



THE CITY OF NEW YORK
OFFICE OF THE COMPTROLLER
MARK D. LEVINE

October 14, 2026

Dario Amodei
Chief Executive Officer
Anthropic, PBC
500 Howard Street
San Francisco, California

Dear Mr. Amodei:

I write to urge Anthropic to pause its planned initial public offering (IPO) until the company fully and transparently addresses the artificial intelligence (AI) safety, governance, and development pace concerns it has raised. Instead of proceeding headlong with its IPO, the company should focus its energies on addressing these unprecedented risks.

Anthropic and five other companies recently signed a “morally binding” accord with the White House, which is insufficient to meet the moment. Before going public, the company should prioritize addressing the myriad safety concerns surrounding AI threats and establishing a governance regime that matches the seriousness of the concerns being raised by you as well as industry researchers, policymakers, and the public. The demands of preparing for a public offering diverts resources and attention from this very necessary work.

As Comptroller of the City of New York, I serve as investment adviser to, and custodian and a trustee of, the New York City pension funds, which have over \$325 billion in assets and constitute one of the largest public pension systems in the United States.

In recent weeks, Anthropic has issued unusually direct public statements acknowledging that the pace of frontier-model capability development may now exceed society’s capacity to understand and control the resulting systems. You have disclosed that unchecked recursive self-improvement could outrun human oversight. In your September 12 essay, “We Must Pace the Frontier,” you stated: “...over the last few months, I have become convinced that fully addressing the risks requires even more prudence ... We must slow the pace at which we improve the capabilities of AI models.”

Despite these heightened safety warnings and broad calls for greater caution across the industry, Anthropic reportedly continues to push forward an IPO that could value the company at approximately \$2 trillion.

Your essay followed the September 9 resignation of Anthropic researcher Jacob Coxon, who publicly said that insiders earnestly believe AI “could kill us all by the end of the decade.” This claim was supported by Anthropic's own Alignment Science Lead, Evan Hubinger, who personally thinks the risk is greater than 10% within the next decade.¹

According to news reports, Anthropic’s draft S-1 filing warns that advanced AI could pose “catastrophic or existential risks to humanity.”² Reporting also notes that the risk disclosures span 80 pages, nearly twice the length of the section describing the business. These disclosures are not routine and indicate risks that are without precedent, let alone at a public offering of this scale.

Anthropic is a Delaware public benefit corporation with a Long-Term Benefit Trust. Press reports indicate that Anthropic intends to adopt a governance structure that may give founders and trustees concentrated control of the company. A gap between control and accountability is especially concerning for investors given the unprecedented risks the company has already acknowledged.

These concerns underscore the seriousness of the risks at issue. Even if these outcomes prove overstated, such dire warnings illustrate the broad economic risk that imprudent AI advances pose. Addressing the Hugging Face incident specifically, you warned 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.”

Considering these concerns, Anthropic should disclose with specificity how it will implement appropriate audits and controls, the scope of what these controls are intended to address, as well as a projected timeline. This should also include an explanation of whether and how these controls will sufficiently address your call to slow the pace at which the capabilities of AI models advance.

In any event, given the gravity of your stated concerns, prudence should not only be shown with respect to the pace of the frontier, it should be also shown with respect to the timing of an IPO.

I therefore urge Anthropic to pause the IPO process so that the company may fully address these key issues, ones which investors must also assess regarding the company’s revised development strategy, governance commitments, and emerging risk landscape.

¹ <https://startupfortune.com/anthropic-researcher-jacob-coxon-quits-warning-ai-could-kill-us-all/>;
<https://www.cbsnews.com/news/ai-kill-humans-anthropic-researcher-more-than-ten-percent-chance/>;
<https://www.bbc.com/news/articles/ckgwy1k42w4o>

² <https://www.reuters.com/business/finance/anthropic-warns-ai-may-pose-existential-risks-humanity-ipo-filing-2026-09-29/>

I request a written response and welcome the opportunity to discuss these matters further.

Sincerely,

A handwritten signature in black ink, appearing to read "Mark D. Levine". The signature is fluid and cursive, with the first name "Mark" being more prominent and the last name "Levine" following in a similar style.

Mark D. Levine
New York City Comptroller