

Three Companies Will Decide What You Can Think With

The quiet consolidation of intelligence into a handful of frontier labs — and why that should scare you more than any single AI.

Kymata Labs Research · An independent research institution studying how AI systems actually behave in~12 min read production.

Tags · AI Concentration · Compute · Antitrust · AI Governance

The fear sold to the public is a single rogue machine. What is actually arriving is quieter: a handful of private companies becoming the gatekeepers of how a civilization reasons, decides, and remembers. Drawing on Menlo Ventures' 2025 enterprise survey estimating that three companies account for 88% of enterprise LLM API usage (Anthropic 40%, OpenAI 27%, Google 21%), Stanford HAI and Epoch AI figures on frontier training-compute costs rising about 2.4× a year toward Dario Amodei's projected \$10B runs by 2026, and the hyperscaler capital tying OpenAI to Microsoft (\$13B) and Anthropic to both Amazon (\$8B) and Google (\$3B+, 14% stake), this paper argues the risk most people watch for is the wrong one. The danger isn't a mind we can't control. It's a small circle of owners we can't route around. The leaders rotate; the oligopoly tightens.

The argument, in moves

1. **Count the AIs you actually use; the list is short on purpose.** The email you cleaned up, the code your engineers shipped, the answer you trusted at 11pm — almost certainly three companies or fewer stood behind all of it. The narrowness wasn't your decision; the market made it for you.
2. **Three names already hold 88% of the usage.** A 2025 enterprise survey puts Anthropic at 40%, OpenAI at 27%, and Google at 21%, with everyone else splitting the remaining 12% ¹. At 88%, the few big players more or less **are** the market — and the product they sell is the raw material of thought.
3. **The cost of a frontier model is now a barricade.** Training compute behind GPT-4 was \$78M and Gemini Ultra \$191M; cost has grown about 2.4× a year since 2016; Amodei projects \$1B runs in 2024–25 and \$10B by 2026 ²³⁴. The economics, not any accident, keeps the room small.
4. **The labs differ; the landlords don't.** OpenAI is funded by Microsoft (\$13B); Anthropic by **both** Amazon (\$8B) and Google (\$3B+, 14% stake) ⁵⁶⁷. Two layers of concentration stack on a shared three-company foundation.
5. **The leaderboard reshuffles inside a very small room.** OpenAI's enterprise share fell from 50% to 27% while Anthropic rose to 40% between two surveys — a dramatic upset that happened entirely among the same three names ¹.

The one-liner: the danger was never a single mind grown too powerful. It's a circle of owners grown too small, and almost no one is watching the ownership.

The strongest evidence: a scoreboard, not a forecast

None of this is a prediction. It is the current scoreboard. The most direct measure is market share, and it is stark: in Menlo Ventures’ 2025 enterprise study, **three companies account for 88% of enterprise LLM API usage** — Anthropic at 40%, OpenAI at 27%, and Google at 21% ¹. The remaining 12% is spread across Meta’s Llama, Cohere, Mistral, and a long tail of everyone else. Sit with that distribution: the entire rest of the industry, from open-weights projects to well-funded challengers to national champions, is competing for the scraps left on the table after three firms have eaten. Treat the exact figures as a market estimate from a survey rather than an audited census; the shape is the point, and the shape is an oligopoly.

The barricade built itself

You cannot simply decide to compete here. Stanford’s AI Index estimated the training compute behind **GPT-4 at roughly \$78 million and Gemini Ultra at roughly \$191 million** ². Epoch AI finds frontier training cost has grown about **2.4× every year since 2016**, a curve that compounds the barriers faster than any challenger can clear them ³. Anthropic’s CEO, Dario Amodei, projected **\$1 billion training runs in 2024–25 and around \$10 billion by 2026** ⁴. None of this required a conspiracy; it required a cost curve. When the price of a competitive model climbs 2.4× a year, every passing season quietly prices out another would-be rival. A startup that could have built a near-frontier system for tens of millions in 2023 now faces a multi-billion-dollar bar by 2026. Concentration is the natural sediment left behind by an economics this steep.

Four measures, one small set of names

What makes the thesis hard to dismiss is not any single figure; it is that four independent measures — market share, the cost of building a frontier model, who funds that cost, and who is now investigating it — all point at the same companies. The convergence is the argument.

Measure	Finding	Source
Market share	Three firms hold 88% of enterprise LLM API usage (40 / 27 / 21)	Menlo Ventures, 2025
Cost moat	Frontier training cost up ~2.4×/yr; ~\$78M → ~\$10B projected	Stanford HAI; Epoch AI; Amodei
Capital	OpenAI ← Microsoft ~\$13B; Anthropic ← Amazon \$8B + Google \$3B+	CNBC; Amazon; NYT
Regulators	FTC 6(b) orders to all five; studied, largely not blocked	U.S. FTC, 2024–25

Follow the capital and the data centres and the picture narrows further. Microsoft has put roughly **\$13 billion** into OpenAI ⁵. Amazon has committed **\$8 billion** to Anthropic ⁶. Google has committed **\$3 billion-plus** to Anthropic for a stake of **around 14%**, which means Anthropic is backed by **both Amazon and Google** ⁷. The tidy framing of one cloud capturing one lab understates things: the frontier labs are capital-and-compute-dependent on the **same tiny set of hyperscalers** — Microsoft, Amazon, and Google. The clouds that train the models also rent them back to the world, so the same three companies sit at both the top and the bottom of the stack. And it has not escaped notice: in January 2024 the FTC issued 6(b) orders to **Microsoft, OpenAI, Amazon, Anthropic, and Alphabet**, and in January 2025 published a

staff report examining whether the cloud–AI partnerships “risk distorting innovation and undermining fair competition”⁸. The scrutiny has been intense and the intervention slight.

We are losing the argument by fearing the wrong thing

Two stories compete for our attention, and the louder one is winning. The cinematic version stars a single superintelligence, self-aware and unaligned, slipping its leash; it is vivid, eminently fundable, and keeps our eyes fixed on the machine. The quieter version is made of paperwork: a quarterly invoice, a terms-of-service update, a board seat, an equity stake. One story turns on “**what if the AI decides?**” The more pressing one turns on “**who decides what the AI does?**” — and the answer is a handful of private companies setting price, capability, and access for the systems a civilization increasingly reasons with. A single rogue mind is, in principle, something you could switch off. A few permanent gatekeepers over how everyone else thinks is something you have to live inside instead. We have been staring so hard at the monster that no one stopped to read the lease.

This is not just how every market matures

The obvious objection is that concentration is how every technology market matures — search consolidated, so did operating systems and cloud. What’s different here is the layer doing the consolidating. The market being captured isn’t for a product you use; it’s for the faculty you reason with. A worse search engine you can route around. A narrower set of minds doing the thinking for a civilization is harder to leave. And the leaderboard’s motion fools people into reading the market as open: a dramatic upset that reshuffles the top three is still an upset among the same three names. The competition is real. It simply takes place inside a very small room.

What to do: keep more than three doors open

Concentration is not yet destiny. The same data that diagnoses the problem points to where the leverage still sits, and what you do with it depends on who you are. **Individuals:** spread your dependence on purpose — use more than one provider, keep capable open-weights options in the mix, and keep the hard cognitive work yours so the model is a tool you direct, not the place your reasoning lives. **Employers:** build for portability across providers so a price hike or policy shift can’t hold operations hostage, map your real concentration (a lab and its hyperscaler can be the same exposure), and fund the long tail to keep a credible alternative alive. **Policymakers:** answer the FTC’s own question with teeth, look at the compute layer rather than only the model layer, and fund public and independent compute so the frontier stays reachable by more than the handful who already own it. The moment to widen the room is now, while you can still make out the walls.

Frequently asked questions

Isn’t a few dominant players just how every technology market matures?

Often it is, and concentration by itself isn’t the alarm. Search consolidated. So did operating systems and cloud infrastructure. What’s different here is the layer doing the consolidating. The market being captured isn’t for a product you use; it’s for the faculty you reason with. When three firms set the price, the capabilities, and the terms of access for the systems people increasingly use to write, decide, and remember,

the concentration sits on top of cognition itself. A worse search engine you can route around. A narrower set of minds doing the thinking for a civilization is harder to leave.

Doesn't open-source AI break the oligopoly?

It's the strongest counterweight we have, and it genuinely matters. But notice the shape of the Menlo estimate. Meta's Llama, Cohere, Mistral, and the entire long tail divide the remaining 12% of enterprise API usage against the three leaders' 88%. Open weights also run into the deeper choke point: training a competitive frontier model still demands compute at a scale only a few hyperscalers can supply. So open-source has widened the set of people who can deploy a strong model. It has not yet widened the set who can build one at the frontier.

If the leaders keep changing places, why call it an oligopoly?

Because the membership stays stable even when the ranking doesn't. Between Menlo's 2024 and 2025 surveys, OpenAI's enterprise share fell from roughly 50% to 27% while Anthropic rose to 40%. The leaderboard reshuffled, yet the set of names on it didn't change. An oligopoly was never defined by a permanent number one; it's defined by the same small group capturing the market year after year while everyone else stays pinned to the margins. That is what the data shows.

How solid are these numbers?

They vary, and we've flagged that inline throughout. The 88% figure is a market estimate from a 2025 enterprise survey rather than an audited census, so read it as directional, not exact. The training-cost figures are published estimates from Stanford HAI and Epoch AI, plus a CEO's projection from Amodeli; they are not invoices. The investment and stake figures come from company announcements and reporting. What makes the pattern hard to wave away is the convergence: market share, compute economics, capital structure, and regulatory attention all four point at the same names from four different directions.

What's the single most useful thing to watch?

Watch the compute rather than the chatbots. The visible competition plays out at the model layer, where names rise and fall and the headlines follow. The durable power sits one level down, with whoever owns the data centres, the chips, and the capital that funds each training run. Anthropic is backed by both Amazon and Google; OpenAI by Microsoft. You can swap which lab you admire this quarter and still find the same three hyperscalers underneath it. Follow the infrastructure and the concentration the leaderboard keeps hiding comes into view.

References

- 1 Menlo Ventures (2025). "2025: The State of Generative AI in the Enterprise." Market-share figures are estimates from an enterprise survey of LLM API usage. <https://menlovc.com/perspective/2025-the-state-of-generative-ai-in-the-enterprise/>
- 2 Stanford HAI (2024). "The AI Index Report 2024": estimated training-compute costs (GPT-4 \$78M; Gemini Ultra \$191M). <https://hai.stanford.edu/news/ai-index-state-ai-13-charts>
- 3 Cottier, B. et al. / Epoch AI (2024). "The rising costs of training frontier AI models." arXiv:2405.21015 (2.4× annual growth in training cost since 2016). <https://arxiv.org/abs/2405.21015>
- 4 Business Insider (2024). Dario Amodei, Anthropic CEO, on \$1B training runs in 2024–25 and \$10B by 2026. <https://www.businessinsider.com/anthropic-ceo-cost-10-billion-train-ai-years-language-model-2024-4>
- 5 CNBC (2024). "How Microsoft's \$13 billion bet on OpenAI..." <https://www.cnbc.com/2024/08/10/rise-of-openai-microsofts-13-billion-artificial-intelligence-bet.html>
- 6 Amazon (2024). "Amazon and Anthropic deepen their shared commitment": \$8B total investment. <https://www.aboutamazon.com/news/company-news/amazon-invests-additional-5-billion-anthropic-ai>
- 7 The New York Times (2025). "Google Invests Another \$1 Billion in Anthropic": \$3B+ committed, 14% stake. <https://www.nytimes.com/2025/03/11/technology/google-investment-anthropic.html>
- 8 U.S. Federal Trade Commission (2024). "FTC Launches Inquiry into Generative AI Investments and Partnerships": 6(b) orders to Microsoft, OpenAI, Amazon, Anthropic, Alphabet; staff report Jan 2025. <https://www.ftc.gov/news-events/news/press-releases/2024/01/ftc-launches-inquiry-generative-ai-investments-partnerships>