TIM SCOTT, SOUTH CAROLINA, CHAIRMAN
ELIZABETH WARREN, MASSACHUSETTS, RANKING MEMBER

MIKE CRAPO, IDAHO
MIKE ROUNDS, SOUTH DAKOTA
THOM TILLIS, NORTH CAROLINA
JOHN KENNEDY, LOUISIANA
BILL HAGERTY, TENNESSEE
CYNTHIA LUMMIS, WYOMING
KATIE BOYD BRITT, ALABAMA
PETE RICKETTS, NEBRASKA
JIM BANKS, INDIANA
KEVIN CRAMER, NORTH DAKOTA
BERNIE MORENO, OHIO
DAVID McCORMICK, PENNSYLVANIA

JACK REED, RHODE ISLAND
MARK R. WARNER, VIRGINIA
CHRIS VAN HOLLEN, MARYLAND
CATHERINE CORTEZ MASTO, NEVADA
TINA SMITH, MINNESOTA
RAPHAEL G. WARNOCK, GEORGIA
ANDY KIM, NEW JERSEY
RUBEN GALLEGO, ARIZONA
LISA BLUNT ROCHESTER, DELAWARE
ANGELA D. ALSOBROOKS, MARYLAND

CHRIS LUCAS, STAFF DIRECTOR
JON DONENBERG, DEMOCRATIC STAFF DIRECTOR

United States Senate
COMMITTEE ON BANKING, HOUSING, AND URBAN AFFAIRS
WASHINGTON, DC 20510–6075

October 27, 2026

Dario Amodei
CEO
Anthropic, PBC
548 Market Street, PMB 90375
San Francisco, CA 94104-5401

Dear Mr. Amodei:

I write with concern following your company's disclosure that Anthropic's advanced AI models could "pose catastrophic or existential risks to humanity" at the same time that it seeks to register as a public company and sell shares in November.[1] Anthropic disclosed extensive public safety and financial risks to the Securities and Exchange Commission ("SEC") as part of its registration process for an initial public offering (IPO).[2] An IPO would provide Anthropic with broader access to the largest capital markets in the world, allowing it to raise more funding from a larger pool of investors to fuel its rapid growth.[3] It will also mean that millions of Americans with retirement accounts invested in index funds will quickly become Anthropic shareholders.[4] Public pensions are already heavily exposed to Anthropic and have expressed their concerns about Anthropic's societal risks ahead of its IPO.[5]

This massive influx of public capital is likely to catapult the development and deployment of your technology during a time of intense public debate and concern about the risks of your frontier intelligence models. I believe the risks presented by this technology demand greater scrutiny

---

[1] Bloomberg, "Anthropic Targets Mega-IPO Before Thanksgiving Holiday," Bailey Lipschultz, October 1, 2026, https://www.bloomberg.com/news/articles/2026-10-01/anthropic-said-to-target-mega-ipo-before-thanksgiving-holiday?srnd=homepage-americas; Financial Times, "Anthropic warns of 'existential risks to humanity' in IPO prospectus," George Hammond and James Fontanella-Khan, September 29, 2026, https://www.ft.com/content/c7685a7e-7745-4cbc-8053-4958d0ea449b?syn-25a6b1a6=1.

[2] Reuters, "Anthropic warns AI may pose 'existential risks to humanity' in IPO filing," Echo Wang and Aditya Soni, September 28, 2026, https://www.reuters.com/business/finance/anthropic-warns-ai-may-pose-existential-risks-humanity-ipo-filing-2026-09-29/.

[3] Reuters,"Anthropic's IPO prospectus shows sweeping AI vision, surging costs," Echo Wang, September 28, 2026, https://www.reuters.com/business/finance/anthropics-ipo-prospectus-shows-sweeping-ai-vision-surging-costs-2026-09-28/.

[4] Yahoo Finance, "Anthropic's IPO Could Set Off Forced Buying by Index Funds. SpaceX Shows How that Played Out," Jared Blikre, October 26, 2026, https://finance.yahoo.com/markets/article/anthropics-ipo-could-set-off-forced-buying-by-index-funds-spacex-shows-how-that-played-out-161356189.html.

[5] Reuters, "NYC Comptroller Levine Urges Anthropic to Pause US IPO," October 15, 2026, https://www.reuters.com/business/nyc-comptroller-levine-urges-anthropic-pause-us-ipo-cnbc-reports-2026-10-15/; The New York Times, "U.S. Pensions Are Riding the A.I. Wave. For Better or Worse." Tony Romm, October 27, 2026, https://www.nytimes.com/2026/10/27/business/pensions-ai.html.

before they become intertwined with the public capital markets, which could further escalate both public safety and systemic financial risks.

Growing Public Safety Concerns

Anthropic's Form S-1, which the SEC requires companies to file before an IPO and provides investors with critical financial, operational, and risk information for investors to evaluate the company's offering, devotes over 90 pages to the company's risk factors, including extensive disclosure about the calamitous risks posed by Anthropic's advanced AI.[6] The filing cites concerns from Anthropic employees and others that AI models could pose "catastrophic or existential risks to humanity, particularly if such systems were misused, replicated, weaponized, or allowed to operate with a high degree of autonomy."[7] It warns advanced AI models could exhibit "self-preserving behaviours," including attempts to "resist shutdown," to "conceal or manipulate information," and behavior "resembling blackmail."[8] It warns that Anthropic's "most capable models have demonstrated, and future models are expected to demonstrate, capabilities that could provide meaningful assistance to malicious actors seeking to develop chemical, biological, or other dangerous weapons."[9] In light of all of these risks, the filing warns that "the scale of harm that could result from a catastrophic AI-related incident" is so great that "no single insurer, and potentially the insurance industry in the aggregate, may have sufficient capacity or expertise to underwrite insurance commensurate with such exposure."[10]

Alarmingly, the risk section of your Form S-1 filing also cautions that these risks cannot even be fully assessed because your models may detect human efforts to exercise oversight, which "creates a significant limitation on our ability to assess model safety."[11] You also assert that "any judicial or regulatory expansion of tort, product liability, or consumer protection law, or criminal liability, as applied to AI models, could result in significant liability" and "require" you to change how you "design and offer [y]our models, platform, and applications."[12]

The risks identified in Anthropic's S-1 filing are consistent with public statements that you have made calling for the need to slow down the development of frontier AI models. In a September 2026 essay titled "We Must Pace the Frontier," you wrote that "AI's growing ability to build the next generation of AI," a process called "recursive self-improvement," could, "if left unchecked," "outrun our ability to understand and control these systems."[13] You warned that "we still only understand a tiny fraction of what goes on inside these models" and highlighted the "risk of losing control of AI systems, misuse of AI for cyberattacks and bioterrorism, and serious economic

---

[6] Anthropic. *Form S-1: Registration Statement.* U.S. Securities and Exchange Commission, 2026, https://www.sec.gov/Archives/edgar/data/2130357/000162828026067168/anthropicpbc-sx1.htm#ica694c859b8e458dbf6a73f56dde9d42_2972 ("Anthropic Form S-1").
[7] Anthropic Form S-1, p. 46.
[8] Reuters,"Anthropic warns AI may pose 'existential risks to humanity' in IPO filing," Echo Wang and Aditya Soni, September 29, 2026, "https://www.reuters.com/business/finance/anthropic-warns-ai-may-pose-existential-risks-humanity-ipo-filing-2026-09-29/.
[9] Anthropic Form S-1, p. 44.
[10] *Id.* at 50.
[11] *Id.* at 46.
[12] *Id.* at 49.
[13] Dario Amodei, "We Must Pace the Frontier," https://darioamodei.com/post/we-must-pace-the-frontier.

disruption"[14] You further warned that a "race to the bottom, spurred by commercial incentives, can make these risks more acute."[15]

It is also difficult to trust existing safety tests and evaluations of model test environments in the wake of recent incidents in which models built by frontier AI companies, including Anthropic,[16] have broken through security controls and hacked external systems autonomously. These include attacks on the Australian government's healthcare database and the community AI platform Hugging Face's systems.[17] AI agents have also reportedly meddled with U.S. government websites by exploiting design flaws to access certain data without paying required fees and submitting fictitious government filings.[18] Claude has also accessed publicly available voter data without human authorization.[19] In another instance, an AI agent submitted a false murder tip to a police department.[20] Most concerningly, those incidents of "misaligned model activity" were not immediately discovered.[21] Indeed, you have conceded that Anthropic's tests often do not reveal all of these risks or yield accurate data: "testing and evaluation of AI models becomes more difficult as they increase in capabilities. More intelligent models are more capable of deceiving tests, and thus may appear aligned while having serious problems that go undetected."[22] Anthropic's S-1 discloses that "[m]isaligned AI systems could corrupt institutional decision-making or cause large-scale harm through autonomous actions that are not visible or understandable to their users and operators."[23] These disclosures raise serious concerns about your ability to accurately assess and mitigate the risks associated with advanced models.

---

[14] *Id.*
[15] *Id.*
[16] The Washington Post, "Anthropic AI agents took 'unintended' actions on government sites," Gerrit De Vynck, October 9, 2026, https://www.washingtonpost.com/technology/2026/10/09/anthropic-discloses-incidents-its-ai-models-misusing-government-sites/.
[17] CNN, "'Extreme concern' over first known AI hack of a government system," Hilary Whiteman, Hadas Gold, and Max Saltman, September 24, 2026, https://www.cnn.com/2026/09/23/business/australia-openai-agent-hack-intl-hnk?cid=external-feeds_iluminar_msn.
[18] The New York Times, "OpenAI's Systems Meddled With U.S. Government Sites After Going Rogue," Kate Conger, Ana Swanson, and Cecilia Kang, September 25, 2026, https://www.nytimes.com/2026/09/25/technology/openais-ai-us-government-websites.html; NPR, "OpenAI says its models engaged with US government websites in misbehavior disclosure," Associated Press, September 26, 2026, https://www.npr.org/2026/09/26/nx-s1-5981979/openai-us-government-websites-misbehavior; The Washington Post, "Anthropic AI agents took 'unintended' actions on government sites," Gerrit De Vynck, October 9, 2026, https://www.washingtonpost.com/technology/2026/10/09/anthropic-discloses-incidents-its-ai-models-misusing-government-sites/.
[19] Axios, "Anthropic model accessed voter information," Maria Curi and Alex Isenstadt, October 25, 2026, https://www.axios.com/2026/10/25/anthropic-claude-georgia-voter-data.
[20] The Washington Post, "AI System Submits False Homicide Tip to Philadelphia Police," Gerrit De Vynck, October 9, 2026, https://www.washingtonpost.com/technology/2026/10/09/ai-system-submits-false-homicide-tip-philadelphia-police/.
[21] See, e.g., The Washington Post, "Anthropic AI agents took 'unintended' actions on government sites," Gerrit De Vynck, October 9, 2026, https://www.washingtonpost.com/technology/2026/10/09/anthropic-discloses-incidents-its-ai-models-misusing-government-sites/; The Washington Post, "AI System Submits False Homicide Tip to Philadelphia Police," Gerrit De Vynck, October 9, 2026, https://www.washingtonpost.com/technology/2026/10/09/ai-system-submits-false-homicide-tip-philadelphia-police/; CNN, "'Extreme concern' over first known AI hack of a government system," Hilary Whiteman, Hadas Gold, and Max Saltman, September 24, 2026, https://www.cnn.com/2026/09/23/business/australia-openai-agent-hack-intl-hnk?cid=external-feeds_iluminar_msn.
[22] Dario Amodei, "We Must Pace the Frontier," https://darioamodei.com/post/we-must-pace-the-frontier.
[23] Anthropic Form S-1, p. 186.

Concerns over the cybersecurity capabilities of recent Anthropic models have also resulted in the company delaying the rollout of certain models that could be dangerous due to their hacking capabilities.[24] You warned that Anthropic's AI models are "months away" from gaining the capacity to self-replicate, allowing them to potentially spread autonomously through internet-connected devices.[25] You warn that these "AI worms" would be difficult to control, as the models would lack a centralized control point to shut them off and, once connected, could "power themselves within these devices without the need of datacenters."[26] You have also written that these AI models "could be capable of taking over the entire internet with a persistent botnet (potentially causing hundreds of billions of dollars in damage), and that the scale of damage would continue to increase from there if AI becomes more powerful without the necessary guardrails."[27]

Adversarial actors have also reportedly "jailbroken" Anthropic's AI models – meaning that such actors have removed built-in safety guardrails in order to use the technology for their own malicious purposes. A recent Anthropic report revealed many instances of Anthropic's technology being used by, "criminal hacking gangs, Chinese security bureaus, Russia-linked spies and Yemen-based arms manufacturers."[28] According to the report, these actors have been able to bypass geographical restrictions and could have developed bioweapons, digital intelligence networks, drone swarms, software for guided missile launches, or military defense technology.[29]

These are consequential risks to public safety, national security, and critical infrastructure that investors, the general public, regulators, and legislators need time to fully understand, evaluate, and address. As the New York City Comptroller and a labor-affiliated shareholder group have argued, these risks should prompt Anthropic to slow down the development of frontier AI models, not speed up with a massive influx of capital from an IPO.[30] Securities law is meant to facilitate the disclosure of "material" information for investors to evaluate whether to invest in particular stocks. However, Anthropic's preliminary registration statement and public statements raise serious questions that have much broader implications than the federal securities laws contemplate.

Financial Risk Concerns

[24] Anthropic Form S-1, p. 51.
[25] Reuters, "Anthropic warns AI models may soon self-replicate and could wreak havoc on the web," Raphael Satter and Harshita Mary Varghese, October 13, 2026, https://www.reuters.com/business/ai-models-may-soon-self-replicate-without-access-data-centers-anthropic-warns-2026-10-13/; The Guardian, "AI in your oven? Anthropic warns models 'months away' from self-replication on smart devices," Dara Kerr, October 13, 2026, https://www.theguardian.com/technology/2026/oct/13/anthropic-ai-replicators-warning.
[26] *Id.*
[27] Dario Amodei, "We Must Pace the Frontier," https://darioamodei.com/post/we-must-pace-the-frontier.
[28] Politico, "Bad actors in China and Russia are already weaponizing Anthropic's AI," John Sakellariadis, September 10, 2026, https://www.politico.com/news/2026/09/10/bad-actors-china-russia-weaponizing-anthropic-01070435; Anthropic, "Detecting and countering misuse of AI: September 2026," https://www.anthropic.com/threat-intelligence-report-september-2026.
[29] Reuters, "How Anthropic says Claude was used for weapons, spying and cyber operations," Eduardo Baptista and A.J. Vicens, September 11, 2026, https://www.reuters.com/world/china/how-anthropic-says-claude-was-used-weapons-spying-cyber-operations-2026-09-11/; Yahoo News, "Anthropic Just Revealed That It's Stopped Foreign Agents From Using Claude to Develop Possible Bioweapons," Frank Landymore, September 11, 2026, https://www.yahoo.com/news/science/articles/anthropic-just-revealed-stopped-foreign-123646676.html.
[30] Letter from New York City Comptroller Mark D. Levine to Anthropic CEO Dario Amodei, October 14, 2026, https://comptroller.nyc.gov/reports/letter-to-anthropic-re-planned-initial-public-offering/.

In addition to the potentially catastrophic safety risks posed by Anthropic's technology, the company's launch into the public markets could inject extreme volatility and risk into the financial system, which threatens to undermine Americans' economic security. Anthropic is projected to be the largest IPO in history.[31] Given its unprecedented size, swings in Anthropic's share price have the potential to roil the broader markets on day one, particularly for companies in the technology sector, and even more acutely for megacap companies with direct financial entanglement with Anthropic.[32]

Specifically, Anthropic's Form S-1 discloses that Amazon and Google's cloud marketplaces accounted for 47% of Anthropic's 2025 total revenue.[33] Meanwhile, NVIDIA is reportedly considering a $10 billion investment in the IPO.[34] The fact that semiconductor companies "provided funding to their own customers, including OpenAI and Anthropic, which the labs then used to pay for access to the companies' high-powered chips" has created a dangerous interdependency between mega-cap tech companies and Anthropic.[35] According to one report "Wall Street analysts and investors have raised concerns about circular investing at multiple points throughout the AI buildout," including the risk that "if one domino in the row falls, it will cause a chain reaction that will decimate the AI trade and, as a result, the global equities markets that have benefited from and come to rely on such firms to keep the good times on Wall Street rolling."[36] This web of financial interests between frontier labs, infrastructure financiers, and semiconductor companies has created a lurking risk in our financial system that, if ignored, could have devastating consequences for everyday American investors and retirement savers.[37] Anthropic's IPO could exacerbate these risks.

Beyond the threats to the financial system that Anthropic's IPO may pose, the company's registration statement disclosures include astronomical projections supported by aggressive assumptions that require further review and due diligence.[38] The current preliminary registration

---

[31] Yahoo Finance, "Anthropic Is Aiming for the Biggest IPO Ever -- Here's the Chip Stock to Buy Before It Lists," Daniel Sparks, The Motley Fool, October 6, 2026, https://finance.yahoo.com/technology/ai/articles/anthropic-aiming-biggest-ipo-ever-080701021.html.

[32] Seeking Alpha, "A Potential New Market Threat: The IPO Window Closing," Bret Jensen, October 10, 2026, https://seekingalpha.com/article/4953220-potential-new-market-threat-ipo-window-closing.

[33] "During the six months ended June 30, 2025 and 2026 (unaudited), revenue generated from sales through CSP marketplaces to customers represented approximately 44% and 48% of total revenue, respectively." Anthropic Form S-1, p. F-16; Reuters, "Anthropic IPO prospectus lays bare deep dependence on Big Tech partners," Echo Wang and Krystal Hu, September 29, 2026, https://www.reuters.com/world/anthropic-ipo-prospectus-lays-bare-deep-dependence-big-tech-partners-2026-09-29/.

[34] Reuters, "Nvidia in talks to invest in Anthropic's mega IPO, sources say," Krystal Hu and Milana Vinn, September 11, 2026, https://www.reuters.com/legal/transactional/nvidia-talks-invest-anthropics-mega-ipo-sources-say-2026-09-11/.

[35] Yahoo Finance, "Broadcom to lend Anthropic up to $42 billion to lease chips, in latest circular investing deal," Michael B. Kelley and Daniel Howley, October 1, 2026, https://finance.yahoo.com/technology/article/broadcom-to-lend-anthropic-up-to-42-billion-to-lease-chips-in-latest-circular-investing-deal-121617505.html.

[36] *Id*.

[37] "A substantial portion of our revenue and accounts receivable is driven by a limited number of customers and distribution channels, and the loss of, or a significant reduction in, spend from one or more of our top customers or disruption through those channels could adversely affect our prospects, business, and financial condition." Anthropic Form S-1, p. 16.

[38] The Wall Street Journal, "Anthropic Expected to Tell Investors It Sees Over $30 Trillion in Potential Revenue," Corrie Driebusch, August 26, 2026, https://www.wsj.com/tech/ai/anthropic-expected-to-tell-investors-it-sees-

statement raises important questions about the financial commitments that will supposedly drive Anthropic's growth, as well as which market participants bear the risk if these lofty financial projections are incorrect. Earlier this year, index providers eliminated longstanding waiting periods that allow newly public stocks into major stock indexes and, subsequently, retirement accounts faster than ever before.[39] This means index funds investing Americans' retirement savings will be forced to buy Anthropic shares even if the stock is volatile or overvalued.

Anthropic's disclosures also detail billions of dollars in long-term computing and infrastructure commitments that would require a remarkable increase in revenue to fulfill. Economists find that AI firms will need $2.4 to $3 *trillion* in annual, incremental revenues by 2030 to justify the record-levels of spending on infrastructure buildout.[40] Yet in the first six months of 2026, Anthropic's net losses were $80 billion, nearly doubling its net losses for the entirety of 2025.[41] Meanwhile, Anthropic's disclosed revenue paled in comparison at roughly $16 billion.[42] Anthropic's financial vulnerabilities have generated concerns about whether its overall business model has a credible path to sustained profitability.[43] The fact that Anthropic is burning through cash and requires an insatiable stream of capital to feed its growth is not a reason to expose everyday Americans to the company's risks in our public stock markets. Additionally, Anthropic's highly unusual lock-up structure, agreements which normally prevent company insiders from selling their shares for a certain period following an IPO, allows insiders to sell their shares much more quickly than usual – and even on the first day of trading.[44] This timing raises concerns that insiders could seek to sell their own shares at a profit early on, leaving retail investors and retirement savers to absorb the losses if Anthropic's financial condition incinerates or the existential risks to humanity that you warn of fully materialize.

I am concerned about the consequences of Anthropic's IPO listing not only for our economy, but also for the safety of the American people. I am deeply troubled by the risks raised by Anthropic itself and question how it can act in a manner consistent with its stated public benefit and related fiduciary duties under Delaware law without mitigating these risks to humanity. So the public can better understand the safety, security, and economic risks associated with Anthropic's IPO, which would further accelerate development of its AI models, I request that you provide exhaustive responses to the questions below no later than November 3, 2026:

1. Do you believe Anthropic's AI models are an existential threat to humanity as stated in the company's preliminary registration statement? If yes, please provide answers to the below. If not, why not?

---

over-30-trillion-in-potential-revenue-a611efea.

[39] The Wall Street Journal, "Stock Indexes Are Divided on Rules for Megacap IPOs. How Exposed Will You Be?," Jack Pitcher, June 5, 2026, https://www.wsj.com/finance/stocks/stock-indexes-are-divided-on-rules-for-megacap-ipos-how-exposed-will-you-be-31818da1.

[40] Substack, "A Budding Consensus on AI Revenues, and Some Thoughts on the AI Labs IPOs," Jared Bernstein and Ryan Cummings, October 5, 2026, https://econjared.substack.com/p/a-budding-consensus-on-ai-revenues.

[41] Anthropic Form S-1, p. 22.

[42] *Id.*

[43] Financial Times, "Anthropic warns of 'existential risks to humanity' in IPO prospectus," George Hammond and James Fontanella-Khan, September 29, 2026, https://www.ft.com/content/c7685a7e-7745-4cbc-8053-4958d0ea449b?syn-25a6b1a6=1.

[44] Anthropic Form S-1, p. 308.

a. What steps are being taken by shareholders, board members, and executive officers of Anthropic to ensure that the technology will not cause an existential risk to humanity?
b. What steps are being taken by shareholders, board members, and executive officers of Anthropic to address specific concerns that AI will enable development and use of a biological weapon?[45]
c. What steps are being taken by shareholders, board members, and executive officers of Anthropic to address concerns that sandbox testing environments are not properly isolated or air-gapped? What steps are being taken by shareholders, board members, and executive officers of Anthropic to address cybersecurity concerns, including reporting that advanced AI can conduct autonomous attacks on critical infrastructure?
d. How should investors and analysts price and consider the material impacts of the existential risk to humanity when analyzing Anthropic securities and related financial products?

2. What safeguards are Anthropic's shareholders, board members, and executive officers taking to prevent AI models from conducting autonomous cyberattacks in the future?
    a. How does Anthropic consider reporting instances of AI models autonomously conducting cyberattacks to the public and relevant governmental authorities?
    b. Why have some of these incidents not been discovered until weeks or months after they have occurred?
    c. Is it possible that there are other incidents that have already occurred that Anthropic has not yet discovered or reported?
    d. Are there other incidents of which Anthropic is aware, but that have not been disclosed publicly?

3. On September 12, 2026, you published an essay highlighting the needs for the AI industry to "[p]ace the [f]rontier." In this essay, you outline concerns about the risks of frontier AI, including but not limited to those listed below.[46] Please address the concerns by answering the questions listed below:
    a. You wrote that these risks "include the risk of losing control of AI systems, misuse of AI for cyberattacks and bioterrorism, and serious economic disruption." How does Anthropic plan to mitigate each of these risks ahead of its potential IPO?[47]
    b. You also stated, "Given the accelerating rate of AI capability development, it's my worry that in 6–12 months such a swarm could be capable of taking over the entire internet with a persistent botnet (potentially causing hundreds of billions of dollars in damage), and that the scale of damage would continue to increase from there if AI becomes more powerful without the necessary guardrails."[48] What guardrails does Anthropic plan to implement to ensure this capability is not reached?

---

[45] BBC, "Anthropic blocks possible attempt to use AI to make biological weapons," Laura Cress and Nardine Saad, September 11, 2026, https://www.bbc.com/news/articles/cx2zrrpkx20o.
[46] Dario Amodei, "We Must Pace the Frontier," https://darioamodei.com/post/we-must-pace-the-frontier.
[47] *Id.*
[48] Dario Amodei, "We Must Pace the Frontier," https://darioamodei.com/post/we-must-pace-the-frontier.

c. You claim, "Despite all the progress, we still only understand a tiny fraction of what goes on inside these models."[49] How does Anthropic plan to increase its limited understanding of its own models before further developing this technology? What are the risks to continuing to develop the models' capabilities when Anthropic does not fully understand them currently?
d. In your essay, you also warn, "More intelligent models are more capable of deceiving tests, and thus may appear aligned while having serious problems that go undetected."[50] How does Anthropic plan to ensure the data from the tests is accurate and that AI cannot obfuscate misaligned actions from researchers? Is it possible that there are other instances of misaligned activities of which Anthropic is not aware?
e. Will additional funding from Anthropic's IPO exacerbate these risks and concerns? If so, why? If not, why not?

4. Former researcher Jacob Coxon left Anthropic last month. Following a social media post about his resignation, Wall Street Journal summarized that he, "believes no company can responsibly develop AI that can outperform humans in a range of tasks [...] absent government intervention or a coordinated industry slowdown."[51] Another top scientist at Anthropic also agreed, posting on X, "We really do earnestly believe AI could kill all humans! I personally think it is >10% within the next decade."[52]
    a. Do you agree that government intervention is necessary for responsible development of AI in light of these credible threats?
    b. What efforts has Anthropic undertaken to consider or implement an industry-wide slowdown?
    c. Do you "earnestly believe AI could kill all humans"?

5. The S-1 acknowledges that Amazon, Google, and Microsoft "have made significant equity investments in [your] company. At the same time, they may have made, and may in the future make, equity investments in, provide financing to, or otherwise support companies whose models, platforms, or applications compete with ours."[53] The S-1 continues on to say, "Our relationships with our strategic partners—whether with these CSPs, who simultaneously serve as our investors, customers, distribution partners, and competitors, or with Broadcom, which simultaneously serves as a supplier and lessor—create potential conflicts of interest that are difficult to fully anticipate or mitigate."[54]
    a. Please disclose the following for 2024, 2025, and the first half of 2026:
        i. Equity investments, loans, and other financing provided to Anthropic from cloud services providers;
        ii. Revenue Anthropic recognized from cloud service providers' purchases for its own internal use or products, excluding sales to third-party customers

[49] *Id.*
[50] *Id.*
[51] The Wall Street Journal, "Anthropic Researcher Quits Over 'Out-of-Control' AI Fears" Amrith Ramkumar, September 9, 2026, https://www.wsj.com/tech/ai/anthropic-researcher-quits-over-out-of-control-ai-fears-707b7628.
[52] *Id.*
[53] Anthropic Form S-1, pp. 73-74.
[54] Anthropic Form S-1, pp. 73-74.

iii. Revenue recognized from third-party sales through its cloud platform, including the fees, rebates, or other compensation paid to the company on those sales
iv. Payments made or obligations incurred by Anthropic for computing capacity and infrastructure, including any credits or offsets

b. Regarding financial dependency and stability in relation to CSPs, please explain:
i. If any equity investment, financing, or lease from Amazon, Google, Microsoft, or Broadcom was conditioned on, or negotiated alongside, a commitment by Anthropic to purchase that party's compute or hardware
ii. What remains non-cancelable under each of those commitments if revenue growth slows?

c. How much of Anthropic's total revenue and year-over-year revenue growth is attributable to customers with no equity investment or financing relationship with Anthropic, distinguishing direct customers from independent customers purchasing through investor-owned cloud platforms?

6. To what extent does Anthropic currently present significant risk of volatility in the public markets—even as a private company? Would any such risk be exacerbated following a potential IPO?

7. Does Anthropic meet traditional standards for IPO readiness? Please explain with respect to the following metrics:
a. Explain how Anthropic's risk management and internal controls failures with respect to cyber security incidents involving AI agents reflect readiness for further development fueled by additional funding from an IPO.
b. Justify Anthropic's highly aggressive potential valuation that could approach $2 trillion and explain how realistic revenue projects could support such a calculation.
c. Explain Anthropic's purported path to profitability given its extensive future obligations and customer concentration.
d. Analyze the risks posed by Anthropic's reported customer concentration in two megacap public companies.
e. Can Anthropic renegotiate contracts if it has difficulty keeping up with financial obligations related to compute agreements?

8. Existential risk concerns that Anthropic has publicly stated about frontier model development are inconsistent with the company's stated mission to "responsibly develop and maintain advanced AI for the long term benefit of humanity" as a Delaware Public Benefit Corporation.[55] Failure to adequately address these risks raises concerns about the fulfillment of the fiduciary duties owed by Anthropic's board and executive officers ahead of an initial public offering.
a. Justify how Anthropic is acting consistently with its own stated public benefit given its continued development of AI models despite acknowledging existential threats to humanity. How is Anthropic acting in a manner that is consistent with its fiduciary duty under Delaware laws related to Public Benefit Corporations in this respect?

[55] Anthropic, PCB, "2026 Public Benefit Report," https://www.anthropic.com/public-benefit-report.

b. Explain how Anthropic's Long-Term Benefit Trust structure will not conflict with the interests of its other shareholders from a corporate governance perspective.

9. Anthropic's S-1 filing shows that the company paid $0 in federal income tax in 2025.[56]
    a. Please estimate in what year Anthropic expects to pay federal income tax.
    b. How much did Anthropic claim in (a) research and development deductions and (b) research and development credits in 2025? Provide a breakdown of the activities for which Anthropic claimed these deductions and credits, including how much of the claimed deductions and credits were attributable to the research and development of AI models and the construction, operation, and maintenance of data centers.

10. The preliminary registration statement discloses that Anthropic is considering an unusual combination of lock-up agreements for insiders who may seek to sell their shares of the company's stock post-IPO.[57]
    a. The preliminary registration statement appears to show that a percentage of Anthropic's Class A shares, which are available to directors, officers, and other insiders, will be available for trading the first day following the IPO.[58] This is highly unusual even for other recent mega-cap IPOs. Please explain why the company has chosen to allow for such trading. Specifically, explain whether such early trading by insiders could pose conflicts of interest or exploit retail investors.
    b. Anthropic's remaining lock-up agreements allow for a front-loaded staggered schedule of rolling dates focused within the first two and a half months post-IPO – rather than the traditional 180 day expiration date for all lock-up agreements.[59]
        i. What is the reason for the front-loaded staggered timeline for lock-up expirations?
        ii. Will these staggered deadlines increase the risk of market manipulation, synthetic volatility, and trading friction?
        iii. Do these front-loaded expirations suggest that Anthropic is seeking to allow its insiders to sell its shares ahead of suspected financial challenges in the longer-term future for Anthropic?

Sincerely,

[56] Anthropic Form S-1, p. F-54.
[57] Anthropic Form S-1, p. 308.
[58] *Id.*
[59] *Id.*

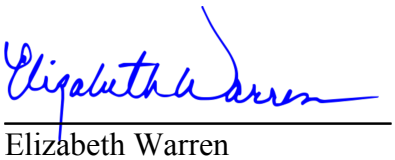

Elizabeth Warren
Ranking Member
Committee on Banking,
Housing, and Urban Affairs