Nicholas C. Rowley (SBN 220036)
Lead Counsel
TRIAL LAWYERS FOR JUSTICE
8235 Santa Monica Blvd.
West Hollywood, CA 90046
Tel.: (310) 855-3727
Fax: (310) 855-3595
ncr@tl4j.com

Andrew T. Tutt (SBN 343206)
R. Stanton Jones (DC Bar No. 987088)*
TRIAL LAWYERS FOR JUSTICE
421 W. Water Street
Decorah, IA 52101
Tel.: (202) 490-7080
Fax: (563) 204-4205 (Tutt)
andrew@tl4j.com; stanton@tl4j.com

Jakob Z. Norman (WY Bar No. 6-3685)*
TRIAL LAWYERS FOR JUSTICE
724 S. Durbin Street
Casper, WY 82601
Tel.: (307) 288-4400 ext. 101
jakob@tl4j.com

Attorneys for Plaintiffs and the Proposed Class

* Application for *Pro Hac Vice* Admission Forthcoming

IN THE UNITED STATES DISTRICT COURT
FOR THE NORTHERN DISTRICT OF CALIFORNIA
SAN FRANCISCO DIVISION

CHARLES BUIST, CHEYENNE HUNT, CHRISTINE BULLOCK, and NICK SPETSAS, individually and on behalf of all others similarly situated,

Plaintiffs,

v.

ANTHROPIC, PBC; OPENAI OPCO, LLC; SPACEXAI LLC; AND GOOGLE LLC,

Defendants.

Case No. 3:26-cv-10693

COMPLAINT

Class Action

Demand for Jury Trial

INTRODUCTION

1. This action challenges an agreement among the four companies at the frontier of artificial intelligence—Anthropic, OpenAI, SpaceXAI, and Google—to slow the pace at which each company improves the competing products it sells to consumers.

2. The agreement was proposed in public, accepted in public, and confirmed in public. On the morning of September 12, 2026, Anthropic's Chief Executive Officer, Dario Amodei, published an essay titled We Must Pace the Frontier. It declared: "We must slow the pace at which we improve the capabilities of AI models." It called for "industry-wide coordination" among frontier laboratories to set "limits on the rate of unchecked AI progress," and it explained that coordinated slowing would let each participant reduce its pace "without sacrificing commercial advantage."

3. Amodei's rivals confirmed their agreement the same day. Within about an hour, Elon Musk, who founded and controls SpaceXAI's Grok business, quote-posted the proposal and wrote: "Dario is right." OpenAI's Chief Executive Officer, Sam Altman, wrote: "I agree with Dario that we need to pace the frontier," and committed OpenAI to the first concrete step of Amodei's plan. Google DeepMind co-founder Demis Hassabis called the essay "the right path forward" and tied it to the "industry-wide standards body for frontier AI" he had proposed two months earlier as a vehicle through which the laboratories could coordinate.

4. Two days later, Altman said what pacing meant: AI progress "should be slower than it otherwise could be." He added that OpenAI would not wait for an antitrust exemption or legislation before beginning that work with its "colleagues across the industry." The following day, OpenAI's Global Policy Chief confirmed that OpenAI, Anthropic, and Google DeepMind had already been working together for weeks, reportedly through a working group of company representatives that had met regularly since July.

5. Defendants themselves have explained why they needed an agreement. Frontier-AI companies compete above all by improving their models faster than their rivals: better reasoning, better coding, more capable agents, larger context windows, new features. A July 2026 statement signed by Amodei, OpenAI's Chief Scientist, a Google DeepMind co-founder, and dozens of their

colleagues acknowledged that each company faces "intense competitive pressure not to unilaterally slow" development. A firm that slows alone loses customers, revenue, and technological leadership to rivals that keep going. Amodei's proposal solved that problem the way cartels always have: by agreeing to slow together, so that no participant bears the competitive cost of restraint.

6. An agreement among competitors to reduce the quality of their products and the rate at which those products improve is an agreement to restrict output. It is among the restraints Section 1 of the Sherman Act, 15 U.S.C. § 1, has always condemned, and it does not become lawful because the products are new.

7. Plaintiffs take AI safety seriously. Plaintiffs, and each of them, believe that AI should have guardrails with ethical considerations and direct liability for any breach being paramount. Those guardrails and consequences should not be decided by the Defendants, but by we the people, through state and federal regulation, and through the voices of juries in our communities.

8. Artificial intelligence is among the most consequential technologies in human history. Researchers have used it to predict the three-dimensional structure of nearly every protein known to science, to identify cancers from imaging studies with accuracy matching or exceeding trained specialists, to accelerate the development of new materials for clean energy, to solve previously unsolved mathematical problems, and to compress years of scientific labor into days. It is undisputed that AI alone could add trillions of dollars annually to the global economy. These are not hypothetical gains. They are the product of the competitive race that Defendants' agreement now threatens to slow.

9. Plaintiffs recognize that powerful technology carries real risks, and they take those risks seriously. But the lesson of every previous transformative industry is that safety and responsibility are properly the work of each company acting on its own, not a cartel license. Automobile manufacturers do not allow ordinary passenger vehicles to accelerate to 300 miles per hour. They do not need a horizontal agreement among the major car manufacturers to make that choice: engineering constraints, internal safety standards, product-liability exposure, government regulation through the National Highway Traffic Safety Administration, and basic competitive self-interest each

independently supply the guardrail. The same logic applies here. Each Defendant is free, and indeed obligated, to set its own safety thresholds, retain its own evaluators, build its own product guardrails, and restrain its own development. None of that requires an agreement with rivals.

10. Defendants are also well-positioned to lead on environmental responsibility, and to do so on their own. AI's energy footprint is substantial and growing: training a single large frontier model can consume tens of millions of kilowatt-hours of electricity, and the International Energy Agency projects that, by 2030, global data-center electricity demand will more than double from 2024 levels, driven in significant part by AI workloads. Each Defendant has pledged environmental commitments. Each has the engineers, the capital, and the competitive incentive to pursue more efficient training architectures, cleaner energy procurement, and smarter compute scheduling. Innovation in AI efficiency—models that do more with less energy—is itself a domain of competitive advantage. An agreement to slow capability development does nothing to make AI more environmentally sound; only engineering competition, applied to efficiency as a prize, will.

11. In short, the path to responsible, safe, and environmentally conscious AI development is each company doing the right thing: investing in safety research, adopting rigorous internal evaluations, reducing its environmental footprint, and welcoming government oversight and regulation. That path is fully open to every Defendant. What the antitrust laws forbid is the shortcut of competitors agreeing among themselves to substitute collective restraint for individual accountability. This action is brought in defense of the competitive process that makes genuine responsibility and genuine progress possible.

12. The antitrust laws do not permit competitors to decide among themselves that competition is too dangerous. Whether frontier AI should develop more slowly is a question for each company acting alone, or for Congress and the agencies. It is not a question that four competitors may answer together. The law leaves each Defendant free to be as cautious, as responsible, and as environmentally conscious as it chooses to be, just not through agreement with the companies it competes against.

13. Plaintiffs support each Defendant's responsibility to consider its conduct carefully and do not challenge any Defendant's unilateral decisions about safety, testing, environmental impacts, or the pace of its own development. Each Defendant remains free to slow itself, to retain outside evaluators, to adopt its own product guardrails, to reduce its environmental impacts, and to ask Congress for regulation or an antitrust exemption. Plaintiffs challenge only what the antitrust laws forbid: an agreement among competitors about how fast their competing products will improve. Congress has granted no exemption for that agreement. Amodei wrote that the government would need to issue a narrow waiver for certain kinds of safety conversations; Altman announced that OpenAI would proceed without it.

14. Plaintiffs are consumers who pay a Defendant directly for subscriptions to one or more of the following services—ChatGPT, Claude, Grok, and Gemini. The value of those subscriptions consists of access to Defendants' most capable models and to the improvements Defendants continually release. Defendants market and price those subscriptions on exactly that basis. An agreement that slows improvement lowers the quality of what subscribers receive for the price they pay. That is an overcharge, and it is an injury of the kind the antitrust laws were enacted to prevent.

15. Plaintiffs bring this action on behalf of a nationwide class of direct purchasers of paid consumer subscriptions to ChatGPT, Claude, Grok, and Gemini for treble damages under Section 4 of the Clayton Act, 15 U.S.C. § 15, and for injunctive relief under Section 16, 15 U.S.C. § 26, to restore the condition the antitrust laws require: that each Defendant decide for itself, in competition with the others, how aggressively to develop and improve its products.

PARTIES

A. Plaintiffs

16. Plaintiff Charles Buist is a resident of Florida. During the Class Period, he personally purchased a paid individual consumer subscription to Claude, ChatGPT, Grok, and Gemini. He continues to subscribe and intends to continue subscribing.

17. Plaintiff Cheyenne Hunt is a resident of California. During the Class Period, she personally purchased a paid individual consumer subscription to Claude, ChatGPT, Grok, and Gemini. She continues to subscribe and intends to continue subscribing.

18. Plaintiff Christine Bullock is a resident of California. During the Class Period, she personally purchased a paid individual consumer subscription to Claude. She continues to subscribe and intends to continue subscribing.

19. Plaintiff Nick Spetsas is a resident of Florida. During the Class Period, he personally purchased a paid individual consumer subscription to Claude, ChatGPT, Grok, and Gemini. He continues to subscribe and intends to continue subscribing.

20. Plaintiffs purchased their subscriptions to obtain access to the respective Defendant's most capable models and to the features and improvements that Defendant continually releases, and each Plaintiff expected that the Defendant would decide independently, in competition with its rivals, how aggressively to develop and improve its product.

21. Plaintiffs have paid, and continue to pay, subscription prices for products whose quality and rate of improvement have been restrained by the agreement alleged herein. Because Plaintiffs continue to subscribe, they face continuing injury from the agreement and have standing to seek injunctive relief.

B. Defendants

22. Defendant Anthropic, PBC ("Anthropic") is a public benefit corporation with its principal place of business in San Francisco, California. Anthropic develops, markets, and sells Claude and related products and services throughout the United States.

23. Anthropic sells consumers paid Claude subscriptions, including Claude Pro and Claude Max, which provide greater usage and access to Anthropic's most advanced models and features.

24. Dario Amodei is Anthropic's Chief Executive Officer and a co-founder. Jared Kaplan is a co-founder and its Chief Science Officer. Each spoke and acted with respect to the matters alleged herein as an authorized representative of Anthropic, and each of their statements alleged herein is attributable to Anthropic.

25. Defendant OpenAI OpCo, LLC ("OpenAI") is a Delaware limited liability company with its principal place of business in San Francisco, California. It provides ChatGPT to individual consumers under OpenAI's Terms of Use.

26. OpenAI sells consumers paid ChatGPT subscriptions, including ChatGPT Plus (currently approximately $20 per month) and the higher-priced ChatGPT Pro, which provide greater access to OpenAI's most advanced models, reasoning, research, coding, image-generation, memory, and other features than OpenAI's free tier.

27. Sam Altman is OpenAI's Chief Executive Officer. Jakub Pachocki is its Chief Scientist. Chris Lehane is its Global Policy Chief. Each spoke and acted with respect to the matters alleged herein as an authorized representative of OpenAI, and each of their statements alleged herein is attributable to OpenAI.

28. Defendant SpaceXAI LLC ("SpaceXAI") is a Nevada limited liability company with a registered office in Austin, Texas. It provides Grok to individual consumers under its Consumer Terms of Service.

29. SpaceXAI develops, markets, and sells Grok throughout the United States, including paid SuperGrok subscriptions that increase usage limits and provide access to its frontier models and additional features.

30. Elon Musk founded the xAI business, controls it, and speaks for it publicly on the development and pacing of frontier AI. His statements alleged herein were made as SpaceXAI's controlling person and authorized representative and are attributable to SpaceXAI.

31. Defendant Google LLC ("Google") is a Delaware limited liability company with its principal place of business in Mountain View, California. Google develops Gemini, including through its Google DeepMind unit, and markets and sells to consumers throughout the United States paid Google AI subscriptions, including Google AI Pro and Google AI Ultra, that provide premium access to Gemini's most advanced models and capabilities.

32. Demis Hassabis is a co-founder and Chair of Google DeepMind and Alphabet's Chief Scientist. Shane Legg is a co-founder of Google DeepMind and its Chief AGI Scientist. Each

spoke and acted with respect to the matters alleged herein as an authorized representative of Google, and each of their statements alleged herein is attributable to Google.

C. Agents and Co-Conspirators

33. Whenever this Complaint refers to an act, statement, or agreement of a Defendant, it means that the Defendant acted through its officers, directors, employees, representatives, or agents within the course and scope of their authority.

34. Various persons and entities not named as Defendants, including affiliates of Defendants and other frontier-AI developers, participated in, facilitated, or knew of aspects of the conduct alleged herein. Plaintiffs reserve the right to amend this Complaint to add them as discovery identifies their roles.

35. Each Defendant acted as a co-conspirator with respect to the agreement alleged herein, and each is liable for the acts of the others in furtherance of it.

JURISDICTION, DIVISIONAL ASSIGNMENT, AND VENUE

A. Jurisdiction

36. This action arises under Section 1 of the Sherman Act, 15 U.S.C. § 1, and Sections 4 and 16 of the Clayton Act, 15 U.S.C. §§ 15 and 26.

37. This Court has subject-matter jurisdiction under 28 U.S.C. §§ 1331 and 1337(a) and 15 U.S.C. §§ 15 and 26.

38. This Court also has jurisdiction under the Class Action Fairness Act, 28 U.S.C. § 1332(d), because the proposed Class has more than 100 members, at least one Class member is a citizen of a state different from at least one Defendant, and the amount in controversy exceeds $5,000,000, exclusive of interest and costs.

39. Defendants' conduct occurred in and substantially affected interstate commerce. Defendants sell subscriptions to consumers in every state, receive subscription payments across state lines, deliver their services over the internet from computing and data-center infrastructure located in multiple states, and compete with one another throughout the United States.

B. Divisional Assignment

40. This action is properly assigned to the San Francisco Division under Civil L.R. 3-2(d) because a substantial part of the events or omissions giving rise to the claims occurred in San Francisco County. Anthropic and OpenAI maintain their principal places of business in San Francisco, and the agreement alleged herein was proposed, accepted, and implemented in substantial part by executives and personnel working in and from this District.

C. Personal Jurisdiction and Venue

41. This Court has personal jurisdiction over each Defendant because Anthropic, OpenAI, and Google are at home in this District; each Defendant markets and sells subscriptions to residents of this District and directs the conduct alleged herein at them; and Section 12 of the Clayton Act, 15 U.S.C. § 22, authorizes nationwide service of process on a corporate defendant in an antitrust action brought in a proper district.

42. Venue is proper in this District under 15 U.S.C. § 22 and 28 U.S.C. § 1391(b) and (c) because Anthropic, OpenAI, and Google inhabit and are found in this District; SpaceXAI transacts substantial business in this District, including the sale of subscriptions to its residents; and a substantial part of the events or omissions giving rise to the claims occurred in this District.

FACTUAL ALLEGATIONS

A. Defendants Compete Through the Pace of Innovation

43. ChatGPT, Claude, Grok, and Gemini are close substitutes for millions of consumers who use them for research, writing, coding, reasoning, mathematical and document analysis, image generation, voice interaction, planning, education, professional work, and automated or agentic tasks.

44. Defendants compete principally on capability: models that reason better, code better, answer more accurately, use tools more effectively, process more information, respond faster, perform more complex tasks, and act more autonomously. They also compete on usage limits, context windows, memory, integrations, voice, image and video generation, deep research, agents, connectors, reliability, and speed of access to newly released frontier models.

45. Defendants' paid subscription tiers monetize those differences directly. OpenAI markets paid ChatGPT subscriptions as providing greater access to advanced models, reasoning, research, coding, image generation, and memory. Anthropic markets Claude Pro and Max through greater usage, access to advanced models and features, priority availability, and early access to improvements. SpaceXAI markets SuperGrok as providing access to frontier models, greater usage, and priority access. Google markets Google AI Pro and Ultra as providing access to advanced Gemini models and capabilities unavailable or limited at its free tier.

46. The pace of improvement is therefore a central dimension of price and quality competition among Defendants. Each major model release or feature improvement causes consumers to subscribe, cancel, switch products, maintain multiple subscriptions, or change which product they use most. Defendants closely monitor one another's releases, benchmark results, features, usage limits, and pricing, and each responds when a rival ships a materially better model.

47. Before the challenged agreement, that rivalry drove each Defendant to invest aggressively in research, compute, training runs, engineering, and deployment. Defendants themselves have repeatedly described the sector as a race and acknowledged that the commercial pressure it generates makes unilateral slowing costly.

48. The stakes compound. OpenAI researchers have explained that a principal research priority is using advanced AI to accelerate AI research itself, and Anthropic describes recursive self-improvement as a principal driver of accelerating capability. An agreement about today's pace of development therefore shapes the trajectory and timing of every subsequent generation of products.

B. Defendants Recognized That No Firm Would Slow Alone

49. In July 2026, senior executives of Anthropic, OpenAI, and Google published a statement titled Pacing the Frontier. Signatories included Amodei; Kaplan; Pachocki; OpenAI Chief Research Officer Mark Chen; Legg; and numerous other employees of Anthropic, OpenAI, and Google.

50. The statement identified the collective-action problem at the heart of this case: each company was under "intense competitive pressure not to unilaterally slow" capability development. It called for tools capable of deliberately pacing frontier-wide development.

51. Altman told reporters in July that OpenAI had helped shape the statement's language and that he had discussed the "need" to pace AI development with White House officials as models became more capable.

52. The statement thus reflected a shared understanding among Defendants' leadership, months before September 12, that slowing was commercially costly for any firm that acted alone, and that the cost could be avoided only if the frontier laboratories slowed together.

C. Defendants Built the Machinery for Coordination Before September 12

53. On July 14, 2026, Hassabis proposed a U.S.-led standards body for frontier AI modeled in part on the Financial Industry Regulatory Authority. Under the proposal, frontier laboratories would initially submit models voluntarily for pre-release review; the body could later impose common requirements before models entered the U.S. market and could escalate those requirements as risks increased. Contemporaneous reporting described the proposal as a mechanism that could coordinate an industry-wide slowdown.

54. Such a body would give competing developers a standing forum in which to exchange information, adopt common requirements, review one another's frontier systems, and coordinate responses to capability thresholds.

55. Beginning in July 2026, representatives of Anthropic, OpenAI, and Google below the CEO level formed a working group that met regularly to work together on creating a standards body for the industry. On September 13, The Information reported that the working group's meetings were ongoing and that it had met during the preceding week. On September 15, OpenAI confirmed that it had been discussing AI safety with Anthropic and Google DeepMind for several weeks.

56. During the same period, Altman told an OpenAI companywide town hall that the major laboratories might need to create a standards body themselves if government did not, and that the

best course could be to act together with the other laboratories. According to published reports, Altman also told OpenAI employees that OpenAI might slow the pace of frontier development, potentially in conjunction with other laboratories.

57. On September 6, 2026, OpenAI published on its website an essay by Pachocki titled An Alien Mind. Pachocki wrote that one of the principal options available to frontier developers was "coordinating to slow down future development"; that no laboratory had solved alignment and monitoring well enough to continue scaling at maximum speed indefinitely; and that he expected voluntary slowdowns to become common until shared safety thresholds were established.

58. On September 10, 2026, WIRED reported that OpenAI had sought guidance from members of Congress on whether coordinating an industry-wide slowdown in frontier-AI development could violate the antitrust laws. OpenAI did not substantively deny the report. The inquiry reflected OpenAI's understanding that the contemplated conduct was coordination among competitors, not unilateral safety decision-making.

59. On September 11, 2026, in an interview with Fortune, Altman was asked why leading AI executives, including Amodei and Musk, did not get together and develop a common plan. He answered: "I think that will happen." He declined to "pre-announce private discussions" that he believed should later be shared "as a group."

60. The day before Amodei's essay, Altman thus referred to private discussions about industry collaboration, declined to disclose their details, and said he thought a common plan would happen.

D. Amodei Proposes the Agreement: "We Must Pace the Frontier"

61. On the morning of September 12, 2026, Amodei published We Must Pace the Frontier and promoted it publicly.

62. The essay's thesis was direct: "We must slow the pace at which we improve the capabilities of AI models."

63. Amodei proposed three steps. The first—embedding third-party evaluators with employee-like access to models and training pipelines—Anthropic adopted unilaterally.

64. The second step, which Amodei called "Democratic Coordination," he expressly distinguished from the first: it "requires industry-wide coordination." Under it, the frontier laboratories would jointly establish common safety standards and "limits on the rate of unchecked AI progress."

65. Amodei identified the inputs through which those limits would operate. He proposed consideration of limits on training compute, on the nature of training runs, and on the internal use of AI to improve AI—the principal means by which each Defendant improves its products. He proposed capability "checkpoints" at which continued advancement would depend on agreed safety certifications or conditions.

66. Amodei also proposed the means of enforcement. Embedded evaluators, he explained, would make "verifiable pacing" practical by giving outsiders access to detailed properties of models and training pipelines—that is, the ability to confirm that each participant was in fact slowing as agreed.

67. Amodei was explicit about the commercial problem the plan solved. A company that slows alone while rivals race ahead loses customers, revenue, talent, and technological leadership; coordinated pacing would let frontier developers slow "without sacrificing commercial advantage." That is the economic function of every output-restricting cartel: to remove the competitive penalty for restraint by ensuring that rivals restrain themselves too.

68. Amodei acknowledged that legislation might take too long and urged AI companies, in parallel with any government action, to "voluntarily work together" on the proposed arrangements.

69. Amodei also acknowledged the antitrust problem. His essay stated that government mediation would be helpful and that the government would need to issue a narrow waiver for certain kinds of safety conversations. No such waiver exists. Congress has enacted no exemption, and no agency has compelled the conduct alleged herein.

70. The proposal thus supplied every element of a horizontal agreement in restraint of trade: the parties (the frontier laboratories); the term (a slower rate of capability advancement than each would choose alone); the mechanisms (shared limits on compute, training runs, and self-

improvement, and capability checkpoints); and the means of policing defection (verifiable pacing through embedded evaluators).

E. Defendants Immediately Accept

71. Within approximately an hour of publication, Musk quote-posted Amodei's announcement of the proposal and responded: "Dario is right."

72. Later that day, Altman wrote: "I agree with Dario that we need to pace the frontier." He added that pacing had been a primary subject of discussion inside OpenAI in recent weeks; committed OpenAI to Amodei's first step, stating that OpenAI "will do the same" in giving independent evaluators employee-like access; and said OpenAI would have more to share soon.

73. Also on September 12, Hassabis stated publicly that Amodei's essay pointed toward "the right path forward" and that "the direction is correct," while noting that implementation details required further work. He tied his endorsement to his own proposed "industry-wide standards body for frontier AI"—the mechanism the Anthropic-OpenAI-Google working group had been developing since July.

74. By the end of September 12, therefore, a senior executive or controlling person of each Defendant had publicly signaled assent to a proposal that the frontier laboratories coordinate to slow the pace of capability development. Each signal was given with knowledge of the others, in sequence, on a single day, in direct response to the proposal.

75. A public offer followed by public acceptances forms an agreement no less than more traditional cartel behavior of private, secretive communications. Conducting part of the exchange in view of the market assured each participant that its rivals were committed.

F. Defendants Confirm the Agreement and Its Object

76. On September 14, Altman published further statements explaining what OpenAI had agreed to. OpenAI favored consistent frontier-AI safety rules and common approaches to safety and monitoring across the industry, and looked forward to "collaborating with our colleagues across the industry" to develop them.

77. Altman then stated the object of the agreement in plain terms: AI progress "should be slower than it otherwise could be."

78. Altman added that OpenAI did not intend to wait for an antitrust exemption or legislation before beginning this work, and that the firms should first do what they could themselves.

79. On September 15, OpenAI Global Policy Chief Chris Lehane confirmed that OpenAI had been working with Anthropic and Google DeepMind on these issues for several weeks—before Amodei's essay—and that OpenAI believed the firms could work together without waiting for an antitrust waiver.

80. Lehane's statement corroborated the reporting that OpenAI, Anthropic, and Google had been communicating privately before the public exchange of September 12, and established that the September 12 statements were the public face of an existing course of dealing among competitors.

81. The coordination continues. Amodei continues to advocate interfirm discussions on the remaining stages of his framework; Altman has said that OpenAI anticipates further collaboration and additional announcements; and the working group and the proposed standards body remain the vehicles through which Defendants intend to implement pacing.

G. Each Defendant Knowingly Joined the Agreement

82. Anthropic proposed the agreement through its Chief Executive Officer, defined its terms, supplied its enforcement mechanism, participated in the working group from July, and has continued to advocate implementation.

83. OpenAI participated in the working group from July; helped draft the July statement; explored through its CEO—with employees and with the White House—whether to coordinate a slowdown with other laboratories, and sought guidance from members of Congress about such coordination; published its Chief Scientist's endorsement of "coordinating to slow down future development" six days before the proposal; accepted the proposal through its CEO on September 12; stated its object on September 14; and confirmed weeks of prior collaboration through its Global Policy Chief on September 15.

84. Google participated in the working group from July; proposed, through Hassabis, the standards body that Defendants have identified as the forum for coordination; had a Google DeepMind co-founder sign the July statement; and, on September 12, endorsed the direction of Amodei's proposal while directing its implementation toward the body Google had proposed.

85. SpaceXAI, through the person who founded and controls its AI business, publicly endorsed the proposal within an hour of its publication, communicating to its three principal rivals that SpaceXAI would not exploit their restraint by racing ahead. Altman's September 11 acknowledgment of private discussions came in direct response to a question that named Musk as a participant in the contemplated common plan.

H. The Agreement Is Not Mere Parallel Conduct

86. The agreement is evidenced by direct communications among Defendants: a written proposal for "industry-wide coordination" to limit the rate of AI progress and innovation; express statements of agreement by the proposer's competitors; and admissions that those competitors had been working together privately for weeks.

87. The circumstances also include every traditional indicator that Defendants' conduct was the product of agreement rather than independent decisionmaking:

(a) Motive. Each Defendant had a strong motive to conspire, because, as Defendants themselves stated, slowing the pace of innovation was seen to be commercially untenable for any firm acting alone but advantageous if all slowed together.

(b) Action against unilateral self-interest. According to Defendants' own statements, absent an agreement, no Defendant would announce that it will improve its products more slowly than it could, because the announcement would invite rivals to take its customers. Each Defendant's announcement is explicable only by its understanding that its rivals had made, or would make, the same commitment.

(c) Interfirm communications and opportunity. Representatives of Anthropic, OpenAI, and Google met regularly from July through September; Altman acknowledged private discussions on September 11; and Lehane confirmed weeks of collaboration on September 15.

(d) Market structure. Frontier AI is highly concentrated among a handful of firms protected by enormous entry barriers, and Defendants have repeatedly described themselves as the "frontier labs" whose joint participation is necessary for pacing to work.

(e) Consciousness of antitrust risk. OpenAI asked members of Congress whether coordinating a slowdown would violate the antitrust laws; Amodei wrote that the government would need to issue a narrow waiver for certain kinds of safety conversations; Altman announced that OpenAI would proceed without one.

(f) Pretext. Every safety objective Defendants have identified can be pursued unilaterally, through independent evaluators, or through government regulation. Only the elimination of competitive pressure requires an agreement among rivals.

I. Plaintiffs Do Not Challenge Defendants' Petitioning of Government

88. Plaintiffs do not seek to impose liability for any Defendant's advocacy to Congress, the White House, or any agency, including advocacy for AI regulation, for a standards body created by statute, or for an antitrust exemption.

89. The restraint challenged here is Defendants' private agreement among themselves, formed and implemented outside any governmental process, to slow the pace of competitive development. Defendants' own statements establish that they are proceeding with that agreement without waiting for, and independent of, any governmental action.

90. Defendants' petitioning statements are alleged only as evidence of Defendants' knowledge, intent, and understanding of the competitive problem the agreement addresses, and of the antitrust risk they knowingly accepted.

J. How the Agreement Restrains Competition

91. In a competitive market, each Defendant independently decides how much to invest, how quickly to train new models, when to begin training runs, how much compute to deploy, when to release improvements, and how aggressively to pursue new capabilities.

92. Under the agreement, each Defendant makes those decisions knowing that its principal rivals have agreed to reduce their own pace. That knowledge eliminates the competitive penalty

for slowing—the risk that a rival will race ahead—and allows each Defendant to reduce its investment and competitive effort without losing ground.

93. A cartel can restrain competition by reducing output, quality, investment, or competitive effort; it need not eliminate them. Here the agreement operates through a common commitment to reduce the rate at which competing products improve, shared checkpoints and standards that delay advancement, reciprocal expectations about development conduct, and monitoring that deters defection.

94. Plaintiffs challenge Defendants' agreement to coordinate those decisions. Each Defendant remains free to make them unilaterally.

MARKET POWER

95. Defendants' agreement is a naked horizontal restraint on a central dimension of competition among direct rivals, and it is unlawful without proof of a relevant market or market power.

96. To the extent market definition is required, the relevant product market is paid consumer subscriptions to general-purpose frontier generative-AI assistants (the "Paid Frontier AI Assistant Subscription Market"). Products in this market provide consumers ongoing access, through a conversational or assistant interface, to highly capable, general-purpose generative-AI systems that perform a wide range of intellectual and creative tasks. The principal products are the paid tiers of ChatGPT, Claude, Grok, and Gemini.

97. These products are reasonably interchangeable. Consumers use them for the same tasks, compare them on the same attributes—model capability, reasoning, coding, accuracy, research functionality, speed, multimodality, agents, context, integrations, usage limits, and price—and switch among them in response to differences in capability and price. Defendants market them as alternatives to one another.

98. Conventional search engines, ordinary productivity software, narrow machine-learning tools, and human professional services are not reasonable substitutes. None provides general-purpose frontier capability through an assistant interface, and consumers would not substitute to them

in response to a small but significant, non-transitory decrease in the quality or rate of improvement, or increase in the price, of paid frontier-AI subscriptions.

99. Free tiers of Defendants' own products do not constrain the paid market. They impose substantially lower usage limits, restrict or exclude access to frontier models, and offer fewer features and lower priority. And because Defendants control their free tiers as well, the agreement restrains the improvement of free and paid tiers alike; a free tier offers no competitive discipline that its owner has agreed to withhold.

100. In the alternative, the agreement restrains competition in an innovation market for the development of general-purpose frontier generative-AI models, in which Defendants are the principal participants and on which the restraint operates directly.

101. The relevant geographic market is the United States. Defendants sell subscriptions nationwide on uniform terms, and consumers cannot practically obtain frontier-AI subscriptions from sellers outside the country on materially different terms.

102. The Paid Frontier AI Assistant Subscription Market is highly concentrated. On information and belief, Defendants collectively account for at least 80% of paid consumer subscriptions to general-purpose frontier generative-AI assistants in the United States.

103. Entry at the frontier requires expenditures on computing infrastructure, specialized semiconductors, electricity, data-center capacity, engineering, research personnel, model training, safety systems, and deployment infrastructure; access to scarce technical talent and large-scale compute; consumer distribution and product reliability; and the ability to fund successive generations of models. These barriers limit the number of firms capable of imposing competitive discipline at the frontier to a handful.

104. Defendants' own statements confirm that they regard one another as the relevant competitive set. Their public discussions refer to a handful of "frontier labs," to coordination among those laboratories, and to the inability of any one of them to slow without competitive consequence. In a market this concentrated, an agreement among Defendants determines market-wide competitive incentives.

ANTICOMPETITIVE EFFECTS

105. The agreement has the purpose, tendency, and effect of restraining competition in the development and improvement of frontier generative-AI products.

106. The agreement reduces the independent competitive pressure that otherwise drives each Defendant to improve its products as rapidly as it can, and it reduces each Defendant's incentive to devote research effort and compute to capability advancement, to begin or expand training runs, to accelerate model development, to deploy innovations promptly, and to respond aggressively to rivals' advances.

107. The agreement causes and threatens delay in the development and delivery of improved capabilities to consumers; reduced product quality relative to what unrestrained competition would produce; reduced investment and competitive effort; and fewer or later improvements in reasoning, coding, research, multimodal capability, reliability, agents, speed, context, memory, and other features.

108. The agreement diminishes the value consumers receive for the subscription prices they pay. Consumers pay the same or higher nominal prices for products that improve more slowly than they would under unrestrained competition; the quality-adjusted price of Defendants' subscriptions is therefore higher than it would be absent the agreement.

109. The agreement reduces consumer choice by foreclosing the possibility that one competitor will break from the pack and pursue technological improvement more aggressively than its rivals, and it harms the competitive process itself by substituting the collective judgment of competitors for independent rivalry.

110. Because the agreement was formed recently and Defendants' development cycles run months, its full effect on released products has not yet manifested. But the agreement has already altered each Defendant's incentives and, on information and belief, its investment, training, and release decisions, and its continuing operation threatens exactly the injury that Section 16 of the Clayton Act exists to prevent.



111. The agreement has no cognizable procompetitive justification. It is not reasonably necessary to permit any Defendant to conduct safety research, test its own models, adopt its own safety standards, retain independent evaluators, disclose safety incidents, or slow its own development. Each Defendant can do all of those things unilaterally, without limitation.

112. To the extent Defendants contend that competition in frontier AI is itself unsafe and should be curtailed, that contention is not a procompetitive justification cognizable under the Sherman Act. The Act does not permit competitors to decide among themselves that competition is undesirable as a policy matter; that judgment belongs to Congress.

113. Nor is the agreement ancillary to any legitimate collaboration among Defendants. They do not integrate productive assets, share risk, or jointly produce anything. They remain separate firms that independently develop, market, price, and sell competing products, and the sole function of their agreement is to reduce the competition among them.

114. No statute exempts Defendants or their agreement to "pace the frontier" from the antitrust laws, and no governmental body has compelled it.

ANTITRUST INJURY

115. Plaintiffs and Class members purchase directly from Defendants the products whose improvement the agreement restrains. They pay Defendants money for access to products whose value depends on technological capability and continued improvement.

116. Absent the agreement, each Defendant would face stronger independent incentives to invest, train, improve, and release capabilities rapidly, because each would face the risk that a rival would do so first, and vigorous competition would force Defendants to provide more value for the subscription prices they charge.

117. As a direct and foreseeable result of the agreement, Plaintiffs and Class members have paid and continue to pay subscription prices that exceed the competitive value of the products they receive. Stated another way, they receive less capability, slower improvement, or both, for the prices they pay. That difference is an overcharge recoverable under Section 4 of the Clayton Act.

118. That injury is measurable on a class-wide basis through common evidence, including Defendants' development and release plans; actual and counterfactual investment, compute, and training decisions; product capabilities and benchmark performance; subscription prices; consumer demand, switching, and usage data; and accepted economic methods for valuing model and feature improvements. The amount of damages will be established through discovery and expert proof.

119. Plaintiffs' injuries are injuries of the type the antitrust laws were designed to prevent, and they flow directly from what makes Defendants' conduct unlawful: horizontal competitors agreed to diminish the competitive pressure to improve the products Plaintiffs buy.

120. Plaintiffs and Class members are direct purchasers and the most efficient enforcers of the antitrust laws with respect to this agreement. No more direct group of purchasers exists.

121. Because Plaintiffs and Class members continue to purchase Defendants' subscriptions, and because the agreement is ongoing, they face threatened loss and damage from its continuation and are entitled to injunctive relief under Section 16 of the Clayton Act.

122. Defendants are jointly and severally liable for all damages caused by the agreement.

CLASS ACTION ALLEGATIONS

123. Plaintiffs bring this action individually and on behalf of the following class under Federal Rules of Civil Procedure 23(a), 23(b)(2), and 23(b)(3):

Nationwide Class: All persons in the United States who, from September 12, 2026 through the date on which the effects of Defendants' unlawful conduct cease (the "Class Period"), purchased directly from one or more Defendants a paid individual consumer subscription providing premium access to ChatGPT, Claude, Grok, or Gemini.

124. Plaintiffs also propose, as appropriate, the following Defendant-specific Subclasses:

OpenAI Subscriber Subclass: All persons in the United States who, during the Class Period, purchased directly from OpenAI a paid individual consumer subscription to ChatGPT.

Anthropic Subscriber Subclass: All persons in the United States who, during the Class Period, purchased directly from Anthropic a paid individual consumer subscription to Claude.

SpaceXAI Subscriber Subclass: All persons in the United States who, during the Class Period, purchased directly from SpaceXAI a paid individual consumer subscription to Grok.

Google Subscriber Subclass: All persons in the United States who, during the Class Period, purchased directly from Google a paid individual consumer subscription providing premium access to Gemini.

125. Plaintiffs reserve the right to modify the Class and Subclass definitions based on facts developed in discovery, including evidence that the agreement was formed before September 12, 2026.

126. Excluded from the Class are Defendants; their parents, subsidiaries, and affiliates; their officers and directors; governmental entities; the Court and its staff assigned to this action; and any person or entity that validly and timely requests exclusion from a Rule 23(b)(3) Class.

Numerosity

127. The members of the Class and each Subclass are so numerous that joinder is impracticable. Each Defendant sells paid subscriptions to large numbers of consumers throughout the United States, and Class members can be identified from Defendants' subscription and billing records using objective criteria.

Commonality

128. This action presents questions of law and fact common to the Class, including:

(a) whether Defendants entered a contract, combination, or conspiracy concerning the pace of frontier-AI development;

(b) the scope, participants, terms, and duration of that agreement;

(c) whether Defendants communicated privately before or after their September 12, 2026 public statements, and what they agreed to;

(d) whether the agreement is a naked horizontal restraint that is unlawful per se, unlawful under abbreviated analysis, or unlawful under the rule of reason;

(e) whether a relevant market must be defined and, if so, its boundaries, and whether Defendants possess market power individually or collectively;

(f) whether the agreement reduced competition in innovation, quality, output, investment, or product development;

(g) whether the agreement caused paid AI subscriptions to provide less value than they would have provided absent the restraint;

(h) whether antitrust impact and damages can be measured using common economic evidence;

(i) whether Defendants' conduct is ongoing or threatens continuing injury; and

(j) the nature and scope of appropriate injunctive relief.

129. These questions will be resolved through evidence common to the Class, including Defendants' communications, meeting records, internal development plans, research budgets, compute allocations, model-training plans, release schedules, competitive analyses, subscription and pricing records, and common economic evidence.

Typicality

130. Plaintiffs' claims are typical of the claims of the Class and any applicable Subclass. Plaintiffs purchased a subscription product affected by the same agreement, suffered the same type of antitrust injury, and will confront the same defenses concerning the existence, legality, and effect of the agreement.

Adequacy

131. Plaintiffs will fairly and adequately protect the interests of the Class. Plaintiffs' interests do not conflict with those of absent Class members, and Plaintiffs have retained counsel experienced in complex antitrust class actions who are capable of prosecuting this action vigorously.

Rule 23(b)(3)

132. Common issues predominate over individual issues. The existence and scope of the agreement, Defendants' common course of conduct, the nature of the restraint, and its competitive effects will be established principally through evidence common to all Class members, and antitrust impact and damages are susceptible to proof through methodologies applicable to the Class or Subclasses as a whole.

133. A class action is superior to individual litigation. Individual consumers' damages are too small to justify separate antitrust suits requiring extensive discovery of some of the world's largest and most sophisticated technology companies, and concentrating the litigation in one proceeding promotes efficiency, avoids inconsistent adjudications, and serves the private-enforcement purposes of the antitrust laws.

Rule 23(b)(2)

134. Defendants have acted and continue to act on grounds generally applicable to the Class, and their ongoing coordination threatens continuing injury to every consumer who purchases their subscriptions. Final injunctive and declaratory relief concerning the agreement is therefore appropriate respecting the Class as a whole.

CLAIMS FOR RELIEF

FIRST CLAIM FOR RELIEF

Violation of Section 1 of the Sherman Act, 15 U.S.C. § 1

(Against All Defendants)

135. Plaintiffs incorporate by reference the preceding paragraphs.

136. Defendants are horizontal competitors in the development and sale of general-purpose frontier generative-AI products and of paid consumer subscriptions to those products, and each is a separate economic actor capable of conspiring with the others.

137. Beginning no later than September 12, 2026, and on information and belief earlier, Defendants entered into a contract, combination, or conspiracy to restrain the pace at which they develop and improve their competing frontier-AI products.

138. Defendants formed and manifested the agreement through, among other things, Amodei's written proposal for industry-wide coordination to limit the rate of AI progress; Musk's, Altman's, and Hassabis's public statements of assent on September 12, 2026; the private discussions and recurring working-group meetings among Anthropic, OpenAI, and Google from July 2026 onward; Altman's September 14 statement that AI progress "should be slower than it otherwise could be"; and Lehane's September 15 confirmation of weeks of collaboration.

139. The agreement restrains competition in model development, product quality, innovation, investment, output, and the timing of new capabilities, and its purpose and effect is to replace each Defendant's independent decisions about the pace of innovation with collective restraint.

140. The agreement is a naked horizontal restraint on output and product quality among competitors. It is not ancillary to any integration of productive assets or legitimate joint venture; Defendants remain separate firms that independently develop, market, price, and sell competing products. The agreement is therefore unlawful per se.

141. Alternatively, the agreement is unlawful under abbreviated, or "quick-look," analysis, because an observer with even a rudimentary understanding of economics would conclude that an agreement among the principal competitors in a market that technological progress in their competing products should be slower than competition would otherwise produce has an anticompetitive effect on consumers.

142. Alternatively, the agreement is unlawful under the rule of reason. Defendants possess substantial collective power in the relevant market alleged above; the agreement has caused and threatens substantial anticompetitive effects, including slower innovation, reduced quality, diminished investment and competitive effort, delayed improvements, and higher quality-adjusted prices; the agreement has no cognizable procompetitive justification; and any legitimate objective Defendants may identify can be achieved through substantially less restrictive means, including unilateral safety decisions, independent third-party evaluation, disclosure, government regulation, safety research, and technical standards that do not limit independent competitive development.

143. Defendants' conduct occurred in and substantially affected interstate commerce.

144. As a direct and proximate result of Defendants' violation, Plaintiffs and Class members have been injured in their business or property, in that they have paid subscription prices exceeding the competitive value of the products they received and have received lower-quality or more slowly improving products than competition would have produced.

145. Plaintiffs and Class members are entitled to recover treble damages, the costs of suit, and reasonable attorneys' fees under Section 4 of the Clayton Act, 15 U.S.C. § 15.

SECOND CLAIM FOR RELIEF

Injunctive Relief Under Section 16 of the Clayton Act, 15 U.S.C. § 26

(Against All Defendants)

146. Plaintiffs incorporate by reference the preceding paragraphs.

147. Defendants' agreement violates Section 1 of the Sherman Act for the reasons alleged above.

148. The agreement is ongoing. Defendants continue to communicate and collaborate concerning frontier-AI standards, safety thresholds, evaluators, and pacing, and they have announced that further collaboration and announcements will follow.

149. Plaintiffs and Class members who continue to purchase Defendants' subscriptions face threatened loss and damage from the continued suppression of independent competition in innovation, and monetary damages alone will not prevent that injury.

150. Plaintiffs therefore seek injunctive relief prohibiting each Defendant from entering into, maintaining, enforcing, or monitoring any agreement with a horizontal competitor concerning:

(a) the rate at which competing AI products are developed, improved, trained, or released;

(b) limits on training compute, or on the nature, frequency, or timing of training runs, adopted pursuant to a horizontal pacing arrangement;

(c) limits on the use of AI to develop improved AI systems, where imposed by agreement among competitors;

(d) coordinated delays in the development or release of models or features;

(e) capability checkpoints that operate as horizontally agreed restraints on competitive development, as distinct from lawful governmental requirements or independent safety measures; and

(f) the exchange of competitively sensitive information for the purpose of policing or enforcing any such restraint.

151. Plaintiffs do not seek to prohibit any Defendant from independently adopting safety measures or product guardrails, independently considering and reducing environmental impacts, independently slowing or suspending its own development, independently retaining evaluators, conducting lawful safety research, complying with governmental requirements, petitioning any governmental body, or participating in legitimate standard-setting activities that do not restrain independent competition over the pace of development.

152. Plaintiffs and Class members are entitled to injunctive relief under 15 U.S.C. § 26.

PRAYER FOR RELIEF

WHEREFORE, Plaintiffs, individually and on behalf of the proposed Class and Sub-classes, respectfully request that the Court:

A. certify this action as a class action under Federal Rule of Civil Procedure 23, appoint Plaintiffs as representatives of the Class and any applicable Subclass, and appoint Plaintiffs' counsel as Class Counsel;

B. adjudge and decree that Defendants entered into and participated in an unlawful contract, combination, or conspiracy in violation of Section 1 of the Sherman Act, 15 U.S.C. § 1;

C. award Plaintiffs and the Class damages in an amount to be established at trial, trebled as provided by Section 4 of the Clayton Act, 15 U.S.C. § 15;

D. enter preliminary and permanent injunctive relief under Section 16 of the Clayton Act, 15 U.S.C. § 26, as set forth above, preventing Defendants from continuing, renewing, enforcing, or implementing any horizontal agreement to restrain competitive innovation;

E. award Plaintiffs and the Class pre-judgment and post-judgment interest to the fullest extent permitted by law;

F. award Plaintiffs reasonable attorneys' fees and costs of suit as provided by law; and

G. grant such other and further relief as the Court deems just and proper.

DEMAND FOR JURY TRIAL

Pursuant to Federal Rule of Civil Procedure 38(b), Plaintiffs demand a trial by jury on all claims and issues so triable.

Dated: September 18, 2026

Respectfully submitted,

/s/ Andrew T. Tutt

Nicholas C. Rowley (SBN 220036)
Lead Counsel
TRIAL LAWYERS FOR JUSTICE
8235 Santa Monica Blvd.
West Hollywood, CA 90046
Tel.: (310) 855-3727
Fax: (310) 855-3595
ncr@tl4j.com

Andrew T. Tutt (SBN 343206)
R. Stanton Jones (DC Bar No. 987088)*
TRIAL LAWYERS FOR JUSTICE
421 W. Water Street
Decorah, IA 52101
Tel.: (202) 490-7080
Fax: (563) 204-4205 (Tutt)
andrew@tl4j.com
stanton@tl4j.com

Jakob Z. Norman (WY Bar No. 6-3685)*
TRIAL LAWYERS FOR JUSTICE
724 S. Durbin Street
Casper, WY 82601
Tel.: (307) 288-4400 ext. 101
jakob@tl4j.com

Attorneys for Plaintiffs and the Proposed Class

* Application for *Pro Hac Vice* Admission Forthcoming