



Insight

Pacing the AI Frontier: An (Allegedly Illegal) Agreement

FRED ASHTON | OCTOBER 7, 2026

Executive Summary

- On September 18, four leading artificial intelligence (AI) firms – Anthropic, OpenAI, SpaceX AI, and Google – were sued in federal court for allegedly agreeing to slow the pace of AI development in violation of federal antitrust law.
- The lawsuit, brought by a class of four paid subscribers, alleges a horizontal agreement stemming from an industry proposal calling for “industry-wide coordination” to slow down development and create safety limits amid the industry’s rapid development of AI capabilities.
- With safety at top of mind and industry moving at breakneck speed, Congress should reject calls for an antitrust exemption; AI firms can create voluntary safety standards and engage in procompetitive collaborations under existing law without slowing product development.

Introduction

On September 18, four companies at the frontier of artificial intelligence (AI) – Anthropic, OpenAI, SpaceX AI, and Google – were sued in federal court for allegedly agreeing to slow the pace of AI development, violating federal antitrust law.

The lawsuit, brought by a class of four paid subscribers, alleges a horizontal conspiracy stemming from an industry proposal calling for “industry-wide coordination” to slow down development and create safety limits amid the industry’s rapid development of AI capabilities.

With safety at top of mind and industry moving at breakneck speed, Congress should resist

creating an antitrust exemption for dominant AI firms. Instead, these companies can create voluntary safety standards and engage in procompetitive collaborations under existing law without slowing product development.

Section 1 of the Sherman Act and the Alleged Agreement

A class of four paid subscribers [filed suit](#) against Anthropic, OpenAI, SpaceX AI, and Google DeepMind under Section 1 of the Sherman Act – which prohibits “every” contract, combination, or conspiracy “in restraint of trade” – alleging an illegal horizontal agreement to suppress competition by slowing product quality improvements.

The complaint stems from a September 12, 2026, [essay](#) by Anthropic’s Chief Executive Officer Dario Amodei titled “We Must Pace the Frontier.” In the piece, Amodei called for “industry-wide coordination” to “establish common safety standards as well as limits on the rate of unchecked AI progress,” declaring that developers “must slow the pace at which we improve the capabilities of AI models.”

That same day, executives at competing firms publicly endorsed the proposal on social media. SpaceX AI founder Elon Musk took to his X platform and [quote-posted](#) Amodei’s proposal, writing “Dario is right.” OpenAI CEO Sam Altman similarly [quote-posted](#), “I agree with Dario that we need to pace the frontier,” while Google DeepMind co-founder Demis Hassabis [called](#) the essay “the right path forward.”

According to the plaintiffs – who hold paid subscriptions to Anthropic’s Claude, OpenAI’s ChatGPT, SpaceX AI’s Grok, and Google’s Gemini – subscribers pay for access to the “most capable models” and to continuous improvements. These features, the plaintiffs assert, are used by the companies to “market and price” the subscriptions and argued that an agreement to slow technological progresses “lowers the quality of what subscribers receive for the price they pay.”

The lawsuit also alleged that this concerted action predated Amodei’s essay and the subsequent responses by months. In July 2026, Amodei, OpenAI’s Chief Scientist, a Google DeepMind co-founder, and other high-ranking employees from leading AI labs [acknowledged](#) the “intense competitive pressure” not to unilaterally slow development.

Amodei’s Ask

In his essay, Amodei framed the need for the slowdown as vital for safety, explaining that today’s AI models “are an almost endless gold mine of insight into both how to build AI well and what can sometimes go wrong with it if it isn’t built well.” Slowing down the pace of development, therefore, would afford the industry “an extra year or two before models

reach critical levels of capability” and “greatly reduce the risk that something goes seriously wrong.”

The following is Amodei’s proposed three-step plan to pace the frontier:

1. **Embedded Evaluators.** Each frontier AI company commits to giving ongoing, employee-like access to a team of embedded third-party evaluators (such as [METR](#)), whose role is to verify adherence to safety practices and commitments, report incidents, and help assess the alignment of not just completed AI models but training pipelines and processes. This is the key step for *verifiability* of any pacing commitments, and has precedent in the banking industry, which sometimes involves regulatory “supervisors” embedded along with employees. **Anthropic is unilaterally committing to this step now.** We intend this to be part of a broader push to redouble efforts on our safety and alignment work.
2. **Democratic Coordination.** Frontier AI companies within democratic countries coordinate to establish common safety standards as well as limits on the rate of unchecked AI progress. Some forms of coordination that would be impactful for pacing are legally challenging, and will require government support.
3. **Global Coordination.** The US and other democratic governments attempt to coordinate with authoritarian governments, to the extent this is possible, while taking seriously the challenges of verifying compliance.

Amodei explicitly acknowledged that the second step “requires industry-wide coordination” in conjunction with “government mediation or waivers of antitrust restrictions.” These conditions are necessary for AI developers to do the work “without sacrificing commercial advantage or the United States’ lead in AI.” In other words, competition to improve models is so intense domestically and internationally – specifically with China – that any firm that unilaterally slows down would be punished by the market.

Forget Exemptions, There Are Other Solutions

Part of Amodei’s proposal calls for AI companies to establish “common safety standards,” which, according to the essay, requires a “narrow [antitrust] waiver for certain kinds of safety conversations.” Yet Congress should disregard this request. Antitrust exemptions inherently require an understanding of the competitive structure of a market in the future, a task that comes with wide, and potentially costly error bands – especially in a market as rapidly changing as AI. Moreover, incumbents can use sweeping antitrust exemptions to create a moat that solidifies their market position. Federal law already accommodates collaborations among competitors in different forms with varying degrees of regulatory

oversight.

Self-Regulatory Organization and Standards Development Organization

In requesting antitrust relief, Amodei pointed to Google DeepMind CEO Demis Hassabis's [proposal](#) of a self-regulatory organization (SRO) to create standards for frontier AI lab safety and security practices. Hassabis's proposal called for an SRO to develop "assessment protocols" and to "conduct testing in areas relevant to national security." Furthermore, the SRO would create metrics and benchmarks to qualify AI as "Frontier-class." Google released a [white paper](#) exploring this option in which the company called for a frontier AI regulatory organization.

Hassabis's proposal and the Google white paper pointed to the Financial Industry Regulatory Authority (FINRA) as a model for an SRO. FINRA writes and enforces rules that govern broker-dealers with oversight from the Securities and Exchange Commission. Google identified the North American Electric Reliability Corporation, the Public Company Accounting Oversight Board, and the American Medical Association as examples of "independent, government-supervised, industry-backed regulators." As Google explains, these "industry-funded bodies...write and enforce binding rules on their members," but "operate under the supervision (and ultimate veto) of a government agency."

Another option would be to create industry-wide safety standards through a formal standards development organization (SDO). [SDOs](#) are entities that bring together industry stakeholders to create, establish, and promote technical standards to ensure product safety, interoperability, and quality. Participants generate these standards through industry consensus and adopt them voluntarily, preventing dominant firms – such as the four named in the lawsuit – from creating barriers to entry.

SDOs can neither mandate firms adopt standards nor fine companies for failing to do so. SROs, by contrast, can levy fines, expel bad actors, or otherwise sanction firms and individuals from working in the industry. Without government oversight, an SRO risks antitrust infractions, especially if it creates barriers to entry that protect incumbents or prevent firms from releasing new models without the approval of competitors.

Collaboration Among Competitors

In the jointly published [2000 Antitrust Guidelines for Collaborations Among Competitors](#), the Department of Justice (DOJ) and Federal Trade Commission (FTC) explained that collaborations involving research and development are often considered procompetitive. Despite support for these guidelines being officially withdrawn by the agencies in December 2024, the evaluation framework outlined in the document is still supported by caselaw.

Replacement guidelines are [currently under consideration](#).

Notably, the 2000 Guidelines explain that such joint ventures – excluding those to fix prices or output, rig bids, or divide markets – are evaluated under the rule of reason framework. Under rule of reason, courts weigh the negative, anticompetitive effects against procompetitive justifications. In other words, security collaborations among competitors are not *per se* illegal.

Furthermore, cybersecurity information sharing among competitors already enjoys an antitrust exemption under the [Cybersecurity Information Sharing Act of 2015](#) (CISA). The act “provides a statutory exemption to Federal antitrust laws for the sharing between and among private entities of cyber threat indicators, defensive measures, or assistance relating to the prevention, investigation, or mitigation of a cybersecurity threat for a cybersecurity purpose.” Moreover, the CISA guidance document points to a DOJ and FTC May 2014 policy statement supporting the [rule of reason framework](#).

These guidelines and the rule of reason framework offer a set of general principles that help guide companies. Of note, they can legally adapt to market evolutions while protecting the competitive process.

Safety as a Mode of Competition

Competition for AI is not limited to the price firms charge their customers for AI subscriptions. Nonprice factors including reliability, speed, accuracy, and safety are all modes of competition. Firms can take unilateral action to set qualitative and quantitative safety standards and use them to differentiate their product.

Regulation and Liability

Another approach [discussed](#) by Neil Chilson, Head of AI Policy at the Abundance Institute, would be a congressionally mandated duty on frontier model developers to “reasonably mitigate the risk of mass casualties and catastrophic property loss.” Chilson explained that

“The statute should not mandate particular model design, testing method, or governance system. It should require covered developers to disclose dangerous capabilities and model-specific mitigations, and to give the government a reasonable way to verify such disclosures. It should also require incident reporting. This framework should replace state efforts at regulating model safety — this is a national issue and needs a national standard.”

Meanwhile, Manhattan Institute Director of Research Judge Glock has [advocated](#) for AI safety to “run[] through the old-fashioned legal system. Torts, or legal claims for damages

against a company, may provide the best mechanism to prevent the dangers of AI while preserving its advantages.” Glock explained that “The benefit of the tort system is that it forces those companies to avoid such harms in their own self-interest and mitigate them step-by-step.”



Conclusion

With four leading AI developers facing a lawsuit over an alleged agreement to slow product development, Congress should reject calls for an antitrust exemption. Existing frameworks, including SDOs, procompetitive collaborations, and market incentives provide firms with paths to innovate that do not violate federal antitrust law.