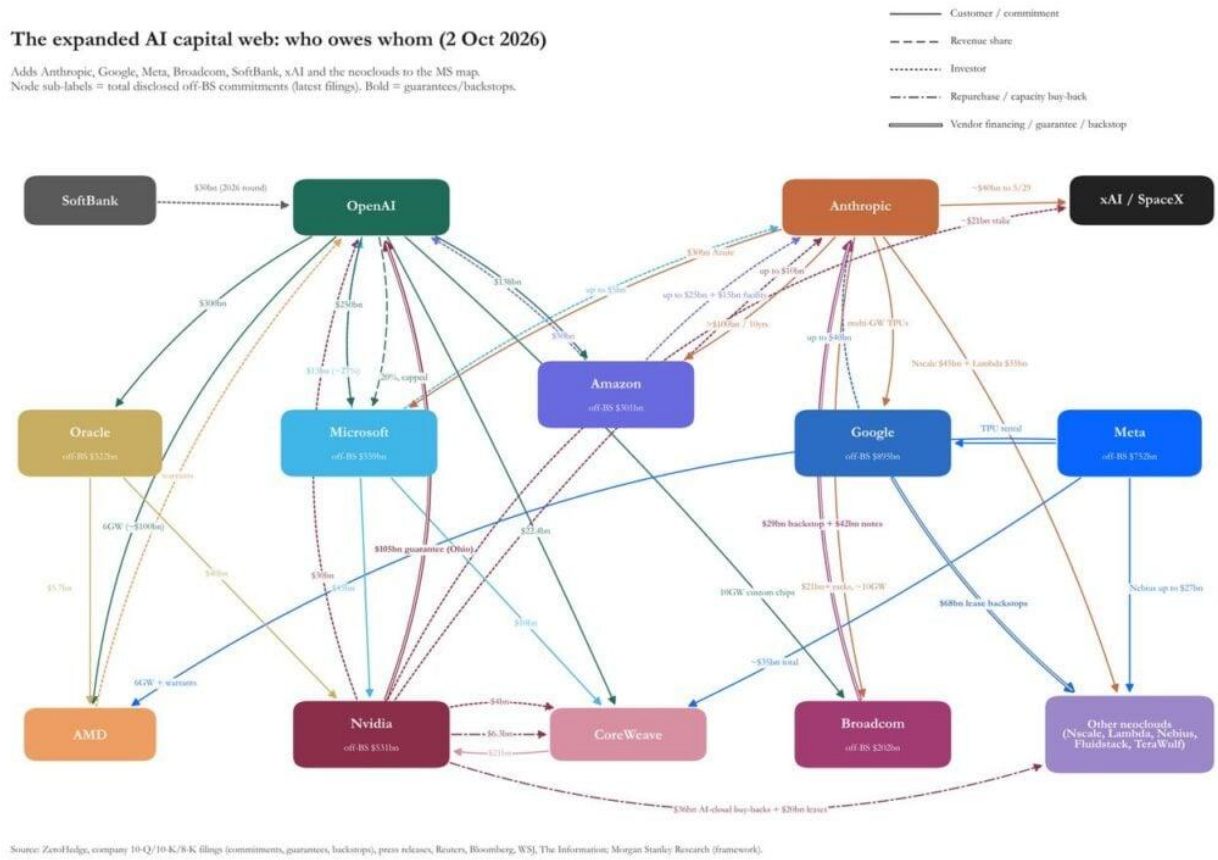

Our AI giants are tied together by commitments, guarantees, backstops and vendor financing, much of it off balance sheet, so one weak link strains them all. Third-party chart: "The expanded AI capital web: who owes whom" (on-chart sources: ZeroHedge, company filings, Morgan Stanley Research framework), used with permission.[56]

That weak link is OpenAI, a cash furnace, IPO delayed [Geomatrix].

[56] **AI capital web source.** *Source:* "The expanded AI capital web: who owes whom (2 Oct 2026)"; on-chart source line: ZeroHedge, company 10-Q/10-K/8-K filings (commitments, guarantees, backstops), press releases, Reuters, Bloomberg, WSJ, The Information; Morgan Stanley Research (framework). The on-chart source line lists ZeroHedge first. From Robert Howe's file 'AI and the 2-Year Teaser Wall chart 9.docx', saved Oct 17, 2026 (HST). Third-party chart; used with permission. *Caveats:* node figures are the chart's own totals of disclosed off-balance-sheet commitments; we have not re-checked them against the filings.

CHART 12

The Reset Wall, Then and Now

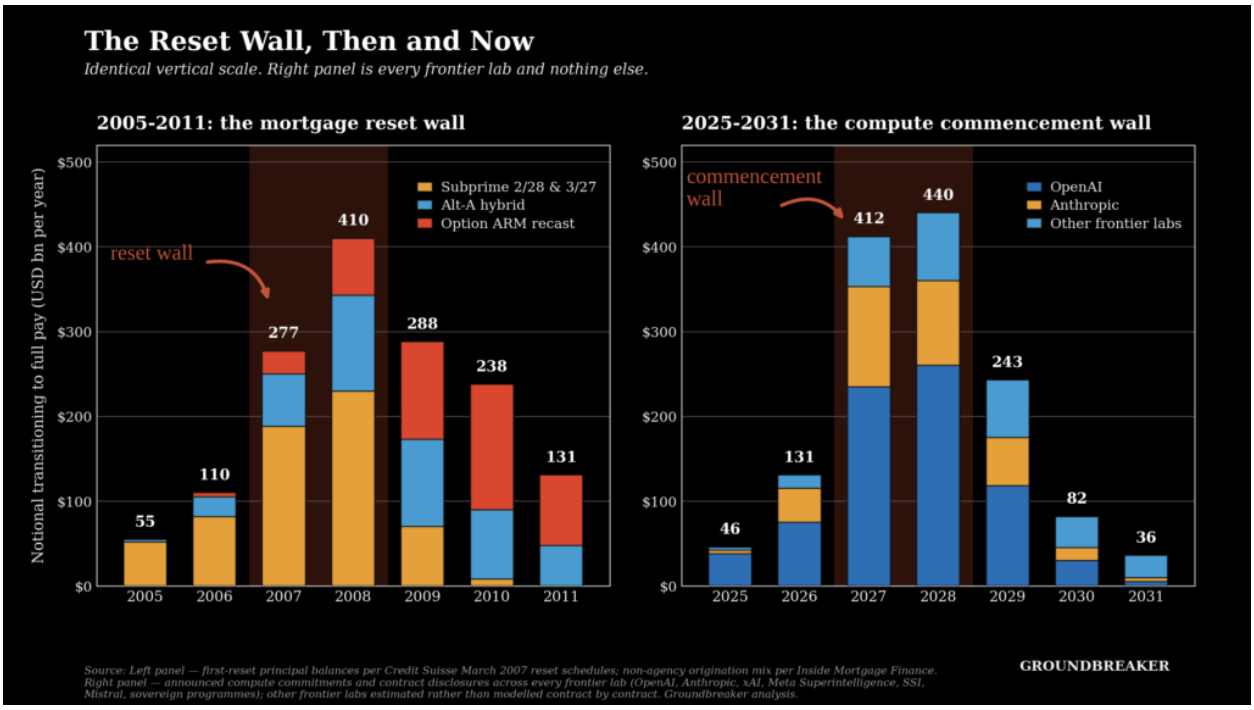

Our frontier labs' compute commencement wall, $412bn in 2027 and $440bn in 2028, matches the 2008 mortgage reset peak of $410bn. Third-party chart: Groundbreaker, "The Reset Wall, Then and Now", used with permission.[57]

Groundbreaker's essay "The Teaser Period: Why the AI Boom Is Hitting a Reset Wall", which we shared on LinkedIn, frames the risk:[58]

> *The bull case wins if - and it is a real if - demand scales into the committed supply before the reset wall lands, and the counterparties stay funded through any air pocket in between. The bear case in this piece is not that artificial intelligence will fail, or that the demand is fake, or that the technology disappoints. It is narrower: that the financing structure can break before the demand arrives, because the obligations are fixed and front-loaded in commencement, while the revenue is variable and back-loaded in adoption - and a fixed obligation meeting a lagging revenue stream is a solvency problem, regardless of how transformative the underlying technology turns out to be.*
>
> *The industry will spend the next eighteen months debating whether artificial intelligence is a bubble, which is the wrong question, asked at the wrong layer. The technology is real; so were the houses. The question is narrower: what happens when instruments underwritten at the teaser meet their reset schedule, and who is holding the paper when the obligations*

[57] **Reset Wall source.** *Source:* Groundbreaker, "The Reset Wall, Then and Now" (Groundbreaker analysis), from Robert Howe's file 'AI and the 2-Year Teaser Wall chart 8.docx', saved Oct 17, 2026 (HST). Third-party chart; used with permission. ***Note:*** on-chart source: left panel, first-reset principal balances per Credit Suisse March 2007 reset schedules; non-agency origination mix per Inside Mortgage Finance. Right panel, announced compute commitments and contract disclosures across every frontier lab (OpenAI, Anthropic, xAI, Meta Superintelligence, SSI, Mistral, sovereign programmes); other frontier labs estimated rather than modelled contract by contract. ***Alternate measure:*** an earlier Groundbreaker version saved Oct 15, 2026 ('Bar chart 2-year teaser loans versus AI Financing.jpeg') shows the whole AI stack, consolidated and net of eliminations: $732bn in 2027 and $820bn in 2028.

[58] **Quoted source.** Groundbreaker, 'The Teaser Period: Why the AI Boom Is Hitting a Reset Wall', Substack, 20 Aug 2026, https://www.groundbrkr.com/p/the-teaser-period-why-the-ai-boom ; shared in our LinkedIn post of 9 Oct 2026, https://www.linkedin.com/feed/update/urn:li:activity:7514224427975213056/ . Quoted with permission; the views quoted are the author's.

cannot be met as written. The reset wall is published above and the AI boom sits in a period of fiction.

AKAMAI Recommended Strategy for AI Bubble Pop

Fallout from the tech bubble pop? From the Nasdaq's March 2000 peak to year-end, the Nasdaq fell some 50%, while the Dow rose some 9%.[59] Money rotated into cheap, neglected value stocks, staples and retailers, some on single-digit P/Es.[60] Today the same stocks sell on 20-50x earnings,[61] and staples and retail are hit 3 ways: store brands taking share,[62] GLP-1 drugs shifting demand from high-margin processed food to low-margin fresh,[63] and the wealthy top of the K-shaped economy, owners of the AI-heavy S&P, spending less as portfolios shrink.[64]

Many investors are swapping equities for bonds, drawn by safe US Treasury yields above 5% (10-year and longer)[65] and the chance of capital gains as rates fall in a recession. We are an equity shop, not a balanced manager allocating this way.

Where is the safe ice to skate to on the cracking market pond? Three areas. First, we are long the USERS of AI: white-collar, repetitive-task industries like insurance and banking, mostly on single-digit to low-teens P/Es and low price-to-book. They underperform whenever the bubble inflates and rebound nicely as it pops.[66] Second, we are long biotech, a huge beneficiary of AI in drug discovery, breakthroughs vital to aging boomers worldwide.

Third, we are long many non-US markets where AI helps and the bubble has not stretched valuations: Greece and Chile and their banks, and, believe it or not, Chinese banks, balance sheets repaired, with

[59] **2000 rotation.** *Data:* price change from the Nasdaq record close (Mar 10, 2000) to Dec 29, 2000: Dow Jones Industrial Average +8.7%, Nasdaq Composite -51.1%; Berkshire Hathaway A +71.9%, Clorox +18.8%, Coca-Cola +33.9%, Walmart +10.8%. To Dec 31, 2001: Dow +0.9%, Nasdaq -61.4%; Berkshire +83%, Clorox +32%, Coca-Cola +4%, Walmart +20%. *Caveats:* over calendar 2000-01 (Dec 31, 1999 to Dec 31, 2001) the Dow fell 12.8% and the Nasdaq 52.1%, so the Dow's gain holds only from the Nasdaq peak. Prices are split-adjusted, before dividends. *Sources:* daily closes via Yahoo Finance (^DJI, ^IXIC, BRK-A, CLX, KO, WMT), https://finance.yahoo.com/ , retrieved Oct 17, 2026 (HST).

[60] **2000 valuations.** *Data:* in early 2000 Philip Morris (Kraft foods, Miller beer, Marlboro) traded at about 5-8x earnings and R.J. Reynolds at about 5x; Kmart traded at 6.5x in May 2000. Not all were cheap: the same May 2000 note put Wal-Mart at 38x and Target at 23x, so we say "some". *Sources:* Floyd Norris, The New York Times, Feb 18, 2000, https://archive.nytimes.com/www.nytimes.com/library/financial/columns/021800norris-column.html ; New York Post, Feb 8, 2000, https://nypost.com/2000/02/08/stock-watch-philip-morris-may-be-bargain/ ; Computerworld, May 29, 2000, https://www.computerworld.com/article/1375233/kmart-belatedly-ready-to-make-online-play.html

[61] **Value-stock valuations.** *Data:* trailing P/E at the Oct 16, 2026 close: Costco 47.6x, Walmart 43.1x, Coca-Cola 28.2x, Procter & Gamble 24.1x, Clorox 18.1x (just under the range). Berkshire is left out because its GAAP earnings include investment gains. *Sources:* stockanalysis.com statistics pages (data: S&P Global Market Intelligence), https://stockanalysis.com/stocks/cost/statistics/ ; /wmt/ ; /ko/ ; /pg/ ; /clx/ , accessed Oct 17, 2026 (HST).

[62] **Store brands.** *Data:* US store-brand dollar share 17.1% in supermarkets in 2007 (PLMA 2008 Private Label Yearbook, Nielsen), against a record 21.3% across all outlets in 2025, with unit share 23.5% and sales of $282.8 billion (PLMA, Circana Unify+). Caveats: channel coverage differs between the two sources. *Sources:* https://www.foodprocessing.com/business-of-food-beverage/business-strategies/article/11368546/business-store-brand-sales-increase-by-54-billion-in-2008 ; https://www.plma.com/about_industry/store_brand_facts

[63] **GLP-1 and food demand.** *Data:* households with a GLP-1 user cut grocery spending 5.3% within six months of adoption (higher-income households 8.2%), with the biggest falls in calorie-dense processed categories (savory snacks -10.1%) and small rises in yogurt, fresh fruit, nutrition bars and meat snacks; 16.3% of US households had a GLP-1 user by July 2024. Hristakeva, Liaukonyte and Feler (Cornell; Numerator panel), Journal of Marketing Research 63(3), 2026. *Sources:* https://doi.org/10.1177/00222437251412834 ; https://news.cornell.edu/stories/2025/12/ozempic-changing-foods-americans-buy

[64] **K-shaped consumption.** *Data:* the top 10% of US households by income accounted for 49.2% of consumer spending in Q2 2025 (Moody's Analytics, Mark Zandi), revised to 45.8% for 2025 after a change in method. The top 10% by wealth held 88.1% of US corporate equities and mutual fund shares in Q2 2026 (top 1% 50.9%, next 9% 37.2%; Federal Reserve Distributional Financial Accounts). *Sources:* Bloomberg, Sep 16, 2025, https://www.bloomberg.com/news/articles/2025-09-16/top-10-of-earners-drive-a-growing-share-of-us-consumer-spending ; Bloomberg Opinion, Apr 7, 2026, https://www.bloomberg.com/opinion/articles/2026-04-07/the-k-shaped-economy-s-defining-statistic-has-some-problems ; FRED series WFRBST01122 and WFRBSN09149, https://fred.stlouisfed.org/series/WFRBSN09149 , retrieved Oct 18, 2026.

[65] **US Treasury yields.** *Data:* Treasury par yield curve on Oct 16, 2026: 10-year 5.09%, 20-year 5.47%, 30-year 5.43%; shorter maturities sit below 5% (3-month 4.26%, 2-year 4.74%, 5-year 4.89%, 7-year 4.98%). Caveats: "above 5%" holds from the 10-year benchmark outward; bills and notes out to 7 years yield less. Long bonds gain most in price when rates fall. *Sources:* U.S. Department of the Treasury, Daily Treasury Par Yield Curve Rates, https://home.treasury.gov/resource-center/data-chart-center/interest-rates/TextView?type=daily_treasury_yield_curve&field_tdr_date_value_month=202610 , accessed Oct 17, 2026 (HST). Rechecked Oct 18, 2026 (HST) against the same Treasury table: Oct 16, 2026 is the latest published date, and the 10-year yield has closed above 5% on every trading day of October 2026 (5.09% to 5.31%).

[66] **Bank and insurer valuations.** *Data:* KBW bank ETF (KBWB, 25 large US banks) trailing P/E 13.0x, forward 11.4x (Oct 5, 2026); US insurance industry P/E 12.0x vs a 3-year average of 15.5x (Oct 18, 2026); ICBC 6.5x trailing P/E and 0.66x book (Oct 16, 2026). Caveats: US banks and insurers sit in the low teens rather than single digits, so we say "single-digit to low-teens"; single digits are found mainly outside the US. *Sources:* https://etf-metrics.com/ticker/kbwb ; https://simplywall.st/markets/us/financials/insurance ; https://stockanalysis.com/quote/hkg/1398/statistics/ (S&P Global Market Intelligence data).

among the world's lowest interest rates to compete for emerging-market lending.[67] Our biggest long is Japan, getting ROE religion,[68] with a few tech companies that are known or emerging dominators in IT hardware.

A correction or a bear market may have started, but who knows until after? Baby Bears come from short-lived events: war, pandemic, elections perceived bad for profits. Mama Bears come from higher inflation and rates, i.e., that 70s show. Papa **B**ears come from bad loans that infect the banking system.[69] The good news: only a few financial institutions, BlackRock and Blackstone, have loan exposure to the AI ecosystem.[70] We don't expect a Papa. Our view, not a forecast, is a Mama Bear of -20% to -30%, the damage mostly in tech but some contagion as ETFs are sold in fear and allocation to bonds. Not straight down, of course, and with the sectors benefiting from AI recovering holding up and recovering as profits prove our case.

Indeed, beneficiaries will be all of us: every tech breakthrough has improved living standards, despite some labor market dislocations. We are not Skynet Armageddon believers, just confident that many data centers will become white elephants like dark fiber[71] and parallel railways,[72] while our health and productivity improve and our portfolios recover soon enough. Human connection and kindness? We'll see and do our part.

Thanks for reading today, and for supporting AKAMAI. Click here to register for our monthly 2-pager with occasional special reports, most far shorter than this one! Cancel any time. Website www.geomatrix-investment.com. Mahalo.

Important information

This document is issued by Geomatrix (H.K.) Limited (SFC CE No. ADJ583), licensed by the Securities and Futures Commission of Hong Kong for Type 4 (advising on securities) and Type 9 (asset management) regulated activities. It is intended solely for professional investors as defined in the Securities and Futures Ordinance (Cap. 571) and must not be distributed to the public. It is for information only and is not an offer or solicitation to buy or sell any security or any interest in the AKAMAI Master Fund, nor investment advice. Figures marked as estimates are our calculations from third-party sources we believe reliable but have not independently verified; all

[67] **Greek, Chilean and Chinese banks.** *Data:* Greece's four systemic banks average 10.2x 2026 P/E and 1.5x P/TBV vs 11.9x and 2.2x for European banks (Beta Securities, Jul 2, 2026). Banco de Chile 8.9x trailing P/E, 1.65x book; Banco Santander-Chile 13.4x, 2.9x book (Oct 16, 2026). ICBC 6.5x P/E, 0.66x book, 4.8% dividend yield. China's 10-year government bond yields 1.65% (Oct 16, 2026), against 0.57% in Switzerland, 2.95% in Japan and 5.09% in the US, so we say "among the lowest". *Caveats:* Chilean banks trade at 1.7-2.9x book, richer than Greek or Chinese banks, on ROE of 23-25%. *Sources:* https://www.euro2day.gr/en/article/2358883/which-listed-companies-combine-earnings-with-a-low ; https://stockanalysis.com/stocks/bch/statistics/ ; https://stockanalysis.com/stocks/bsac/statistics/ ; https://stockanalysis.com/quote/hkg/1398/statistics/ ; https://fxmacrodata.com/rates/gov-bond-10y (national debt offices), accessed Oct 18, 2026.

[68] **Japan ROE push.** *Data:* in March 2023 the Tokyo Stock Exchange asked Prime and Standard companies to manage with cost of capital and share price in mind. By March 2026, Prime companies trading below 1x book fell to 27% (-23 points since July 2022), and 93% of Prime companies had disclosed plans. *Sources:* TSE, Mar 31, 2023, https://www.jpx.co.jp/english/equities/follow-up/uorii50000004sse-att/uorii50000004tcv.pdf ; TSE, Apr 7, 2026, https://www.jpx.co.jp/english/equities/follow-up/b5b4pj000004yqcc-att/vk0khi00000151mi.pdf

[69] **Types of bear market.** *Data:* Goldman Sachs sorts US bear markets since 1835 into event-driven (one-off shocks; average fall 29%, about 9 months), cyclical (rising rates, recession, falling profits; 31%, about 27 months) and structural (imbalances and financial bubbles; 57%, about 42 months). Our baby, mama and papa bears map loosely onto these. *Sources:* Goldman Sachs, Bear Market Anatomy, Apr 8, 2025, https://www.gspublishing.com/content/research/en/reports/2025/04/08/0bf285f5-8d4a-478c-843f-4b4ea81256d5.html ; MarketWatch, Mar 11, 2020, https://www.marketwatch.com/story/goldman-sachs-analyzed-bear-markets-back-to-1835-and-heres-the-bad-news-and-the-good-about-the-current-slump-2020-03-11

[70] **Blackstone and BlackRock.** *Data:* Blackstone owns QTS (US$10bn take-private, 2021) and AirTrunk (A$24bn), with $70bn of data centers and over $100bn in its development pipeline, and led a large private debt financing for CoreWeave. BlackRock's Global Infrastructure Partners and the AI Infrastructure Partnership (with Microsoft, Nvidia and MGX) closed the $40bn purchase of Aligned Data Centers in July 2026. The Chicago Fed finds direct bank exposure generally modest. *Caveats:* much of this is equity and private credit rather than bank loans, and banks still arrange AI debt (JPMorgan and MUFG took months to place $38bn of Oracle-linked data-center debt). *Sources:* https://www.blackstone.com/insights/article/behind-the-deal-airtrunk/ ; https://www.esgdive.com/news/blackrocks-gip-microsoft-backed-ai-group-buy-aligned-data-centers-for-40/825920/ ; https://www.chicagofed.org/publications/chicago-fed-insights/2026/ai-tail-risk-for-banks ; https://finance.yahoo.com/technology/ai/articles/banks-arranged-38-billion-ai-152113770.html

[71] **Dark fiber.** *Data:* in 2002 only 2.7% of installed US fiber was in use (TeleGeography), and Salomon Smith Barney put lit fiber at about 5%. Bandwidth prices fell an average of 65% a year in 2000-02 and most long-haul carriers filed for Chapter 11. Caveats: some intercity routes ran at 70%+ of lit capacity (TeleChoice, 2001). *Sources:* The Wall Street Journal, Sep 26, 2002, https://www.wsj.com/articles/SB1032982764442483713 ; The New York Times via https://moglen.law.columbia.edu/CPC/archive/barbecue/18MELT.html ; Light Reading, https://www.lightreading.com/cable-technology/fiber-utilization-figures-challenged

[72] **Parallel railways.** *Data:* Mark Casson estimates that Victorian Britain's 1914 network could have given the same service with about two-thirds of the track, the rest being duplicated routes and excess branches built partly to block rivals in and after the 1840s Railway Mania. *Sources:* Mark Casson, The World's First Railway System (Oxford University Press, 2009), https://doi.org/10.1093/acprof:oso/9780199213979.001.0001 ; Andrew Odlyzko, https://www-users.cse.umn.edu/~odlyzko/doc/mania08.pdf

figures and opinions are as of 17 October 2026 and may change without notice. Quoted third-party material reflects its authors' views. Forecasts are not guarantees. Past performance is not indicative of future results, and the value of investments can go down as well as up.