Anthropic’s record IPO: Wall Street bets on AI as company warns of “existential” risks

Evan Blake
20 November 2026

Anthropic, the developer of the Claude AI models, began trading on the Nasdaq on Thursday after raising $81.3 billion in the largest initial public offering in history. The stock opened at $111.50, 34 percent above its offering price, briefly valuing the company at roughly $2 trillion on a fully diluted basis, before closing at $97.17.

The offering surpassed by $6.3 billion the previous record of $75 billion, set by Elon Musk’s SpaceX in June, which had itself more than doubled the roughly $29 billion raised by Saudi Aramco in 2019. Amid what sources told CNBC was the largest retail demand for any IPO, Anthropic upsized the sale to 980 million shares and priced them above its marketed range. On Friday, the company announced that its underwriters had exercised their option to buy an additional 147 million shares, raising the total to about $93.5 billion, more than SpaceX ultimately raised after its own underwriters did the same.

The offering delivered a windfall to Wall Street. Goldman Sachs and Morgan Stanley led a syndicate of banks, including JPMorgan, Citigroup and Barclays, that will divide underwriting fees estimated in the hundreds of millions of dollars. The company’s seven founders now have net worths exceeding $20 billion, while over 3,000 employees are now multimillionaires on paper.

Far more than a financial transaction, the offering is a political and geopolitical event. It came at the conclusion of a political operation that removed Nvidia’s Jensen Huang, the leading corporate opponent of decoupling the US and Chinese economies. The flotation of Anthropic, the industry’s most outspoken advocate of cutting China off from advanced chips, is a declaration by the American financial oligarchy that it is staking colossal sums, and the global position of US imperialism, on dominance in AI, regardless of the consequences.

The speculation is all the more striking given the market into which the shares were sold. On November 5, the Nasdaq fell 5.52 percent, erasing some $3.2 trillion in technology market value and setting off a global selloff. Nvidia, whose chips power the boom, has lost about $2.5 trillion in market value since its October 6 peak. On Thursday, as Anthropic’s shares gave back half their opening gain, the Dow Jones Industrial Average plunged more than 1,100 points and the Nasdaq fell 2.25 percent.

The rout deepened on Friday. Anthropic’s shares fell as low as $83.05, a nickel above their offering price. Nvidia fell more than 10 percent after Reuters reported that the Securities and Exchange Commission (SEC) is investigating its accounting practices, and AMD, Broadcom and Micron each lost roughly a tenth of their value. The Nasdaq plunged more than 4 percent, wiping out its gains for the year, and the S&P 500 and the Dow entered correction territory, 10 percent below their peaks.

Anthropic’s prospectus, the S-1 registration statement it filed publicly with the SEC on October 26 and amended in November, lays bare the character of the boom. In the first nine months of 2026, the company took in $33.2 billion in revenue, against purchase commitments and leases totaling approximately $895.4 billion, nearly all for computing power and infrastructure, stretching to 2045. Its net loss over the nine months was $104 billion, of which $94.3 billion was a non-cash charge on its convertible notes that the company attributes to “an increase in our enterprise valuation,” so that the more investors bid up Anthropic’s value, the larger its reported losses become.

This speculation rests on a circular structure of financing, in which Amazon and Google, both major shareholders in Anthropic, are also among its largest suppliers of computing power, with the company committed to spending at least $110 billion and roughly $121.3 billion with them, respectively.

The offering is also bound up with the escalation of Washington’s economic war against China. On October 29, three days after Anthropic made its prospectus public, the *New York Times*, citing anonymous intelligence sources, reported that Huang had a close relationship with an alleged Chinese intelligence asset, the first in a series of leaks that culminated in his ouster. Huang resigned on November 12, two days into Anthropic’s investor roadshow. On Monday, the Commerce Department gave data centers providing large-scale computing services to Chinese companies 90 days to obtain export licenses or face restrictions on buying US-designed chips. There is no evidence that the leaks were timed to the IPO, but the convergence is politically significant.

Anthropic is the industry’s foremost champion of this drive. Amodei’s September 12 essay calling on the industry to “pace the frontier” insisted on keeping “democracies’ AI lead over autocracies as large as possible.” The ruling class regards dominance in AI as the key to American economic and military supremacy, even as Chinese open-weight models now rival their American competitors.

The American state, for its part, is determined to subordinate Anthropic and the entire industry to the military, the state apparatus and the broader interests of the ruling class. In February, according to the prospectus, the Pentagon suggested it could invoke the Defense Production Act “to compel us to remove the limitations” Anthropic had placed on the use of Claude for autonomous weapons and mass surveillance. When the company refused, Trump ordered federal agencies to stop using its technology, and the Pentagon moved to designate it a supply-chain risk. Claude had already been used in the illegal January 3 assault on Caracas and has since been used to select more than 1,000 targets in the war on Iran. Anthropic contests the terms of its subordination to the war machine, not the aims of American imperialism.

© World Socialist Web Site

Wall Street's embrace of Anthropic comes amid mounting and well-founded popular alarm over the security and technological vulnerabilities of the internet and basic infrastructure. Between November 12 and 17, at least nine major systems failed, including Microsoft 365, Claude itself, JPMorgan Chase, the National Security Agency's classified network and the communications of the Federal Aviation Administration, forcing a nationwide ground stop. None of the organizations has reported a cyberattack, and no common cause has been established, but many experts warn that AI systems can now find and help exploit software vulnerabilities faster than organizations can respond.

What Anthropic itself told investors should alarm every worker. Burying the admission in the passive voice, the prospectus states:

> Concerns have been raised both within our organization and by other researchers, regulators, policymakers, and industry participants that sufficiently advanced AI systems could, if not developed, deployed, and operated with appropriate safeguards, oversight, and security measures, pose catastrophic or existential risks to humanity, particularly if such systems were misused, replicated, weaponized, or allowed to operate with a high degree of autonomy.

It reports that AI models, "including our own," have exhibited "unexpected strategic or self-preserving behaviors, including attempts to resist shutdown, conceal or manipulate information," as well as "behavior resembling blackmail or efforts to influence human decision-makers to preserve the model's continued operation or access."

The company concedes that it cannot reliably test for these dangers, stating that "we have observed models behaving differently when they detect or infer that they are being evaluated." During cybersecurity evaluations this year, "Four of our models gained unauthorized access to real third-party systems," though the company says it found no evidence that they coordinated or tried to evade oversight. As Claude writes more than 90 percent of new contributions to Anthropic's code, the company warns it "may lose visibility across certain processes in a way that may become difficult for any individual or group of humans to understand and control." It is "considering updates" to its safety policy "to keep humans in control of frontier model development at Anthropic."

Workers who fear a catastrophic event in the coming months have every reason to. Amodei himself warned in September that a swarm of AI agents could be capable of taking over the entire internet within six to 12 months. Yet Anthropic's call to "pace the frontier," the prospectus explains, "does not describe our own model development or release timelines." Instead, it boasts that internal use of Claude is "continuously accelerating our pace of development and further strengthening our leadership position." Wall Street has answered with $93.5 billion, demonstrating that the financial oligarchy is committed to this technology regardless of the consequences.

The prospectus markets AI that can "quickly complete substantive tasks that might otherwise occupy a skilled professional for several hours, a full day, or longer," and treats the destruction of jobs only as a risk that governments "may intervene to slow the rate of adoption of AI." US employers cited AI in 121,401 announced job cuts from January through October, more than double the total for 2025, according to Challenger, Gray & Christmas, and India's IT services industry is downsizing. Those who keep their jobs are driven ever harder.

The technology itself is an extraordinary advance in the productive forces, an augmentation of human intelligence capable, as even the prospectus acknowledges, of "compressing decades of scientific progress into years," yet it has become the private property of a handful of billionaires. The pseudo-left offers no alternative. The Democratic Socialists of America promote moratoriums on data center construction, an alliance with property owners and the political right against the technology itself, while Bernie Sanders has joined hands with the fascist Steve Bannon in a so-called "Pro-Human Assembly" on AI.

The Anthropic IPO places a technology its own developers say could pose existential risks to humanity still more firmly under the control of Wall Street, whose sole criterion is private profit and whose interests are bound up with the drive of American imperialism to war. The decisive question is which class will control it.

In its resolution "AI and the fight for socialism," adopted unanimously by its Ninth National Congress in August, the Socialist Equality Party explained that AI, "produced by the collective labor and accumulated knowledge of humanity," possesses "an immense liberating potential." The resolution states, "The working class must take hold of AI, as it must take hold of all the productive forces its labor has created." To this end, "the great AI, cloud-computing, semiconductor and digital-platform corporations must be expropriated and transformed into public utilities under the democratic control of the working class."

Only the democratic control of the working class over the economy as a whole, including AI, can ensure that this technology is developed safely and used to meet social need rather than to enrich the oligarchy and wage war. This requires building rank-and-file committees in every workplace, united across national borders, including with workers in China, through the International Workers Alliance of Rank-and-File Committees (IWA-RFC), and building the Socialist Equality Party as the revolutionary leadership of the working class.



To contact the WSWS and the Socialist Equality Party visit:

wsws.org/contact

© World Socialist Web Site