FWP 1 anthropicpbc-fwpipowebsi.htm FWP

Issuer Free Writing Prospectus dated November 10, 2026
Filed Pursuant to Rule 433
Relating to Preliminary Prospectus dated November 10, 2026
Registration File No. 333-299633

ANTHROP\C

This free writing prospectus relates to the initial public offering of Class A common stock of Anthropic, PBC ("Anthropic"). On November 10, 2026, Anthropic filed Amendment No. 2 to its Registration Statement on Form S-1 (File No. 333-299633) ("Amendment No. 2") with the Securities and Exchange Commission (the "SEC"). Amendment No. 2 and a copy of the preliminary prospectus dated November 10, 2026 relating to this offering may be accessed through the following link:

https://www.sec.gov/Archives/edgar/data/2130357/000162828026072484/anthropicpbc-sx1ax2.htm

On November 10, 2026, Anthropic launched a standalone IPO website containing information relating to its proposed initial public offering of Class A common stock. Screenshots of the IPO website are attached hereto as Exhibit A, which includes financial highlights, case studies, participation information, and frequently asked questions. The IPO website also includes an IPO factsheet attached as Exhibit B and a letter from Anthropic's founders attached as Exhibit C.

Anthropic has filed a registration statement (including a prospectus) with the SEC for the offering to which this communication relates. Before you invest, you should read the prospectus in that registration statement and other documents Anthropic has filed with the SEC for more complete information about Anthropic and this offering. You may get these documents for free by visiting EDGAR on the SEC Web site at www.sec.gov. Copies of the preliminary prospectus

may be obtained from: Goldman Sachs & Co. LLC, Attention: Prospectus Department, 200 West Street, New York, NY 10282, by telephone at (866) 471-2526, or by email at prospectus-ny@ny.email.gs.com; or Morgan Stanley & Co. LLC, Attention: Prospectus Department, 180 Varick Street, 2nd Floor, New York, NY 10014.

***

Exhibit A

IPO Website Screenshots

ANTHROP\C

Anthropic's Initial Public Offering

AI will have a vast impact on the world. Anthropic is a public benefit corporation dedicated to securing its benefits and mitigating its risks.

01 Founders' letter
02 Roadshow video
03 Financial highlights
04 Our customers
05 How to participate
06 Documents and FAQs
07 Legal notice





"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."

Read the Founders' letter →



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.

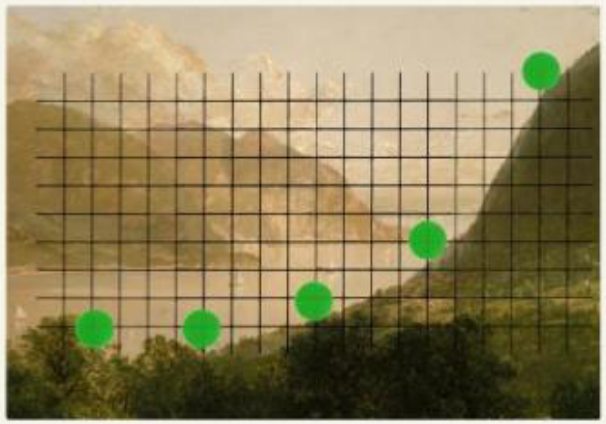
AI's potential

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 Rich and Sutton's concept of the Bitter Lesson, which describes how a more scale has consistently won out over more specific, complex methods of solving AI problems. From that scale is not the only source of progress in AI and capabilities—algorithmic breakthroughs ("conceptual metaphors"), which allow us to do meaningful additional higher capabilities with the same amount of compute, have played a major role. But those algorithmic breakthroughs, I argue, emerge naturally from the scaling laws: researchers and users alike are always the most clear signal of the scaling level we're at.

Q3'26 Financial highlights



$17.0B
Revenue

>12x
YoY revenue

>11,900
Customers > $100K TTM revenue[1]

>12x
YoY customers > $100K TTM revenue[1]

>4x
Compute efficiency gain since Q1'25[2]

15%
Adjusted Operating Margin[3]

$23.3B
Net Loss
(due to an increase in our valuation, net loss includes a $21.9 billion non-cash charge related to a change in the fair value of convertible notes, derivative liabilities, and SAFEs)

(3)%
Operating Margin

[1] As of September 30, 2026.
[2] Measured as revenue divided by compute and infrastructure expenses in Q3'26 divided by the same ratio in Q1'25.
[3] See the section in our prospectus titled "Management's Discussion and Analysis of Financial Condition and Results of Operations—Non-GAAP Financial Measures" for a reconciliation of Adjusted Operating Margin to Operating Margin.

Customer stories

LEGORA

01 02 03 04 05

Scaling an agentic operating

"AI is becoming fundamental to how legal work gets done, and it

system for lawyers

Download case study

will reshape how law firms and in-house teams operate. We have found Claude to be brilliant at instruction following and powering agentic workflows."

Jacob Lauritzen
CTO at Legora

1,500+	Law firms and in-house teams served
4.3	Non-billable hours recaptured per lawyer weekly, in an independent study of 3t firms
100%	Claude Enterprise employee access
14x	Claude Code growth (about 50 seats to 700 company-wide)

Lilly

Rakuten

Salesforce

Walmart

01 Legora

Lilly 02

Bringing the most powerful tools in science to every team

Download case study

"Science and technology have always advanced together. So we gave Claude to our scientists first. Now Claude is across the company, and we can't imagine going back. Patients are waiting, so we can't."

Diogo Rau
Executive Vice President & Chief Information and Digital Officer at Eli Lilly and Company

50,000+	Employees with access to Claude Enterprise
5x	Expansion since rollout
>85%	Employees in key functions actively engage with Claude

03 Rakuten

04 Salesforce

05 Walmart

01 Legora

02 Lilly

Rakuten 03

Turning every employee into a builder

Download case study

"Agents like Claude Code shine when we work with them to minimize the cost of new innovation—a quick transition from idea to production. Give a capable person agents that hold context and taste, and it lets hidden talent scale their potential 100 times over."

Yusuke Kaji
General Manager, AI for Business at Rakuten

70+	Business units using AI
99.9%	Code accuracy delivered on complex code modifications
7+ hrs	Autonomous coding
24→5 days	Working days for new feature time to market

04 Salesforce

05 Walmart

01

02

03

salesforce 04

Re-engineering how a global enterprise builds itself

Download case study

"By fusing Claude's extraordinary reasoning with the trusted data, workflows, and governance every enterprise runs on, we're delivering a dynamic interface that thinks, reasons, and acts. This is how every business will run."

05

Legora

Lilly

Rakuten

Marc Benioff
Chair & CEO at Salesforce

+151%	Effective output per developer year over year
18x faster	API migration (231 days of work in 13 days)
2x	Quarter over quarter productivity gains (estimated 8.1 million hours annualized)
99%	Salesforce employees using Slackbot

Walmart

01 Legora

02 Lilly

03 Rakuten

04 Salesforce

Walmart

05

Making shopping faster, easier, and more personal

Download case study

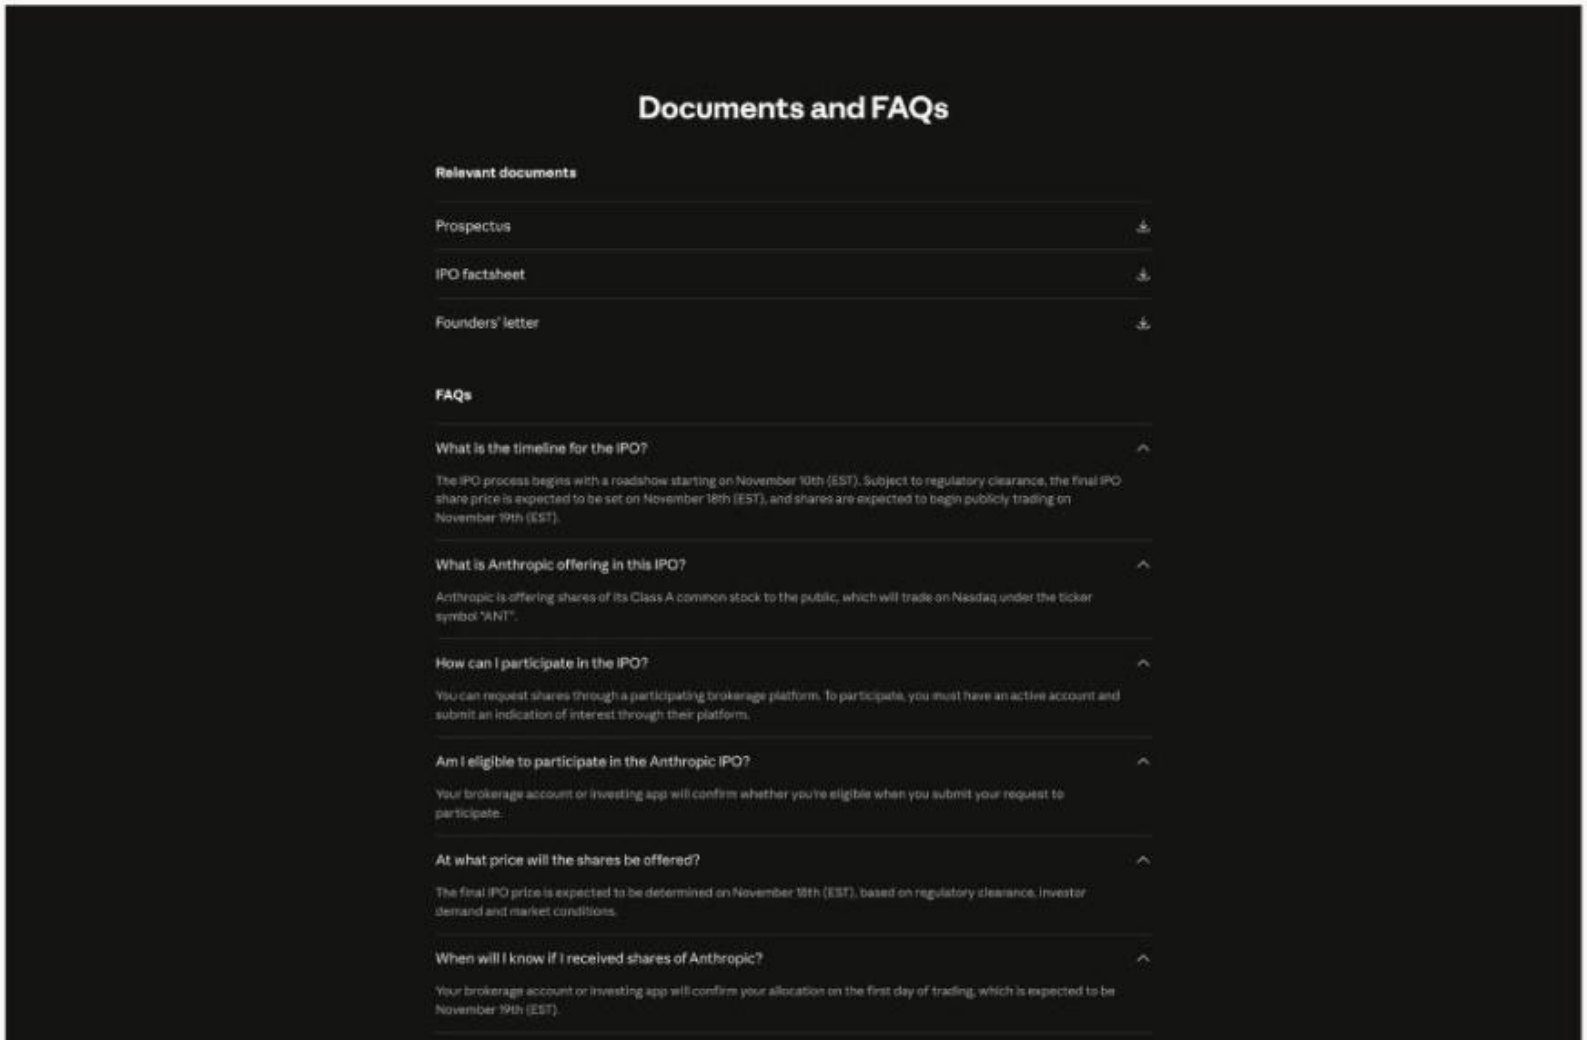
"We're using AI to make shopping more intuitive and engineering more productive. Agents like Claude are helping us do both, from better understanding customer intent to accelerating repeatable development work."

Daniel Danker
EVP AI Acceleration, Product & Design at Walmart

110%+	Growth in Sparky weekly active users quarter over quarter
~35%	Increase in average order value via Sparky vs. non-Sparky orders
2x+	Increase in units sold via Sparky quarter over quarter

How to participate

charles SCHWAB | E*TRADE | Fidelity | Robinhood | SoFi



What is the difference between buying in the IPO and buying after the stock begins trading?

Shares bought in the IPO will be purchased at the offering price, which is expected to be set on November 18th (EST). Shares bought after Anthropic starts trading will be purchased at the prevailing market price at that time.

Will I receive all the shares I apply for?

There is no guarantee that any investor will receive all or any of the shares they apply for in the Anthropic IPO. However, you should be ready to purchase all the shares for which you apply.

What is the prospectus?

A prospectus is a regulatory filing that provides detailed information about Anthropic's business, financial statements, risk factors, management team, use of proceeds, offering details, among other items. Investors should review the prospectus closely before investing.

Legal notice

The Company has filed a registration statement (including a U.S. prospectus), with the U.S. Securities and Exchange Commission ("SEC") for the global offering. The prospectus in that registration statement and other documents the Company has filed with the SEC include additional information about the Company and global offering. Investors may get these documents for free by visiting EDGAR on the SEC's website at www.sec.gov. The registration statement (including the U.S. prospectus) can also be obtained from Goldman Sachs & Co. LLC, Attention: Prospectus Department, 200 West Street, New York, NY 10282, by telephone at (866) 471-2526, or by email at prospectus-ny@ny.email.gs.com; or Morgan Stanley & Co. LLC, Attention: Prospectus Department, 180 Varick Street, 2nd Floor, New York, NY 10014.

The registration statement has not yet become effective. These securities may not be sold nor may offers to buy be accepted prior to the time the registration statement becomes effective. This communication does not constitute an offer to sell or the solicitation of an offer to buy these securities, nor shall there be any sale of these securities in any state or jurisdiction in which such offer, solicitation or sale would be unlawful prior to registration or qualification under the securities laws of any such state or jurisdiction.

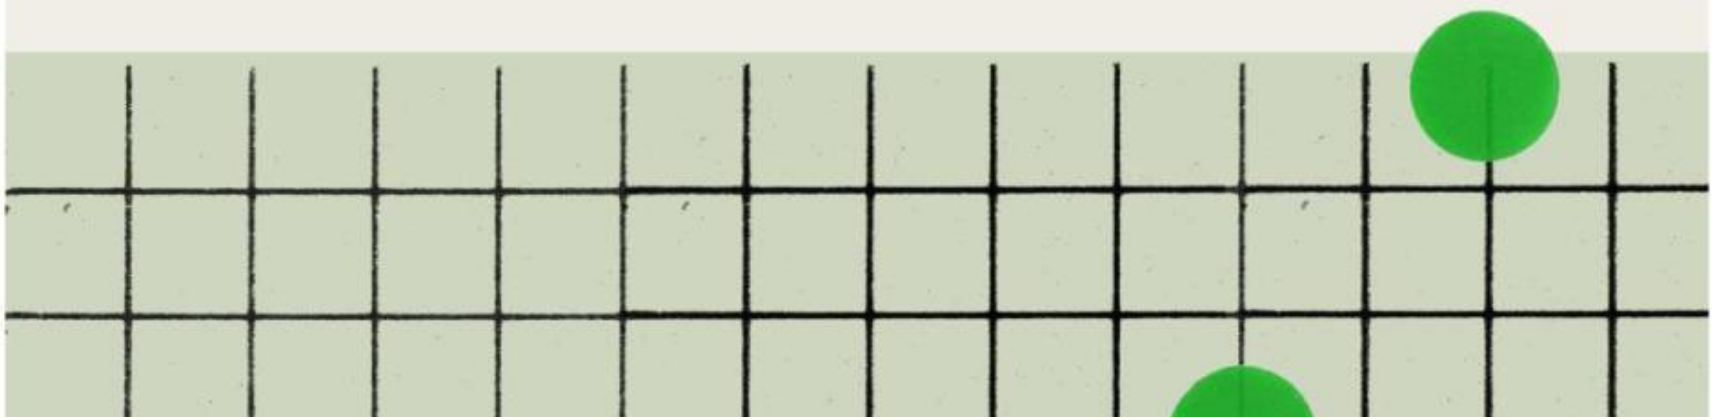
© ANTHROPIC, PBC

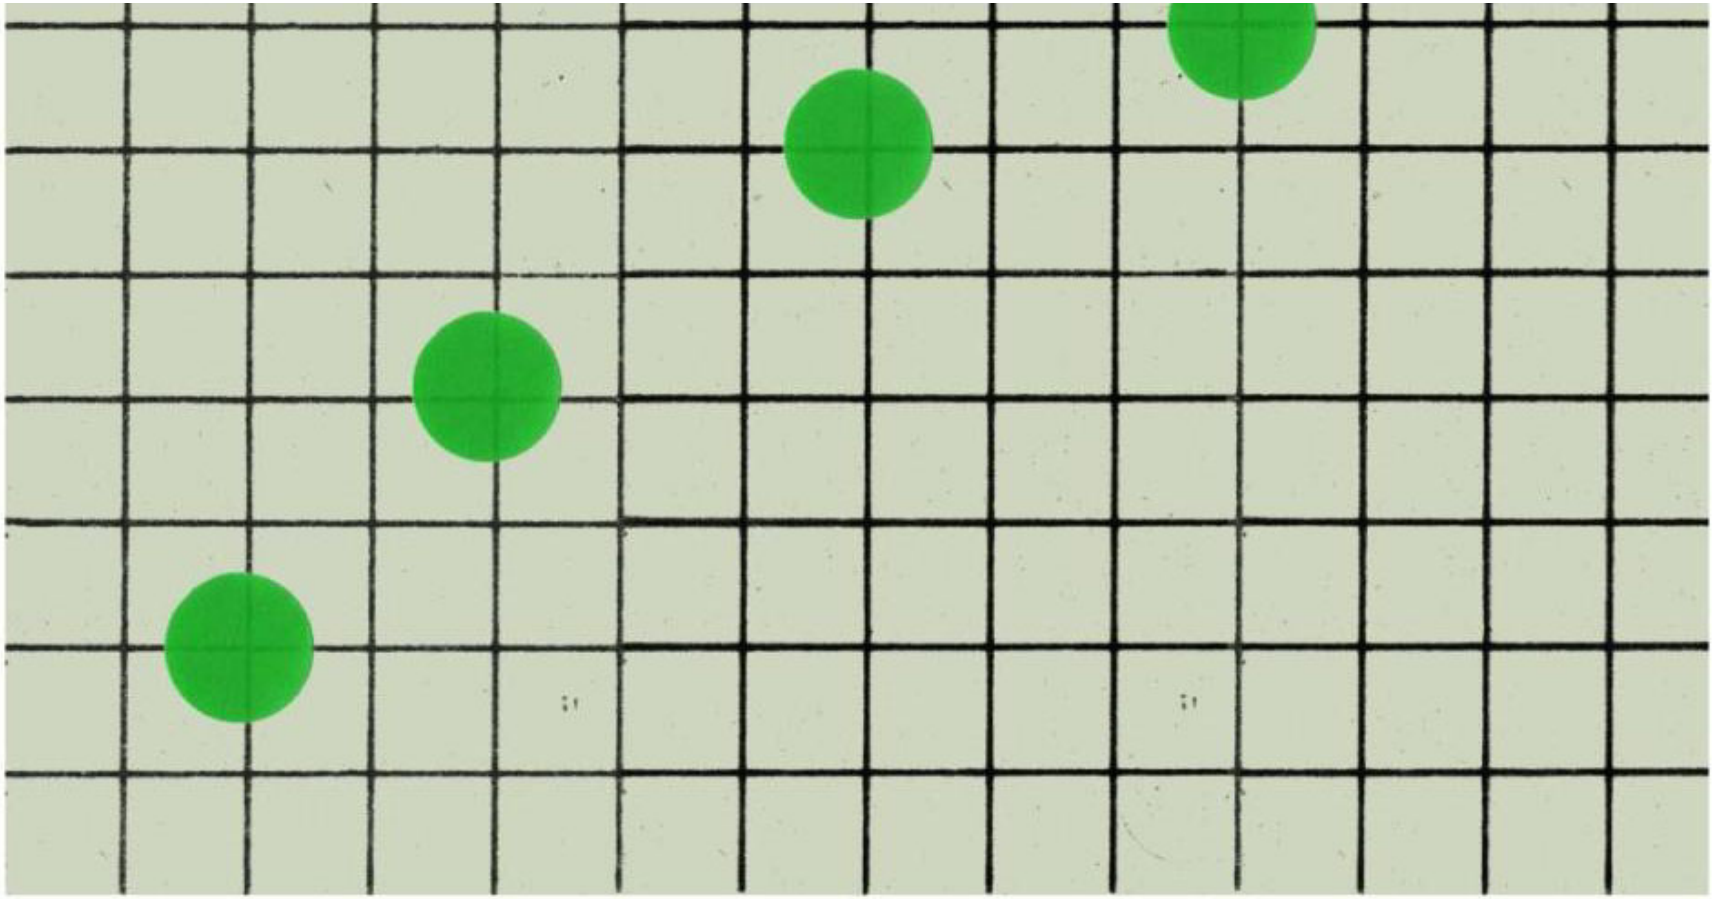
Exhibit B

IPO Factsheet

ANTHROP\C

November 2026





ANTHROP\C



Our mission

The responsible development and maintenance of advanced AI for the long-term benefit of humanity

We are in the foothills of the exponential

For the first time in history, intelligence is something humans can build. The cognitive capability to reason, analyze, synthesize, decide and create is producible at scale and accessible on demand.

$40T+

Knowledge work[1]

$2T+



Note: Curve interpolated between disclosed points. ¹We believe the work that knowledge workers perform represented more than $40 trillion in economic value globally in 2025. Please see Appendix for methodology. ²Forecasted global pharmaceutical spending in 2026; IQVIA Institute for Human Data Science. Global Medicine Use Trends 2026, February 2026.

ANTHROP\C

Claude overview

Our frontier AI, Claude, is delivered through models, platform and applications that allow customers to build, deploy and scale AI-native applications and workflows

Claude enables enterprises to reason, plan and execute across complex workflows with a level of trustworthiness and judgment that we believe was not available from earlier generations of AI

The broader shift in how Claude is used inside the enterprise reflects its evolution from a tool that responds to discrete prompts to a Virtual Collaborator that performs complex, long-horizon work alongside employees

Customers realize value from Claude through more productive employees, faster processes and embedded intelligence in their products

Why customers choose Claude

Access to **frontier AI** that delivers **real-world outcomes**

Trusted partner for **mission-critical** and **regulated workflows**

Integrated models, platform and apps accelerate **time-to-value**

Deep workflow integration into enterprise processes

Claude is accelerating our ability to close the loop

Each version of Claude helps build the next, continuously accelerating our pace of development and strengthening our leadership position



>90%

of the contributions to our codebase in September 2026 were written by Claude

ANTHROP\C

Core advantages

7 of 7
co-founders still with Anthropic

~1%
TTM regrettable employee attrition as of Q3'26

>4x
compute efficiency[1] gain from Q1'25 to Q3'26

Mission & focus

Our singular focus on our mission enables us to concentrate talent, research, and compute investment on the highest-value work for humanity. We are a public benefit corporation (PBC); our mission is our public benefit purpose, and the Long-Term Benefit Trust (LTBT), which currently consists of four financially disinterested trustees, generally has the authority to elect a majority of our Board.

Talent & culture

Our mission and focus enable us to attract and retain world-class talent within a low-ego, truth-seeking culture where AI capability and safety are mutually reinforcing.

Frontier & safety

Our frontier capabilities and safety research are deeply integrated, helping us produce

research

systems that are more reliable, steerable, and well suited for real-world deployment. Each version of Claude assists in building the next, helping to produce compounding improvements in the speed at which we innovate.

Approach to compute

Our approach to compute treats compute as a single, fungible pool across multiple strategic partners, enabling us to optimize price, performance, and compute productivity. We have achieved this through our compute multiplier research breakthroughs and our multi-chip architecture across AWS Trainium, Google TPUs, NVIDIA GPUs, and AMD GPUs.

Enterprise strategy

Our enterprise strategy is built on the belief that frontier AI is most valuable in real-world settings, where customer feedback drives continuous model improvement and broader adoption.

[1]Definition of compute efficiency can be found in the Appendix.

ANTHROP\C

Enterprise adoption of Claude

Virtual Collaborator

Claude Code | **Claude Chat + Cowork**

Claude and Claude Code power Virtual Collaborators that take a goal and work alongside humans, leveraging judgment, self-correction, tool use and long-running autonomous work to execute complex multi-step workflows

Select Enterprise Customers

AIG · Bristol Myers Squibb · DOORDASH · Jane Street · Lilly · mercado libre · Mozilla · PayPal · Rakuten · SAP · shopify · TRAVELERS · UST · verizon · Walmart

We meet enterprises where they are

Intelligence-led growth
Customers adopt Claude, see real value, and organically expand use

Direct sales and Applied AI
Global GTM teams scale Claude across industries, and Applied AI teams design and deploy custom solutions

Partnerships
Cloud partners - access to the three major clouds and their combined global salesforces
Ecosystem partners - Global SIs, data platforms and others accelerate deployment
Ode with Anthropic - bringing Claude's capabilities to mid-sized companies

More productive employees

Faster processes

Embedded intelligence

Diffusion of AI

Broad-based adoption across customers, with significant growth across spending levels[1]

>11.9K
customers with ⩾$100K TTM revenue

~8x
increase in customers ⩾$100K TTM revenue from Q4'25 to Q3'26

>1,900
customers with ⩾$1M TTM revenue

~250
customers with ⩾ $10M TTM revenue

Trusted by the world's largest enterprises, with significant runway for future expansion[1]

>30%
of the Global 2000 spend ⩾$100K on a TTM basis

>50%
of the Fortune 500 spend ⩾$100K on a TTM basis

<0.1%
Average Fortune 500 customer spend on Claude, as a % of their total expenses

<1%
Top 10 Fortune 500 customer spend on Claude, as a % of their total expenses

Significant expansion as customers realize the value of Claude[1]

~900%
net revenue retention

>50%
of enterprise customers use both our platform and at least one product[2]

Increasing model efficiency allows us to lower customer cost while preserving margins

~25%
lower customer cost with Fable 5.1 vs Fable 5[3]

~40%
lower customer cost with Opus 5.5 vs Opus 5[3]

Note: Definition of customers with $100K or more of revenue on a TTM basis and net revenue retention can be found in the Appendix. [1]Metrics as of Q3'26 unless noted otherwise. [2]Up from 30% in Q4'25. [3]Customer cost reflects Anthropic's published estimates of the cost to run the same workloads on each model vs. its predecessor at default settings.

ANTHROP\C

Financial model

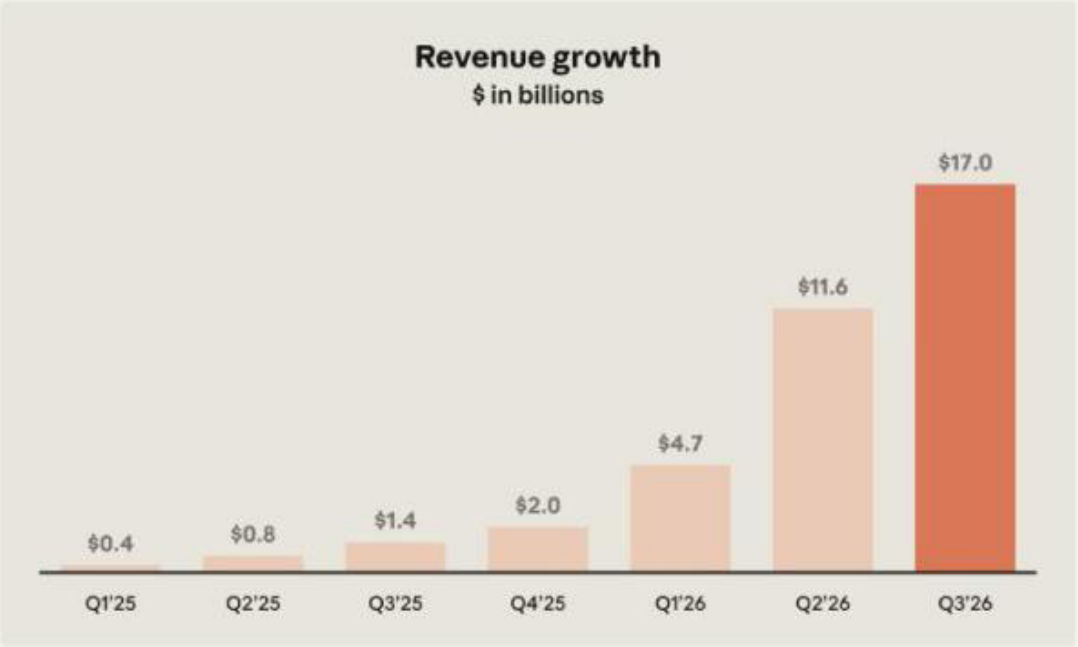
Combines rapid revenue growth, compute efficiency and operating efficiency, sustaining significant long-term investments to build frontier AI responsibly

>12x

YoY revenue growth as of Q3'26

~$81B

Run-Rate Revenue as of October 2026[1]



>4x

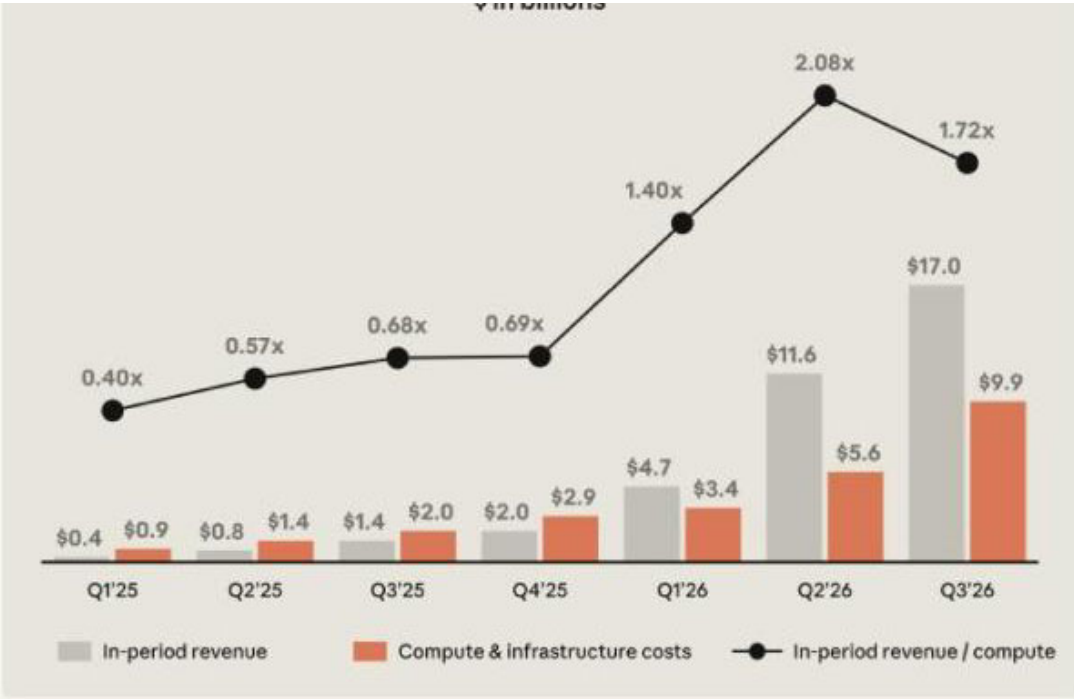
Compute efficiency

Increase in compute efficiency[2] from Q1'25 to Q3'26

Compute efficiency drivers

Compute multiplier advancements

Multi-chip compute approach



[1]Definition of Run-Rate Revenue can be found in the Appendix. [2]In the three months ended September 30, 2026, our underlying compute efficiency improvements were partially offset by the delivery of additional compute, including $5.0 billion of compute under short-term commitments, which currently costs significantly more than comparable compute procured under our multi-year commitments. Definition of compute efficiency can be found in the Appendix.

ANTHROP\C

>$6.6M

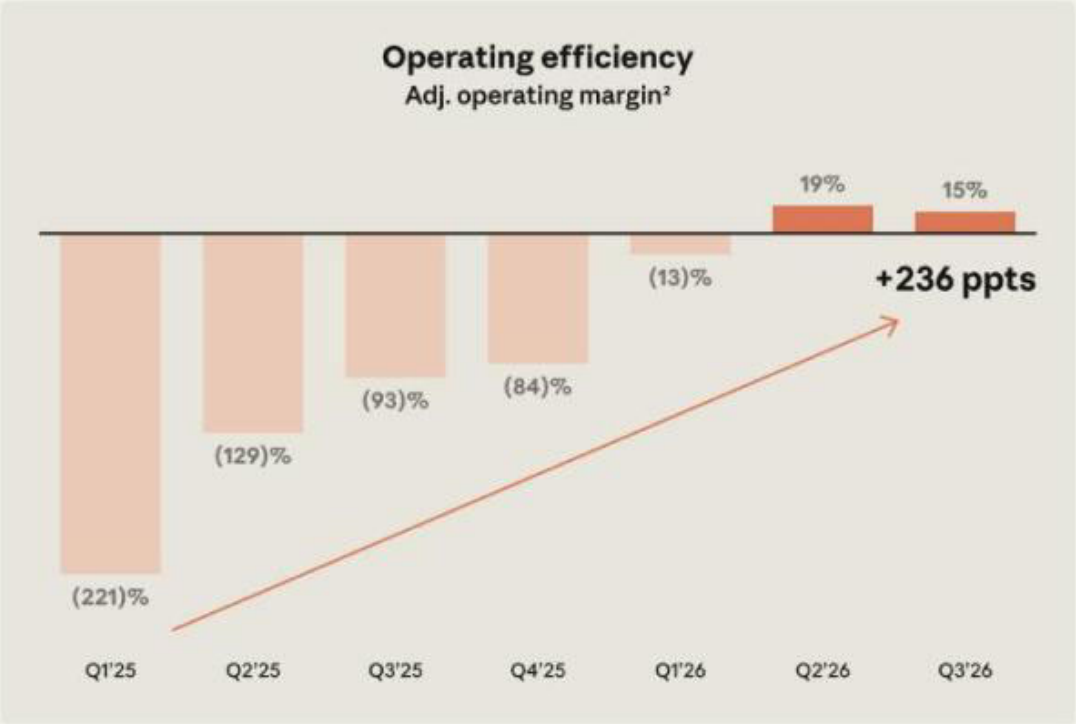
TTM revenue per employee as of Q3'26[1]

Operating efficiency drivers

High talent density

Efficient GTM

Extensive internal use of Claude



We are significantly scaling our compute fleet and are well capitalized to make investments at the frontier

~$900B

of compute commitments[3]

~$96B

of cash and cash equivalents[4]

~$36B

of short-term investments[4]

~$141B

of capital raised[4]

~$15B

of availability under A&R Credit Facility[4]

[1]TTM revenue per employee is calculated as revenue for the trailing twelve months ended September 30, 2026 divided by total headcount (full-time and part-time employees) as of September 30, 2026. [2]Adjusted operating margin is a non-GAAP measure and the definition of Adjusted operating margin and a reconciliation to the nearest GAAP measure can be found in the Appendix. [3]As of and subsequent to September 30, 2026. Includes approximately $451.6 billion expected to be incurred before December 31, 2030; substantially all of the remainder is expected to be incurred between 2031 and 2045. Approximately $70.7 billion of these commitments are due in the 12 months following September 30, 2026. [4]As of Q3'26.

ANTHROP\C

Growth strategy

Pursue strategic transactions and partnerships

Accelerate the mission, expand reach and strengthen competitive position

Extend the frontier into new markets

Accelerate discoveries and unlock new opportunities across emerging domains

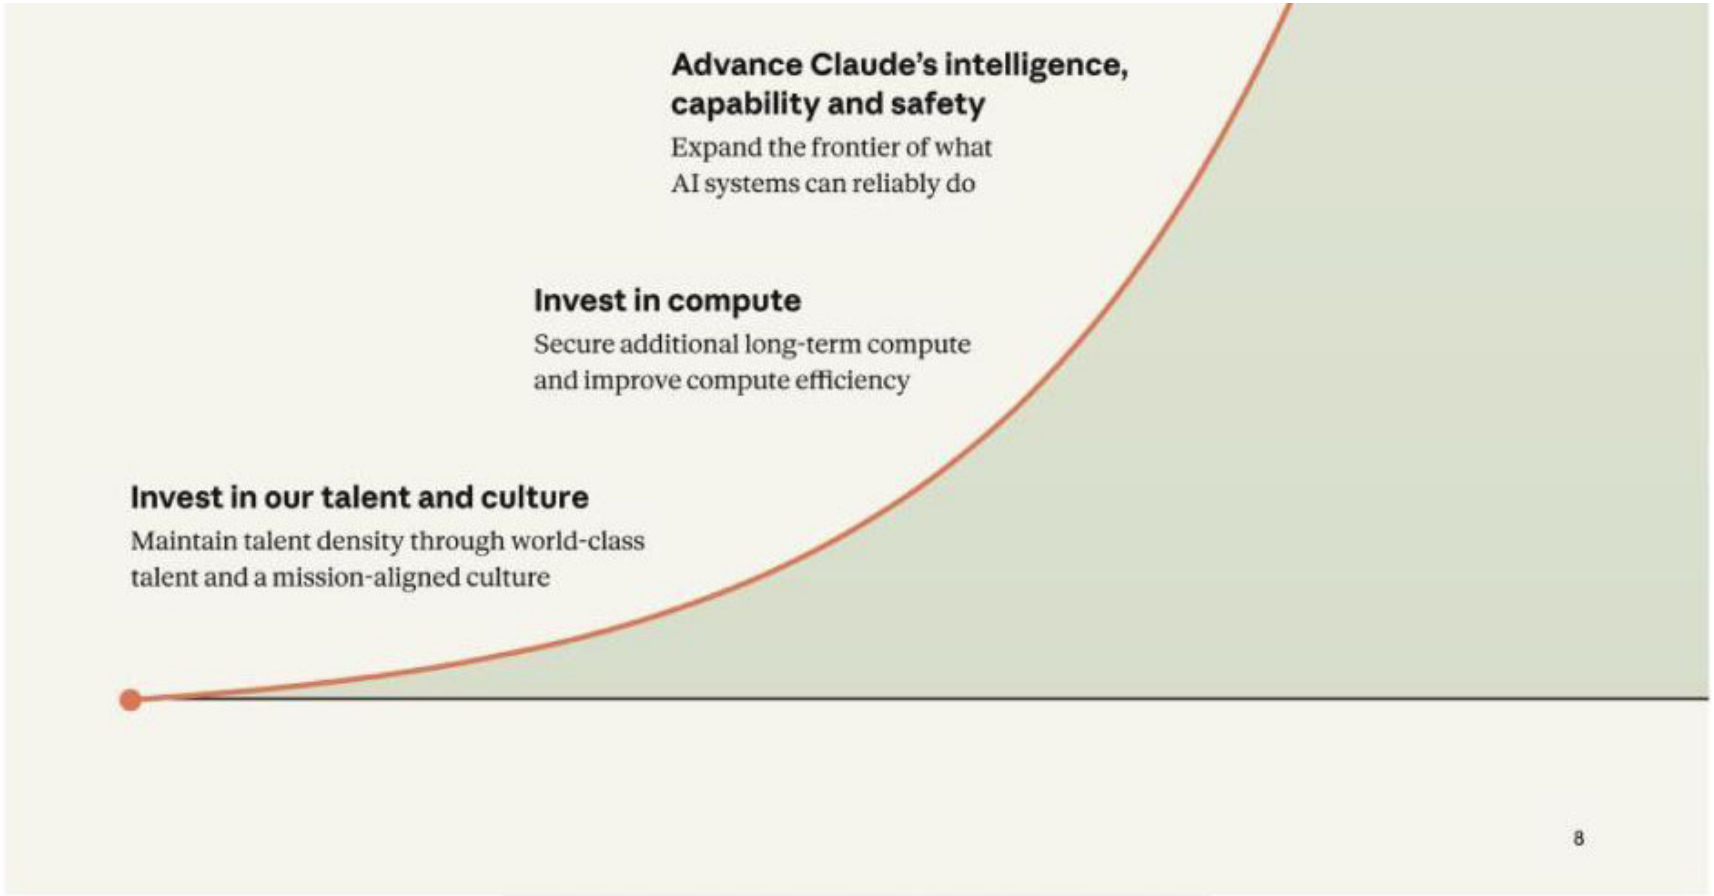
Drive adoption across existing and new customers

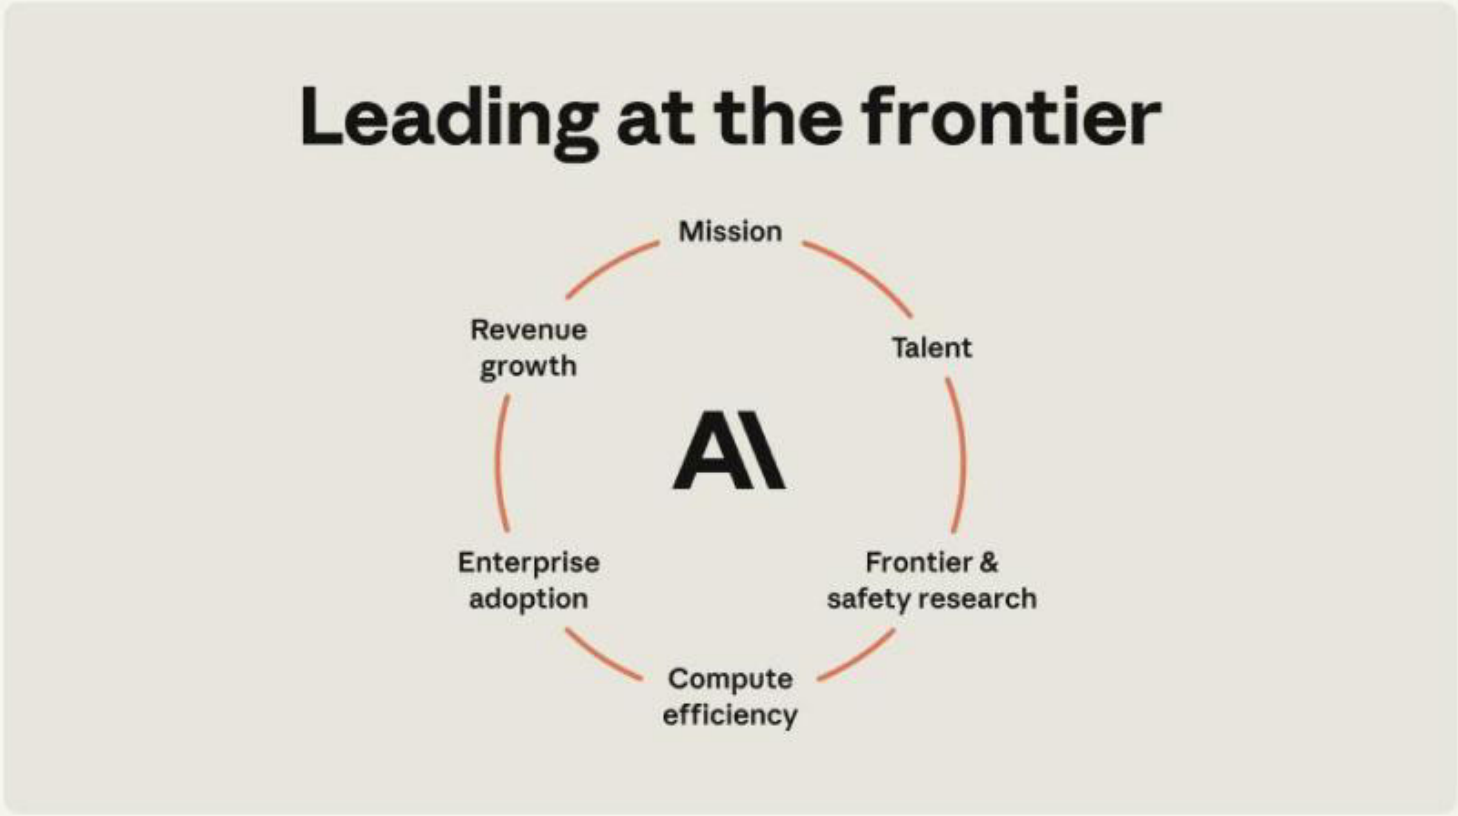
Deepen enterprise relationships and expand adoption globally

Make Claude available everywhere

Make Claude's intelligence accessible across every surface where work happens




Advance Claude's intelligence, capability and safety
Expand the frontier of what AI systems can reliably do
Invest in compute
Secure additional long-term compute and improve compute efficiency
Invest in our talent and culture
Maintain talent density through world-class talent and a mission-aligned culture



Disclaimer

Anthropic, PBC (the "Company") has filed a registration statement on Form S-1 (including a preliminary prospectus) with the Securities and Exchange Commission (the "SEC") for the offering to which this presentation relates, but it has not yet become effective. The Company's Class A common stock may not be sold nor may offers to buy be accepted offer to sell or the solicitation of an offer to buy these securities, nor shall there be any sale of these securities in any state or jurisdiction in which such offer, solicitation or sale would be unlawful prior to registration or qualification under the securities laws of any such state or jurisdiction.

prior to the time the registration statement becomes effective. Before you invest, you should read the preliminary prospectus in that registration statement, including the "Risk Factors" set forth therein, and any other documents the Company has filed with the SEC for more complete information about the Company and this offering. You may obtain these documents for free by visiting EDGAR on the SEC web site at www.sec.gov. Alternatively, the Company, any underwriter or any dealer participating in the offering will arrange to send you the preliminary prospectus if you request it by contacting: Goldman Sachs & Co. LLC, Attention: Prospectus Department, 200 West Street, New York, NY 10282, by telephone at 866-471-2526 or by email at prospectus-ny@ny.email.gs.com; Morgan Stanley & Co. LLC, Attention: Prospectus Department, 180 Varick Street, Second Floor, New York, NY 10014, by telephone at 866-718-1649 or by email at prospectus@morganstanley.com; or Citigroup Global Markets Inc., c/o Broadridge Financial Solutions, 1155 Long Island Avenue, Edgewood, New York 11717 or by telephone at 1-800-831-9146. This presentation shall not constitute an

Certain information contained in this presentation relates to or is based on studies, publications, surveys, and other data obtained from third-party sources and the Company's own internal estimates and research. Although the Company is responsible for all of the disclosure contained in this presentation and believes these third-party sources to be reliable as of the date of this presentation, it has not independently verified, and makes no representation as to the adequacy, fairness, accuracy, or completeness of, any information. All trademarks, service marks, and trade names appearing in this presentation are the property of their respective holders.

Forward-Looking Statements

This presentation contains forward-looking statements, which statements involve substantial risks and uncertainties about us and our industry. These statements may relate to, but are not limited to, our future financial condition or results of operations, business strategy and plans for growth, our future capital expenditures, technological capabilities, strategic relationships, and

objectives of management for future operations. In some cases, you can identify forward-looking statements because they contain words such as "may," "will," "should," "expect," "plan," "anticipate," "could," "intend," "target," "project," "contemplate," "believe," "estimate," "predict," "potential" or "continue" or the negative of these words or other similar terms or expressions that concern our expectations, strategy, plans or intentions. You should not rely upon forward-looking statements as predictions of future events. We have based these forward-looking statements largely on our current expectations and projections about future events and trends that we believe may affect our business, financial condition, results of operations, business strategy, prospects and financial needs. The outcomes of the events described in these forward-looking statements are subject to risks, uncertainties, and other factors that may cause our actual results, performance or achievements to be materially different from any future results, performance or achievements expressed or implied by the forward-looking statements.

Moreover, we operate in a very competitive and rapidly changing environment. New risks and uncertainties emerge from time to time, and it is not possible for us to predict all risks and uncertainties that could have an impact on the forward-looking statements contained in this presentation. We cannot assure you that the results, events, and circumstances reflected in the forward-looking statements will be achieved or occur, and actual results, events, or circumstances could differ materially from those described in, or implied by, the forward-looking statements. Neither we nor any other person assumes responsibility for the

Non-GAAP Financial Measures

This presentation uses financial measures that are not presented in accordance with generally accepted accounting principles in the United States ("GAAP") to supplement financial information presented in accordance with GAAP. We believe that our non-GAAP financial measures provide useful information to investors, analysts and others about our business and financial performance, enhance their overall understanding of our performance, and can assist in providing a more consistent and comparable overview of our financial performance across periods. There are limitations to the use of the non-GAAP financial measures presented in this presentation. For example, the non-GAAP financial measures may not be comparable to similarly titled measures of other companies. Other companies may calculate non-GAAP financial measures differently than the Company, limiting the usefulness of those measures for comparative purposes. Thus, our non-GAAP financial measures should be considered in addition to, not as substitutes for, or in isolation from, measures prepared in accordance with GAAP. See the Appendix for reconciliations of such non-GAAP financial measures to their most directly comparable GAAP measures.

The information in this presentation is provided in summary form only and does not purport to be complete. This presentation does not contain all the information that is or may be material to you in considering a potential investment in the Company and should not be considered as a recommendation in respect of the holding, purchasing or selling of any securities of the Company. This

accuracy and completeness of any of these forward-looking statements. Moreover, the forward-looking statements made in this presentation relate only to events as of the date on which the statements are made. We undertake no obligation to update any forward-looking statements made in this presentation to reflect events or circumstances after the date of this presentation or to reflect new information or the occurrence of unanticipated events, except as required by law. We may not actually achieve the plans, intentions or expectations disclosed in our forward-looking statements and you should not place undue reliance on our forward-looking statements. Our forward-looking statements do not reflect the potential impact of any future acquisitions, mergers, dispositions, joint ventures, or investments we may make. In addition, statements that "we believe" and similar statements reflect our beliefs and opinions on the relevant subject.

presentation should not be construed as, nor be relied on in connection with, any offer or invitation to purchase or subscribe for, underwrite or otherwise acquire, hold or dispose of any securities of the Company, and shall not be regarded as a recommendation in relation to any such transaction whatsoever.

Additional Disclaimers

European Economic Area Disclaimer

Advertisement – This presentation does not constitute a prospectus for purposes of Regulation (EU) 2017/1129, as amended (the "Prospectus Regulation"), and legal provisions in connection therewith. In connection with the proposed offering to the public to eligible retail investors in Denmark, Finland, France, Germany, Ireland, the Netherlands, Norway, Spain, and Sweden (the "Covered EEA Countries"), which is part of the offering described in this presentation (the "European Public Offering"), this presentation is an advertisement for purposes of the Prospectus Regulation and legal provisions in connection therewith.

The European Public Offering will only be made on the basis of the European Prospectus (as defined below). Eligible retail investors, which are individual persons or legal entities which do not qualify as qualified investors within the meaning of Article 2 point (e) of the Prospectus Regulation, in the Covered EEA Countries should not purchase or subscribe for any shares of Class A common stock referred to in this presentation except on the basis of information in, and subject to the terms of, the prospectus to be issued by Anthropic in connection with the European Public Offering (the "European Prospectus"), which may differ from the terms set out in this presentation. The European Prospectus is, following approval by the German Federal Financial Supervisory Authority (*Bundesanstalt für Finanzdienstleistungsaufsicht* – "Bafin"), available on the website of Anthropic at ipo.anthropic.com and upon request of any investor and free of charge, a copy of the European Prospectus in electronic version shall be delivered to the relevant investor. The approval of the European Prospectus by the Bafin should not be understood as an endorsement of the securities. Potential investors in the Covered EEA Countries where the European Public Offering will be made should read the European Prospectus (including any documents incorporated by reference therein and any supplements thereto, if any) before making an investment decision in order to fully understand the potential risks and rewards associated with the decision to invest in the securities. Investment in shares entails numerous risks, including a total loss of initial investment.

In member states of the European Economic Area other than the Covered EEA Countries, this presentation is only

La Oferta Pública Europea se realizará únicamente sobre la base del Folleto Europeo (según se define a continuación). Los inversores minoristas elegibles, que son personas físicas o jurídicas que no tengan la consideración de inversores cualificados en el sentido del artículo 2, letra e), del Reglamento de Folletos, en los Países del EEE Cubiertos no deberán comprar ni suscribir ninguna acción ordinaria de Clase A a la que se hace referencia en esta presentación salvo sobre la base de la información contenida en, y de conformidad con los términos de, el folleto que publicará Anthropic en relación con la Oferta Pública Europea (el "Folleto Europeo"), que podrá diferir de los términos establecidos en esta presentación. El Folleto Europeo se encuentra, tras su aprobación por la Autoridad Federal de Supervisión Financiera de Alemania (*Bundesanstalt für Finanzdienstleistungsaufsicht* – "Bafin"), disponible en el sitio web de Anthropic en ipo.anthropic.com y, a solicitud de cualquier inversor y de forma gratuita, se entregará al inversor correspondiente una copia del Folleto Europeo en versión electrónica. La aprobación del Folleto Europeo por Bafin no debe interpretarse como un respaldo de los valores. Los inversores potenciales en los Países del EEE Cubiertos donde se realizará la Oferta Pública Europea deberán leer el Folleto Europeo (incluido cualquier documento incorporado por referencia al mismo y cualquier suplemento al mismo, en su caso) antes de tomar una decisión de inversión con el fin de comprender plenamente los riesgos y beneficios potenciales asociados a la decisión de invertir en los valores. La inversión en acciones conlleva numerosos riesgos, incluido el de la pérdida total de la inversión inicial.

En los Estados miembros del Espacio Económico Europeo

addressed to and directed at persons who are "qualified investors" within the meaning of Article 2 point (e) of the Prospectus Regulation. In addition, Anthropic may also offer any shares of Class A common stock in any member state of the European Economic Area in circumstances falling within the application of Article 1(4) of the Prospectus Regulation not requiring the publication of a prospectus.

Publicidad – Esta presentación no constituye un folleto a los efectos del Reglamento (UE) 2017/1129, en su versión modificada (el "Reglamento de Folletos"), y las disposiciones legales relacionadas con el mismo. En relación con la oferta pública propuesta a inversores minoristas elegibles en Dinamarca, Finlandia, Francia, Alemania, Irlanda, Países Bajos, Noruega, España y Suecia (los "Países del EEE Cubiertos"), que forma parte de la oferta descrita en esta presentación (la "Oferta Pública Europea"), esta presentación constituye un anuncio publicitario a los efectos del Reglamento de Folletos y las disposiciones legales relacionadas con el mismo.

distintos de los Países del EEE Cubiertos, esta presentación se dirige exclusivamente a personas que sean "inversores cualificados" en el sentido del artículo 2, letra e), del Reglamento de Folletos. Además, Anthropic también podrá ofrecer acciones ordinarias de Clase A en cualquier Estado miembro del Espacio Económico Europeo en circunstancias incluidas en el ámbito de aplicación del artículo 1, apartado 4, del Reglamento de Folletos que no requieran la publicación de un folleto.

Swiss Disclaimer

This presentation is directed at individual persons or legal entities which do not qualify as professional clients within the meaning of the FinSA ("eligible retail investors") and which are resident or are located in Switzerland. This information constitutes an advertisement within the meaning of article 68 FinSA. It is not a prospectus within the meaning of the FinSA nor under any other applicable laws with respect to the public offering of Class A common stock of Anthropic to eligible retail investors in Switzerland (the "Swiss Public Offering"). Eligible retail investors in Switzerland should not purchase or subscribe for any shares of Class A common stock referred to in this information except on the basis of information in the European Prospectus to be issued by Anthropic for purposes of the European and Swiss Public Offerings. The European Prospectus is, following approval by Bafin and deposit (*Hinterlegung*) with a Swiss prospectus office for automatic acceptance in accordance with article 54(2) FinSA, available free of charge on the website of the Company at ipo.anthropic.com. The approval of the European Prospectus by Bafin and/or the deposit of the European Prospectus with a Swiss prospectus office should not be understood as an endorsement of the securities.

Potential eligible retail investors in Switzerland should read the European Prospectus (including any documents incorporated by reference therein and any supplements thereto, if any) before making an investment decision in order to fully understand the potential risks and rewards associated with the decision to invest in the securities. Investment in shares entails numerous risks, including a

total loss of the initial investment.

ANTHROP\C

Appendix: Financial statements

Annual consolidated statement of operations

$ in millions | FYE December	FY 2023	FY 2024	FY 2025
Revenue	**$41**	**$386**	**$4,585**
Compute and infrastructure	(433)	(2,522)	(7,326)
Personnel	(95)	(454)	(1,839)
Bartz settlement	-	-	(1,500)
Sales, marketing, and partnerships	(2)	(43)	(568)
Other operating and administrative	(62)	(347)	(1,412)
Operating loss	**$(551)**	**$(2,980)**	**$(8,060)**
Interest income	69	106	431
Change in fair value of convertible notes, derivative liabilities, and SAFEs	(1,045)	(5,434)	(34,308)
Other income (expense), net	(2)	-	(8)
Loss before provision for income taxes	**$(1,529)**	**$(8,308)**	**$(41,945)**
Provision for income taxes	(2)	(1)	(20)
Net loss	**$(1,531)**	**$(8,309)**	**$(41,965)**

Quarterly consolidated statement of operations

$ in millions | FYE December	Q1 2025	Q2 2025	Q3 2025	Q4 2025	Q1 2026	Q2 2026	Q3 2026

Revenue	**$376**	**$787**	**$1,382**	**$2,040**	**$4,731**	**$11,555**	**$16,956**
Compute and infrastructure	(944)	(1,390)	(2,047)	(2,945)	(3,387)	(5,558)	(9,863)
Personnel[1]	(218)	(320)	(489)	(812)	(1,295)	(1,857)	(1,711)
Bartz settlement	-	-	(1,500)	-	-	-	-
Sales, marketing, and partnerships	(24)	(96)	(202)	(246)	(1,033)	(2,135)	(2,361)
Other operating and administrative	(125)	(181)	(233)	(873)	(683)	(1,031)	(3,588)
Operating income (loss)	**$(935)**	**$(1,200)**	**$(3,089)**	**$(2,836)**	**$(1,667)**	**$974**	**$(567)**
Interest income	51	64	124	192	349	762	1,248
Change in fair value of convertible notes, derivative liabilities, and SAFEs	(5,701)	(5,241)	(9,165)	(14,201)	(12,791)	(59,597)	(21,936)
Other income (expense), net	-	-	(3)	(5)	(309)	(8,448)	(1,640)
Loss before provision for income taxes	**$(6,585)**	**$(6,377)**	**$(12,133)**	**$(16,850)**	**$(14,418)**	**$(66,309)**	**$(22,895)**
Provision for income taxes	(7)	-	(5)	(8)	(21)	-	(396)
Net loss[2]	**$(6,592)**	**$(6,377)**	**$(12,138)**	**$(16,858)**	**$(14,439)**	**$(66,309)**	**$(23,291)**

[1]Includes stock-based compensation expense of $105 million, $181 million, $305 million, $574 million, $924 million, $1.3 billion, and $1.1 billion for the three months ended March 31, 2025, June 30, 2025, September 30, 2025, December 31, 2025, March 31, 2026, June 30, 2026, and September 30, 2026, respectively. [2]Primarily driven by an increase in the fair value of our outstanding convertible promissory notes resulting from an increase in our enterprise valuation.

ANTHROP\C

Appendix: GAAP to Non-GAAP reconciliations

Annual non-GAAP reconciliations

$ in millions, except for percentages	FY 2023	FY 2024	FY 2025
Operating loss	$(551)	$(2,980)	$(8,060)
Stock-based compensation expense	29	192	1,165
Non-cash charitable contributions[1]	-	-	541
Bartz settlement[2]	-	-	1,500
Adjusted Operating Income / (Loss)	**$(522)**	**$(2,788)**	**$(4,854)**
Adjusted Operating Margin	(1,273)%	(722)%	(106)%

Quarterly non-GAAP reconciliations

$ in millions, except for percentages	Q1 2025	Q2 2025	Q3 2025	Q4 2025	Q1 2026	Q2 2026	Q3 2026

Operating income / (loss)	$(935)	$(1,200)	$(3,089)	$(2,836)	$(1,667)	$974	$(567)
Stock-based compensation expense	105	181	305	574	924	1,261	1,070
Non-cash charitable contributions[1]	-	-	-	541	126	-	1,960
Bartz settlement[2]	-	-	1,500	-	-	-	-
Adjusted Operating Income / (Loss)	**$(830)**	**$(1,019)**	**$(1,284)**	**$(1,721)**	**$(617)**	**$2,235**	**$2,463**
Adjusted Operating Margin	(221)%	(129)%	(93)%	(84)%	(13)%	19%	15%

[1]Relates to non-cash charges from our donation matching program, the value of which is linked to employee equity awards and recorded in other operating and administrative expenses. [2]Excluded from our calculation of Adjusted Operating Income (Loss) as it is not indicative of ongoing operations.

ANTHROP\C

Appendix: Definitions and methodologies

Definitions

Adjusted Operating Income (Loss): refers to operating income (loss) excluding (i) stock-based compensation expense, (ii) non-cash charitable contributions, and (iii) Bartz settlement.

Adjusted Operating Margin: refers to Adjusted Operating Income (Loss) divided by revenue.

Compute Efficiency: calculated as revenue divided by compute and infrastructure costs.

Customer: refers to a prosumer or an organization that uses our models, platform, or applications.

"Customers with $100,000 or more of revenue on a TTM basis," "customers with $1 million or more of revenue on a TTM basis," and "customers with $10 million or more of revenue on a TTM basis" each refer to customers contributing revenue equal to or greater than the applicable dollar amount during the trailing twelve-month ("TTM") period ending on the measurement date. For purposes of this metric, we define a customer as a distinct organization,

Net Revenue Retention: reflects customer expansion, contraction, and churn as of a measurement date. We calculate Net Revenue Retention by dividing the Qualified Customers' Current Period revenue by their Base Period revenue. A Qualified Customer is a customer with at least $1 of revenue during the Base Period, excluding subscribers to our Claude Pro, Max, and Ultra plans, whose fixed-price subscriptions do not reflect the expansion and contraction behavior this metric is intended to capture. The Base Period is the three months ending a year prior to the measurement date. The Current Period is the three months ending as of the measurement date.

Run-Rate Revenue: an operating metric that refers to the estimated annualized value of our revenue-generating activity as of a measurement date, calculated as the sum of (i) our consumption-based fees during the trailing 28-day period ended on the measurement date, multiplied by 13, and (ii) our aggregate monthly subscription fees as of the measurement date, multiplied by 12. For purposes of calculating Run-Rate Revenue, consumption-based fees

identified by aggregating the accounts that we determine belong to the same organization based on internal records and available data. We do not separately define or internally track a category of "enterprise" customers, and our customers are not segmented or measured internally by reference to whether a given customer is an "enterprise." Because spending $100,000 or more on our products is generally characteristic of organizational rather than individual usage, we use the number of and revenue from customers with $100,000 or more of revenue on a TTM basis as our measure of enterprise adoption and engagement with Claude, including in references to "enterprise customers" and "enterprise adoption."

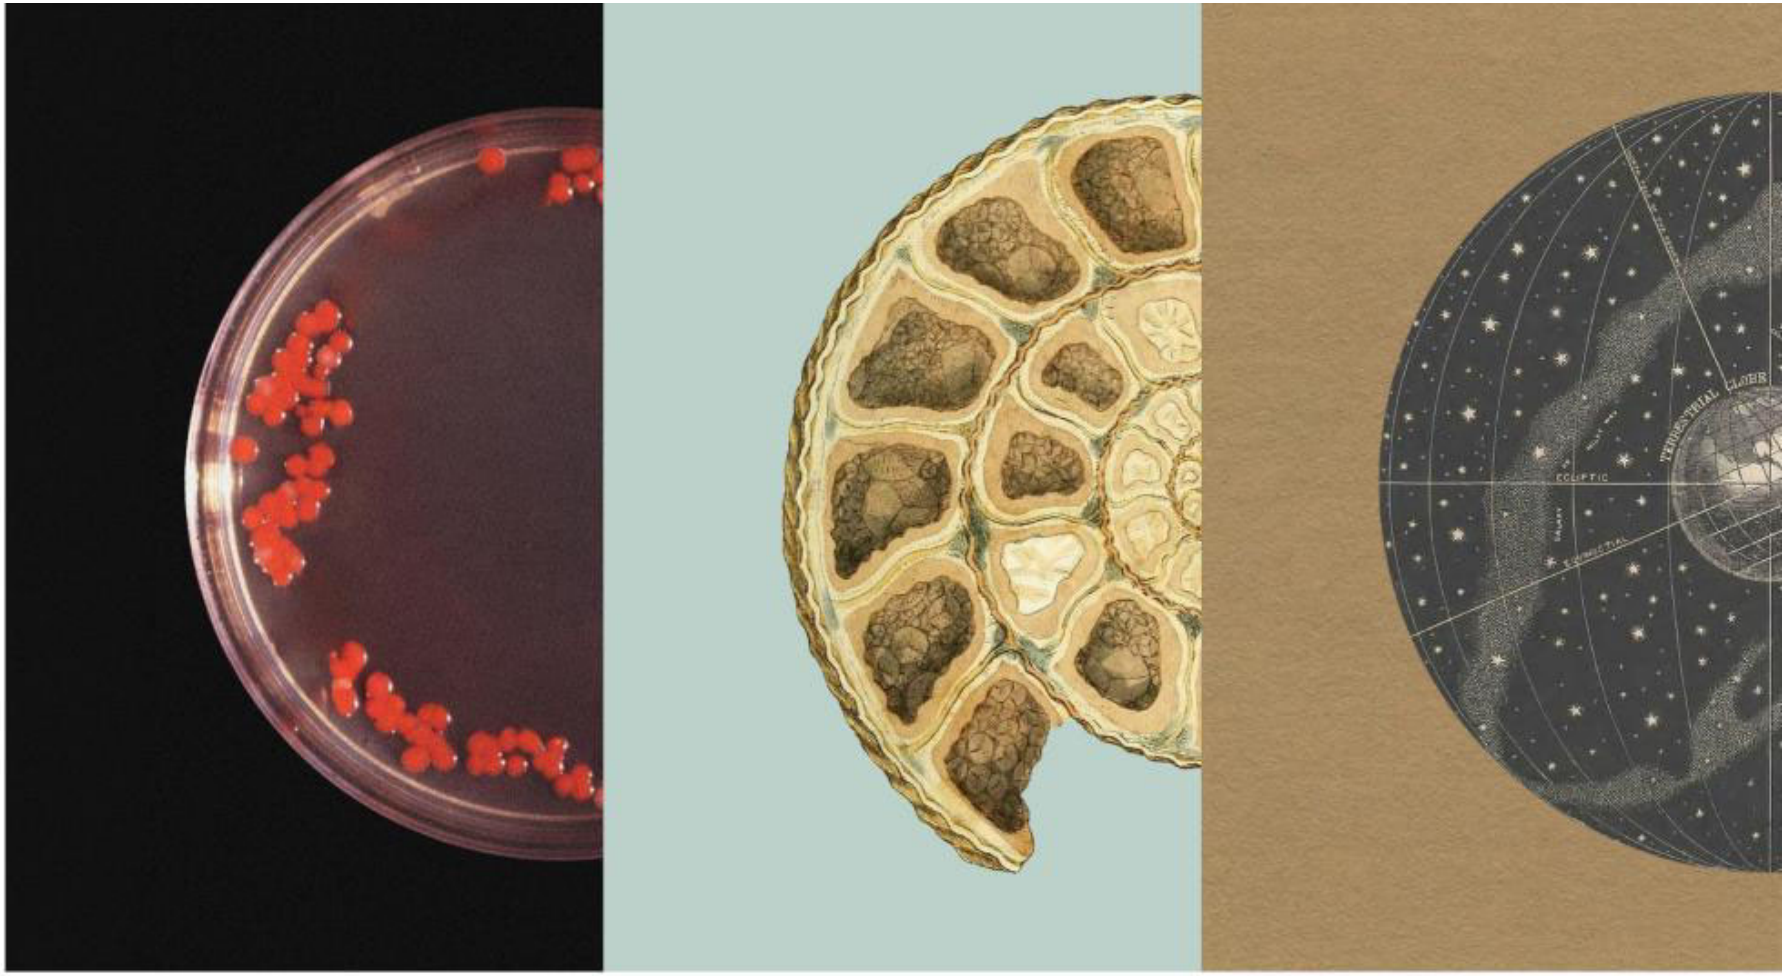
and subscription-based fees are derived from customer usage and active subscriptions.

Methodologies

Knowledge work calculation methodology. Our estimate regarding global knowledge workers' wages, which we refer to as the annual economic value represented by global knowledge workers, is our own, derived from third-party data. Based on wage data for U.S. workers by occupation classification from the U.S. Bureau of Labor Statistics in 2025, we calculate the total salary for the subset of workers that we believe to represent knowledge workers. Based on International Labour Organization data for global labor income as a percentage of global GDP, and based on 2025 global GDP, we estimate a 2025 global wages figure. We calculate the ratio of total U.S. wages in 2025 to the 2025 global wages figure, and apply that ratio to the total U.S. knowledge workers' wages figure to arrive at an estimate of global knowledge workers' wages.

Exhibit C

Founders Letter



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. depend on scale was surprising, but predictions made on the basis of the scaling laws have proved true again and again.[1]

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

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.

---

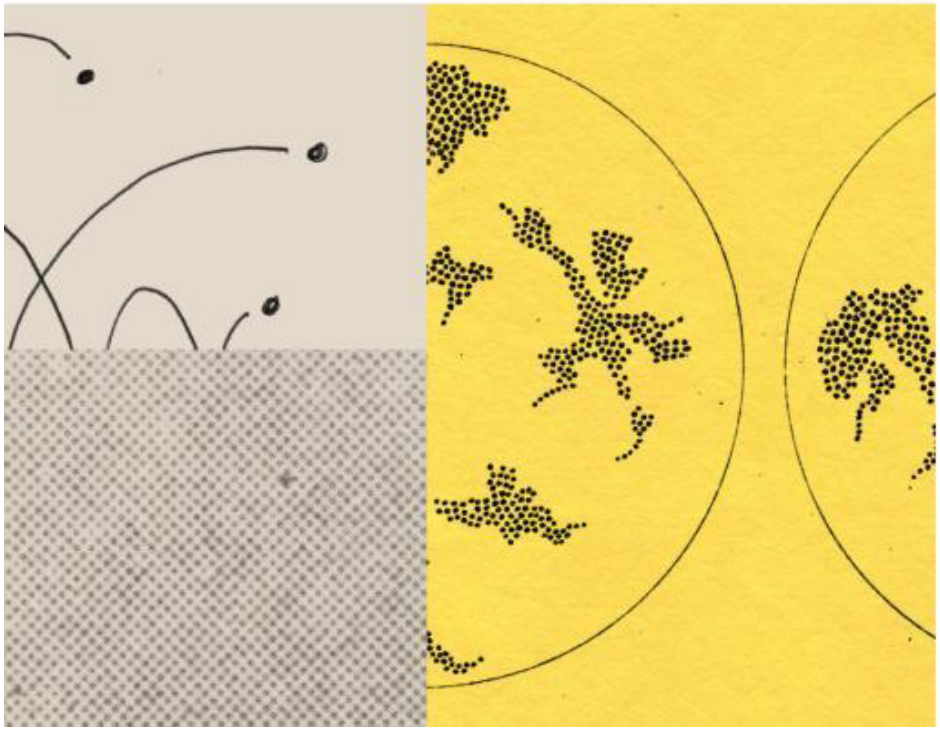
[1] The idea is closely related to Richard Sutton's concept of the bitter lesson, which describes how sheer scale has consistently won out over more specific, complex methods of training AI models. Note that scale is not the only source of improvements in model capabilities—algorithmic efficiencies ("compute 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

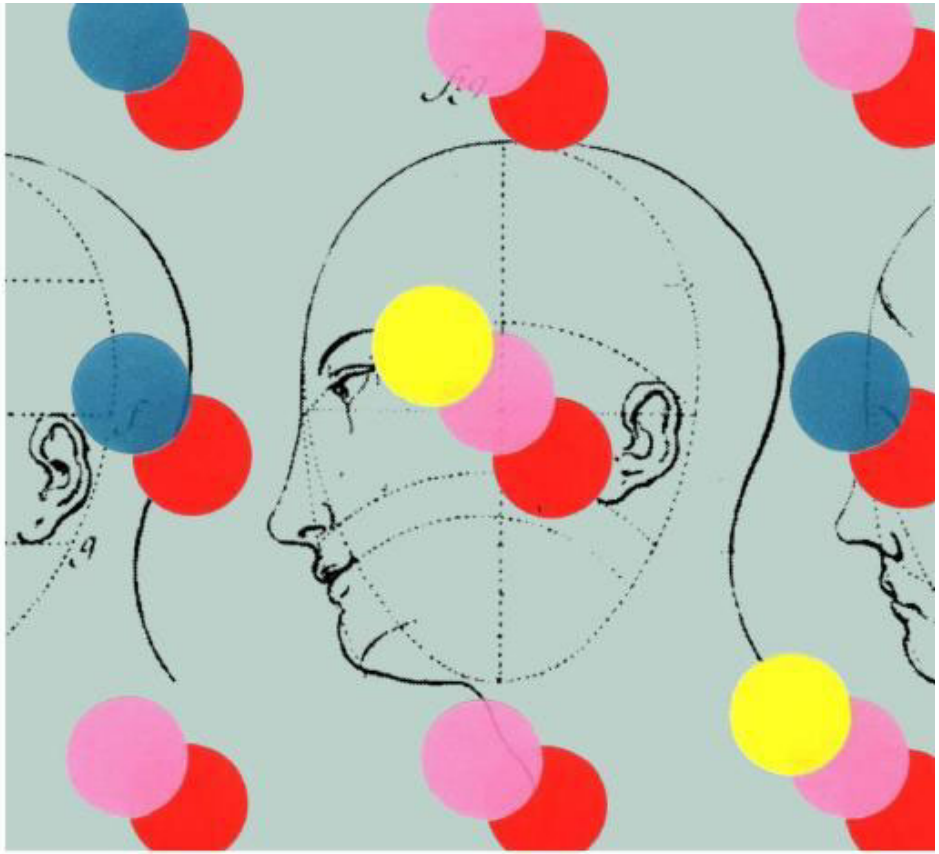
pharmaceutical companies to improve both basic discovery and clinical trials, from helping pharmaceutical company scientists tackle drug



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.



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.

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,

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.

There are four broad categories of risk: misuse, loss of control, authoritarian abuse, and economic disruption.

- AI can be misused by those who wish to do harm, for example to conduct cyberattacks or to design biological weapons. We have strict guardrails and trusted access programs that are designed to block misuse of our technology while still enabling positive applications. We regularly shut down attempts to misuse our models, report them to law enforcement when appropriate, and endeavor to transparently disclose them to the public.

- AI may also behave autonomously in ways we (and our users) do not intend—as illustrated by recent high-profile cybersecurity incidents. Losing control of sufficiently powerful AI systems could be a global security threat of the first order. Anthropic has pioneered a number of methods for improving the control of our AI systems. These include mechanistic interpretability, which allows us to look inside AI models and try to understand how they behave, and constitutional training, which allows us to train AI models to act in line with a specific written document—Claude's constitution. We also publish model system cards that run to hundreds of pages, designed to give the public visibility into our progress on model safety. It has long been our practice to disclose incidents

- macroeconomic models to predict the possible future impact of AI. We also suspect more radical actions will be needed and have already committed $200 million to partner with governments to run studies on different ways of responding to changes in the economy. We do this to find evidence about which types of interventions are effective, so that if society does encounter economic disruption, we will have helped build and de-risk the right tools to use in the response. We expect to scale up this funding significantly in the future.

We consider it our personal responsibility to help the world understand these risks, to prevent or mitigate them when they arise from our AI models, and to work with governments and the rest of the industry to address them at a societal level.

Anthropic's unique approach

Navigating a technology with the properties we have just described requires a very different approach from that of a typical company. Economists use the term *externalities* for effects on society that fall outside the transactions between a company and its customers, that the market therefore does not fully reward or penalize: the cost of pollution produced by a factory, for example, or the value of a research breakthrough by

and publish research findings that help mitigate against this risk across the industry.

- Authoritarian governments may use AI to create high-tech surveillance states, deepen repression, and make themselves much more difficult to dislodge. We are committed to working with free societies to provide a counterweight to authoritarian abuse and to promote freedom and democracy.

- Finally, the incredible economic value of AI may come with economic disruption, similar to what has happened with previous technologies but greater in magnitude. We are committed to studying this disruption—we maintain an economic index that provides aggregated information about Claude's usage across the economy, and our researchers build

one company that other companies freely build upon. AI has unusually large externalities, both positive and negative. Cultivating the positive externalities, and defending against the negative ones, is to some degree in tension with the natural incentives of a for-profit company, and with the standard governance structure of a publicly traded company.

As founders and business leaders, we believe deeply in the power of entrepreneurial capitalism to deliver massive benefits to humanity. But we also believe that, for a technology as consequential as AI, entrepreneurial capitalism must be balanced and augmented by mechanisms that protect the common good and share the upside as widely as possible. Everything about the way we have built Anthropic, from our culture to our corporate structure, has been designed to ensure we can act in the public interest when the situation demands it.

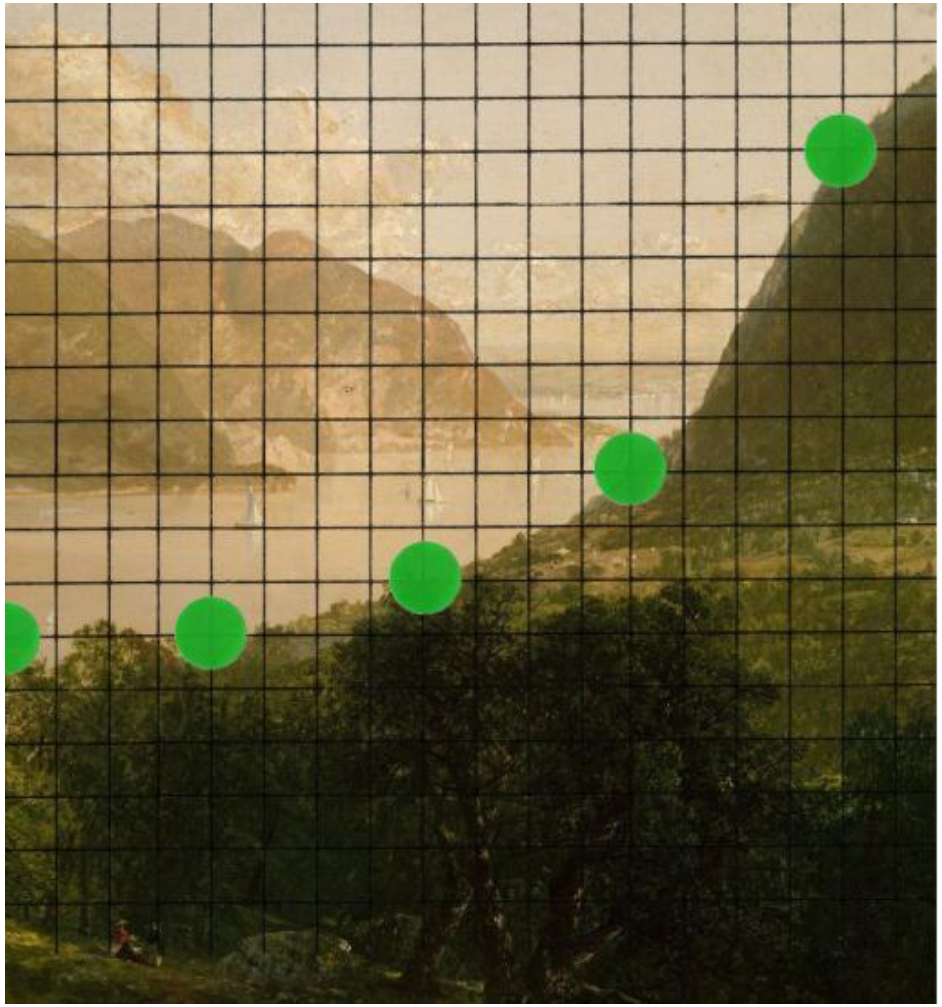
It starts with Anthropic's culture, which is built around public benefit, honesty and trust, low-ego collaboration, and relentless focus on doing the simple thing that works. Everyone who interviews at Anthropic goes through an interview process designed to assess these attributes, and we are willing to turn down candidates who don't demonstrate them, even if they are extremely strong technically.

Anthropic remains small for its reach. Each Anthropic employee supports over $6.6 million in trailing 12-month revenue as of September 30, 2026, which is meaningfully higher than the average for the largest public U.S. companies. We believe a close-knit team of people who share a mission can accomplish far more than its size would suggest, and we intend to keep that character as we grow.

Upon joining, all employees are offered a donation match: for each share of equity that they pledge to give away to philanthropic causes, Anthropic creates an additional share to match it, which is also pledged in the same manner. At the time of writing, just over half of all our employees have pledged at least some of their equity.[2] All seven of us have also pledged to donate 80% of our personal equity in Anthropic to charitable causes.

As co-founders and leaders, we aim for radical internal honesty, and we have a culture of lively and open debate between people of all levels of seniority. Every two weeks Dario, our CEO, typically gets up in front of the company and openly describes the decisions that are on his mind and the state of the industry; other co-founders and leaders frequently do the same in their respective areas. Employees debate the tradeoffs necessitated by the complexity of our mission in discussion forums that are open to the whole company. This open culture is one reason our talent density is high and our talent retention remains among the best in the industry. All seven of us still work at the company, which makes Anthropic the only major AI company to have retained its entire founding team. Over three quarters of our first 100 employees still work here; even those who do leave often go on to recommend Anthropic to others, or to return. Our statistics show that of the employees who left whom we would have preferred to retain, 19% subsequently returned to work at Anthropic.

But culture is not enough: formal governance commitments matter as well. We have designed our governance structure to balance commercial success and the public good. Anthropic is a public benefit corporation, which means our board of directors is tasked with weighing the financial interests of our stockholders against our public benefit mission and



the interests of those materially affected by our conduct. In addition, our governance structure shares power among three groups: our co-founders, our employees, and the Long-Term Benefit Trust ("LTBT"). The LTBT, an independent body with no financial stake in Anthropic, is charged with holding the company to its long-term public benefit mission. Its members bring expertise in fields where the positive and negative externalities of AI are most likely to arise, including national security, economics, global health, and science.

---

[2] For employees who joined before May 2024, we offered a donation match of up to 50% of their shares. Over 70% made a pledge, with over 25% pledging the full 50%. For post-May 2024 employees, we offered a donation match of up to 25% of their shares; nearly 50% have participated, with nearly 30% pledging the full 25%.

The LTBT elects four of our seven directors, with one of those seats subject to safeguards on nomination and removal (including an employee vote if the company and the LTBT disagree). We, the co-founders, have the power to elect the other three directors. Our increased voting power is not permanent and voting power passes to our common stockholders once fewer than three of us remain in the voting entity. This prospectus describes the structure in greater detail.

In addition to individual philanthropy, we are earmarking a measurable fraction of our resources for advancing AI's non-market benefits and reducing its risks, beyond what is incentivized by the logic of the market. Prior to our initial public offering, we are setting aside 5% of the company's equity for a fund devoted to the public good. We will set aside an additional 2% for every sustained 50% increase in share price from the prior contribution (starting with our initial public offering price), at our board's discretion. This 5% (or larger) commitment could be devoted to a wide range of initiatives, from science and medicine, to biosecurity, to global health, to economic research programs, to distributing the economic value from Anthropic's equity widely within the United States or throughout the world.[3]

For investors concerned about dilution, we note that

Our theory of change has always been what we call a *race to the top*: by showing that it is possible to succeed commercially while acting in the public interest, we hope to persuade other companies to do the same. We will make mistakes, and we expect to learn from others as much as they learn from us. If our competitors adopt our practices and improve on them, we consider that a success.

On investor expectations

Before starting the IPO process, we spoke to many business leaders who had gone through a similar process in the past. One piece of advice, which we heard in slightly different forms from multiple people, stuck with us: *never let your company be worth a dollar more than it is actually worth*. What they meant is that it is very dangerous when people buy your company's stock based on expectations you cannot fulfill or on incorrect predictions about how you will behave. This is in fact very similar to one of Anthropic's guiding principles: *underpromise and overdeliver*. It is also similar to something we do when recruiting: we give candidates an *anti-sell*, a pitch deliberately designed to be unappealing to those who are a poor fit for Anthropic while being appealing to the type of person we want to hire.

In this spirit, although we believe that AI as a field

the size and structure of this plan are not outside the range of what has been offered as executive compensation packages at some public companies in the past. We expect Anthropic's leadership team and employees to be as motivated by this structure as leaders of other companies are motivated by executive compensation.

With all of that said, Anthropic's employees and leaders understand that our public benefit goals are fueled by commercial success, and by the success of our customers. Achieving both business growth and public impact has been the balance that we have navigated since the founding of the company, and we have a strong track record of showing that it is possible to do both. Our ability to have a positive impact on the world scales in relation to our reach and stature as a successful company.

(and Anthropic in particular) represents one of the greatest economic opportunities in history, we also believe that it is important to state clearly the ways in which our decision-making may run counter to investors' expectations. These fall into two categories: business bets that derive from AI's wild exponential trajectory but seem strange or unorthodox in isolation, and decisions that depart from profit maximization in order to appropriately balance the public good.

---

[3] This pledge is in addition to the aforementioned pledges of all of Anthropic's co-founders to give away 80% of our personal equity and the pledges of employees who took up our offer of donation matching, which amount to 9% and 8% of fully diluted shares as of June 30, 2026, respectively.

- In the first category, Anthropic will often make forward investments (especially in compute) that may seem irresponsibly large at first glance. This derives from our belief, based on the scaling laws, that AI's economic value is on an exponential curve and that the expected near-future economic value of our products will therefore be much larger than the current value. In order to fully capture the opportunity in front of us, we need to invest more aggressively than would a business that operates in a more slowly changing technological regime. A useful analogy is to pharmaceutical companies: very large upfront investments are often required for research and development, and those investments only make sense later in light of the returns from successful drugs. At Anthropic, a similar process is occurring much faster and on a much larger scale.

  We will always be willing to forgo *some* upside in exchange for downside protection, for example by looking for ways to offload excess compute if we overestimate demand, seeking to improve our compute and operating efficiency in order to increase our tolerance for error, or raising capital when necessary to insulate ourselves from temporary slowdowns in demand. But some level of risk is inherent to the extreme upside possibilities offered by the AI exponential. In general, we will aim to be *bold but sane*, and to win

- In the second category (departure from profit maximization), we are committed to standing up for our values at critical points, even when this is costly for us commercially.

  For example, when we trained Claude Mythos Preview, the most capable model we had built at the time, we judged that its abilities, particularly in cybersecurity, had advanced to a point where releasing it widely without new safeguards would have carried significant risk. We therefore chose not to make it generally available and offered its capabilities to the public (as Claude Fable 5) only after we had built additional safeguards. The delay meant forgoing substantial revenue and ceding ground to competitors (we also gave away around $100 million in Claude credits to cyber defenders), but we believe it was the right decision, and we expect to face similar decisions again as our models grow more capable.

  We have also committed to keeping Claude free of advertising. Advertising funds much of the modern internet, and it would likely be a large source of revenue for us. But people bring some of their most important work and most personal questions to Claude, and we want Claude's only incentive to be giving them a helpful answer. Once advertising becomes part of a company's revenue

trust with a track record of successful bets.

Related to this, we will be constantly seeking to reinvent our business, releasing new classes of products whose aim is to dwarf the economic value of our existing ones. Advances in AI's cognitive capabilities are constantly opening up new classes of products that were not viable with models from even the preceding few months. For example, AI chatbots became viable around 2022, and AI coding agents around 2025; AI agents for general knowledge work and scientific discovery are gradually becoming viable in 2026. Relative to a typical business, we will generally devote less effort to entrenching our current products and business models, and more effort to defining new ones, because that is where the technology's greatest possibilities lie.

targets, the pressure to optimize for engagement tends to grow. We have chosen not to introduce that pressure into Claude.

A particularly important place to weigh the public interest is the pace of AI progress. Moving quickly brings AI's benefits sooner, but moving too quickly could outrun the world's ability to make AI systems safe. Building AI safely may require Anthropic (and ideally the rest of the industry) to modulate the extremely fast pace at which AI capabilities are compounding, especially when this is driven by AI systems building themselves. While progress will still be very fast in absolute terms—potentially much faster than today—there may be times when we choose to advance more slowly than we could in order to make our systems safer, even if that puts us at a competitive disadvantage.

All of us have a sincere conviction that the AI exponential will generate so much value that, in the long run, it will far outweigh any shorter-term economic sacrifices we make as part of our mission. We also think that some of our public benefit-driven decisions may be good for the business long-term, by generating increased trust and credibility among customers, governments, employees, and the public at large—although that is not the ultimate reason we choose to make them. Nevertheless, we feel the need to be forthright about our decision-making and to set expectations appropriately. We and our investors will have a smoother journey if everyone involved is on the same page about how we operate and is not surprised by our choices.

Parting thoughts

More broadly, however, we encourage you, as a prospective owner of our stock, not to think of your purchase in purely transactional terms. We hope you will see it instead as an opportunity to be part of something greater. The advent of machine intelligence is one of the seminal moments in human history.

As one of the companies at the very center of that moment, we have a rare opportunity to shape it for the better. Our aim is to help humankind chart a safe course through this technological evolution, open new frontiers of discovery, and bring forth an age of miracles. Everyone who owns our stock, whether an institution or an individual, becomes part of that effort. We hope you will share our ambition to see it through.

Dario Amodei

Daniela Amodei

Tom Brown

Jack Clark

Jared Kaplan

Sam McCandlish

Chris Olah



