

A severe case of COVIDIA: prognosis for an AI-driven US equity market

NVIDIA and its GPU customers are now a large driver of equity market returns, earnings growth, earnings revisions, industrial production and capital spending. NVIDIA’s financial results are extraordinary (it beat on revenues and earnings again last week, and announced a \$50 billion stock buyback). As shown below, NVIDIA has also experienced the fastest road to being the market’s largest stock in the post-war era.

But for investors, the more important questions look past the economics of selling GPUs and focus on the ability of hyperscalers (Google, Amazon, Microsoft, Meta, etc) and other AI infrastructure users/providers to earn adequate returns on hundreds of billions in AI-related capital spending. The level of this spending now rivals the mainframe era of the late 1960’s and the fiber optic deployment of the late 1990’s. For adequate returns on AI infrastructure to materialize, within the next 12-18 months we will need to see a greater shift in favor of “inference” tasks (AI used to run production models for corporate customers) rather than GPU capacity primarily being used to train foundational models and chatbots. In this Eye on the Market, we take a closer look.

Also: some brief comments on the Third Circuit’s decision to strike down Section 230 protections for certain social media company activities.

Michael Cembalest
JP Morgan Asset Management

Table of contents

COVIDIA: GPU orbit now driving equity markets and capital spending 2

NVIDIA financials: phenomenal, and the antithesis of the dot-com era 4

From A (Altman) to Y (Combinator): the transformational view of generative AI 5

And yet...there are a lot of questions on spending, revenues, productivity benefits, energy and more 6

Tracking corporate sector AI adoption so far: mostly positive but some mixed signals 8

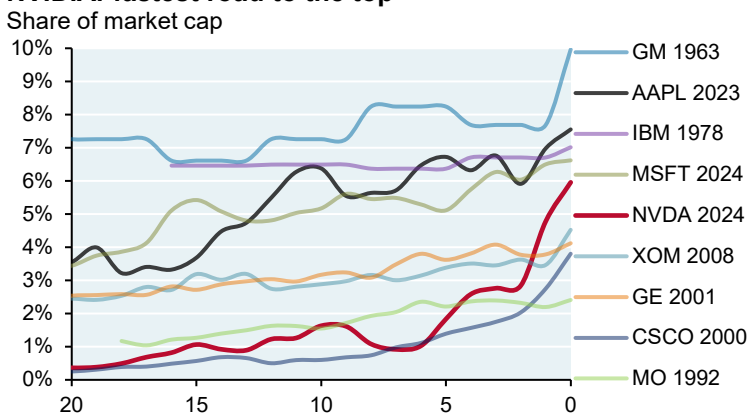
What are equity markets pricing in? 9

COVIDIA prognosis: judgment day on AI capital spending is probably 2-3 years away 10

Appendix I: Current and prospective NVIDIA competitors 11

Appendix II: The fate of prior equity market leaders 12

NVIDIA: fastest road to the top



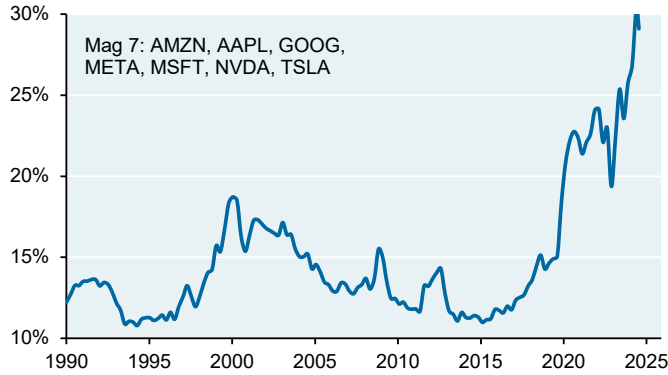
Source: Bloomberg, Factset, GS, JPMAM, August 2024

COVIDIA: GPU orbit now driving equity markets and capital spending

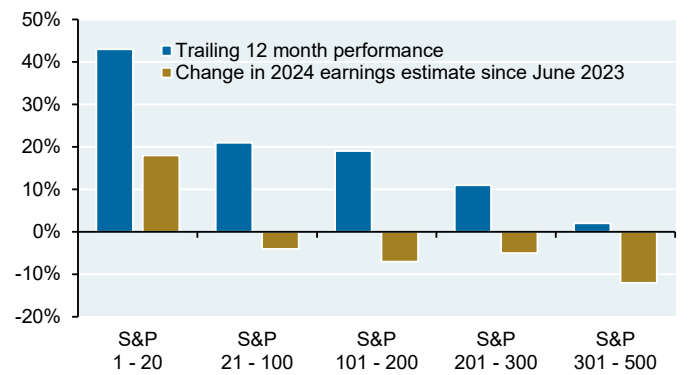
NVIDIA and its GPU customers are an increasingly large driver of equity market returns, earnings growth, earnings revisions, industrial production and capital spending. You have probably seen the first chart before on the Mag 7 share of market cap. The others are just as notable, highlighting how earnings, industrial production and capital spending are increasingly impacted by AI related spending. Some colleagues pointed out recently that Mag 7 capital spending now exceeds that of the entire energy sector.

Magnificent 7 domination

% of total large-cap market cap



Top quintile of stocks are driving earnings revisions and market returns, Percent



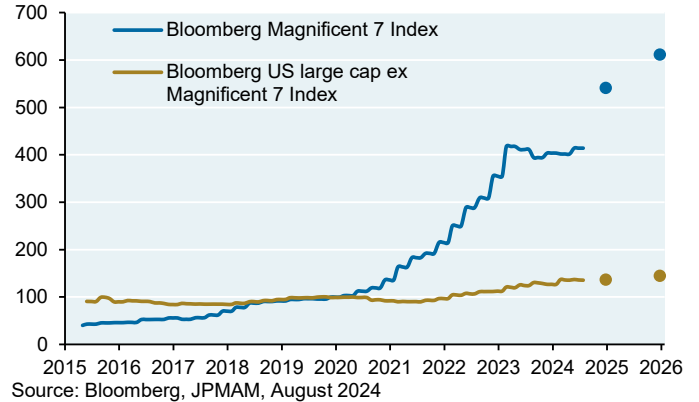
US industrial production of high tech industries

Index (100 = January 2020)



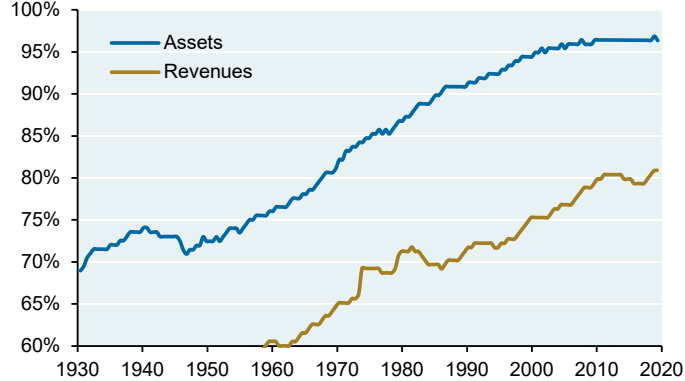
Capex + R&D, Magnificent 7 vs rest of market

Index (100 = January 2020)



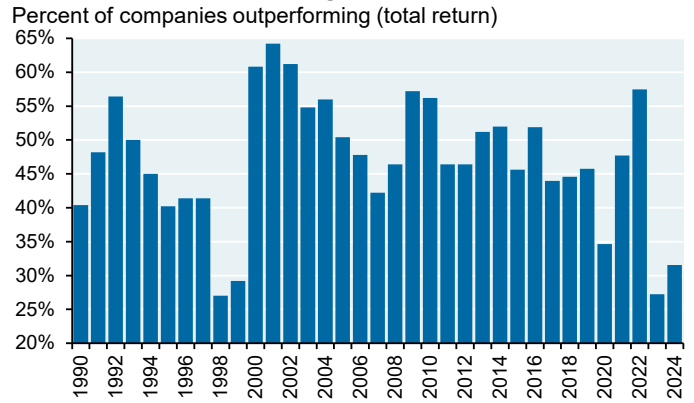
US corporate concentration has been steadily increasing for the last 100 years; see the chart on the share of assets and revenues controlled by the largest 1% of companies. But we have now hit new extremes, exemplified by the collapse in the share of companies outperforming the S&P 500, and the outperformance of the market cap weighted S&P 500 vs the equal weighted version of the same index (last reached in the late 1990's).

Largest 1% of US corporations share of assets and revenues, Percent



Source: "100 years of corporate concentration", Kwon et al., August 2024

Constituents outperforming the S&P 500



Source: Bloomberg, JPMAM, August 16, 2024

Rolling 2 year performance of market cap vs equal weighted S&P 500



Source: Bloomberg, JPMAM, August 2024

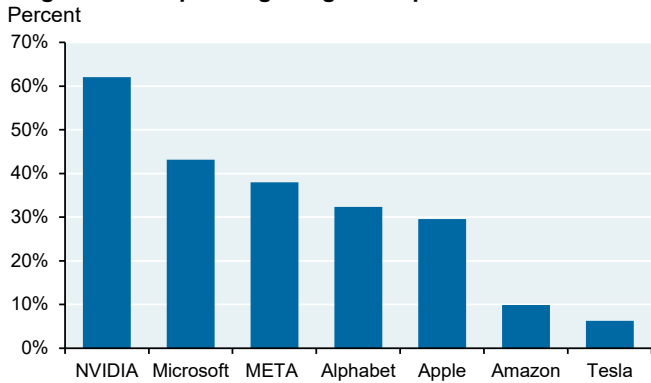
Third Circuit removes some Section 230 protections from social media companies

- This year, in *Moody v. NetChoice LLC*, the Supreme Court ruled that the use of algorithms by social media or internet companies to automatically prioritize or otherwise curate content that is posted by third parties reflects that company’s editorial judgment and constitutes “expressive activity” which is protected by the First Amendment. In other words, your TikTok, Twitter or LinkedIn feed constitutes that company’s own speech, even though the content was originally created and uploaded by third parties. So...the Court essentially concluded that social media companies have a First Amendment right to use algorithms to boost or downplay posts based on whatever political, social or economic standards they wish. The decision implied that a recently passed law in Texas may be unconstitutional, and remanded the case back to lower Texas courts
- Fine, said the Third Circuit: **if curated feeds are first-party speech, then Section 230 of the Communications Decency Act does not shield these companies from liability for harm caused by content that their algorithms recommend to users.** Thus, in *Anderson v. TikTok*, the court allowed a lawsuit against TikTok to proceed. Its algorithm is alleged by plaintiffs to have recommended a video featuring the “blackout challenge” to a child who died while performing it
- This all makes perfect sense to me; let the lawsuits begin

NVIDIA financials: phenomenal, and the antithesis of the dot-com era

- NVIDIA’s operating margins are high and rising. The company has seen a step-change in projected revenue growth over the last year (third chart, from gold dots to red dots). Last week, the company beat expectations on revenues and earnings, and announced a \$50 billion stock buyback program
- NVIDIA bears no resemblance to dot-com market leaders like Cisco (see below) whose P/E multiple also soared but without earnings to go with it
- Consensus analyst reports project that NVIDIA will maintain 90%+ market share in AI chips for the next two years. Even with a possible delay in NVIDIA’s next generation “Blackwell” chip shipments, few if any firms are positioned to take immediate advantage...but competition is coming. **See Appendix I on current and prospective NVIDIA competitors, and Appendix II for history since the 1960’s on what happened to stocks once they represented the largest market cap, as NVIDIA did for one brief day in June 2024**

Magnificent 7 operating margin comparison



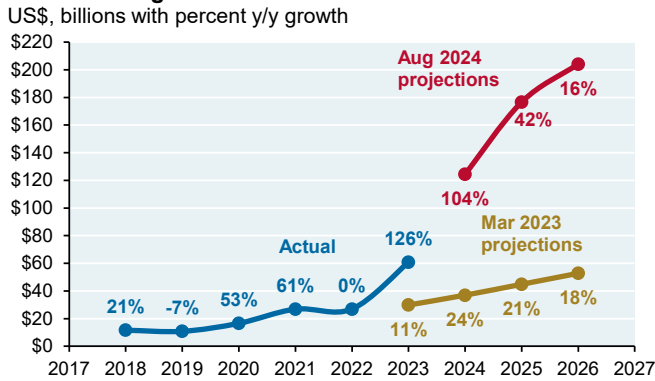
Source: Bloomberg, JPMAM, July 2024

NVIDIA net income margins



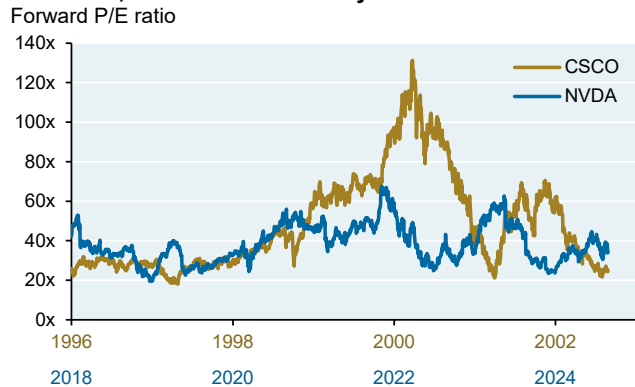
Source: Bloomberg, JPMAM, July 2024

Nvidia sales growth: actuals and estimate revisions



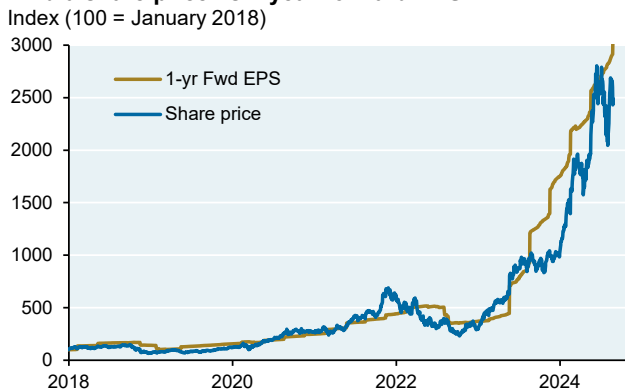
Source: Bloomberg, JPMAM, August 30, 2024

P/E ratios, dot-com era vs today



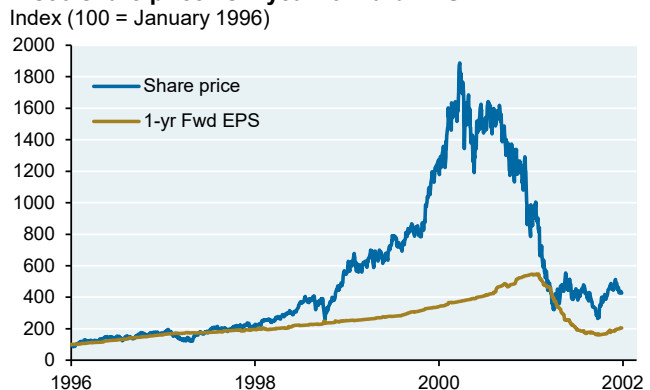
Source: Factset, JPMAM, August 30, 2024

Nvidia share price vs 1-year forward EPS



Source: Factset, Bloomberg, JPMAM, August 30, 2024

Cisco share price vs 1-year forward EPS



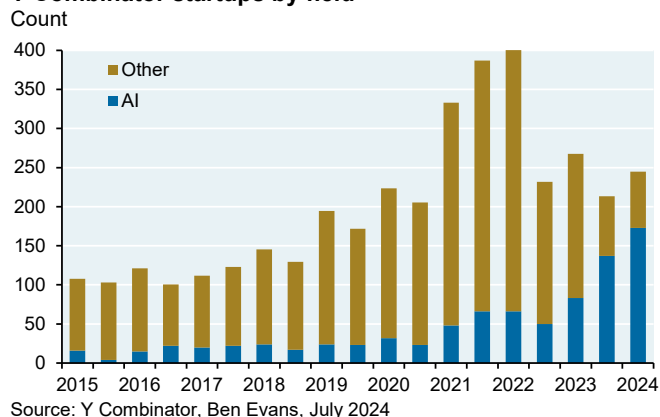
Source: Factset, Bloomberg, JPMAM, June 25, 2024

From A (Altman) to Y (Combinator): the transformational view of generative AI

The prognosis for COVIDIA is less related to NVIDIA financials and more related to whether/when the AI transformation markets are pricing in will actually happen. AI adherents make powerful arguments:

- Sam Altman at OpenAI considers AI the “biggest, best and most important of all technology revolutions”, and believes it will become more ubiquitous over time¹
- According to McKinsey (never the types to shy away from seismic predictions against which they’re never measured), generative AI could eventually add ~\$8 trillion to the global economy each year.²
- VC accelerator Y Combinator is “all-in” with respect to AI, and its track record is impressive. Jared Heyman’s Rebel Fund estimated returns on every Y Combinator deal since 2005. If you had invested in each (which is practically impossible), your annual average return would have been 176% net of dilution³
- AI is driving the VC ecosystem, accounting for 40%+ of new unicorns in the first half of 2024 and 60%+ of the increase in VC-backed valuations; US unicorn value grew by \$162 billion so far in 2024 (Pitchbook)
- AI is advancing so rapidly that performance benchmarks for reading comprehension, image classification and advanced math are becoming obsolete, with new ones replacing them (see box on GPQA benchmark)
- Generative AI may not just improve productivity of call centers, coders and professional writers; a Google study projects that ~80% of jobs could see at least 10% of tasks done twice as quickly⁴
- Elon Musk’s new xAI startup raised \$6 billion at a valuation of \$24 billion this year, with plans to build the world’s largest supercomputer in Tennessee to support AI training and inference efforts. Musk has been further burnishing his libertarian credentials by sharing fake AI-generated videos of VP Kamala Harris⁵
- Then there’s the mind-altering 165-page “*Situational Awareness*” paper by Aschenbrenner which takes AI predictions to another level (superintelligence, 20% of US power supply, \$1 trillion and full AGI by 2030)⁶
- Rob Toews at Radical Ventures describes AlphaFold and other protein folding algorithms as the most important achievements in the history of AI (see April 2 2024 EoTM, page 3 for more details)

Y Combinator startups by field



The Google-proof Q&A benchmark (GPQA)

- Consists of 400+ multiple choice questions
- Even PhD-level scholars correctly answer questions in their field only 65% of the time. The same scholars scored 34% outside their field even with unrestricted access to the internet
- In 2023, most AI models scored 30%-40% on GPQA. Claude 3 (the latest Anthropic chatbot) scored ~60%
- HOWEVER...this is still just an example of excellence in memorization and the problem of “contamination”, since models are first trained on data they’re tested on

¹ MIT Artificial Intelligence, May 2024

² “The economic potential of generative AI: The next productivity frontier”, McKinsey & Company, June 2023

³ “On the 176% annual return of a YC startup index...and why one will never exist”, Jared Heyman, July 2023. Findings: 5%-6% of YC startups became \$ 1 bn unicorns, and of those 10-12% became decacorns (10x more than unicorns). Unfortunately, it isn’t possible to invest in every Y Combinator deal since companies hand-pick investors in oversubscribed start-up rounds. In addition: those 176% average annual returns are concentrated in a subset of deals. Y Combinator’s deal failure rate is high, ~40% during its first 17 batches, and that excludes companies that are struggling or were sold at a loss

⁴ “The economic impact of generative AI”, Andrew McAfee, Visiting Fellow / Technology & Society, Google, April 2024

⁵ “Recipe for disaster: Experts issue warning after Musk shares fake AI video of Harris”, Salon.com, July 2024

⁶ In “*Situational Awareness: the Decade Ahead*” (June 2024), Leopold Aschenbrenner argues that he is among a few hundred people in the world who truly understand the scope and magnitude of the AI transformation. He explains how hundreds of millions of AI systems can automate AI research, resulting in “superintelligence” which would be the most powerful weapon that mankind has ever built...and that the democratic world would not survive China getting true AGI first

And yet...there are a lot of questions on spending, revenues, productivity benefits, energy and more

- Will OpenAI lose \$5 billion this year as reported by The Information?⁷
- Will it really cost \$100 billion to train a single AI model by 2027, a figure cited by Anthropic’s CEO?⁸
- Will the largest tech companies be able to come up with \$500 billion in annual “missing revenues” to break even vs their AI data center spend, a figure estimated by David Cahn at Sequoia?

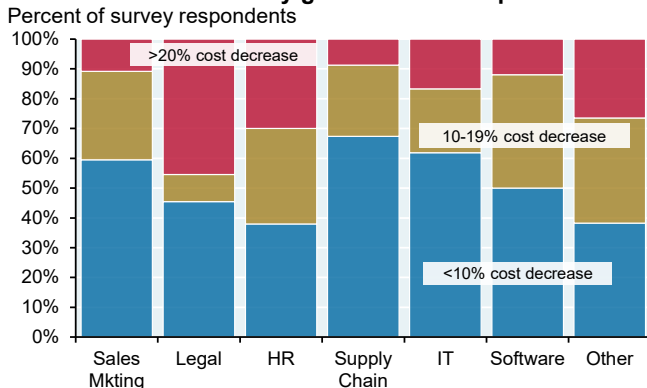
Sequoia: Will hyperscalers and enterprise software companies earn enough to pay for AI infrastructure?

Cahn takes Nvidia’s revenues and multiplies by 2x to obtain total cost of data centers, since GPUs are roughly half the cost of data center ownership (the rest is energy, buildings, generators, etc). Cahn then multiplies by 2x again to reflect the expected 50% gross margins of major hyperscalers (i.e., Microsoft, Amazon and Google which provide cloud services via Azure, Amazon Web Services and Google Cloud Platform). Cahn assumes that Google, Apple, Microsoft and Meta will generate \$10 billion each year from new AI-related revenue, and that Oracle, ByteDance, Alibaba, Tencent, X and Tesla earn another \$5 billion each. The gap between new AI revenues and AI data center spend: \$500 billion per year

“AI’s \$600 billion question”, David Cahn, Sequoia Capital, June 2024

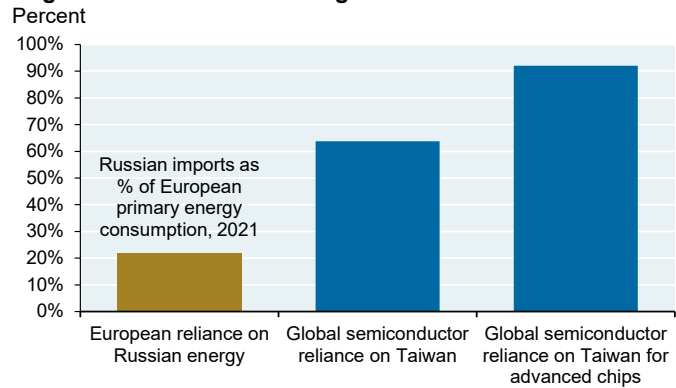
- Are hyperscalers and other cloud/enterprise companies really building enough AI compute infrastructure to power 12,000 different ChatGPTs, assuming its current size of 200 mm daily active users asking 1.25 queries per day, as estimated by Barclays?⁹ Barclays might be substantially underestimating the query intensity and user base of widely adopted generative AI applications, particularly if such applications entail query threads that require much more computing power...but you get the point
- Barclays estimates that in 2024, enough GPUs were built to generate ~\$100 bn of revenues at maximum utilization rates. Actual 2024 payments by GPU users: ~\$10 bn. How long will it take for this gap to close?
- Why do most respondents in a NewStreet Research survey cite cost decreases of just 10% or less after generative AI adoption (see chart below)?
- Why do some large enterprises whose engineers use AI coding tools still not see improvements in cycle time end-to-end software delivery? Note: only ~30% of a software developer’s time is spent coding
- When do markets start thinking more about geopolitical risk given TSMC’s current role as sole supplier of advanced chips to NVIDIA?¹⁰ Remember Trump’s comments in July: “Taiwan should pay us for defense, I wouldn’t feel so good right now if I were them”...which caused an immediate decline in semi stocks

Cost decreases driven by generative AI adoption



Source: NewStreet Research, August 2024

Regional reliance for critical goods



Source: BP Statistical Review, ROC Taiwan, Global Guardian, 2024

⁷ "Why OpenAI Could Lose \$5 Billion This Year", The Information, Amir Efrati and Aaron Holmes, July 2024

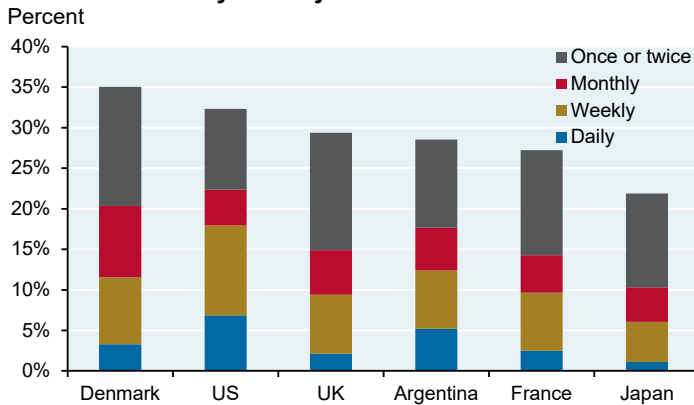
⁸ "In Good Company with guest Dario Amodei", a podcast by Norges Bank Investment Management, June 2024

⁹ "Cloud AI Capex: FOMO or Field-Of-Dreams?", Ross Sandler, Barclays Equity Research, June 2024

¹⁰ As of August 2023, NVIDIA was sourcing all advanced AI chips from TSMC. In July 2024, NVIDIA announced that it would start diversifying and begin sourcing advanced memory chips from Samsung as well

- What if Daron Acemoglu at MIT is right that that the annual boost to total factor productivity from AI will be just 0.06% per year, which is quite small given ~0.6% average annual TFP growth over the last 20 years?¹¹
- Why do so many Chat GPT users access it infrequently, like me?

Use of ChatGPT by country

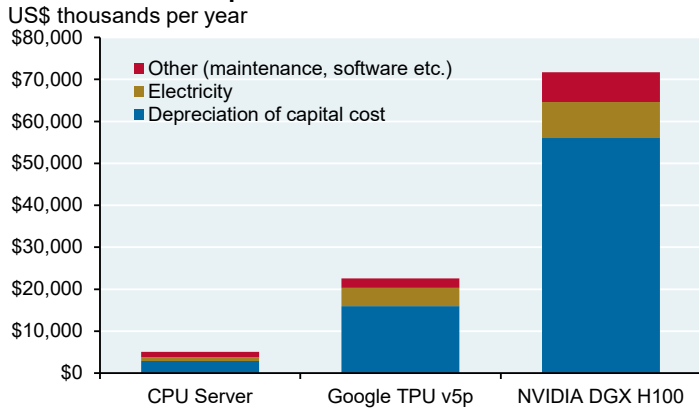


Source: Reuters, Ben Evans, May 2024

“What was I made for?” GPT and the EOTM
 In the September 2023 Eye on the Market (“*What Was I Made For*”), I explained why I’m not a frequent GPT user. I asked GPT-4 70 questions from the Eye on the Market archives on markets, economics, politics and energy. GPT-4’s aggregate grade: C+. It got half the questions right but the other half wrong despite being able to web crawl, open pdfs, etc, and its mistakes were not predictable, requiring me to check every answer. Limited if any productivity benefits for me if that’s the case

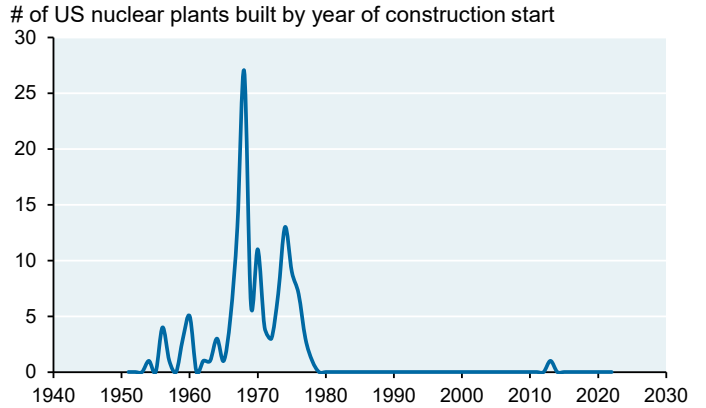
- What about energy consumption? There’s already a dispute over Talen Energy’s partial sale of its nuclear generation to Amazon. AEP and Exelon filed a complaint with FERC citing \$140 million of transmission costs now borne by PJM electricity customers. And remember: the advent of cloud computing barely changed electricity demand since data centers replaced on-premises workloads which use 2x-3x more electricity per unit of compute than the cloud. No such luck with the rise of generative AI, which needs new energy
- By 2024, shouldn’t AI/machine learning do better than its recent poor performance detecting breast cancer? Human doctors identified 18 breast cancers that AI missed, while AI only identified 2 that human doctors missed along with generating of hundreds of false positives¹²
- Will NVIDIA’s next chip rollout bring down AI server costs? The company claims that a large model will be able to be trained using 2,000 Blackwell GPUs @\$35k-\$40k each using 4 MW of power instead of 8,000 H100 GPUs @\$25k each using 15 MW (i.e., 50%+ declines in cost and energy). And what if a lot of inference tasks can be performed by cheaper CPUs rather than GPUs?

AI server cost comparisons



Source: NewStreet Research, August 2024

How realistic is a nuclear renaissance in the US?



Source: Power Reactor System Database, JPMAM, December 2023

¹¹ Acemoglu (MIT) starts by assuming that generative AI could transform ~20% of tasks in the production process, and then adjusts for estimates that only 25% of that figure could be cost-effectively automated within 10 years. He estimates AI-related labor cost savings at ~25% to arrive at an estimated total factor productivity boost of just 0.064% each year. “*The Simple Macroeconomics of AI*”, Daron Acemoglu, MIT, May 2024

¹² “*A Prospective Analysis of Screen-Detected Cancers Recalled and Not Recalled by Artificial Intelligence*”, Smith et al, Journal of Breast Imaging, May 2024

Tracking corporate sector AI adoption so far: mostly positive but some mixed signals

The Census tracks AI adoption in its Business Trends and Outlook Survey. The results are positive given steady AI adoption increases from Oct 2023 to Aug 2024, and projections through 2025 (first chart). CEO surveys also point to rising adoption: 50% of CEOs surveyed say they're hiring for gen AI-related roles that didn't exist last year, 40% plan to increase AI spend in the next year and 70% say generative AI is their top priority¹³.

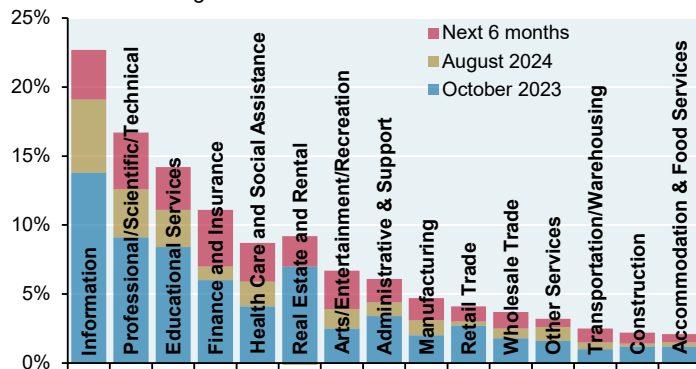
But when looking at enterprise use data from Bain, AI adoption is still concentrated in development/pilot stages; actual production use cases *declined* a little from October 2023 to February 2024. And according to Fed data¹⁴, AI job skill penetration is still low, tracking the early days of 3D printing rather than the faster adoption rates of cloud computing and smart devices (third chart).

Is the decline in tech sector employment at a time of rising payrolls a sign of AI adoption reducing tech labor intensity? Unclear; we compared changes in job listings by type against estimates of each job's AI disruptability and found only a modest correlation.

Perhaps I'm too impatient; it took 20 years for 20% of addressable retail (excluding cars and car parts) to move online as e-commerce, and AI adoption appears faster than that. It also took three years after the 2007 release of the iPhone for Uber to appear, and another five years for Uber to reach 20 mm-30 mm monthly active users. One thing is clear: AI is great for consultants. BCG now earns 20% of its revenue by helping large companies figure out what to do with generative AI¹⁵.

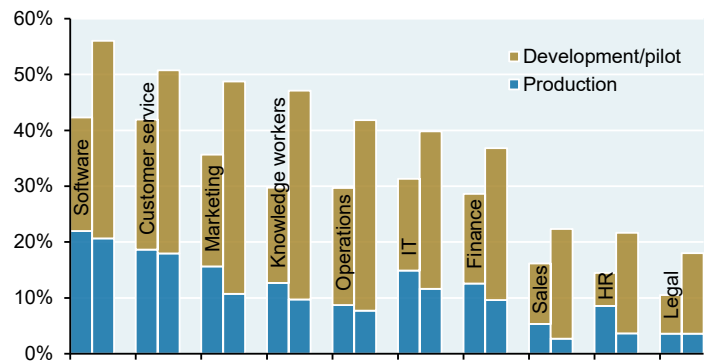
Census: AI adoption rates by sector and date

Share of firms using AI



Source: US Census Bureau, JPMAM, August 2024

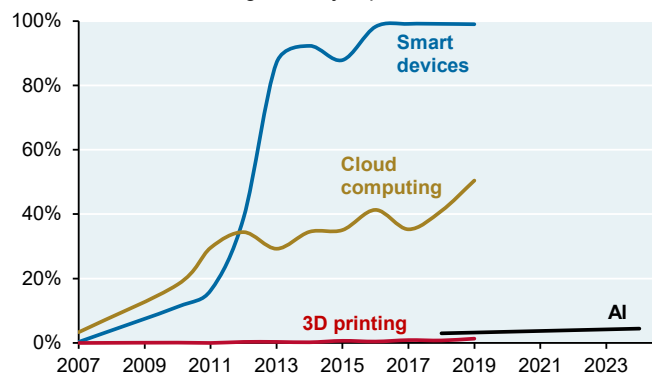
Bain: generative AI adoption rates by stage and sector, Oct 2023 vs Feb 2024



Source: Bain, Ben Evans, July 2024

Fed: technology penetration in job openings

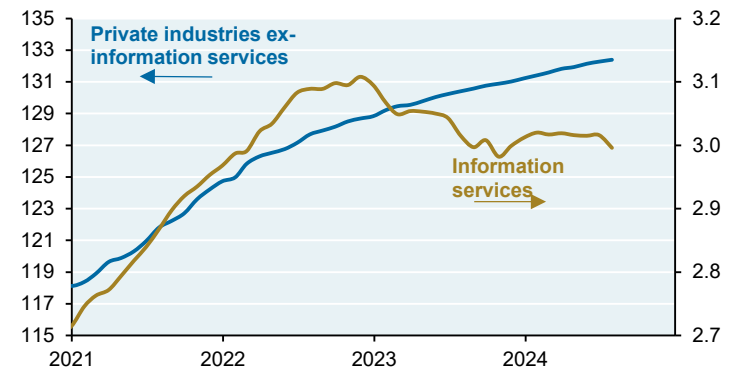
% of urban areas with significant job presence



Source: St. Louis Fed, Kalyani & Hogan, April 2024

US private employment

Millions



Source: BLS, Haver, JPMAM, July 2024

¹³ IBM AI CEO Survey, May 2024; KPMG AI CEO Survey, April 2024

¹⁴ "AI and Productivity Growth", Federal Reserve Bank of St Louis, Kalyani and Hogan, April 2024

¹⁵ "BCG says AI consulting will supply 20% of revenues this year", Financial Times, April 2024

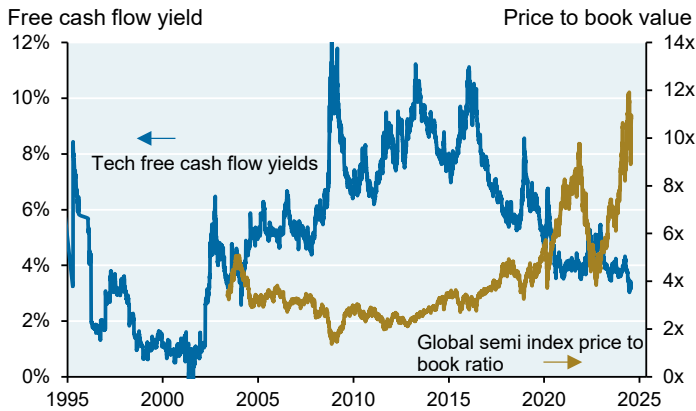
What are equity markets pricing in?

While the surge in NVIDIA’s stock has been accompanied by rising earnings, it’s 38x P/E multiple is consistent with the rest of the technology sector which is priced at the upper end of historical ranges. Investors are bullish: short interest on tech stocks has disappeared, and P/E multiples that markets apply to AI beneficiaries are not that far below multiples applied to AI suppliers. In other words, a lot of the AI revolution appears to be priced into the equity markets. Another look: the aggregate P/E ratio on an indicative generative AI ETF.¹⁶ has risen to 28x this year, although it’s still below its liquidity-induced peak of 38x in 2022.

At least P/E multiples are not as outlandish as in prior tech booms: current tech leaders trade at 1.8x the broad market compared to peak relative P/E multiples of 6x during the dot-com era, 2.8x during the mainframe era (1969) and 2.4x during the PC era (1983).¹⁷

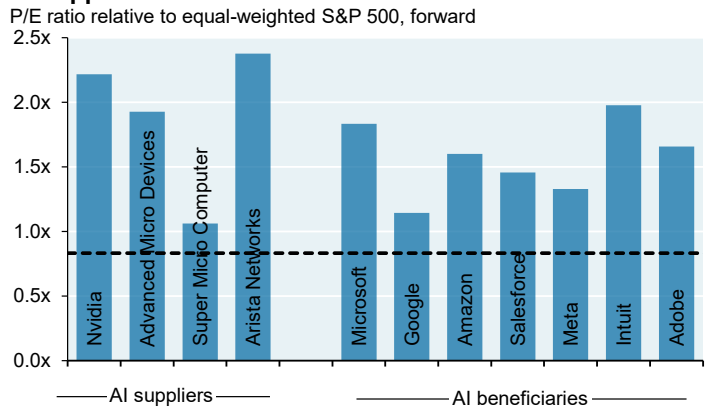
Good news for investors: tech sector capex and R&D is not elevated relative to cash flow, at least when compared to the pre-2000 era (fourth chart). According to Empirical Research, the ratio of Mag 7 capital spending to revenues is actually *below* that of the broad market, and expected to stay that way through 2026.

Tech & semiconductor valuations at extremes



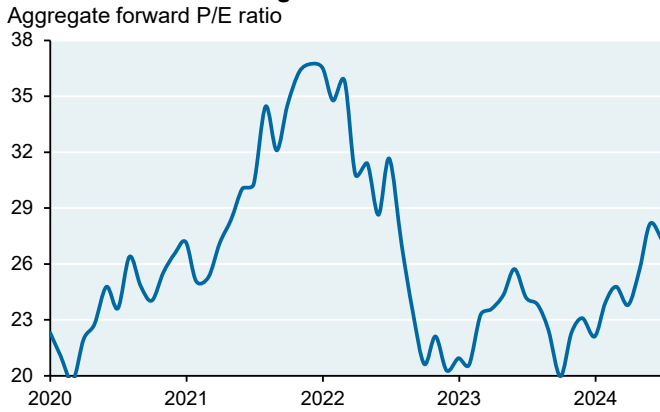
Source: Bloomberg, JPMAM, August 15, 2024

AI suppliers vs beneficiaries



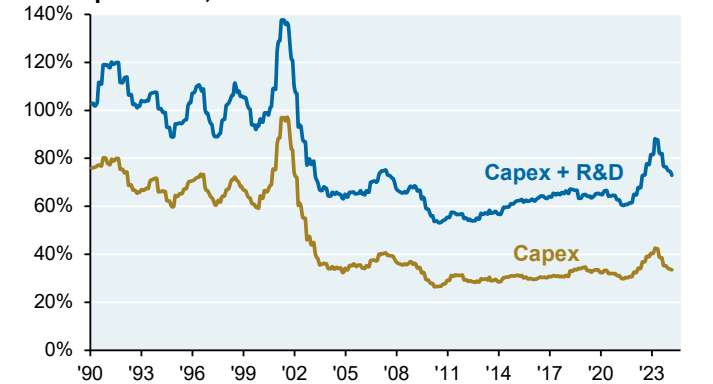
Source: Bloomberg, JPMAM, August 16, 2024

Valuation of indicative generative AI ETF



Source: Bloomberg, JPMAM, July 2024

S&P 500 TMT investment spending as share of cash flows from operations, Percent



Source: Compustat, JPMAM, June 2024. TMT = Info tech, Comm services, AMZN

¹⁶ We selected an AI ETF that has a large weight to NVIDIA (~7%). Some other AI ETFs have performed poorly this year since their NVIDIA weights are lower than the NVIDIA weight in the S&P 500

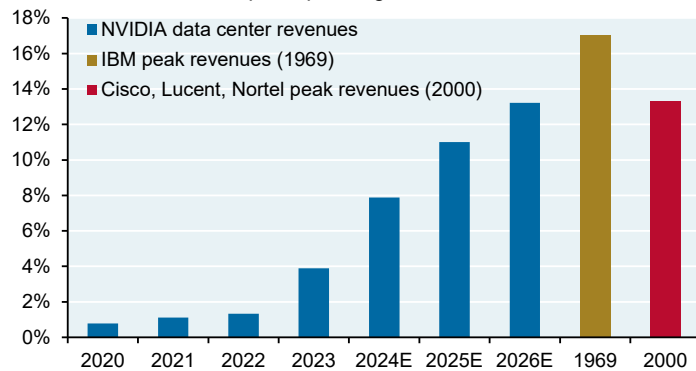
¹⁷ “The tech leaders: is capex killing the CHAT”, Empirical Research, August 26 2024

COVIDIA prognosis: judgment day on AI capital spending is probably 2-3 years away

- Every computing cycle entails infrastructure first, then platforms and then applications. It’s probably too soon to worry that there’s no killer generative AI app yet akin to the enterprise resource planning software of the 1990s or the search and e-commerce applications of the 2000s
- But the clock is ticking: NVIDIA data center revenues as a share of market-wide capital spending are projected to reach heights seen at the peak of the mainframe era in 1969 and the dot-com boom (see chart). As a result, stakes for investors are high
- The table shows how hyperscalers and other data center users now spend more on model training than on inference¹⁸. **In English: these companies spend more than twice as much *training* complex foundational models than they do on *running completed production models* on a day-to-day basis for their end-use customers.** That will probably continue for another year or so: highly profitable hyperscalers will keep the AI capital spending machine moving on their own. Also: some AI demand will be sourced from the military which is price inelastic due to national security reasons
- **Within the next two years, corporate AI adoption trends shown on page 6 need to move higher (i.e., a lot more inference activity) to avoid a “metaverse” outcome for all the capital deployed.** As cited earlier, there needs to be several hundred billion dollars a year of AI-related demand from the corporate sector to pay for AI infrastructure

AI capital spending in context

Share of market-wide capital spending



Source: Empirical Research, August 2024

\$ in billions	2024 AI data center spend			2024 data center operating costs		
	GPUs and other chips	Other AI spend	Total AI capex	Training and R&D	Inference	Total operating costs
Microsoft	\$20	\$20	\$40	\$3	\$3	\$6
Meta	\$11	\$12	\$23	\$2	\$2	\$4
Google	\$14	\$15	\$29	\$3	\$1	\$4
Amazon	\$8	\$8	\$16	\$2	\$1	\$3
Tier 2 cloud	\$26	\$26	\$52	\$8	\$3	\$11
Enterprises & govt	\$26	\$26	\$52	\$8	\$2	\$10
Total	\$105	\$105	\$210	\$27	\$12	\$39

Source: New Street Research, August 2024. Operating costs include cash operating expenses, software, depreciation and electricity

¹⁸ Industry rule of thumb: the cost of inference is roughly the square root of the cost of training

Appendix I: Current and prospective NVIDIA competitors

While NVIDIA currently has 90%+ market share in advanced AI chips, there are several competitors which currently offer products that compete with NVIDIA GPUs, or plan to do so.

Semiconductor companies

- AMD has already developed chips that compete with NVIDIA, and has developed its own ROCm software to compete with NVIDIA’s CUDA
- Intel is designing Gaudi 3 to be 1.5x faster than NVIDIA’s H100 chip and 30%-60% cheaper. Intel also is working with Google, Qualcomm and others to develop a competitor to NVIDIA’s CUDA software. Intel’s upcoming Granite Rapids chip will feature inference processing units on every core (suitable for smaller models), and AMD may follow suit
- Softbank’s Arm has focused on designing and selling chip IP but is reportedly in discussions with TSMC to secure capacity to begin producing AI chips sometime in 2025

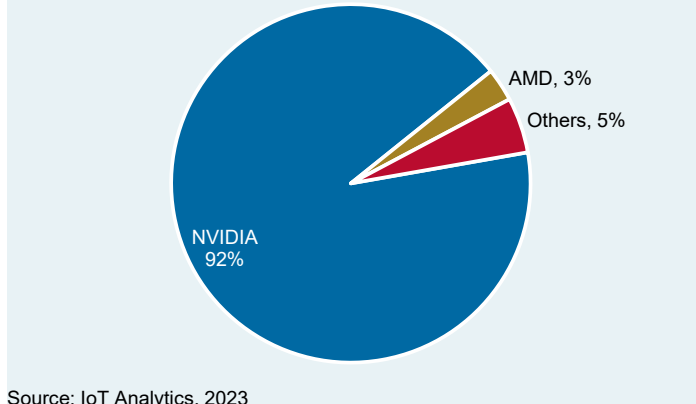
Cloud service providers

- Amazon already offers Trainium and Inferentia chips to existing AWS customers alongside NVIDIA chips. AWS estimates the cost of running models on their own chips to be roughly half that of NVIDIA chips
- Google already offers TPU v5p chips to GCP customers, at roughly 1/3rd the annual running cost of its NVIDIA counterpart (H100 chip). Google is NVIDIA’s closest competitor among cloud service providers
- Microsoft offers the Maia 100, a cheaper alternative to the H100 for Azure customers

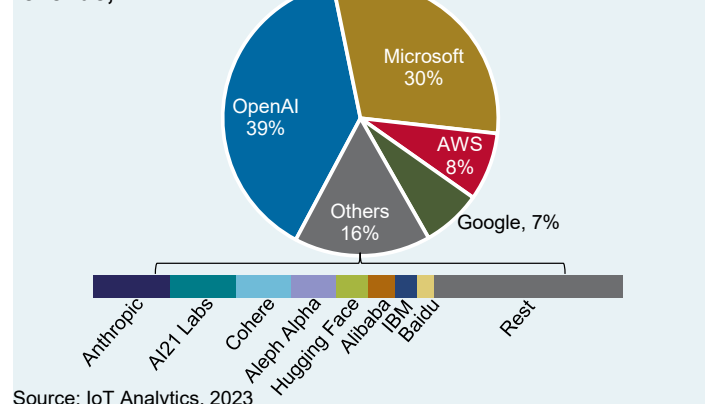
Other AI application developers and internal users

- Apple is working on an AI chip project dubbed Project ACDC for internal inference computations
- Cerebras partnered with Dell to market its WSE-3 chip to compete with NVIDIA’s H100 chip. The WSE-3 chip has 50x the compute power of the H100 and can train a model 10x the size of GPT-4, but can only be purchased as part of Cerebras’s computing platform, similar to buying computing capacity from AWS
- Groq, founded by a former Google TPU designer, is working on a chip optimized for ultra-fast inference for less computationally intensive models. Groq partners with Global Foundries rather than TSMC
- Meta has developed its own chip, the Meta Training and Inference Accelerator (MTIA), for internal AI projects. MTIA is primarily being used to train Meta’s recommendation and ranking algorithms, but the goal is for MTIA to eventually train generative AI like Meta’s Llama
- Tesla has designed its own D1 chip for self-driving car projects on their Dojo supercomputing platform
- OpenAI has reportedly begun discussions with Broadcom on chip development, has hired employees from the Google TPU team and is raising capital for fab construction in partnership with Intel/TSMC/Samsung

Market share of data center GPUs based on 2023 revenue
Percent



Market share of AI models and platforms based on 2023 revenue, Percent

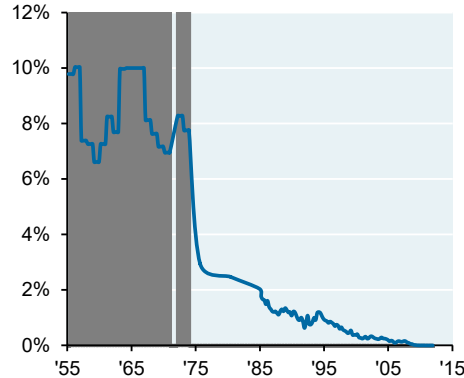


Appendix II: The fate of prior equity market leaders

For one brief single day in June 2024, NVIDIA attained the largest market cap in the S&P 500 at \$3.33 trillion, surpassing Microsoft. Market leaders of decades past typically reached an inflection point of some kind, after which their market cap declined as a share of the market (GM, IBM, Altria, Cisco, GE and Exxon). Only Microsoft and Apple continued to thrive. Notes: gray bars denote periods when each stock was the largest in the S&P 500; market cap shares computed on an end-of-quarter basis

General Motors

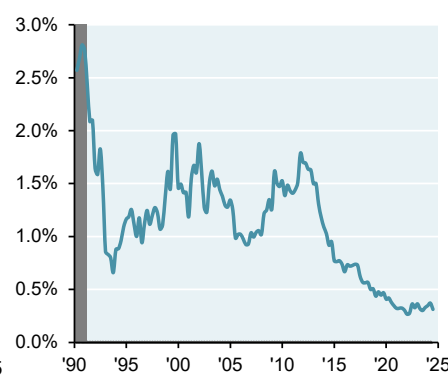
Share of total market cap



Source: Factset, Datastream, JPMAM, June, 2024

IBM

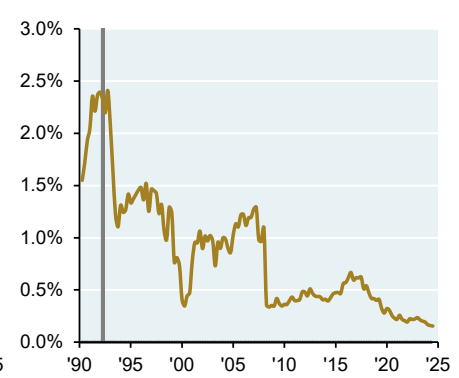
Share of total market cap



Source: Factset, JPMAM, June, 2024

Altria

Share of total market cap



Source: Factset, JPMAM, June, 2024

Cisco

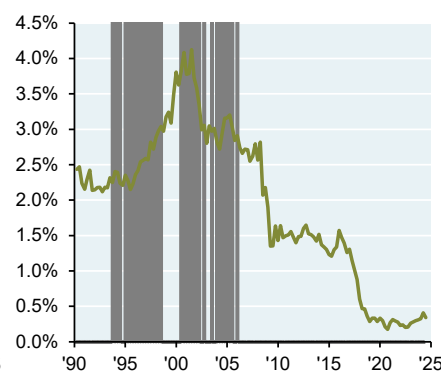
Share of total market cap



Source: Factset, JPMAM, June, 2024

General Electric

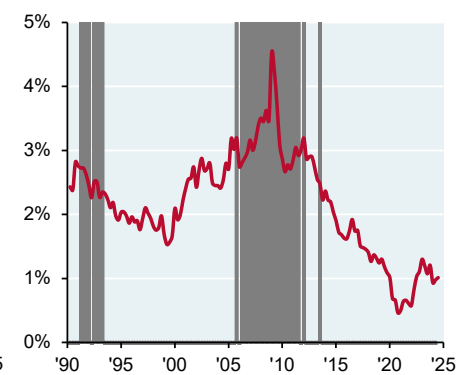
Share of total market cap



Source: Factset, JPMAM, June, 2024

Exxon

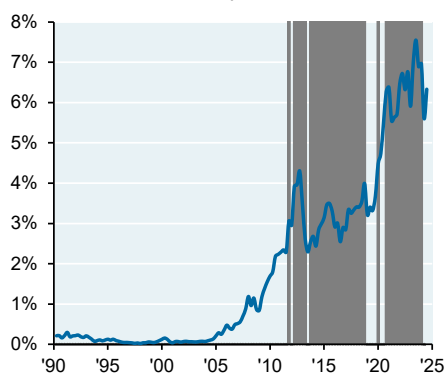
Share of total market cap



Source: Factset, JPMAM, June, 2024

Apple

Share of total market cap



Source: Factset, JPMAM, June, 2024

Microsoft

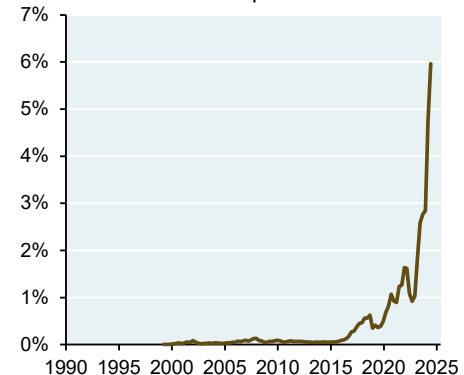
Share of total market cap



Source: Factset, JPMAM, June, 2024

Nvidia

Share of total market cap



Source: Factset, JPMAM, June, 2024

IMPORTANT INFORMATION

This material is for information purposes only. The views, opinions, estimates and strategies expressed herein constitutes Michael Cembalest’s judgment based on current market conditions and are subject to change without notice and may differ from those expressed by other areas of JPMorgan Chase & Co. (“JPM”). **This information in no way constitutes J.P. Morgan Research and should not be treated as such.** Any companies referenced are shown for illustrative purposes only and are not intended as a recommendation or endorsement by J.P. Morgan in this context.

ACCESSIBILITY

J.P. Morgan is committed to making our products and services accessible to meet the financial services needs of all our clients. If you are a person with a disability and need additional support, please contact your J.P. Morgan representative or email us at accessibility.support@jpmorgan.com for assistance

GENERAL RISKS & CONSIDERATIONS Any views, strategies or products discussed in this material may not be appropriate for all individuals and are subject to risks. Investors may get back less than they invested, and **past performance is not a reliable indicator of future results.** Asset allocation/diversification does not guarantee a profit or protect against loss. Nothing in this material should be relied upon in isolation for the purpose of making an investment decision.

NON-RELIANCE Certain information contained in this material is believed to be reliable; however, JPM does not represent or warrant its accuracy, reliability or completeness, or accept any liability for any loss or damage (whether direct or indirect) arising out of the use of all or any part of this material. No representation or warranty should be made with regard to any computations, graphs, tables, diagrams or commentary in this material, which are provided for illustration / reference purposes only. Any projected results and risks are based solely on hypothetical examples cited, and actual results and risks will vary depending on specific circumstances. Forward-looking statements should not be considered as guarantees or predictions of future events. Nothing in this document shall be construed as giving rise to any duty of care owed to, or advisory relationship with, you or any third party. Nothing in this document shall be regarded as an offer, solicitation, recommendation or advice (whether financial, accounting, legal, tax or other) given by J.P. Morgan and/or its officers or employees. J.P. Morgan and its affiliates and employees do not provide tax, legal or accounting advice. You should consult your own tax, legal and accounting advisors before engaging in any financial transactions.

J.P. Morgan Wealth Management is a business of JPMorgan Chase & Co., which offers investment products and services through J.P. Morgan Securities LLC (JPMS), a registered broker-dealer and investment adviser, member [FINRA, opens overlay](#) and [SIPC, opens overlay](#). Insurance products are made available through Chase Insurance Agency, Inc. (CIA), a licensed insurance agency, doing business as Chase Insurance Agency Services, Inc. in Florida. Certain custody and other services are provided by JPMorgan Chase Bank, N.A. (JPMCB). JPMS, CIA and JPMCB are affiliated companies under the common control of JPMorgan Chase & Co.

© 2024 JPMorgan Chase & Co. All rights reserved.