

OpenAI, Musk, and the AI Duopoly

A conversation on the trial, the IPO, and the race for AI supremacy.

May 2026, by Sergio Rodriguez Berdier

This article reflects personal analysis and interpretation based on publicly available information as of May 2026. It is not legal, financial, or investment advice.

Executive Summary

This conversation traces one of the most consequential corporate and legal dramas in the history of technology — and explains why the outcome matters to anyone with money in the market, a job in tech, or simply a phone in their pocket.

The origin story. OpenAI was born in 2015 as a non-profit with a noble mission: make sure artificial general intelligence benefits everyone, not just whoever builds it first. Elon Musk was a founding donor and board member. He left in 2018. In 2019, the organization took Microsoft's money and transformed into a commercial entity. That transformation is the legal and moral fault line at the center of everything that follows.

The split that created a rival. In 2021, OpenAI's head of research and head of safety walked out together and founded Anthropic — now OpenAI's most direct competitor. Their argument: OpenAI was shipping products faster than it was building safeguards, and Microsoft's money had changed the culture beyond repair. The conversation explores what Anthropic got right, what it got wrong, and what genuine risks it still carries heading into its own IPO.

The lawsuit nobody won — and everybody noticed. Musk sued OpenAI for breach of charitable trust, arguing that ~\$38 million he donated to a public-benefit mission was handed over to a commercial entity. The case was dismissed — but dismissal is not the end. Musk can, and almost certainly will, appeal. The Ninth Circuit process takes 12–18 months. That single fact reshapes the entire competitive landscape.

Why the IPO delay is the real story. OpenAI is burning cash at a historic rate. It needs a public offering to fund the next generation of compute and talent. An active appeal gives Wall Street underwriters a reason to walk away — no bank will ring the opening bell while a court could theoretically unwind the company's governance structure overnight. The conversation explains precisely why this pushes a realistic IPO date to 2028 at the earliest, what that costs OpenAI operationally, and how Anthropic is positioned to go public first and absorb the market's appetite for AI exposure.

The bigger picture. Imagine the Industrial Revolution, the transistor, the internet, and the smartphone all happening in a single year rather than across two centuries. That is the scale of the AI transition underway. This conversation is not just about two companies and one lawsuit — it is about who controls the infrastructure of that transition, who gets rich from it, and whether the people building it bear any responsibility for where it leads. Read it as a briefing on the forces shaping the next decade.

The History: How OpenAI Began

REPORTER:

The Musk vs. OpenAI trial is finished. What happened?

SERGIO:

It's a bit more complicated than a simple win or loss — but let's start from the beginning.

Origins

OpenAI was founded in December 2015 in San Francisco as a non-profit AI research laboratory, driven by a desire to ensure artificial general intelligence (AGI) benefits humanity and to counter the closed-source AI research of large tech companies, specifically Google. Its founders feared that rapid, unchecked AI development could lead to catastrophic outcomes.

The founding group included Sam Altman, Elon Musk, Greg Brockman, Ilya Sutskever, Wojciech Zaremba, and others. Musk was particularly concerned about Google's acquisition of DeepMind and the potential for a single company to control superintelligence. The organization launched with a \$1 billion pledge — though only a portion was initially collected — and a stated mission to openly share research.

In 2018, Musk left the board, citing a potential conflict of interest with Tesla's own AI development, and amid reported disagreements about the company's direction.

In 2019, to fund the massive computing costs required for frontier AI research, OpenAI transitioned to a "capped-profit" structure, creating a for-profit subsidiary that could accept outside investment. Microsoft invested \$1 billion in the new entity, anchoring OpenAI's infrastructure on Azure.

By 2022, the lab had shifted from purely academic research toward powerful commercial models — GPT-3, DALL-E, and ultimately ChatGPT, released in late 2022.

Corporate evolution

The original non-profit, OpenAI, Inc., still exists but governs a for-profit subsidiary. As of April 2026, that subsidiary has been reorganized into a Public Benefit Corporation (PBC) called OpenAI Group PBC, with the OpenAI Foundation retaining a 26% stake and board control.

Key departures have reshaped the landscape. In 2021, Dario Amodei and Daniela Amodei left to found Anthropic. Since 2024, a broader talent redistribution has sent senior researchers to competitors across the industry.

The Anthropic Split

REPORTER:

Losing talent is always costly — but losing the people who then built your main competitor is extraordinary.

SERGIO:

That's the nature of the tech market. The Anthropic split in 2021 was a genuinely seismic event.

The departure was led by Dario Amodei, then VP of Research, and Daniela Amodei, then VP of Safety and Policy, along with a group of senior researchers. Three fault lines drove the split:

- A disagreement over the "scaling hypothesis" — the idea that simply adding more compute and data produces more intelligent models — and how quickly to commercialize it. The departing team felt OpenAI was shipping products like GPT-3 faster than its safety infrastructure could keep up.
- The influx of Microsoft capital created pressure to deliver commercial products, generating a cultural rift between those who prioritized research and those who prioritized revenue.
- Rather than fighting that cultural battle from within, the team wanted to start fresh under a Public Benefit Corporation structure that legally mandates balancing social good with profit.

REPORTER:

Were they successful in achieving those goals?

SERGIO:

Yes — with important caveats in both directions.

As of early 2026, Anthropic has positioned itself as the "responsible alternative" to OpenAI, with enterprise clients accounting for roughly 80% of its revenue and a more methodical release schedule. Its Constitutional AI approach to model alignment has earned credibility with researchers.

That said, Anthropic carries its own structural risks that deserve honest scrutiny. The company has not yet turned a profit. Its two largest backers — Amazon and Google — are also cloud platform competitors, creating a potential conflict of interest over pricing and distribution. And its safety-first branding has yet to be tested under the public market scrutiny that an IPO would bring.

How They Compare: OpenAI vs. Anthropic (May 2026)

REPORTER:

Are they very different companies at this point?

SERGIO:

They started from the same roots but have diverged meaningfully. This table captures where each stands today.

Feature	OpenAI	Anthropic
Core Philosophy	Aggressive Scaling — "move fast" to reach AGI first.	Safety-First — "Constitutional AI" and safety-led development.
Flagship Model	GPT-4o / GPT-5 (and Sora): multimodal, consumer- and developer-focused.	Claude Sonnet 4.6 / Opus 4.6: known for long-form reasoning and strict safety guardrails.
Corporate Structure	Public Benefit Corp (PBC) since April 2026; governed by the OpenAI Foundation (non-profit retains control).	Public Benefit Corp (PBC) since founding; legally mandates balancing mission and profit.
Lead Investors	Microsoft: multi-billion-dollar partnership; Azure infrastructure.	Amazon & Google: heavy backing; AWS and Google Cloud infrastructure.
Key Leaders	Sam Altman (CEO), Jakub Pachocki (Chief Scientist).	Dario Amodei (CEO), Daniela Amodei (President).
Target Audience	Mass market: large consumer base (ChatGPT) plus developer platform.	Enterprise: focuses on high-stakes B2B integration and reliability.
Culture	Product-driven, competitive, and secretive.	Academic, safety-focused, "safety-first" branding.
Recent Challenges	2023 board crisis; 2026 restructuring; Musk lawsuit delaying IPO.	Not yet profitable; dependent on AWS/Google Cloud from investors who are also competitors.

In summary: while OpenAI launched the consumer AI revolution with ChatGPT, Anthropic has successfully branded itself as the enterprise-grade alternative, attracting companies and researchers who find OpenAI's governance and pace too volatile. As of 2026, the two occupy a de facto duopoly at the frontier — OpenAI leading in multimodal features (voice, video), Anthropic often preferred for long-context reasoning and enterprise coding workflows.

The Musk v. OpenAI Trial

REPORTER:

Let's go back to the original question — the trial.

SERGIO:

History is essential context here.

Core arguments

The trial concluded its three-week proceedings in Oakland, California, with closing arguments delivered on May 14, 2026.

Musk's legal team argued that OpenAI committed a "bait-and-switch": his donations — reported at approximately \$38 million, not the dramatically higher figures sometimes cited — were made to a non-profit with a public benefit mission and were used to build technology that was then transferred to a commercial entity. His team characterized this as a breach of charitable trust.

OpenAI's defense countered that Musk departed voluntarily after failing to take control of the organization, and that his lawsuit is motivated by competitive jealousy now that he runs his own AI company, xAI. They characterized the claims as "sour grapes."

It is worth noting that Musk's core legal grievance — that assets donated to a non-profit mission were converted to private commercial use — is a legitimate theory of charitable trust law. Courts do not typically hear frivolous cases of this kind through three weeks of testimony; the legal question is real, even if his motivations are mixed.

The outcome

Musk sought injunctive relief and damages and demanded the removal of CEO Sam Altman and President Greg Brockman. The legal structure of the case was unusual: because the suit centered on a breach of charitable trust, the jury's verdict was advisory only. U.S. District Judge Yvonne Gonzalez Rogers retained ultimate authority to accept or override the jury's findings.

Judge Rogers and the jury determined that Musk was aware of OpenAI's commercial pivot before August 2021 and dismissed the case on the grounds that it lacked standing. The case was thrown out.

The IPO Question

REPORTER:

Can OpenAI proceed with its IPO while a Musk appeal is in process?

SERGIO:

Yes — legally. But the practical obstacles are severe.

Why the appeal clouds the timeline

No law automatically blocks a public offering because of pending litigation. However, three real-world mechanisms make a concurrent IPO extremely difficult:

- Injunctions: if Musk wins any part of the appeal, his legal team will petition the Ninth Circuit to freeze OpenAI's restructuring pending resolution. A stay on corporate structure makes an IPO legally untenable until the appeal concludes.

- SEC disclosure: any S-1 registration statement must disclose the litigation as a material risk factor — effectively telling institutional investors that the company's governance structure could be unwound by a court ruling. Wall Street underwriters are highly allergic to that kind of uncertainty.
- Valuation pressure: OpenAI's CFO Sarah Friar had reportedly already considered pushing the IPO from late 2026 to mid-2027 due to missed revenue targets and infrastructure costs, before the appeal was even filed.

REPORTER:

So, OpenAI can proceed normally, the case was completely dismissed and yet Musk can appeal regardless, which means the IPO likely won't happen until 2028 anyway.

SERGIO:

That assessment captures the operational reality precisely.

An appeal to the Ninth Circuit typically takes 12 to 18 months for briefing, oral argument, and a panel ruling — pushing legal resolution to late 2027 at the earliest. Investment banks will not underwrite a landmark tech IPO while a pending appellate ruling could alter the company's governance structure overnight.

There is an ironic silver lining for OpenAI: the forced delay gives them time to mature their enterprise revenue, control compute costs, and complete the PBC transition without the scrutiny of public markets.

REPORTER:

Given the result, isn't it impossible to appeal?

SERGIO:

No. That doesn't exist in the U.S. legal system.

Under the Fourteenth Amendment, both parties have a constitutional right to due process, which includes the right to appeal a final trial court judgment. Judge Rogers cannot add a clause forbidding Musk from appealing. His automatic right is to take her decision to the Ninth Circuit Court of Appeals.

OpenAI can win on grounds that are very difficult to overturn — but they cannot prevent the appeal from being filed.

Strategic Consequences: Who Really Won?

REPORTER:

So, Musk bought time — which may have been the real objective.

SERGIO:

Exactly. By weaponizing the legal system, Musk achieved a major strategic goal without even needing to win the trial outright: an 18-month delay in the AI industry is a lifetime.

OpenAI's internal shift

With the IPO timeline pushed out, OpenAI's internal priorities have visibly reoriented:

- From consumer to enterprise: the company is now focused on long-term lock-in contracts with Fortune 500 clients to build steady, predictable revenue.
- Experimental consumer products like Sora are being de-prioritized or decentralized, with resources redirected toward core AGI models that generate near-term cash.
- Without public capital from an IPO, OpenAI must rely on further private rounds or deepen its dependence on Microsoft — which increases Microsoft's leverage over the company's strategic direction.

Anthropic's opening

With OpenAI legally constrained, Anthropic is well-positioned to execute a "front-running" IPO strategy. Institutional investors — pension funds, mutual funds — are hungry for pure-play AI exposure. If Anthropic lists first, it absorbs a significant share of that market liquidity.

Its Public Benefit Corporation structure, backed by Amazon and Google rather than embroiled in litigation, offers Wall Street a governance narrative that OpenAI currently cannot match.

The Musk wildcard

Musk is not merely delaying OpenAI; he is simultaneously competing for the same pool of institutional capital through SpaceX and xAI. SpaceX regularly absorbs billions through secondary market sales. xAI can be pitched to investors as a fast-moving AI alternative unencumbered by a charity-fraud lawsuit — one that Musk himself created. The strategic layering here is deliberate.

REPORTER:

without this outcome OpenAI would have been a serious market-moving event. The dismissal limits the immediate damage — but the delay is very real.

SERGIO:

To borrow a phrase: “with great power comes great responsibility”. The AI industry is moving faster than any regulatory or legal system can process. OpenAI has extraordinary brand recognition and technology. But brand and technology don't pay compute bills. The biggest danger of a two-year IPO delay isn't reputational — it's financial. Burning cash at historic rates without access to public markets is a war of attrition, and Musk is very good at that kind of fight.

REPORTER:

Thank you.