Filed Pursuant to Rule 424(b)(4)
Registration Nos. 333-299633 and 333-300148

Class A Common Stock 980,000,000 Shares Anthropic, PBC

This is an initial public offering of shares of Class A common stock of Anthropic, PBC. We are selling 956,351,354 shares of Class A common stock and the selling stockholders identified in this prospectus are selling an additional 23,648,646 shares of Class A common stock. We will not receive any proceeds from the sale of shares of our Class A common stock by any of the selling stockholders.

Prior to this offering, there has been no public market for our Class A common stock. The initial public offering price per share is $83.00. We have been approved to list our Class A common stock on the Nasdaq Global Select Market under the symbol "ANT."

We have a novel capital and corporate governance structure. We are incorporated as a public benefit corporation under Delaware law, a structure consistent with our commitment to artificial intelligence ("AI") safety and responsibility. As a public benefit corporation, we are required to balance the pecuniary interests of our stockholders, the best interests of those materially affected by our conduct, and the specific public benefit purpose identified in our amended and restated certificate of incorporation: the responsible development and maintenance of advanced AI for the long-term benefit of humanity. We created our Long-Term Benefit Trust (the "LTBT") to help ensure that Anthropic, PBC is managed in a manner that responsibly balances its goals as a public benefit corporation. The LTBT is intended to operate as an independent body of financially disinterested trustees with expertise in AI safety, national security, public policy, and social enterprise. Through a special class of common stock, Class T common stock, the LTBT generally has the authority to elect and remove the Class T directors, who together will constitute a majority of the authorized members of our board of directors. Our Class T common stock does not otherwise have voting rights with respect to matters submitted to stockholders generally and has no economic rights. Immediately prior to the completion of this offering, we will also issue a single share of Class F common stock, a special class of our common stock, to a limited liability company organized by our co-founders (the "Founder LLC"). Dario Amodei, our Chief Executive Officer and a member of our board of directors, and Daniela Amodei, our President and Chairperson of our board of directors, together with our five other co-founders, will initially be the sole members of the Founder LLC. Our Class F common stock, which will be held by the Founder LLC, will be entitled to an aggregate number of votes that equals 50.1% of the total voting power of our outstanding shares of capital stock with respect to matters submitted to stockholders generally, including the election of our directors other than the Class T directors. Our Class F common stock will not have any economic rights. Immediately prior to the completion of this offering, we will also authorize a special class of our common stock, Class E common stock, which will be issued to our employees, excluding our co-founders. Our Class E common stock will be entitled to vote only on certain limited matters and will not have any economic rights. In addition, we have authorized Class A common stock and Class N common stock. The rights of the holders of Class A common stock and Class N common stock are identical, except with respect to voting and conversion. Each share of Class A common stock is entitled to one vote and shares of Class N common stock have no voting rights, except as otherwise required by law. The holders of our Class N common stock are primarily strategic business partners, and shares of Class N common stock will generally convert to Class A common stock upon a sale or transfer to an unaffiliated third party. Our Class A common stock is not convertible. See the sections titled "Risk Factors—Risks Related to our Governance Structure," "Our Governance Structure and Its Goals," and "Description of Capital Stock" for more information.

We are considered a "controlled company" within the meaning of the Nasdaq Global Select Market corporate governance requirements. See the sections titled "Our Governance Structure and Its Goals," "Management," and "Description of Capital Stock" for additional information.

See the section titled "Risk Factors" beginning on page 30 to read about factors you should consider before buying shares of our Class A common stock.

Neither the Securities and Exchange Commission nor any other regulatory body has approved or disapproved of these securities or passed upon the accuracy or adequacy of this prospectus. Any representation to the contrary is a criminal offense.

	Per share	Total
Initial public offering price	$ 83.00	$ 81,340,000,000
Underwriting discount[1]	$ 0.56122	$ 550,000,000
Proceeds, before expenses, to Anthropic, PBC	$ 82.43878	$ 78,840,438,875
Proceeds, before expenses, to selling stockholders	$ 82.43878	$ 1,949,565,525

(1) The underwriters will not receive any discount or commission on any shares of our Class A common stock sold pursuant to their option to purchase additional shares of Class A common stock from us. See the section titled "Underwriting" for a description of the compensation payable to the underwriters.

To the extent that the underwriters sell more than 980,000,000 shares of Class A common stock, the underwriters have the option to purchase up to an additional 147,000,000 shares of Class A common stock from Anthropic, PBC at the initial public offering price.

The underwriters expect to deliver the shares against payment in New York, New York, on or about November 20, 2026.

We have entered into an agreement with each of Advanced Micro Devices, Inc., Broadcom Inc., Google LLC (funding a portion of its previously committed investment obligations under the Google 2026 NPA (as defined below)), Micron Technology, Inc., Microsoft Corporation, NVIDIA Corporation, and TSMC Partners, Ltd., pursuant to which they have agreed, severally and not jointly, to purchase an aggregate of $26.0 billion of our Class N common stock in a private placement at a price per share equal to the initial public offering price. The private placement is contingent upon, and is scheduled to close immediately subsequent to, the completion of this offering as well as the satisfaction of certain conditions as further described in the section titled "Concurrent Private Placement." The underwriters will not receive any underwriting discounts or commissions on the shares of Class N common stock sold in the concurrent private placement. The shares of Class N common stock sold in the concurrent private placement will not be registered in this offering and will not be subject to lock-up agreements with the underwriters, but will be "restricted securities" under the Securities Act of 1933, as amended (the "Securities Act"), and may be resold only pursuant to an effective registration statement or an available exemption from registration.

Certain members of our board of directors and entities affiliated or associated with them have indicated an interest in purchasing shares of our Class A common stock in this offering at the initial public offering price at an aggregate purchase price of up to $305 million. Because these indications of interest are not binding agreements or commitments to purchase, the underwriters could determine to sell more, fewer, or no shares to any of these individuals or entities and these individuals or entities could determine to purchase more, fewer, or no shares in this offering. The underwriters will receive the same underwriting discounts and commissions on shares of our Class A common stock purchased by these individuals or entities as they will from other shares of our Class A common stock sold to the public in this offering. Any shares purchased by these individuals or entities will be subject to lock-up restrictions described in the section titled "Shares Eligible for Future Sale."

Goldman Sachs & Co. LLC **Morgan Stanley**

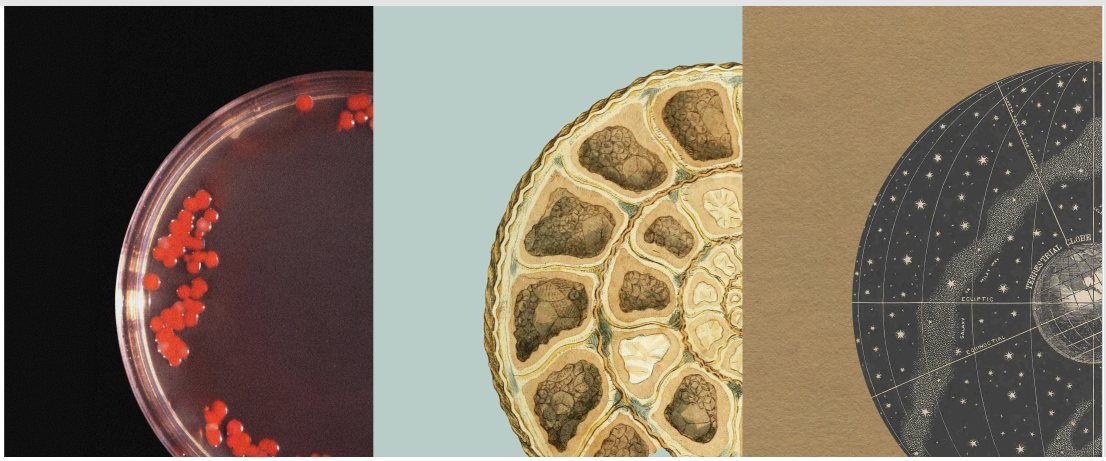
J.P. Morgan Citigroup Barclays Wells Fargo Securities

BofA Securities Deutsche Bank Securities RBC Capital Markets UBS Investment Bank

BMO Capital Markets BNP PARIBAS Credit Agricole CIB Jefferies Mizuho MUFG TD Securities

BTIG Cantor Capital One Securities Macquarie Capital Societe Generale

Prospectus dated November 18, 2026





We believe that AI is one of the most consequential technologies in human history. If managed well, it could rapidly accelerate scientific innovation, greatly increase economic growth rates, bring cures for major diseases within reach, and solve problems that have resisted generations of human effort.

Anthropic aims to realize these benefits while navigating and preventing the risks that come with a technological revolution this fast and this powerful.

There are four things that we would like any investor in Anthropic to understand: the incredible potential of AI, the risks we must address, Anthropic's unique approach to building a company and operating in the world, and the ways in which our behavior may differ from that of a purely profit-maximizing company.

AI's potential

In 2020, four of our founders co-wrote a paper that documented the scaling laws of AI. The scaling laws describe how increasing compute, data, model size, and training environments in combination produces a rapid increase in AI's general cognitive capabilities and, in turn, its economic value. To many at the time, the idea that machine intelligence would so strongly depend on scale was surprising, but predictions made on the basis of the scaling laws have proved true again and again.[1]

The phenomenon is somewhat analogous to Moore's law for computer chips, under which the computing power of these chips doubled every 12 to 18 months for several decades. Moore's law has been the basis of a succession of industries—from the personal computer to the internet—that have revolutionized the modern economy, but we believe that the impact of an exponential in machine cognition will be vastly greater even than that of the exponential in computer chips.

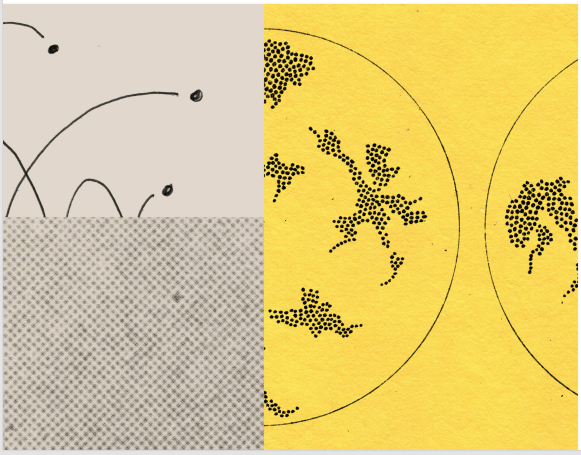
[1] The idea is closely related to Richard Sutton's concept of the bitter lesson, which describes how sheer scale has consistently won out over more specific, complex methods of training AI models. Note that scale is not the only source of improvements in model capabilities—algorithmic efficiencies ("compute multipliers"), which allow us to (for example) achieve higher capabilities with the same amount of compute, have played a major role. But these efficiencies are themselves found using the scaling laws: researchers test new ideas and keep the ones that improve the scaling trend the most.

It is Anthropic's thesis that the exponential in cognition is likely to continue, eventually reaching the level of what we have called "a country of geniuses in a data center"—millions of virtual agents with the ability to revolutionize every area of human endeavor: software, physics, math, biology, medicine, finance, economics, commerce, robotics, military and intelligence capabilities, and much more. Anthropic has always focused first and foremost on working with other businesses, because we understand that the extraordinary benefits of AI will mostly derive from the way our models are deployed by enterprises and other organizations across the world.

It is worth highlighting a few applications with particularly great potential to improve the quality of human life:

- First, we believe that AI will profoundly accelerate the rate of progress in biology and medicine, potentially curing diseases that have plagued humanity for millennia, including cancer, Alzheimer's disease, heart disease, rare genetic diseases, and the major infectious diseases. AI enables entirely new ways of doing scientific research and drug development, and we hope to play a leading role in developing tools to benefit the broader scientific community. Anthropic already works with many biotechnology and discovery challenges to cutting the time it takes to write clinical trial reports from months to minutes. To build the best tools for science, our teams need firsthand experience doing scientific research and drug discovery. We have started to directly make biological discoveries with Claude: among them is the discovery of a molecular machine that might have the potential to be developed into a novel biotechnology tool, perhaps for gene editing. We expect Claude to enable many more novel biological discoveries in the near future, and we will work with our partners to validate them, working together to raise the rate of discovery across the entire field.
- Second, AI can be a powerful lever for economic growth. As a general-purpose technology, it has the potential to raise productivity across every sector of the economy. Productivity growth in advanced economies has stagnated in recent decades, partly for structural reasons and partly because new ideas seem to be harder to find. AI could reverse this trend. AI can directly make workers and firms more productive in tasks that were previously hard to automate (some analyses suggest this alone could undo the recent productivity slowdown). But AI is also an "innovation in the method of innovation"—it can improve the very process by which businesses come up with new products and new ways of

pharmaceutical companies to improve both basic discovery and clinical trials, from helping pharmaceutical company scientists tackle drug

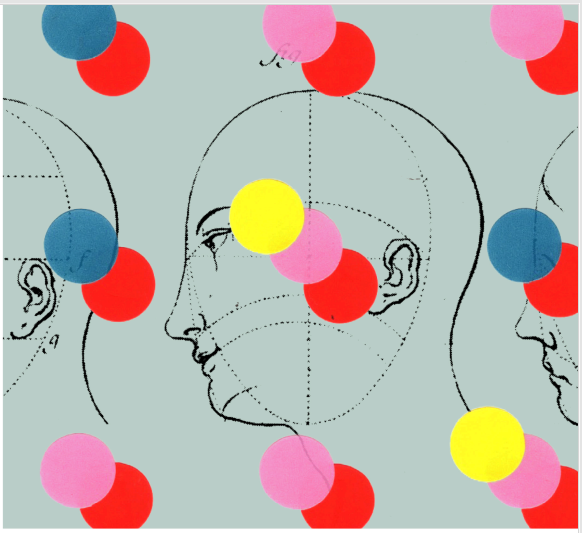


working. That could bring about a sustained acceleration in economic growth, with all the benefits for human well-being that growth can bring.

- Third, AI can diffuse these benefits throughout the world, helping to raise *global* economic growth rates. Today, even when medicines for terrible diseases are developed, they do not always reach everyone who needs them, due to failures of logistics, staffing, or access in poorer parts of the world. An estimated 4.6 billion people, more than half the world's population, still lack access to essential health services, according to the World Health Organization ("WHO") and the World Bank. We believe AI can help here as well, by extending scarce expertise, strengthening the data that health systems run on, and helping get treatments to where they are most needed.

We already have active collaborations on epidemiological modeling for malaria and other diseases in Africa. For example, during the 2026 Ebola outbreak in the Democratic Republic of the Congo, the WHO's Regional Office for Africa used Claude to cut the time needed to prepare its daily situation report from a full day to under an hour, allowing its team to run several disease models instead of one when deciding where to place treatment centers. With the Gates Foundation, we are helping health ministries use their data to deploy health workers, manage supply chains, and detect outbreaks, and we support free AI-powered clinical search for doctors in roughly 100 low- and middle-income countries, from Uganda to Haiti to Mongolia. With the International Rescue Committee, we are exploring how Claude can be used to digitize and analyze children's health records in Nigeria, helping workers identify children at risk of malnutrition and make limited resources go further. More broadly, AI can help to strengthen markets where they are currently weak. Mobile phones did this in the developing world in the 2000s, giving people access to critical information on prices that allowed trade to flow. AI could do the same for expertise, giving anyone access to the knowledge and skills of a lawyer, an accountant, or an agronomist, improving economic efficiency and accelerating growth across the world.

- Fourth, AI could raise the quality of government services and build technological state capacity. Tax authorities, social services, medical programs, planning and permitting—these and similar functions are often frustrating for citizens to interact with, and they could all be radically improved and modernized with sufficiently powerful AI. We already work with the U.S. federal government, many state governments, and U.S. allies in areas such as child services, social security administration, cybersecurity, and intelligence collection. In Utah, developers across ten state agencies are using Claude to modernize the software behind public services; in California,



Claude is helping to shorten the lines at the DMV and to better serve people on Medicaid. In the United Kingdom, it is helping job seekers find their way through government services. Similarly, AI can be used to promote and strengthen democracies and bend the arc of geopolitical history in the direction of freedom, by helping democratic governments deliver on their promises to citizens, defending democratic institutions against cyberattacks and foreign interference, and ensuring that free societies lead in the world's most powerful technology.

These are far from the only applications of AI, but they illustrate the remarkable potential of the technology to build a better world. Alongside our customers, we intend to apply our AI models to revolutionize these and dozens of other areas of human endeavor. Doing so may well represent the greatest economic opportunity in history.

AI's risks

To build AI is to harness cognition itself, and anything so powerful can be used for ill as well as good. Anthropic is a deeply values-driven organization, founded in large part to help the world realize AI's benefits and avoid its dangers.