

Economically Pacing the Frontier

Dhilan Shah

Contents

1 Introduction and Class Discussion

2 Pacing the Frontier

3 The Economy

4 Pacing Implications for the Economy

This article examines the concept of “pacing the frontier” and its economic implications.

1 Introduction and Class Discussion

On Monday, September 14th, our class discussed Dario Amodei’s article “We Must Pace the Frontier.” After a brief overview of the article, we dived into a more general discussion about Amodei’s concept of “pacing.” The first topic of conversation concerned the role of embedded evaluators. As I explain later in my essay, an embedded evaluator is someone with “employee-level access” at a lab who monitors and reports on internal model development. Students were frustrated by the apparent “vagueness” of the role: What protections do evaluators have? Who can become an embedded evaluator? Are they appointed by the labs?

“External reviewers should have the right to publish key findings about risk levels, incidents, practices, and the access they received or didn’t receive — without editorial control by Anthropic. We will have the narrow ability to redact security-sensitive, legally privileged, commercially sensitive, or third-party confidential information, but we can’t redact findings just because they are unfavorable. The reviewers can say publicly if a redaction removed something important to their conclusions.” [1]

This statement in particular prompted some interesting discussion. What is a “commercially sensitive” finding? What is the bar for censorship? Does the report of a redaction have different effects (publically) than a report of an incident? We also discussed the apparent need for better protection for whistleblowers, leading to a conversation about the state of current legislation (specifically S. 1792 and H.R. 3460) [30, 31]. In addition to discussing embedded evaluators, the class also debated the feasibility of global coordination in pacing the frontier. The class analyzed China’s response to calls for slower frontier AI development: Specifically, *a September 14, 2026 press briefing, where Foreign Ministry spokesperson Guo Jiakun criticized fear-mongering and confrontation as obstacles to global AI governance and called for open, inclusive AI development [22]*. The consensus was that pacing the frontier globally will require unprecedented collaboration across governments, and that the realistic first step is to pursue a domestic solution.

The rest of this essay expands on these discussions and examines an issue that received less attention in class: how pacing the frontier could affect economic and financial stability.

2 Pacing the Frontier (summary)

Dario Amodei, CEO and co-founder of Anthropic, formalized the concept of pacing the frontier in early September. In his article, “We Must Pace the Frontier,” Amodei shares his concerns about the state of AI advancement and his framework for safer development. He begins by describing two reasons for his growing fear of AI development: First, he notes that models are becoming increasingly capable of developing and training future models [1], which he argues could allow development to outpace human oversight. Second, he discusses the OAI-HF attack, claiming that it wasn’t merely a failure of the company, but an extreme example of a class of incident that is becoming increasingly common across the industry.

After outlining his concerns, he proposes a framework to address future AI development, which he describes as “Pacing the Frontier” (PTF). He clarifies that PTF doesn’t involve halting model training or technical progress, but rather consists of “ensuring companies take adequate time to align and safeguard their models, and for third-party evaluators to confirm this” [1].

He outlines three steps to PTF:

Embedded evaluators:

Frontier AI companies give evaluators employee-level access to internal data, and the evaluators can actively report their findings to the public. Note that under this proposal, evaluators would still face limits on publishing information that might legally or commercially hurt Anthropic [1]. However, Amodei clarifies that if information were redacted from a public release, evaluators could discuss the reasons for those restrictions (including publicly challenging those redactions) [1].

Democratic coordination:

Frontier AI companies within democratic countries should coordinate to establish common safety standards as well as limits on the rate of unchecked AI progress. Because legislation can take time, Amodei also calls for direct collaboration between labs [1]. He notes that certain forms of coordination would require a narrow antitrust waiver [1].

Global coordination:

The US government and other democracies should coordinate with authoritarian governments to mitigate AI risks. Recently, Amodei reiterated these claims in his address to the UN Security Council on September 23, 2026, where he claimed that unchecked artificial intelligence could pose an existential risk to humanity (specifically citing bioterrorism and lack of control).

Anthropic also published an “Advanced AI Framework,” which outlines (in more detail) the obligations for AI companies that train models requiring more than 10^{25} training FLOPs and earn more than \$500 million in annual AI-derived revenue (or spend more than \$1 billion per year on AI research and development) [2]. In the document, Anthropic defines some basic concepts, such as “catastrophic risk” (a foreseeable and material risk that will materially contribute to significant death, injury, or damage) and “periodic review” (a government agency should review criteria across all developers annually). The document also outlines what developers should publish (transparency) and standards for independent evaluation and security. The document also formalizes a government agency for review and sets “societal resilience measures” to describe how society should prepare for more capable models. Read more here:

<https://www-cdn.anthropic.com/files/4zrzovbb/website/0a58d567024a8b448ff15158ebc3625328dfcc1f.pdf>

A Different Perspective: What Critics Are Saying

This section examines the framing of “pacing” through the lens of Anthropic as a company. Common critics who view labs primarily as commercial enterprises question whether their support for PTF reflects virtue signaling or an effort to secure regulatory advantages.

Recently, Anthropic filed an IPO prospectus, details of which were subsequently reported by media outlets [3, 4]. Perhaps more fascinating than the lab’s proposed valuation (\$2 trillion) and reported 2025 revenue (\$4.6 billion) were the company’s statements about safety. In the prospectus, Anthropic focused heavily on the risk factor of AI, explicitly mentioning things like “self-preserving behaviors” and efforts to “conceal or manipulate

information” [4]. I find this framing interesting; personally, I’m very curious about how Amodei plans on presenting the “pacing” argument to investors. I imagine it would be something like this, where he presents existential risk to investors in terms of its potential commercial and legal implications.


Some people claim that Anthropic’s filing of its S-1, as well as the release of Opus 5.5, is somewhat inconsistent with its stated commitment to pacing. In “Pacing the Frontier is not the actual goal for AI Labs” (LessWrong), rahulxyz argues that “posturing from the labs about slowing down and deeply caring about safety is done in order to retain and calm the employees who they are dependent on to keep pushing capabilities to get to AGI. If it were not for a big contingent of employees pressing them (increasingly publicly), they would make zero public acknowledgements of risks at all” [5]. He reaches this conclusion after 1) pointing to the seemingly unchanged roadmap for model releases from labs and 2) the relationship between top talent in the AI industry and safety (safety’s importance in attracting and retaining researchers). Now, it’s important to understand that there is evidence that labs aren’t merely virtue signaling. However, it’s important to understand that from the perspective of “a rational company”, this would be a reasonable argument for why appearing to support PTF could offer a competitive advantage.

Another common rationale for the push for AI safety is regulatory capture. A statement from Michael Best & Friedrich LLP (an Am Law 200 firm, according to the firm’s own announcement [24]) captures the essence of this idea: “the big established labs are happily joining the hysteria bandwagon to help protect their turf. By convincing regulators and the public that AI is an uncontrollable, existential monster, these companies can successfully lobby for heavy compliance frameworks and regulatory bottlenecks. Those rules would disproportionately burden smaller competitors and open-source projects.” [6]. By viewing Anthropic as a venture-backed business that has investors and must become profitable, the idea that PTF is a strategy for retaining talent or securing regulatory advantages becomes plausible.

3 The Economy


We now consider pacing with respect to the US economy, focusing on general spending, foreign investment, and the stock market.

First, we are spending a lot on AI. The total investment in data centers and related artificial intelligence infrastructure is projected to total nearly \$10.3 trillion from 2025 to 2032 (according to new estimates from economist Stijn van Nieuwerburgh, published by the Brookings Institution). The projection is equivalent to an average of approximately 3.6% of U.S. GDP annually [29].



[8]

Additionally, the data center buildout is mainly responsible for keeping the private construction industry afloat, as hyperscalers continue to commit substantial capital to building out infrastructure [29]. Finally, foreign investment into US equities is at a record high [9]. Countries like Saudi Arabia are investing billions of dollars into US companies, largely for AI exposure [10].




[9]

Here’s a statement from the government about a strategic AI partnership with Saudi Arabia:

The two sides noted that this strategic partnership encompasses the supply of advanced semiconductors, the development of AI applications, building and developing advanced AI infrastructure, building national capabilities, and expanding high-value investments between the two countries. This partnership will contribute to boosting productivity and innovation, growth and prosperity, and achieving economic and social returns for both countries. [10]

The dollar is largely supported by oil and military force. With both being challenged by continued conflict in the Middle East, AI remains the sole bulwark of American spending power. Comparing shareholder returns with capital spending reveals a similarity between Big Oil and Big Tech.



The Stock Market

10-year bond yields reached 5.1% and are heading toward 5.2% this year [11]. Typically, higher yields lower stock valuations due to higher discount rates, but the market continues to surge. The irrationality of the market can be attributed to a few factors:

- Expectations of a peace agreement in the Middle East (lower oil prices)
- Expectations about Federal Reserve policy: “the market has largely priced in expectations for future interest rate hikes” (Warsh) [11]
- AI infrastructure investment

We’ll focus on this third part. The Fed and the government view AI as “too big to fail” and believe it will ultimately boost US productivity and serve as a disinflationary force in the long run (AI stocks are macro-immune).

The following statements illustrate Federal Reserve officials’ views on AI investment:

Speaker & Date	Quote
Gov. Michelle Bowman, Jan 30, 2026	“strength in AI-related investment ... will continue to boost productivity gains” [12]
Gov. Michael Barr, Mar 26, 2026	“investments in AI are likely in the future to contribute to strong productivity growth” [13]
Chairman Kevin Warsh, Jul 29, 2026 (FOMC press conference)	“The surge in high-tech capex has been remarkable.” [14]
Gov. Christopher Waller, Sep 3, 2026	“AI investment is a legitimate part of GDP today” [15]
Gov. Michael Barr, Sep 29, 2026	“the surge of investment, and related demand from the AI buildout, is having a measurable effect on prices” [16]

Pacing Implications for the Economy

To understand how PTF can directly impact the economy, we must examine the expectations surrounding investment in AI. Governor Lisa Cook notes that companies have spent only a small fraction of the \$2 trillion in announced data-center plans [21]. Investors, labs, and hyperscalers are betting on continued scaling. Revenue multiples provide an excellent starting point for examining how much future growth those valuations require.

Company	Valuation	Revenue*	Multiple
OpenAI · Mar 2026	\$852B	\$24B	35.5×
Anthropic · May 2026	\$965B	>\$47B	≤20.5×
Cisco · Mar 2000	\$569B	\$17.4B	32.7×

OpenAI and Anthropic disclosed these valuations alongside their revenue figures in their respective funding announcements [25, 26]. Cisco provides a historical comparison: its March 2000 market capitalization divided by its latest reported quarterly revenue, annualized, yields a multiple comparable to those of these AI companies [27, 28]. Wider adoption, enterprise sales, and continued model development all contribute to the commercial thesis. OpenAI’s explanation to investors explicitly links additional compute to more intelligent models, better products, adoption, revenue, and cash flow [25].

Continued capability improvement is therefore an explicit part of the investment thesis, although this statement alone does not establish how much of the valuation depends on it. Pacing the frontier could slow one mechanism through which these companies expect to increase revenue. To the extent that expected returns depend on capabilities beyond those available today, a slowdown could delay cash flows, reduce valuations, and make some planned investments less profitable than expected. Given the scale of capital committed to the AI buildout, a sufficiently large reassessment could affect financial conditions

4 Conclusion

“We Must Pace the Frontier” presents the moral imperative to slow development of new models at Frontier AI labs. It presents a loose framework for what should be organized (within companies, domestically, and internationally) to prevent threats from a misaligned superintelligent model. However, data suggests that the current US economy *depends* on the continued development of the frontier. Thus, we are presented with a difficult challenge: how do we properly address the problems of AI safety (slow down) while also keeping the United States economy from collapsing? I think the natural next step is to assess the actual risk of “pacing” at each layer of the AI stack: ie., how does the data center and grid buildout get impacted by this? We also need to better understand the GPU market: why hasn’t compute become less expensive, why has demand for compute continued to rise (as inference gets even cheaper), how much compute is used for inference/training (and prices for each).

These are all questions that must be answered as we continue to solve the problem of alignment. I do believe that there is a solution to the AI safety problem that preserves the prosperity of the American People, but it begins with an acceptance: *the US economy is at the mercy of the AI buildout, and if it slows, we risk disaster.*

Sources

[1] Dario Amodei. “We Must Pace the Frontier.”

<https://darioamodei.com/post/we-must-pace-the-frontier>

[2] Anthropic. “Advanced AI Framework.”

<https://www-cdn.anthropic.com/files/4zrzovbb/website/0a58d567024a8b448ff15158ebc3625328dfcc1f.pdf>

[3] The New York Times. Anthropic IPO filing.

<https://www.nytimes.com/2026/09/29/business/dealbook/anthropic-ipo-filing-s1.html>

[4] Reuters. Anthropic warns AI may pose existential risks to humanity in IPO filing.

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

[5] rahulxyz. “Pacing the Frontier is not the actual goal for AI Labs.” LessWrong.

<https://www.lesswrong.com/posts/Nm4ewbYovtjq69dvH/pacing-the-frontier-is-not-the-actual-goal-for-ai-labs>

[6] Michael Best. “We should be afraid of regulatory capture, not AI.”

<https://www.michaelbest.com/insights/we-should-be-afraid-of-regulatory-capture-not-ai-102o365/>

[7] LessWrong. “Timeline of AI safety.”

<https://www.lesswrong.com/posts/SEfjw57Qw8mCzy36n/timeline-of-ai-safety>

[8] The Wall Street Journal. “The AI Build-Out Is Becoming the Biggest Economic Bet in U.S. History.”

<https://www.wsj.com/economy/the-ai-build-out-is-becoming-the-biggest-economic-bet-in-u-s-history-c60716dd>

[9] Apollo. “The Dollar’s Hidden Dependence on the AI Trade.”

<https://www.apollo.com/wealth/insights-news/insights/daily-spark/the-dollars-hidden-dependence-on-the-ai-trade>

[10] U.S. Department of State. “Joint Statement: The Strategic Artificial Intelligence Partnership.”

<https://www.state.gov/releases/office-of-the-spokesperson/2025/11/joint-statement-the-strategic-artificial-intelligence-partnership>

[11] MishTalk. “Dot Plot of Fed Expectations Plus an Excellent Press Conference by Warsh.”

- <https://mishtalk.com/economics/dot-plot-of-fed-expectations-plus-an-excellent-press-conference-by-warsh/>
- [12] Michelle Bowman. January 30, 2026. Federal Reserve Board.
<https://www.federalreserve.gov/newsevents/speech/bowman20260130a.htm>
- [13] Michael Barr. March 26, 2026. Federal Reserve Board.
<https://www.federalreserve.gov/newsevents/speech/barr20260326a.htm>
- [14] Kevin Warsh. FOMC press conference transcript. July 29, 2026.
<https://www.federalreserve.gov/mediacenter/files/FOMCpresconf20260729.pdf>
- [15] Christopher Waller. September 3, 2026. Federal Reserve Board.
<https://www.federalreserve.gov/newsevents/speech/files/waller20260903a.pdf>
- [16] Michael Barr. September 29, 2026. Federal Reserve Board.
<https://www.federalreserve.gov/newsevents/speech/barr20260929a.htm>
- [17] Bank for International Settlements. Basel III global regulatory framework. Revised June 2011.
<https://www.bis.org/publications/201106-standards-basel-iii-global-regulatory-framework-more-resilient-banks-and-banking-systems-revised-version-june-2011>
- [18] OpenAI. “Hugging Face Incident and the Road Ahead.”
<https://openai.com/index/hugging-face-incident-and-the-road-ahead/>
- [19] Michael Barr. February 17, 2026. Federal Reserve Board.
<https://www.federalreserve.gov/newsevents/speech/barr20260217a.htm>
- [20] Kevin Warsh. Jackson Hole. August 28, 2026. Federal Reserve Board.
<https://www.federalreserve.gov/newsevents/speech/warsh20260828a.htm>
- [21] Lisa Cook. September 28, 2026. Federal Reserve Board.
<https://www.federalreserve.gov/newsevents/speech/cook20260928a.htm>
- [22] Ministry of Foreign Affairs of the People’s Republic of China. “Foreign Ministry Spokesperson Guo Jiakun’s Regular Press Conference on September 14, 2026.”
https://www.fmprc.gov.cn/eng/xw/fyrbt/202609/t20260914_12021997.html
- [23] AFP, via CNA. “OpenAI cancels release of newest model due to safety concerns.” September 28, 2026.
<https://www.channelnewsasia.com/business/open-ai-new-model-safety-astra-6416906>
- [24] Michael Best & Friedrich LLP. “Michael Best Appoints Jeff Mayer as Firm’s Next Chief Operating Officer.” May 19, 2026.
<https://www.michaelbest.com/insights/michael-best-appoints-jeff-mayer-as-firms-next-chief-operating-officer/>
- [25] OpenAI. “Accelerating the Next Phase of AI.” March 31, 2026.

<https://openai.com/index/accelerating-the-next-phase-ai/>

[26] Anthropic. “Anthropic Raises \$65B in Series H Funding at \$965B Post-Money Valuation.” May 28, 2026.

<https://www.anthropic.com/news/series-h>

[27] Cisco. “Market Capitalization.” March 27, 2000 historical milestone.

https://heritage.cisco.com/2000_marketcap/

[28] Cisco. “Cisco Systems Reports Second Quarter Earnings and Two-for-One Stock Split.” February 8, 2000.

<https://newsroom.cisco.com/c/r/newsroom/en/us/a/y2000/m02/cisco-systems-reports-second-quarter-earnings-and-two-for-one-stock-split.html>

[29] Stijn Van Nieuwerburgh. “Financing the AI Buildout.” Brookings Papers on Economic Activity, Fall 2026 conference draft. September 23, 2026.

<https://www.brookings.edu/articles/financing-the-ai-buildout/>

[30] U.S. Congress. S. 1792, AI Whistleblower Protection Act, 119th Congress, introduced May 15, 2025.

<https://www.congress.gov/119/bills/s1792/BILLS-119s1792is.pdf>

[31] U.S. Congress. H.R. 3460, AI Whistleblower Protection Act, 119th Congress, introduced May 15, 2025.

<https://www.govinfo.gov/app/details/BILLS-119hr3460ih>