



AI IPOs 101: why they'll change the boom forever

July 01, 2026

Four reasons why the expected IPOs of OpenAI and Anthropic matter so much

The expected wave of initial public offerings of the world's leading AI companies will change the AI boom forever.

It will bring unprecedented **transparency**, **funding** and **accountability** to the industry – as well as finally giving investors a **pure-play benchmark** to value the likely winners and losers of a revolutionary technology.

The [news](#) at the end of last week that [OpenAI](#) is considering delaying its own share sale, originally expected for later this year, indicates just how much is at stake. The company, which lit the fuse for the AI frenzy with its launch of ChatGPT in November 2022, as well as its now slightly larger rival [Anthropic](#), which makes Claude, confidentially filed for IPOs in early June. Those share sales would reportedly aim to raise about \$60bn apiece and value each company at more than \$1 trillion.

They and their peers need to **strike while the iron is hot** to secure computing power, distribution and capital. But there is no guarantee that there will be enough investor appetite to absorb hundreds of billions of dollars in rapid-fire demand.

The AI boom is fast approaching a **moment of truth**, as rapid growth and soaring valuations collide with ballooning capital expenditure, a public backlash and the challenges of real-life adoption. Challenges abound, including Chinese open-source models that promise similar performance for a fraction of the price. Whoever comes out on top could define the era.

In this report, we look at:

- Why have an IPO?
 - Which AI-related companies are likely to have IPOs?
 - OpenAI and Anthropic in a nutshell
 - Four ways that IPOs will change the AI boom forever
1. Greater transparency for financials – and business models
 2. More diverse funding sources – and more spending
 3. Greater accountability for their actions – and the industry
 4. New era for comparing valuations to peers – and the past

Authors

[Adrian Cox](#)
Research Analyst



Why have an IPO?


1. **Access new sources of funding** to invest, reduce debt and improve financial flexibility.
2. **Create liquidity** to give existing investors and employees a way to cash out and enable acquisitions with publicly traded shares.
3. **Improve public credibility and profile**, lending additional weight with consumers, businesses and lawmakers.

Which AI-related companies are likely to have IPOs?

The two leading contenders to sell shares to the public this year are Anthropic, given a 78 percent chance, according to bets on Polymarket, and OpenAI, now down to 28 percent.

Other AI-related companies that may go public include data analytics platform Databricks, enterprise data platform Glean, and compliance tool Vanta, according to Polymarket bets. Others include business process platform Celonis; graphic design tool Canva; finance operations platform Ramp; defence company Anduril; and Europe's only notable frontier-model maker Mistral.

Figure 1: Implied probability of select AI-related IPOs by 2027



Source : Polymarket, Deutsche Bank Research

OpenAI in a nutshell

- San Francisco-based OpenAI was **founded in 2015** as a non-profit research laboratory led by **Co-Chairs Elon Musk and now-CEO Sam Altman**, the former head of Y-Combinator. It was backed by some of the biggest investors in Silicon Valley, including Peter Thiel and Reid Hoffman. (Musk severed ties with OpenAI in 2018, triggering a legal dispute that concluded in May in OpenAI's favour.)
- Its **mission** was "to ensure that artificial general intelligence – AI systems that are generally smarter than humans – benefits all of humanity".
- On November 30, 2022, OpenAI launched **ChatGPT**, a generative AI chatbot based on its GPT-3.5 large language model, which became a



runaway success with 100 million users in three months. GPT-4, released in March 2023, was so far ahead of the competition that it took a year or more for serious challengers to gain a foothold. However, GPT-5, launched last August, was quickly matched amid stiffer competition.

- The non-profit board voted to remove **Altman** as CEO in November 2023. He was reinstated four days later, after a revolt by investors and employees.
- Its **most important strategic relationship has been with Microsoft**. When OpenAI transitioned into a “capped-profit” entity in 2019, it was able to obtain capital investment including \$1bn from Microsoft, which later added another \$10bn in early 2023. Their exclusive partnership and revenue-sharing agreements enabled Microsoft to be an early leader in AI and secured funding and computing power for OpenAI.
- The companies **scaled back their partnership** in late April 2026, though Microsoft retains a 27 percent stake in OpenAI. (OpenAI’s for-profit LLC became a public benefit corporation controlled by the non-profit in May 2025.)
- Today, **more than one in ten people on the planet use ChatGPT every week**. The name is to chatbots what PostIts are to sticky labels. It reportedly surpassed \$25bn in annualised revenue in March and has indicated to investors that it will achieve its first annual profit in 2030.
- Although the market is increasingly crowded, OpenAI is still producing **cutting-edge models**, with GPT-5.5 excelling at agentic coding. Brand new GPT-5.6 won the same badge of honour as the latest Claude model last Friday by being restricted to a small group of users over US government concerns that it was so powerful that it could pose a security risk. It also has a strong developer ecosystem building on top of its Application Programming Interface (API).

Anthropic in a nutshell

- San Francisco-based Anthropic was **founded in 2021** by seven former employees of OpenAI, including **Dario Amodei, its CEO**, and his sister Daniela Amodei, its president. Anthropic is a public benefit corporation whose [mission](#) is “to ensure that the world safely makes the transition through transformative AI”.
- Led by its **Claude models**, the company has focused heavily on the enterprise market, positioning itself as a bastion of high standards of safety and reliability.
- **Google parent Alphabet and Amazon** have invested billions of dollars in Anthropic. Like OpenAI, it has struck complex agreements that commit it to use chips and data centres from its investors in return, including a commitment [reported](#) by The Information to spend about \$200bn with Google Cloud over five years.
- This has been a **breakout year** for Anthropic. The release of its Cework and Claude Code agent tools in January and February prompted a “SaaSocalypse” slump in stocks of software-as-a-service companies, amid concerns that their utility-like business models could be largely disintermediated.
- The **Claude website had almost one billion visits in May**, still behind ChatGPT at 5.6 billion and Google’s Gemini at nearly three billion, but way up from 200 million in January, according to data from Similarweb. Meanwhile, Claude Code now accounts for slightly more than half of the business-focused coding market.
- **Anthropic’s revenue** reached an annual recurring rate of \$47bn in May,



overtaking OpenAI. CEO Amodei [said](#) that an 80-fold increase in usage in the first quarter on an annualised basis was “just crazy” and “too hard to handle,” forcing it to limit access for some users and sign up new sources of compute to keep up with demand.

- A **disagreement with the US Administration** came out into the open in February when the Department of Defense announced that it was a supply-chain risk. The company had demanded restrictions on the use of its technology for domestic surveillance and autonomous weapons.
- Then two weeks ago, Anthropic [abruptly switched off its most powerful](#) **Table 5 and Mythos 5 models** worldwide after the US Department of Commerce ordered it to restrict access by foreigners, due to concerns that it could be weaponised by bad actors. It won approval to restore some access at the end of last week.

Four ways that upcoming IPOs will change the AI boom forever

1. Greater transparency for financials and business models

Nothing sharpens an entrepreneur’s mind – and business model – like the prospect of an IPO.

- It is a decision that they may take with some trepidation, given the obligation to open the doors and explain their finances, business model and outlook in detail.

To prepare for an IPO, AI companies will have to produce public-company-grade reporting, credible performance indicators, risk analyses and forecasts.

- They will also have to make any relevant disclosures about a growing number of lawsuits over copyright, alleged negligence regarding safety, and corporate governance.
- Remuneration may also have to be revamped, with public metrics around pay for performance for management and provisions for employees to vest stock.


Greater transparency may help investors take a view on how much they consider currently lossmaking AI models to be winner-takes-all, like Amazon or Meta, where network effects create a virtuous circle, or commodities, where one model is viewed as interchangeable with another.

- While OpenAI has the leading chatbot by far with around 900 million monthly users, it has been rapidly pivoting its business model this year to try to attract enterprise users who will pay for it.
- That includes creating the [OpenAI Deployment Company](#) in May with 19 private-equity groups led by TPG to distribute its enterprise products, spearheaded by so-called “forward deployed engineers”. It said previously that enterprise revenue was on track to match consumer revenue by the end of the year.
- Since CEO Sam Altman declared a “code red” in December, OpenAI has scrapped what it called “side quests” such as its video generating tool Sora and is [working](#) on a “unified superapp” that brings its chatbot, coding, browsing and agentic capabilities together, aiming to convert satisfied home users to a similar business offering.



- Anthropic overtook it in business subscriptions in May, according to data from Ramp. And Similarweb data shows monthly visits to ChatGPT fell below a majority of the generative AI market for the first time in May, suggesting consumers are increasingly willing to switch between models.

Figure 2: Top four generative AI platforms: monthly visits worldwide (desktop and mobile web)



Source : Similarweb, Deutsche Bank Research

- OpenAI has indicated that it is generating about \$25bn to \$33bn in annualised revenue, up from a \$4bn annual rate at the end of 2024. By comparison, Anthropic's annualised run rate shot up from \$1bn a year at the end of 2024 to \$47bn in late May. The company previously said that its first profitable full year will be in 2029, one year earlier than OpenAI.

Greater disclosure around IPOs will finally illuminate the business models of the frontier model makers.

- While chips and data centres are easy to measure, actual monetisation of AI is opaque, given that private model makers make only limited disclosures and public hyperscalers rely on AI models for just a small part of their AI revenues.
- The key questions are how intense and rapid adoption will be; is the return on investment justified once businesses have to make hard choices to scale beyond pilots; and will premium all-singing, all-dancing frontier models win out over cheap, stripped-down, task-specific alternatives?

The stakes for pure-play model makers are higher now than ever before, with pressure from above and below.

- On the one hand, the hyperscalers have unsurpassed access to data, distribution and compute, as well as the ability to offer competing models on their platforms and still take a cut.
- On the other hand, ever-smarter open-source models may be good enough for most tasks. GLM 5.2, an open-source model from China's Z.ai, has become a viral hit in the US after rocketing up the leaderboards over the past couple of weeks. It scores 51 in Artificial Analysis's Intelligence Index, not




far behind Claude's (currently restricted) Fable 5 and OpenAI's GPT-5.5, yet costs just 48 cents per task in the index, versus \$2.75 and \$1.03 respectively.

Price is becoming a battleground as enterprise customers begin to balk at the true costs of rolling AI out at scale.

- Yes, cost per token, the building block of AI, is falling by between nine and 900 times a year, depending on the task, according to Epoch AI. But that is being offset by ever more compute-hungry agentic applications – as well as by well-meaning initiatives to encourage employees to be creative that have incentivised budget-busting “token-maxxing”.
- The economics of AI suddenly look less attractive once token usage is no longer being subsidised by investors. Anthropic has changed from a flat fee to a usage-based model, and OpenAI's Altman said last month that cost had come out of the blue to become a “huge issue” for customers this year. Both are [reported](#) to be considering price cuts that could erode their profit margins.

Figure 3: Artificial Analysis Intelligence Index of selected AI models vs cost per tasks




Source : Artificial Analysis Intelligence Index. “Fallback” mechanism routes safety-flagged messages to Opus 4.8 model. Meta Muse Spark price not available. US models in blue; Chinese models in red

- OpenRouter data showing weekly token usage of the top models on its platform indicates that the number of tokens processed by US models has risen fivefold this year to 20 trillion, convincingly outpaced by a thirteenfold increase by Chinese models to 26 trillion. The platform allows customers to access hundreds of LLMs from different providers through a single account.



Figure 4: The number of tokens processed by OpenRouter using Chinese models has overtaken tokens processed using US models this year (trillion)



Source : OpenRouter, Deutsche Bank Research

- Amazon, Walmart and Cisco are among companies that have reined in token use. Uber burned through its AI budget for the year in four months and has now capped spending, with its Chief Operating Officer [saying](#) that “the link is not there yet” between the use of Claude Code and customer innovation.
- Demand is growing for new model architectures that can divert queries to the most efficient model for any given prompt, such as OpenRouter’s new Fusion plugin.

2. More diverse funding sources – and more spending

AI is an arms race for compute, pitting startups against established hyperscalers.

- Specialist model-makers are on the back foot compared with a company like Alphabet parent Google that has ample cloud compute facilities and is generating roughly \$50bn in cash from operations every quarter that it can use to buy even more.
- Private companies cannot tap the same pool of investors forever. Abundant public market funding would break the cycle of relying on the piecemeal multi-billion-dollar investments and complicated financing arrangements among the leading firms of the past few months.

To be sure, this has not been a bad time to be raising private money.

- Three lean years for venture investment reversed last year as investors injected \$425bn, an increase of almost a third from the previous year, into more than 24,000 companies, [Crunchbase](#) data show. AI funding was particularly healthy, exceeding \$225bn globally, albeit concentrated in few



companies, according to [CB Insights](#).

Figure 5: Global AI funding surged in recent quarters, albeit concentrated in fewer companies: quarterly equity funding (\$bn) and deal count



Source : CB Insights, Deutsche Bank Research

- Anthropic said at the end of May that it [raised](#) \$65bn at a post-money valuation of \$965bn, more than double its valuation in February. In one leap, it overtook OpenAI, which [raised](#) \$122bn at an \$852bn post-money valuation in March. If either achieved the mooted market value of \$1trn, it would be comfortably within the world's top 20 most valuable companies.

But growth does not come cheap.

- OpenAI told investors in February that it was targeting some \$600bn in infrastructure investments by 2030, less than half of its earlier estimate of \$1.4tn amid concerns about whether future revenue could justify the outlay. Its emphasis has changed from building its own data centres to just renting them.
- OpenAI's funding round in March was anchored by long-term partners Amazon, Nvidia and SoftBank. Amazon will invest \$50bn in OpenAI in return for OpenAI using 2GW of its Trainium chip capacity and cooperating on new joint products.
- Long-term shareholder and partner Microsoft still took part even though the two have grown more distant and [unwound](#) some exclusive arrangements in April. The startup also took investments from at least 19 institutional investors and funds, as well as some retail investors via banks and exchange-traded funds, and expanded a revolving credit facility to about \$4.7bn.
- Meanwhile, Anthropic's latest round looped in at least two dozen investors and included \$15bn of previously committed investments from hyperscalers including \$5bn from Amazon – after the startup agreed to buy up to 5GW of compute capacity from it and a similar amount from Google and Broadcom.
- In a sign of just how complicated the AI funding org chart is becoming, Google, which already [reportedly](#) owned a 14 percent stake in Anthropic



worth \$3bn, pledged to invest another \$40bn in it in April and more recently was said by Bloomberg to have agreed to backstop lease payments worth \$35bn on five data centres.

IPOs would not just unlock new equity finance.

- Public listings would give the companies and their backers leverage capacity: the ability to raise debt against liquid, transparently valued collateral. Those funds could then be recycled into further expenditure on data centres, acquisitions and other ventures.
- A publicly funded OpenAI or Anthropic would likely not just build bigger models but also seek to become the default foundational platform layer similar to what Windows was for PCs and to Amazon's AWS for the cloud. They could use larger war chests to expand their API services, developer tools and enterprise partnerships, commoditising the layers where rivals (and perhaps some customers) operate.

Financial flexibility matters for all of the players in the land grab for capacity in the AI boom.


- Capital expenditure by hyperscalers looks set to rise by two-thirds to more than \$750bn this year. Even cash-rich Google turned to the debt markets this year, raising almost \$32bn in less than 24 hours in February, then raising \$84.75bn in the biggest equity capital markets transaction ever on June 1. Nvidia sold \$25bn of high-grade bonds in mid-June.
- And smaller companies are turning to the public markets, most recently with South Korean memory chip maker SK Hynix announcing that it is planning to raise more than 45 trillion won, or almost \$30bn, in a landmark US listing.

An opening of the IPO floodgates after a five-year drought would give companies the currency they need to embark on a spending spree.

- It could even break the inertia of an upstream drought in fundraising for early investments by unblocking trapped liquidity. Global venture capital fundraising fell last year to \$122 billion, little more than a quarter of the peak in 2021 and 2022, as investors shied away from illiquid investments in favour of less risky strategies, PitchBook [data](#) show.



Figure 6: Global private-equity fundraising dropped significantly after 2022, trapping upstream liquidity



Source : [PitchBook](#), Deutsche Bank Research. Data as of Dec 31, 2025

- That could mean a surge in mergers and acquisitions, turning other private companies into targets and giving an exit route to impatient investors. This year has already seen more than 5,000 technology acquisitions or investments pending or completed with a value of about \$850bn, more than double the level for the same period last year, Bloomberg data show.
- Elevated valuations and regulatory hurdles for the biggest companies would put them out of reach of all but a tiny number of strategic buyers (the hyperscalers) or sovereigns. An alternative approach has been “acquihires”, pioneered by Microsoft two years ago when it hired most of the staff of Inflection, including co-founder Mustafa Suleyman, and licensed the use of its AI models for \$650 million.
- But smaller companies have a much greater chance of being acquired.
- Most recently, US chipmaker Onsemi agreed to buy Synaptics, which makes semiconductors for smart devices, for about \$6.2bn in stock late last week. Indeed, Bending Spoons, which buys up struggling software companies, including AOL and Vimeo, itself filed for an IPO last week. The Italian company aims to raise as much as \$1.6bn at a \$19bn valuation early this month.

Last but not least, IPOs will give companies a potent new weapon in the battle for talent: liquid, publicly traded stock options.

- For top-tier researchers and engineers, this would be a far more tangible and immediate form of wealth than illiquid shares in a private rival that might seek to lure them away.
- Such a liquidity shock could force private competitors into responding with ever larger funding rounds at higher valuations on the promise of “jam tomorrow” for star employees – or into going public themselves. They would also have to be on high alert to prevent their existing core teams from being poached by a newly public OpenAI or Anthropic armed with a trillion-dollar market capitalisation.
- Meta’s Mark Zuckerberg sparked a bidding war for top AI researchers last



year, offering compensation packages worth hundreds of millions of dollars in some cases. OpenAI and Anthropic were the most recent players to fish in the small pond of genuine AI experts a couple of weeks ago, hooking two legends of the industry. Shares of Google, where they worked, fell more than 5 percent on the day.

3. Greater accountability for their actions – and for the industry

The AI industry has lost its lustre in recent months.

- Warm initial perceptions of it as a slightly magical composer of wedding speeches have faded and may vanish once Mom and Pop own shares in pure-play AI companies for the first time. Those companies risk becoming a lightning rod for broader fears about the impact of AI on trust, jobs and vulnerable industries.

At the most basic level, public companies are held to higher standards around consistency and accountability.

- For example, OpenAI has attracted criticism for reversing its earlier opposition to advertising in an attempt to create a more sustainable business model. Critics have said conversations could be biased by sales imperatives.

Beyond that, OpenAI and Anthropic have already had a taste of the political challenges in store as public companies.

- Anthropic was blacklisted by the US Department of Defense in February for refusing to let Claude be used for autonomous weapons and mass surveillance. The dispute cost Anthropic a \$200m [contract](#) and is now in the courts. OpenAI was then criticised from the other side after it struck a deal to replace Anthropic. Altman later described the move as being “complex but the right decision with extremely difficult brand consequences and very negative PR for us in the short term”.
- More recently, the US Department of Commerce issued an export control directive late on June 12 to suspend all access by foreign nationals to Anthropic’s Claude Mythos 5 and less powerful Fable 5 models because of security concerns.
- Although it had to temporarily disable access to all customers to ensure compliance, its reputation for having the most powerful models will have benefitted (see our report [AI twist in the tale for Anthropic’s Fable and Mythos](#)). The silver lining may eventually eclipse the cloud amid a branding boost that money can’t buy. However, alternatives to US models are also likely to benefit from concerns that access could be switched off without warning.

Most broadly, the makers of the best-known chatbots may pay a price for being seen as proxies for all of AI when controversy hits the industry.

- Hot topics include cybersecurity. The leaders of the Five Eyes cybersecurity agencies of the US, UK, Australia, Canada and New Zealand warned in a joint [statement](#) last week that frontier models are likely to supercharge cyber threats: “The timeline is not years, it is months”.

The biggest underlying popular concern is what AI will mean for jobs.

- The leaders of both Anthropic and OpenAI have moderated their



predictions of severe disruption in recent weeks, with Amodeli pivoting to say automation would dramatically increase productivity and Altman saying the impact on jobs has been less than he expected.

- One challenge is that AI has become a scapegoat for job cuts at companies ranging from Block (4,000 employees, or nearly half the workforce) to Atlassian (10 percent of roles) and even British American Tobacco (with AI cited as one of a number of factors in a planned one-fifth reduction in its workforce).
- However, the tide has begun to turn on predictions of catastrophic job losses. A slump in entry-level graduate jobs may have been more due to rising interest rates and remote work than AI – and may be reversing. And business leaders like Amazon Web Services' CEO Matt Garman have [said](#) eliminating the pipeline of talent makes no sense.

Nowhere has the AI backlash become more focused than over data centres.

- As fast as these “AI factories” are mushrooming around the US and the world, so are concerns – both justified and unjustified – about them overwhelming the electricity grid, forcing up power prices, and sucking rivers dry for water to cool their chips.
- How the leading model makers, hyperscalers and infrastructure companies handle the issue could determine whether it remains contained or unleashes moratoriums, regulation or swingeing taxes that put the entire boom at risk.
- In the US, home to half of the world’s data centres, opposition to “AI factories” is becoming a rare area of bipartisan consensus, uniting prominent Democratic Representative Alexandria Ocasio-Cortez and independent Senator Bernie Sanders on the left, and populist strategist Steve Bannon on the right.
- Maine recently passed a moratorium on larger data centres while, across the Atlantic, data-centre hub Ireland has introduced a “Bring Your Own Power” policy that requires new facilities to generate their own electricity or secure renewable energy contracts.

Data centers are – literally – close to home for voters.

- That makes them ripe fodder for targeted lobbying and grassroots campaigns in a way that vague fears about long-term job losses are not. The US midterm elections in November will be an important litmus test of the political consequences of local pressure.
- A Gallup poll in May found that 71 percent of Americans would oppose a data centre being built in their area. At least 75 projects worth a record \$130bn were blocked or delayed in the first quarter, [according](#) to Data Center Watch.
- Tech companies have done their bit to respond. Microsoft, for example, said in its most recent data centre announcement last week that it would create thousands of jobs at the facility in Pecos, Texas; pay to ensure it does not cause a rise in electricity prices; and replenish more water than it uses. “We are approaching this project as a new neighbour,” it said.

The IPOs of the biggest AI companies could also have repercussions in the geopolitical contest for technological supremacy, raising their status to that of strategic national assets.




- That could encourage other powers to accelerate support for their own national champions, like China's "AI tigers" and France's Mistral, to avoid becoming dependent on US-controlled foundation models. Other countries may find themselves having to choose to align with one centre of power or the other.

4. New era for comparing valuations to peers – and the past

Part of the anxiety about an AI bubble is the lack of pure-play, public, investable companies to act as a benchmark. A listed OpenAI or Anthropic would fill a huge gap.

- Sentiment about AI has driven much of the market since the launch of ChatGPT, with anticipated gains from investment and productivity helping to explain the two-thirds rise in the S&P 500 Index during that period.
- That rise has been bumpy, soaring and swooning from week to week, as one area after another, such as software, semiconductors or infrastructure companies, has come into favour and then fallen out again. Investors have crowded into trades on a concentrated and shifting list of potential AI enablers and adopters, and away from prospective casualties.
- To give one example, the Magnificent Seven biggest technology stocks have whipsawed over the past month amid concerns about how much of their AI capital expenditure will pay off, shedding as much as 12 percent by late last week. They have tripled over the past three-and-a-half years to account for almost a third of the S&P 500.

Figure 7: S&P 500 Index, Magnificent Seven tech stocks and other select AI-related sectors performance YTD to June 30 (indexed to 100)



Source : Bloomberg Finance LP, Deutsche Bank Research

- That makes investing in AI right now feel a bit like going to a horse race and having to bet on the jockeys and the condition of the track because the horses themselves are unavailable.

Of course, there is no guarantee that any filing ends up in a successful IPO.

- The San Francisco-based co-working company WeWork was valued at



\$47bn in early 2019 but failed to gain traction when it filed its prospectus in August that year. It eventually went public in 2021 at a heavily reduced valuation and ended up filing for bankruptcy in late 2023.


Another way that major IPOs could remove some distortions is by acting as a release valve for private markets, enabling long-awaited exits from investors.

- Private-equity companies are sitting on a record backlog of unsold companies, thanks in part to geopolitical disruption and fears of a “SaaSocalypse”, according to [data](#) from Pitchbook and others. The median age of still-held assets continued to trend upwards from three years in 2022 to four years in 2025 as exit activity stalled.

Despite concerns that the boom turns to bust, there are differences in the IPO markets from the height of the dot-com bubble.

- Overall valuations are clearly elevated by historical standards, with the Shiller Cyclically Adjusted Price-Earnings (CAPE) ratio for the S&P 500 recently surpassing 40, nearing the record 44 it reached at the peak of the dot-com bubble.
- However, the IPOs coming to market are more mature and more profitable.
- Back in 1999 to 2000, the median age of venture-capital-backed tech startups going public shrank to about 4.5 years, according to University of Florida data. Since then, the median age has been relatively stable at ten-and-a-half years and was 12 years last year.
- Only 6 percent of those tech startups going public at the height of the dot-com bubble were profitable, compared with 22 percent since 2010, and they had a median opening price-to-sales ratio of over 40, four times higher than the average since 2010.

Figure 8: Median age and profitability of technology IPOs (1980-2025)



Source : University of Florida, Jay Ritter, Deutsche Bank Research

Perhaps the greatest insurance against a bubble is the fixation of today’s senior investment managers on the dot-com bust that scarred them in their formative years in the market.



- For them, mega IPOs bring back memories of the IPO of Netscape in August 1995, widely regarded as the start of the dot-com bubble. However, there is just as important a lesson in how long it took for that bubble to go bad despite prophecies of doom and gloom right from day one.
- Bloomberg published an article two days after the IPO on August 9, 1995, entitled “Why investors say Netscape shows a market poised for a fall” and warns “Be worried, very worried”. “Netscape was a little insane,” says a trader quoted in the article, after the stock doubled on its first day of trading, giving the lossmaking company a \$2.9bn market value. “It certainly hints at more frothiness in the market and that’s usually the sign of the beginning of the end of a bull market”.
- Yet in the end, the Nasdaq shot up from 1,005 on the day of Netscape’s IPO, not peaking until almost five years later, at over 5,000 on March 10, 2000. It was only then that it dramatically fell back, reaching a trough of about 1,100 in October 2002.
- As for Netscape, it was bought by America Online (AOL) for \$10bn in 1999, and then vanished in the fallout from AOL’s record-breaking \$124bn purchase of Time Warner in 2001.

Yet Amazon, another star of the times, shows the power of patience – and picking a winner.

- Adjusting for stock splits, the shares debuted at 7.5 cents in the company’s \$54m IPO on May 15, 1997. After the roller coaster of the dot-com bubble and crash, the shares took a decade to recapture their early peak – and then continued to rise, reaching almost \$275 a month ago.
- Yet on the day of its IPO, when the stock rose 31 percent, a Pennsylvania-based fund manager told Bloomberg: “I personally am going to stay away from it. I don’t think it’s been proven that one can make a lot of money over the web.”

Recent relevant research

[Assessing AI in finance with Evident Co-CEO Alexandra Mousavizadeh](#) (June 9, 2026)

[ChatIPO: Quick take on OpenAI’s record-breaking IPO plans](#) (May 21, 2026)

[Will Claude’s AI finance agents change your life?](#) (May 6, 2026)

[AI, Iran and the Gulf 101](#) (April 13, 2026)

[What AI says about AI eating itself and the world...](#) (Feb 16, 2026)

[AI Q&A: Why have market vibes suddenly shifted on AI?](#) (Feb 16, 2026)

[Three AI themes for 2026: the honeymoon is over](#) (Jan 20, 2026)



Appendix 1

This material has been prepared by the Deutsche Bank Research Institute and is provided to you for general information purposes only. The Institute leverages the views, opinions, and research from Deutsche Bank Research, economists, strategists, and research analysts. Accordingly, you should assume that content in this document is based on or was previously published and provided to Deutsche Bank clients who may have already traded on the basis of it.

Any views or estimates expressed in this material reflect the current views of the author(s) and may differ from the views and estimates of Deutsche Bank AG, its affiliates, other Deutsche Bank personnel, and other materials published by Deutsche Bank. The content in this material is valid as of the date shown on the first page and may change without notice. Deutsche Bank has no obligation to provide any updates or changes to the information herein.

This material should not be used as a basis for trading securities or other financial products and should not be considered to be a recommendation or individual investment advice for any particular person. It does not constitute an offer, solicitation, or an invitation by or on behalf of Deutsche Bank to any person to buy or sell any security or financial instrument. Nothing in this material constitutes investment, legal, accounting or tax advice. Deutsche Bank engages in securities transactions, including on a proprietary basis, and may do so in a manner inconsistent with the views or information expressed in this material.

While information in this material has been obtained from sources believed to be reliable, neither Deutsche Bank AG nor any of its affiliates makes any representation or warranty, express or implied, as to the accuracy or completeness of the statements or any information contained in this material and therefore any liability is expressly disclaimed. This material is provided without any obligation, whether contractual or otherwise. Information regarding past transactions or performance is not indicative of future results.

In the US this report is approved and/or distributed by Deutsche Bank Securities Inc., a member of FINRA. In Germany this information is approved and/or communicated by Deutsche Bank AG Frankfurt, licensed to carry on banking business and to provide financial services under the supervision of the European Central Bank (ECB) and the German Federal Financial Supervisory Authority (BaFin). In the United Kingdom this information is approved and/or communicated by Deutsche Bank AG, London Branch, a member of the London Stock Exchange, authorized by UK's Prudential Regulation Authority (PRA) and subject to limited regulation by the UK's Financial Conduct Authority (FCA) (under number 150018) and by the PRA. This information is distributed in Hong Kong by Deutsche Bank AG, Hong Kong Branch, in Korea by Deutsche Securities Korea Co. and in Singapore by Deutsche Bank AG, Singapore Branch. In Japan this information is approved and/or distributed by Deutsche Securities Limited, Tokyo Branch. In Australia, retail clients should obtain a copy of a Product Disclosure Statement (PDS) relating to any financial product referred to in this report and consider the PDS before making any decision about whether to acquire the product.

By accessing this material, you agree that its content may not be reproduced, distributed or published by any person for any purpose, in whole or part, without Deutsche Bank's prior written consent. You also agree that you shall not scrape, extract, download, upload, sell or offer for sale any of the content in this material, and you agree not to use, or cause to be used, any computerized or other manual or automated program or mechanism, tool, or process, including any scraper or spider robot, to access, extract, download, scrape, data mine, display, transmit, or publish, any of the content in this material.