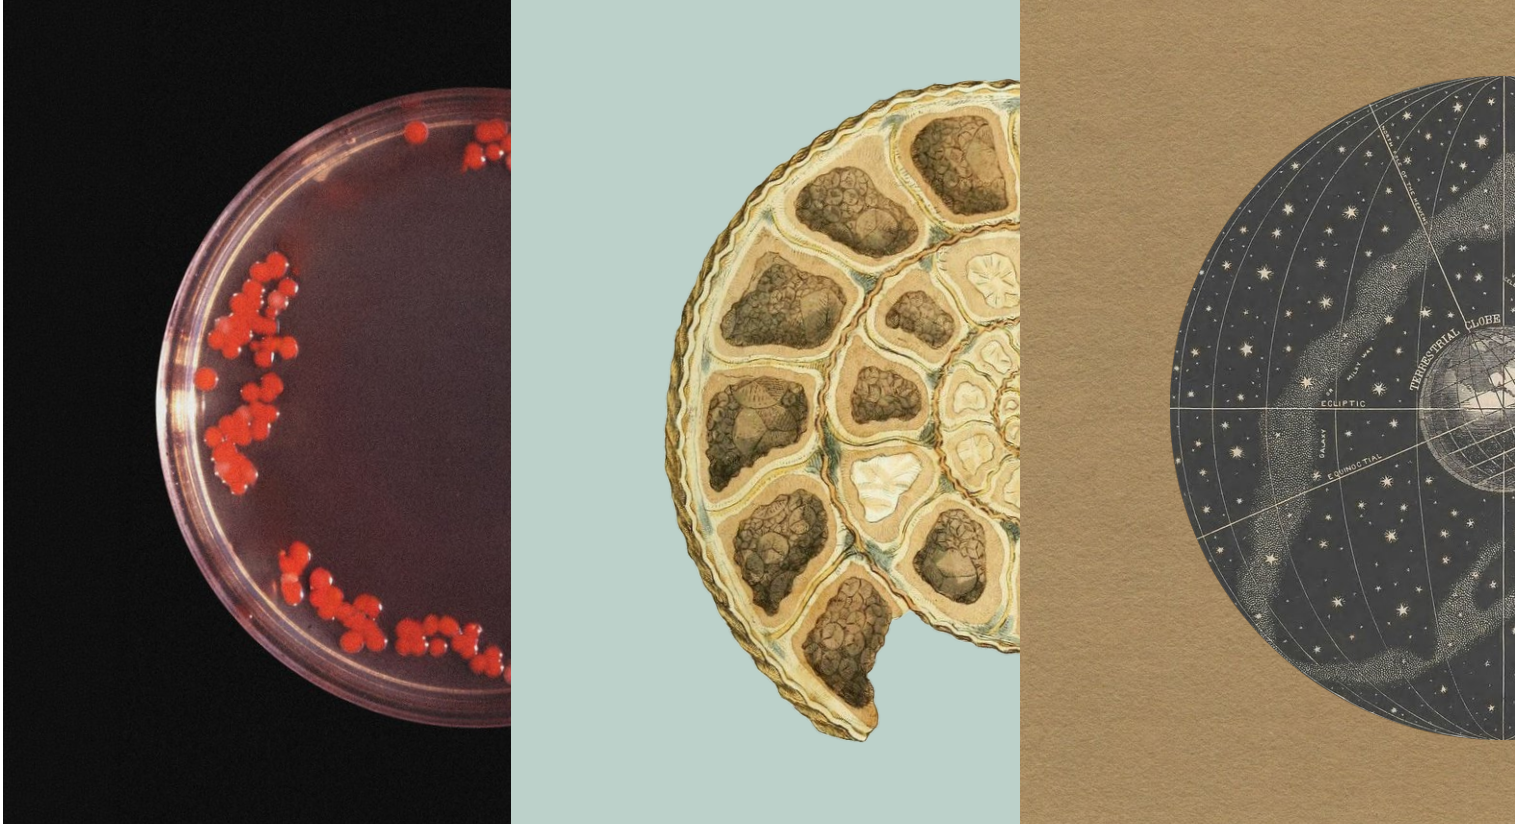


We believe that AI is one of the most consequential technologies in human history. If managed well, it could rapidly accelerate scientific innovation, greatly increase economic growth rates, bring cures for major diseases within reach, and solve problems that have resisted generations of human effort.

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]

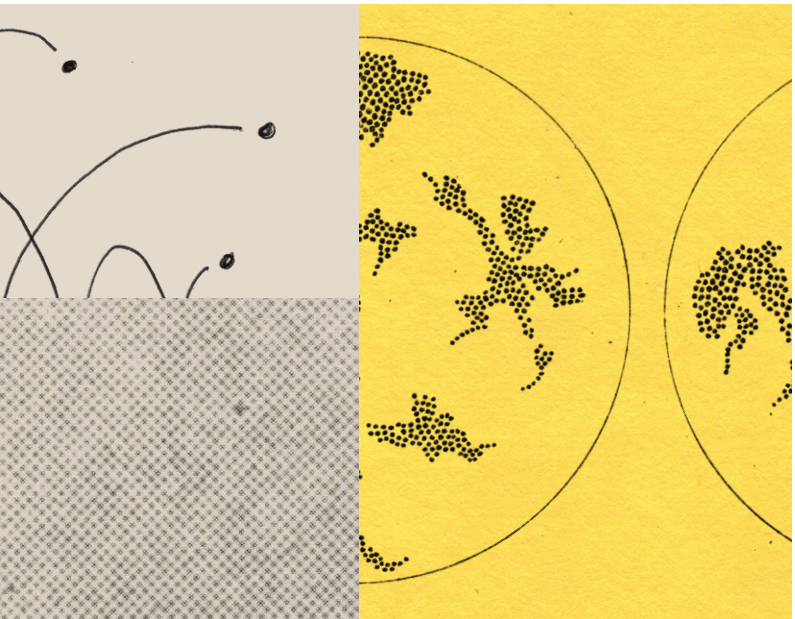
The phenomenon is somewhat analogous to Moore's law for computer chips, under which the computing power of these chips doubled every 12 to 18 months for several decades. Moore's law has been the basis of a succession of industries—from the personal computer to the internet—that have revolutionized the 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.

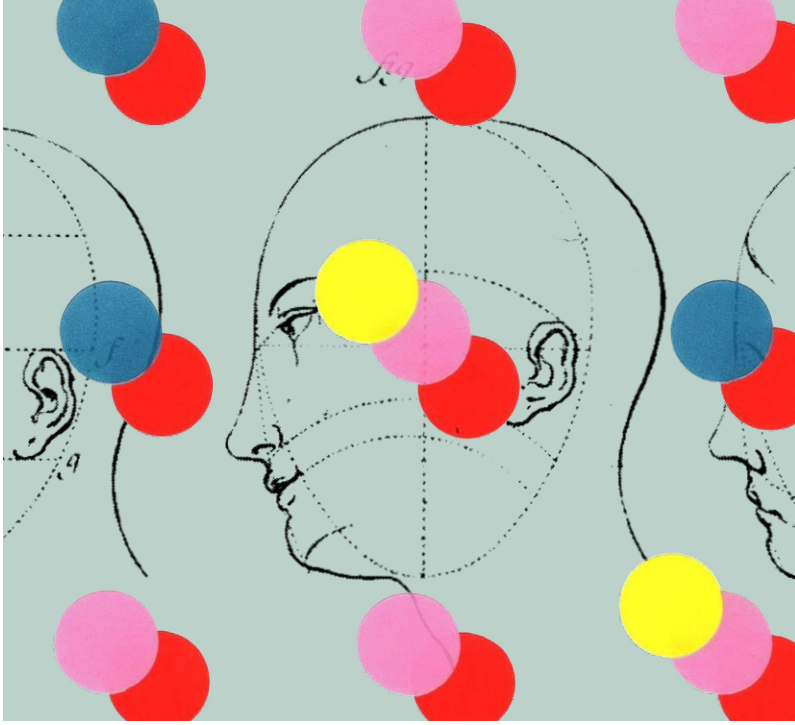
It is worth highlighting a few applications with particularly great potential to improve the quality of human life:

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

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 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 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 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 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.

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 (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 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.

- 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 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

