

Essay

AI Is not a Natural Monopoly

Simon Goldstein & Peter N. Salib[†]

Economists and antitrust scholars have recently warned that the AI industry may be a natural monopoly. In support of this claim, they have argued that the AI industry shares key features with natural monopolies of the past: First, like railroads, AI has high fixed and low marginal costs. That is, training a frontier AI is expensive, but asking it a question is cheap. Next, like social media, AI companies will benefit from network effects. The more users a company has, the more training data they can collect. This leads to better models, then more users, and so on. Finally, some antitrust scholars say, the AI industry is already too concentrated. Today, the market for frontier AI systems contains only three—maybe three and a half—players.

Appearances here are, however, deceiving. This essay argues that the AI industry is not a natural monopoly. Nor is it plagued by the problems of market concentration today.

To show why, the essay identifies three structural features of AI essential for understanding the industry's competitive dynamics in the long run. First, power-law capabilities scaling and fast-following dynamics mean that training costs are not a barrier to competition. Although training the world's best AI model is enormously expensive, training one that is just as good—but six months later—is cheap. This is the story of OpenAI's GPT-4o and DeepSeek v3. Second, recent breakthroughs in reinforcement learning mean that user data—and thus network effects—are no longer central to improving AI systems. Today's AI

[†] Equal contribution. Simon Goldstein is an Associate Professor of Philosophy at The University of Hong Kong, Principal Investigator at the HKU AI and Humanity Lab, and a Research Affiliate at the Center for AI Safety. Peter N. Salib is an Assistant Professor of Law at The University of Houston Law Center, Executive Co-Director of the Center for Law and AI Risk, and Law and Policy Advisor to the Center for AI Safety. Copyright © 2026 by Simon Goldstein & Peter Salib.

companies are not competing to amass the largest mountain of training data, but to engineer the best virtual environments in which their models can “learn by doing.” Finally, and most subtly, we argue that a bit of market power in the AI industry might be a good thing. Today’s AI industry is highly innovative, despite having only a few players. Monopoly power can be bad if it raises prices or degrades quality too much. But as Aghion and Howitt’s 2025 Nobel Prize-winning work argued—and as every patent lawyer knows—monopoly power can also be essential for incentivizing innovation. Thus, antimonopoly interventions for the AI industry could, paradoxically, increase prices and reduce quality in the long run.

INTRODUCTION

People are beginning to worry about AI companies' market power. As a small handful of AI companies suck down a firehose of investment capital, scale their user bases with historic speed, and push their valuations above the half-trillion mark, these firms seem poised to monopolize, or oligopolize, the market for "intelligence on tap."¹

The parties warning of monopolization include prominent legal scholars from the neo-Brandeisian antimonopoly movement, along with mainstream think tanks and economists.² In these thinkers' telling, AI may well be a *natural* monopoly,

1. *Data-Center AI Chip Market – Q1 2024 Update*, TECH INSIGHTS (May 28, 2024), <https://www.techinsights.com/blog/data-center-ai-chip-market-q1-2024-update> [<https://perma.cc/MB7R-XRCB>]; Jared Spataro, *AI at Work: "Intelligence on Tap" Will Reshape Knowledge Work*, MICROSOFT (Apr. 23, 2025), <https://www.microsoft.com/en-us/worklab/ai-at-work-intelligence-on-tap-will-reshape-knowledge-work> [<https://perma.cc/QLU9-P6SS>].

2. Tejas N. Narechania, *Machine Learning as Natural Monopoly*, 107 IOWA L. REV. 1543 (2022); Tejas N. Narechania & Ganesh Sitaraman, *An Antimonopoly Approach to Governing Artificial Intelligence*, 43 YALE L. & POLY REV. 95 (2024). For other analyses from outside the broader neo-Brandeisian sphere, see Anton Korinek & Jai Vipra, *Concentrating Intelligence: Scaling and Market Structure in Artificial Intelligence*, 40 ECON. POLY 225 (2025); Jai Vipra & Anton Korinek, *Market Concentration Implications of Foundation Models* (Brookings Ctr. on Regulations & Markets, Working Paper No. 9, Sept. 2023), https://cdn.governance.ai/Market_Concentration_Implications_of_Foundation_Models.pdf [<https://perma.cc/H3A4-PJUW>]; Joshua S. Gans, *Market Power in Artificial Intelligence* (Nat'l Bureau of Econ. Rsch., Working Paper No. 32270, Mar. 2024); Thibault Schrepel & Alex Pentland, *Competition Between AI Foundation Models: Dynamics and Policy Recommendations*, 34 INDUS. & CORP. CHANGE (2025); Fasheng Xu, Xiaoyu Wang, Wei Chen & Karen Xie, *The Economics of AI Foundation Models: Openness, Competition, and Governance* (Oct. 2025) (unpublished manuscript), <https://arxiv.org/abs/2510.15200>; Christophe Carugati, *Competition in Generative Artificial Intelligence Foundation Models* (Bruegel, Working Paper No. 14/2023, July 2023); Jon Schmid, Tobias Sytsma & Anton Shenk, *Evaluating Natural Monopoly Conditions in the AI Foundation Model Market*, RAND (2024), https://www.rand.org/pubs/research_reports/RRA3415-1.html [<https://perma.cc/3SAM-793M>].

In this essay, we focus on Narechania and Sitaraman's contributions for two reasons. First, because their economic analyses are in many ways representative of the other works. All of these analyses predate the reasoning revolution, and thus the transition to reinforcement learning in bespoke environments as a primary mode of scaling model capabilities. They likewise predate the DeepSeek release and the fast-following open-source model revolution. Both are important to our arguments against AI as a natural monopoly. Second, Narechania and Sitaraman's work are the clearest in calling for regulatory action now. Our analyses of which antimonopoly approaches would, and would not, be appropriate for a concentrated AI market are thus most germane to their work.

meaning that, absent intervention, market power will inexorably concentrate in the hands of a few firms.³ The leading AI companies will leverage their access to deep capital markets to finance AI training runs so expensive that competitors will not dare to follow. Then, they will release the models those training runs create cheaply, or for free, gaining huge user bases. From those huge user bases, the leading companies will collect *even more* training data and use it to make their systems even better. This will increase the user base, and the network effects will compound until a small initial lead becomes unassailable.

As a result, some of these thinkers argue, the government must act now.⁴ It should consider regulating AI companies' rates and products, financing competitors or public options, or imposing democratic governance requirements.⁵ Regulatory delay would, the argument goes, ignore anticompetitive harms accruing today and invite opposition to regulation from a truly consolidated monopolist tomorrow.⁶

This essay challenges that story. AI is not a monopoly, natural or otherwise. It is true that, facially, the AI industry does have some features that hint of monopoly. Like the railroad industry of the nineteenth century, the fixed costs of training a frontier model are enormous, while the marginal cost of each prompt to the trained model is low.⁷ And like the social media networks of the twenty first century, AI does involve software served to user bases that scale exponentially.

But first appearances here are deceiving. We uncover three deep structures of the AI industry that cut against the AI monopoly story: fast-following competition, the reinforcement learning revolution, and continuous innovation pressures that may in fact make modest accumulation of market power a net-positive phenomenon.

Fast-following is simply the observation that the market for *frontier* AI systems does not comprise the entire market for

3. See Narechania, *supra* note 3.

4. *Id.* at 1596.

5. *Id.*

6. *Id.*

7. See Ben Cottier, *Trends in the Dollar Training Cost of Machine Learning Systems*, EPOCH AI (Jan. 31, 2023), <https://epoch.ai/blog/trends-in-the-dollar-training-cost-of-machine-learning-systems> [https://perma.cc/J7TF-RW3M] (describing training costs); *Pricing*, CLAUDE, <https://platform.claude.com/docs/en/about-claude/pricing> [https://perma.cc/2RGG-F44Z] (accessed on Mar. 6, 2026) (describing the cost to prompt).

advanced AI systems.⁸ It is true that just a few firms—OpenAI, Anthropic, Google DeepMind, and maybe xAI—operate at the bleeding edge of AI progress.⁹ And claiming the title of world’s best AI model, even temporarily, does require sinking huge fixed costs, in the form of massive training runs.¹⁰

The bleeding edge, however, advances rapidly. And a gaggle of open-source AI providers—Meta, DeepSeek, Qwen, and more—consistently provide models that would have been bleeding edge, had they been released six months earlier.¹¹ Such near-frontier models can be trained for a fraction of the cost of true frontier models.¹² The savings are passed on to consumers; many fast-following models are available for free.¹³

We argue that deep features of AI technology—including the power-law relationship of training compute to model capability and steady improvements to algorithmic efficiency—will likely enable fast-following competition indefinitely. As a result, advances in model outputs will continue to be rapidly commoditized, initial capital expenditures notwithstanding.

This is not the structure of a natural monopoly; it is the opposite.

The case against network effects in the AI industry likewise turns on technical aspects of AI progress. It is true that, until late 2024, the AI frontier was pushed mostly by training on ever-greater quantities of human-produced data.¹⁴ But the advent of

8. Brian Eastwood, *AI Open Models Have Benefits*, MIT (Jan. 20, 2026), <https://mitsloan.mit.edu/ideas-made-to-matter/ai-open-models-have-benefits-so-why-arent-they-more-widely-used> [<https://perma.cc/DAQ6-MLK5>].

9. Anton Korinek & Jai Vipra, *Concentrating Intelligence: Scaling and Market Structure in Artificial Intelligence*, (Working Paper No. 228, 2024), https://www.ineteconomics.org/uploads/papers/WP_228-Korinek-and-Vipra.pdf [<https://perma.cc/H6T7-SVC9>].

10. Cottier, *supra* note 8.

11. Tharin Pillay, *The Gap Between Open and Closed AI Models Might Be Shrinking. Here’s Why That Matters*, TIME (Nov. 5, 2024), <https://time.com/7171962/open-closed-ai-models-epoch> [<https://perma.cc/AKZ3-PXYZ>].

12. Eduardo Baptista, *China’s Deepseek Says Its Hit AI Model Cost Just \$294,000 to Train*, REUTERS (Sept. 18, 2025), <https://www.reuters.com/world/china/chinas-deepseek-says-its-hit-ai-model-cost-just-294000-train-2025-09-18> [<https://perma.cc/HNC8-48QP>].

13. *Download DeepSeek AI*, SOFTONIC, <https://deepseek-ai.en.softonic.com/chrome/extension/download> [<https://perma.cc/239M-LVGF>] (accessed on Mar. 6, 2026) (providing a link to download DeepSeek AI for free).

14. Kylie Robison, *OpenAI Cofounder Ilya Sutskever Says the Way AI is Built is About to Change*, THE VERGE (Dec. 13, 2024), <https://www.theverge.com>

“reasoning” models like OpenAI’s o1 and o3 changed things dramatically.¹⁵

Today, the most important method of advancing AI capabilities is “reinforcement learning.”¹⁶ Reinforcement learning is a kind of learning by doing.¹⁷ Here, the main determinant of progress is not how much human-produced text you have: It is how many high quality virtual environments your engineers have built in which your AI can work on hard problems, verify correct answers, and update to reinforce the behaviors leading to those correct answers.¹⁸ The centrality of these environments to future AI progress means that AI is not currently, and probably will not become, social media 2.0.

What will the AI economy look like if we are right, and the AI industry will not monopolize? Many scenarios are possible. We sketch one for illustrative purposes, chosen for its unexpected combination of features. In our scenario, AI companies have no monopoly power, have huge revenues, have razor thin profit margins, have large profits by today’s standards but small ones by tomorrow’s, act as net reducers of wealth inequality, and nonetheless are immensely powerful.

The essay’s final contribution is an analysis of what the optimal policy response would be if there *were* an accumulation of market power in the AI industry. Antimonopoly commentators have proposed a wide array of interventions ranging from NPU-law-inspired price controls to public options to novel governance structures.¹⁹

We suggest another remedy may be optimal for AI market power, should it exist: nothing, at least for now. Monopolies are

.com/2024/12/13/24320811/what-ilya-sutskever-sees-openai-model-data-training? [https://perma.cc/6RZF-4UW7].

15. Maxwell Zeff, *OpenAI Trained o1 and o3 to ‘Think’ About Its Safety Policy*, TECHCRUNCH (Dec. 22, 2024), <https://techcrunch.com/2024/12/22/openai-trained-o1-and-o3-to-think-about-its-safety-policy> [https://perma.cc/JNY6-HSFA].

16. Timothy B. Lee, *How a Big Shift in Training LLMs Led to a Capability Explosion*, ARS TECHNICA (Jul. 7, 2025), <https://arstechnica.com/ai/2025/07/how-a-big-shift-in-training-llms-led-to-a-capability-explosion> [https://perma.cc/7NR7-JF8Q].

17. *Id.*

18. Aiko Tanaka, *Simulation Environments: The Secret Sauce of Effective RL Training*, SMARTCR (Aug. 26, 2025), <https://smarter.org/ai-technologies/reinforcement-learning/rl-simulation-environments> [https://perma.cc/8USZ-W89M].

19. See Narechania & Sitaraman, *supra* note 3, at 156–67.

classically harmful because they cause prices to increase and quality to decline.²⁰ But, as 2025 Nobel Prize Winners Philippe Aghion and Peter Howitt have shown, a bit of monopoly power can also be essential to innovation.²¹ Every patent attorney knows this. Intellectual property is a limited monopoly designed to incentivize investment in invention.²²

We apply this framework—which trades off deadweight loss from monopoly pricing against innovation incentives—to the AI industry. Today, that industry is fiercely competitive and highly innovative. Thus, applying antimonopoly regulation to the AI industry could, perhaps paradoxically, increase prices and reduce quality in the long run. We argue that NPU-style price and quality regulations would be especially pernicious in this regard.

One more note before proceeding. This essay's analysis is of the market for general-purpose AI *models*, like GPT-5, Claude Sonnet 4.5, or DeepSeek R2. There are, of course, other markets relevant to the growth and progress of AI as a technology. Among them, the markets for: electricity, training data, frontier training chips, inference chips, EUV lithography machines, foundries, cloud computing, machine learning engineers, hardware engineers, laptops, cell phones, AI-enabled wearables, and more. Some of these markets—especially for training data, chips, and cloud computing—have been subjected to scrutiny by antimonopoly thinkers alongside the market for AI models.²³

We prefer to separate the analyses. These other markets mostly predate the AI revolution.²⁴ Thus, much more has been written about them, and much more of what was true of them ten years ago remains true today.²⁵ Powerful, generalist AI

20. Greg Brailsford, *Monopolies: How Limited Competition Affects Everyday Life*, UPRISE RI (July 25, 2023), <https://upriseri.com/monopolies-limited-competition-affects-you> [<https://perma.cc/E32M-PW7H>].

21. See generally Philippe Aghion and Peter Howitt, *A Model of Growth Through Creative Destruction*, 60 *ECONOMETRICA* 323 (1992).

22. Eric A. Posner and E. Glen Weyl, *Property is Only Another Name for Monopoly*, 9 *J. L. ANALYSIS* 51, 52 (2017). Indeed, as property theorists will tell you, all private property has a monopolistic character. *Id.*

23. See generally Daniel McIntosh, *We Need to Talk About Data: How Digital Monopolies Arise and Why They Have Power and Influence*, 23 *J. TECH. L. & POL'Y* 185 (2019).

24. *Id.*

25. See, e.g., Dina Srinivasan, *The Antitrust Case Against Facebook: A Monopolist's Journey Towards Pervasive Surveillance in Spite of Consumers' Preference for Privacy*, 16 *BERKELEY BUS. L.J.* 39 (2019); Herbert Hovenkamp, *Antitrust and Platform Monopoly*, 130 *YALE L. J.* 1952 (2021); Tim Wu, *Blind Spot:*

models, by contrast, are a genuinely new and, we think, revolutionary technology. Understanding what is new, or not, about the specific market for this new technology seems worthwhile.²⁶

I. THE CAUSES AND EFFECTS OF MARKET POWER

Market power occurs when one, or a small handful, of firms controls the entire market.²⁷ Firms with market power can, in turn, harm various other participants in the market in various ways. They can charge prices that are too high or sell goods that are too low in quality.²⁸ They can impose these harms on consumers, on workers, or on other firms that supply or buy from them.²⁹

There are two basic ways of detecting, or forecasting, monopoly power: with direct evidence of a market's current characteristics, or with indirect evidence rooted in a market's structure.³⁰

Antitrust lawsuits and agency enforcement actions mostly take the first approach, turning on data about the market as it exists today.³¹ They might, for example, examine whether the prices of the relevant goods increased following a merger.³² An

The Attention Economy and the Law, 82 ANTITRUST L. J. 771 (2019); C. Scott Hemphill & Tim Wu, *Nascent Competitors*, 168 U. PA. L. REV. 1879 (2020).

26. One disadvantage of disaggregating the analyses here is that it excludes arguments about how vertical integration might affect market power. *Cf.* Narechania & Sitaraman, *supra* note 3, at 104. We think it is worth sacrificing that discussion here for the sake of increased clarity at the model layer. In general, we think cross-layer analyses are valuable, but the resulting market power is hard to predict. Consider, e.g., that Google's vertical integration to the chip layer—its production of proprietary TPUs—is a market entry challenging NVIDIA's prior near-monopoly. *See* Tensor Processing Units (TPUs), GOOGLE, <https://cloud.google.com/tpu> [<https://perma.cc/C34E-KZDZ>].

27. James A. Schmitz, *The Costs of Monopoly: A New View*, FED. RSRV. BANK MINNEAPOLIS (2016), <https://www.minneapolisfed.org/article/2016/the-costs-of-monopoly-a-new-view> [<https://perma.cc/J4JV-PQDQ>].

28. *Id.*

29. *Id.*

30. *See generally* U.S. DEP'T OF JUST. & FED. TRADE COMM'N, HORIZONTAL MERGER GUIDELINES (1997), <https://www.justice.gov/atr/horizontal-merger-guidelines-0> [<https://perma.cc/CC2P-MDAG>]; THOMAS PHILIPPON, THE GREAT REVERSAL (2019).

31. U.S. DEP'T OF JUST. & FED. TRADE COMM'N, HORIZONTAL MERGER GUIDELINES (1997), <https://www.justice.gov/atr/horizontal-merger-guidelines-0> [<https://perma.cc/CC2P-MDAG>]; PHILIPPON, *supra* note 31.

32. U.S. DEP'T OF JUST. & FED. TRADE COMM'N, HORIZONTAL MERGER GUIDELINES (1997), <https://www.justice.gov/atr/horizontal-merger-guidelines-0> [<https://perma.cc/CC2P-MDAG>].

answer in the affirmative suggests that the merging firms gained sufficient market power to charge uncompetitive prices.³³

Quality is difficult to measure directly. Courts sometimes use barriers to market entry or consumer choice as a proxy.³⁴ The idea is that, when it is hard to switch between products, or for new products to come to the market, quality will suffer.

The second way of assessing, or predicting, market power proceeds from economic theory. This approach involves thinking from first principles about the structure of a market, to see whether it has features that will promote monopoly.

Legal scholars Ganesh Sitaraman and Tejas Narechania have recently argued, individually and together, that the AI industry has structural features of a *natural monopoly*.³⁵

Natural monopolies arise when it is cheaper for a single firm to service a market than for multiple firms to do so.³⁶ More carefully, natural monopolies are characterized by subadditive cost functions.³⁷ Intuitively, a subadditive cost function is a situation where, for any quantity of goods, it would cost a single firm less to produce those goods than it would for two separate firms to each produce part of the quantity.³⁸

Subadditive cost functions can arise when economies of scale increase indefinitely, so that every new good a firm produce is cheaper than the last.³⁹ They can also arise when the production of a good has high fixed and low marginal costs.⁴⁰ Consider, for example, if a single factory could supply all of the ping pong balls the world demanded. Then, once a single firm invested in that factory, its marginal cost to produce balls would be cheap. Here, the building of a second factory would be societally wasteful—and other firms would have little reason to do it. Absent an (illegal) agreement not to compete, the two factories would be forced to sell balls at exactly the marginal cost. Neither would ever recoup the cost of building the physical plant.

33. *Id.* See also PHILIPPON, *supra* note 31.

34. See, e.g., *United States v. Microsoft Corp.*, 253 F.3d 34 (D.C. Cir. 2001).

35. See Narechania, *supra* note 3; Narechania & Sitaraman, *supra* note 3.

36. See William J. Baumol, *On the Proper Cost Tests for Natural Monopoly in a Multiproduct Industry*, 67 AM. ECON. REV. 809, 809 (1977).

37. *Id.* at 810.

38. So, with C denoting the cost function and $Q = Q_1 + Q_2$ denoting quantity, we would have $C(Q) < C(Q_1) + C(Q_2)$.

39. Baumol, *supra* note 37, at 814.

40. *Id.*

In addition to the naturally monopolistic force of subadditivity, *network effects* can also produce market power.⁴¹ A product benefits from network effects when, the more people who use or buy it, the more valuable it is to everyone.⁴² Consider social media. The more people who use Facebook, the more people each user can connect with, and the more valuable Facebook becomes.⁴³ Here, we can get winner-take-all effects, where the market converges to a single provider that connects everyone with everyone.

This does not always happen, nor persist. If switching and startup costs are low, rival networks can enter, provide a differentiated product, and gain market share. Consider that TikTok, X, Pinterest, and YouTube all exist, despite Facebook's massive user base.⁴⁴ And if the new entrant is successful enough, the entire market may migrate rapidly from one large network to another. Think here of MySpace's rapid collapse following Facebook's entry to the marketplace.⁴⁵

II. ARGUMENTS FOR AI MARKET POWER

What are the arguments that AI is, or will become, a monopoly?

Antimonopoly thinkers start with direct evidence from today's marketplace.⁴⁶ They point out that, circa 2025, only a small handful of companies in the world reliably produce the very most capable models.⁴⁷ The top players in the United States are OpenAI, Anthropic, DeepMind, xAI, and Meta.⁴⁸

But antimonopoly scholars' arguments for AI market power have largely been indirect and forward-looking, based on the

41. Michael L. Katz & Carl Shapiro, *Network Externalities, Competition, and Compatibility*, 75 AM. ECON. REV. 424, 438 (1985).

42. *Id.*

43. See Guy Aridor, Rafael Jiménez-Durán, Ro'ee Levy & Lena Song, *The Economics of Social Media*, 62 J. ECON. LIT. 1422, 1438 (2024).

44. *Social Media Fact Sheet*, PEW RSCH. CTR. (Nov. 20, 2025), <https://www.pewresearch.org/internet/fact-sheet/social-media> [https://perma.cc/C8L7-8DDL].

45. Stuart Dredge, *MySpace – What Went Wrong: ‘The Site Was a Massive Spaghetti-Ball Mess’*, GUARDIAN (Mar. 6, 2015), <https://www.theguardian.com/technology/2015/mar/06/myspace-what-went-wrong-sean-percival-spotify> [https://perma.cc/2ZY4-SKYU].

46. See *supra* note 3.

47. *Id.*

48. *Id.*

application of economic theory.⁴⁹ Here, the claim has been that AI has structural features that will lead to natural monopoly—subadditive costs or network effects.⁵⁰

Let's begin with subadditive cost. Here, scholars have pointed out that software in general has high fixed and low marginal cost.⁵¹ It is expensive to develop but cheap to run: all the more so for AI foundation models.⁵² The key distinction is between the compute costs associated with *training* (making the AI model) versus *inference* (running the model).⁵³ Training costs are fixed, while inference is charged bit by bit, with each user's prompt.⁵⁴

This, it is argued, creates barriers to entry.⁵⁵ It cost over \$100 million to train GPT-4.⁵⁶ This cost grows as models get bigger and more powerful, with estimates of the biggest near-future training runs costing \$1 billion.⁵⁷

This cost structure could lead to a natural monopoly. As Narechania puts it, while “the fixed computational costs of training the application will be comparatively high . . . the marginal costs of responding to queries after training will be comparatively low. Hence . . . the average marginal computational cost of each query to the system seems likely to decrease continually. In

49. Narechania, *supra* note 3, at 1550.

50. *Id.*

51. *Id.* at 1574.

52. *Id.* (“Developing software can be an exceptionally complex task, requiring significant resources in terms of expertise and labor (i.e., software developers), as well as computing power. But once a program is built, it is easy to burn copies onto discs or, more commonly (and more cheaply), transmit them over the internet.”).

53. *The Difference Between AI Training and Inference*, NEBIUS (July 25, 2025), <https://nebius.com/blog/posts/difference-between-ai-training-and-inference> [https://perma.cc/J35B-4T2P].

54. Kyle Aubrey, *How the Economics of Inference Can Maximize AI Value*, NVIDIA (Apr. 23, 2025), <https://blogs.nvidia.com/blog/ai-inference-economics> [https://perma.cc/N4BS-5PFN].

55. Narechania & Sitaraman, *supra* note 3, at 122.

56. Will Knight, *OpenAI's CEO Says the Age of Giant AI Models Is Already Over*, WIRED (Apr. 17, 2023), <https://www.wired.com/story/openai-ceo-sam-altman-the-age-of-giant-ai-models-is-already-over> [https://perma.cc/SA3B-BVVE].

57. Ben Cottier et al., *The Rising Costs of Training Frontier AI Models* (May 31, 2024) (unpublished manuscript), <https://arxiv.org/pdf/2405.21015> [https://perma.cc/DSP9-FRG7].

short, machine-learning-based applications seem computationally subadditive.”⁵⁸

The argument from network effects has to do with training data.⁵⁹ Training data is sometimes called the crude oil of AI—the essential commodity that allows the entire industry to exist.⁶⁰ When an AI product becomes more popular, it attracts more users.⁶¹ And each use of the product creates data in the form of text and interactions with the product, which can in theory be used to improve the product.⁶² As with social media companies, this could cause the product to gain more users, more data, more value, and so on indefinitely.⁶³

Network effects create market power from the demand side, rather than the supply side. The product’s *quality gets better* with scale, rather than its *costs falling* with scale.⁶⁴ But the end result is the same: one model may serve the entire market better than two.⁶⁵

III. AGAINST AI MARKET POWER

What is wrong with the arguments above? We have four main points. First, today’s AI industry is characterized not by consolidation, oligopoly, and stasis, but by fierce competition.⁶⁶ Second, that very competition is causing entry costs to fall rapidly, undercutting the possibility of future consolidation.⁶⁷ Third, even absent falling costs of entry, the fixed costs of training an AI system are not actually large, when compared against

58. Narechania, *supra* note 3, at 1583.

59. *See id.* at 1585.

60. Marco Iansiti & Andres V. Lerner, *Assessing the Strength of Network Effects in Social Network Platforms* (Harv. Bus. Sch., Working Paper No. 21-086, 2021), https://www.hbs.edu/ris/Publication%20Files/21-086_a5189999-6335-4890-b050-a59a4b665198.pdf [<https://perma.cc/XRH2-TF6T>].

61. *Id.*

62. *Id.*

63. Narechania, *supra* note 3, at 1584 (“the first application in the market gains access to more recent and more relevant training data—data from in-market consumers—before any competitor.”).

64. Iansiti & Lerner, *supra* note 61.

65. We note here that in classic economic models of market power; network effects alone do not produce monopolies. If, for example, switching costs are low, multihoming is easy, or market segmentation is possible, then even industries with strong network effects can be highly competitive. *See* Iansiti & Lerner, *supra* note 61.

66. *Infra* Section II.A.

67. *Infra* Section II.B.

the relevant marginal costs.⁶⁸ Fourth, network effects are unlikely to matter very much in the market for AI—and they are especially unlikely to matter for reasons having to do with training data.⁶⁹

A. TODAY'S AI MARKET IS HIGHLY COMPETITIVE

We begin with direct evidence from today's market for AI models. Today, the AI industry is highly competitive, and that competition is fueling high levels of innovation. As a result, there is currently little direct evidence of the price or quality effects one would expect from a concentrated market.

Above, we saw that there are only a few companies operating at the bleeding edge of general-purpose AI.⁷⁰ Nonetheless, those companies compete vigorously with one another on both price and quality.⁷¹ This can be seen, for example, on the leaderboards for benchmarks that test AI systems' capabilities.⁷² At various times in the past two to three years, colorable claims to the world's best AI model have been held by: OpenAI, Anthropic, Google DeepMind, and xAI.⁷³

It is not just that different AI companies have held the lead. The *absolute quality* of the leading model has also improved substantially over time.⁷⁴ Across essentially every benchmark heretofore devised, frontier AI models are making rapid progress.⁷⁵ Indeed, progress has been so quick that benchmarks are regularly “saturated”—with AIs achieving near perfect scores—and thus rendered useless.⁷⁶

As the quality has improved, prices have fallen precipitously.⁷⁷ Consider: In January 2025, OpenAI's most capable

68. *Infra* Section II.C.

69. *Infra* Section II.B.II.

70. *Supra* Section II.

71. *LLM Leaderboard - Comparison of over 100 AI Models from OpenAI, Google, DeepSeek & Others*, ARTIFICIAL ANALYSIS, <https://artificialanalysis.ai/models> [<https://perma.cc/JY2R-XLMM>].

72. *Id.*

73. *Id.*

74. *See, ARC-AGI-1 Leaderboard*, ARC PRIZE, <https://arcprize.org/leaderboard> [<https://perma.cc/LY3R-UXZT>].

75. *See id.*

76. Shana Lynch, *AI Benchmarks Hit Saturation*, STAN. HAI (Apr. 3, 2023), <https://hai.stanford.edu/news/ai-benchmarks-hit-saturation> [<https://perma.cc/WYD9-SXNK>].

77. Sean Michael Kerner, *OpenAI o3 and o4 Explained: Everything You Need to Know*, TECHTARGET (June 13, 2025), <https://www.techtarget.com/>

publicly available model, o3, cost \$10 per million input tokens.⁷⁸ In early August 2025, OpenAI released o3's more capable successor, GPT-5, priced at \$1.25 per million input tokens.⁷⁹ That's an eightfold reduction in cost-to-consumer over about eight months, on top of a substantial increase in quality.⁸⁰

And that is just how things look at the bleeding edge. While it is true that the costs of training a world-leading model remain high, the costs of producing a *near*-frontier model are far lower and falling fast.⁸¹ OpenAI has reported that it cost over \$100M to train GPT-4, which was the world's best LLM as of 2023.⁸² In late 2024, DeepSeek published a paper showing that their V3 model—whose performance far outstripped GPT-4's—was trained for just over \$5M.⁸³ This cheap-to-produce, cheap-to-run, and free-to-copy model was slightly behind OpenAI's best model *at the time*.⁸⁴ But it was far ahead of where the state of the art had stood just a year before.

Finally, the AI industry has already seen multiple examples of successful entry by new firms.⁸⁵ AI development is not

whatis/feature/OpenAI-o3-explained-Everything-you-need-to-know [https://perma.cc/C4J3-DEQ5].

78. *Id.*

79. *API Pricing*, OPENAI, <https://openai.com/api/pricing> [https://perma.cc/8JQV-US82].

80. *Introducing GPT-5*, OPENAI, <https://openai.com/index/introducing-gpt-5> [https://perma.cc/7D2Y-VH9E].

81. Knight, *supra* note 57.

82. *Id.*

83. DeepSeek-AI et al., *DeepSeek-V3 Technical Report (Feb. 18, 2025)* (unpublished manuscript), <https://arxiv.org/pdf/2412.19437> [https://perma.cc/CC25-8G48]; Thomas Kwa et al., *Measuring AI Ability to Complete Long Tasks (Mar. 18, 2025)* (unpublished manuscript), <https://arxiv.org/abs/2503.14499> [https://perma.cc/MRV5-ERJ4]. Some commenters have challenged this characterization, arguing that it excludes costs like the R&D investment required to figure out how to create such a cheap, capable model. Dylan Patel et al., *DeepSeek Debates: Chinese Leadership on Cost, True Training Cost, Closed Model Margin Impacts*, SEMIANALYSIS (Jan. 30, 2025), <https://newsletter.semi-analysis.com/p/deepseek-debates> [https://perma.cc/KMG3-33XW]. This is true, but it doesn't change the story very much. OpenAI also has such costs, which could be added to their published training costs. Moreover, now that the R&D has been done, anyone may replicate a V3-level model for around \$5M. DeepSeek-AI et al., *supra* note 84.

84. *See Models & Pricing*, DEEPSEEK, https://api-docs.deepseek.com/quick_start/pricing [https://perma.cc/2Y83-9QCG] (explaining features of the DeepSeek model).

85. Trevor Wagener, *Intense Competition Across the AI Stack*, CCIA (Mar. 13, 2025), <https://ccianet.org/articles/intense-competition-across-the-ai-stack> [https://perma.cc/6ULJ-WVML].

dominated by incumbent firms but instead has focused on startups.⁸⁶ In 2017, Google was the obvious candidate for AI monopolist.⁸⁷ It had just invented the transformer architecture that underlies modern LLMs.⁸⁸ It had deep pockets.⁸⁹ It had industrial scale computing architecture.⁹⁰ Nonetheless, entrants like OpenAI and Anthropic were able to race ahead to become the two most important AI companies in the world.⁹¹ Established financiers were happy to fund their endeavors, despite Google's obvious advantage.⁹²

True, these companies received investments from big tech incumbents like Microsoft.⁹³ But that is how capital works—it comes from prior successful ventures. Neither of these large incumbents owns or controls these AI startups outright.⁹⁴ And indeed, the relationship between, for example, Microsoft and OpenAI is as often conflictual as collaborative.⁹⁵

86. *See id.*

87. *See* Ashish Vaswani et al., *Attention Is All You Need*, 31 ADVANCES NEURAL INFO. PROCESSING SYS. (2017), <https://arxiv.org/abs/1706.03762> [<https://perma.cc/GW2G-XE39>].

88. *Id.*

89. Glenn Chapman, *Google's Annual Revenue Tops \$400bn for First Time, AI Investments Rise*, YAHOO FIN. (Feb. 4, 2026), <https://finance.yahoo.com/news/googles-annual-revenue-tops-400-002026416.html> [<https://perma.cc/KF29-RSMT>].

90. Alex Barrett & Barrett Williams, *What a Year! Google Cloud Platform in 2017*, GOOGLE CLOUD (Dec. 26, 2017), <https://cloud.google.com/blog/products/gcp/what-year-google-cloud-platform-in-2017> [<https://perma.cc/9UPU-KH6B>].

91. Sri Muppidi, *OpenAI Tops \$25 Billion in Annualized Revenue as Anthropic Narrows Gap*, INFO. (Mar. 4, 2026), <https://www.theinformation.com/articles/openai-tops-25-billion-annualized-revenue-anthropic-narrows-gap> [<https://perma.cc/DF43-L98B>].

92. *OpenAI's \$110 Billion Funding Round Draws Investment from Amazon, Nvidia, SoftBank*, REUTERS (Feb. 27, 2026), <https://www.reuters.com/business/retail-consumer/openais-110-billion-funding-round-draws-investment-amazon-nvidia-softbank-2026-02-27> [<https://perma.cc/RD5M-EDQT>].

93. *Anthropic Raises \$30 Billion in Series G Funding at \$380 Billion Post-Money Valuation*, ANTHROPIC (Feb. 12, 2026), <https://www.anthropic.com/news/anthropic-raises-30-billion-series-g-funding-380-billion-post-money-valuation> [<https://perma.cc/P5AU-FCZY>].

94. *Id.*

95. Liz Ticong, *'All Bets Are Off': Inside Microsoft's AGI Clause Fight with OpenAI*, TECHREPUBLIC (July 11, 2025), <https://www.techrepublic.com/article/news-openai-agi-clause-threatens-microsoft-deal> [<https://perma.cc/E7V5-QG8V>].

None of this looks very much like an industry plagued by market power. Yet.

B. TOMORROW'S AI MARKET SEEMS LIKELY TO BE COMPETITIVE, TOO

Maybe all of this competition will be short lived. Narechania has argued that, whatever the state of the AI market today, it plausibly has the structural features of a *natural monopoly*.⁹⁶ If so, a competitive market today would likely devolve into a monopolistic one tomorrow.

1. AI Is Not a Natural Monopoly

The future is inherently uncertain, but we are inclined to doubt these gloomy prognostications. Possibly, the AI market will someday develop features of a natural monopoly. But today, we think, it lacks them.

Classically, natural monopolies arise when a single firm can serve the entire market more cheaply than two firms could.⁹⁷ As described above, this can happen when a product has high fixed and low marginal costs—which, in some sense, AI certainly does.⁹⁸ Training a frontier LLM is expensive; asking it a question is cheap.

But this alone is not enough to produce a natural monopoly.⁹⁹ Our ping pong ball factory, discussed above, has high fixed and low marginal costs.¹⁰⁰ But if the first factory reaches capacity, and the market still wants more ping pong balls, it may be cheaper for a second firm to build a factory than for the initial firm to build a second one.¹⁰¹

96. *See generally* Narechania, *supra* note 3.

97. *See* Baumol, *supra* note 37, at 809.

98. *See id.* at 813.

99. Narechania likewise concedes that high fixed costs alone aren't sufficient for natural monopoly: "though the fixed costs of software development are much higher than the marginal costs of software distribution, those fixed costs are, in absolute terms, not all that high. Compared, say, to the fixed costs of building a power plant and an electric grid, developing a software program may seem trivial. And so competition in these software markets, even if properly characterized as natural monopolies, yields less waste than in traditionally regulated markets." Narechania, *supra* note 3 at 1576–77.

100. *See supra* Section I.

101. *See generally* Paul L. Joskow, *Regulation of Natural Monopolies*, in HANDBOOK OF LAW & ECONOMICS (A. Mitchell Polinsky & Steven Shavell eds., 2007).

These factors are subtle and hard to predict. Many industries with high fixed and low marginal costs—airlines, power plants—look facially like natural monopolies but turn out not to be.¹⁰² The market equilibrium for AI is thus difficult to forecast *ex ante*.

Nonetheless, at least today, the industry's cost structure does not look like a natural monopoly.

To begin, the fixed costs of creating a frontier AI are actually *not* that high, given the correct comparator. Training costs certainly appear large in absolute dollar terms.¹⁰³ But they are likely to be small compared to the total market for powerful AI. And this ratio, between fixed costs and addressable demand, is what really matters for natural monopoly.

Consider, for example, that GPT-4 cost roughly \$100 million to train.¹⁰⁴ Some have speculated that training GPT-5 may cost more like \$1 billion.¹⁰⁵ But ChatGPT has over 900 million users.¹⁰⁶ The cost of inference for all of these users is significantly more than \$100 million, perhaps closer to the order of \$5 billion.¹⁰⁷ And in 2025, OpenAI is expected to take in more than \$13 billion in revenue.¹⁰⁸ A market wherein a \$1 billion upfront

102. Lawrence J. White, *Economies of Scale and the Question of Natural Monopoly in the Airline Industry*, 44 J. AIR L. & COM. 545, 571 (1978); Paul L. Joskow, *Deregulation and Regulatory Reform in the U.S. Electric Power Sector*, in DEREGULATION OF NETWORK INDUSTRIES 113 (Sam Peltzman & Clifford Winston eds., 2000).

103. Ben Cottier & Robi Rahman, *Training Compute Costs are Doubling Every Eight Months for the Largest AI Models*, EPOCH AI (June 19, 2024), <https://epoch.ai/data-insights/cost-trend-large-scale> [https://perma.cc/J8KR-NMQA].

104. Will Knight, *supra* note 57.

105. Matthew Griffin, *OpenAI GPT-5 Is Costing \$500 Million per Training Run and Still Failing*, FANATICAL FUTURIST (May 27, 2025), <https://www.fanaticalfuturist.com/2025/05/openai-gpt-5-is-costing-500-million-per-training-run-and-still-failing> [https://perma.cc/KU3Z-SX47].

106. Aisha Malik, *ChatGPT reaches 900M weekly active users*, TECHCRUNCH (Feb. 27, 2026), <https://techcrunch.com/2026/02/27/chatgpt-reaches-900m-weekly-active-users> [https://perma.cc/K6CE-D3BD].

107. Sebastian Moss, *OpenAI Training and Inference Costs Could Reach \$7bn for 2024, AI Startup Set to Lose \$5bn - Report*, DATACENTER DYNAMICS (July 24, 2024), <https://www.datacenterdynamics.com/en/news/openai-training-and-inference-costs-could-reach-7bn-for-2024-ai-startup-set-to-lose-5bn-report> [https://perma.cc/EVQ2-KUZP].

108. Niccolo Conte, *Charted: The Soaring Revenues of AI Companies (2023–2025)*, VISUAL CAPITALIST (Nov. 25, 2025), <https://www.visualcapitalist.com/charted-the-soaring-revenues-of-ai-companies-2023-2025> [https://perma.cc/RKG3-CMD3].

investment, plus \$5 billion in marginal cost, unlocks over \$10 billion in revenue is not a market characterized by exclusionary fixed costs.¹⁰⁹

Furthermore, the fixed barriers to entry in AI are lower than they initially seem. In the AI market, competition at the absolute *frontier* is enormously expensive.¹¹⁰ But fast-following competition is both cheap and highly effective.¹¹¹

There are two reasons for this. First, as discussed above, rapid innovations in training techniques have caused training costs to fall quickly over time.¹¹² Researchers at Epoch AI recently estimated that they halve every eight months.¹¹³

Second, even without continued reductions in training costs, AI capabilities scale as a power law of training compute.¹¹⁴ This means that doubling the performance of the model requires increasing training compute by several orders of magnitude.¹¹⁵ So long as this “scaling law” holds, it will remain relatively easy for new entrants to offer AI systems nearly as good as frontier models, but at much lower prices.

These two facts suggest that any window to extract monopoly rents by sinking large upfront costs will be short.¹¹⁶

Consider: just a year after GPT-4’s release, DeepSeek entered the market with a superior model that could be trained for 1/10th of GPT-4’s cost.¹¹⁷ Since then, the range of tasks for which *anyone* could charge monopoly prices has been limited to the set of things which GPT-5 could do, but which no open weight model can do.

Thus, once one considers the *entire* market for generalist AI models, it looks quite different from classic natural monopolies.

109. Note that frontier AI companies currently operate at a significant loss. However, this is probably driven more by massive investments in research, development, and growth than by the failure of previous investments to pay off.

110. See Knight, *supra* note 57.

111. Anson Ho et al., Algorithmic Progress in Language Models (Mar. 9, 2024) (unpublished manuscript), <https://arxiv.org/abs/2403.05812> [<https://perma.cc/BY5C-ZXP8>].

112. *Id.*

113. *Id.*

114. Jared Kaplan et al., Scaling Laws for Neural Language Models 7–9 (Jan. 23, 2020) (unpublished manuscript), <https://arxiv.org/abs/2001.08361> [<https://perma.cc/Q4FP-Y8C6>].

115. *Id.*

116. Anson Ho et al., *supra* note 112.

117. See Knight, *supra* note 57 and accompanying text.

Imagine, for example, that in the 19th century the cost of laying rail had cut in half every eight months. In that case, the railroad industry would not have been an especially pernicious natural monopoly. A few years after the first company laid their rail, a competitor could do the same at much lower cost and then compete on price. True, the new investment would in some sense be deadweight loss, as it would be more efficient for the first company to keep servicing all customers. But the size of this loss would fall by half every 8 months, and so over time would be insignificant. If new entrants could offer even modestly higher quality—or even differentiated—service, the upfront investment would easily be worth it.

The history of (apparent) natural monopolies is full of such cases—where falling costs or follow-on innovation undermine the monopoly.¹¹⁸ Narechania gives some examples: “[N]ewer technological advances proved these markets capable of sustaining (evermore) competition on their own—once, for example, the railroads faced competition from an array of long-haul truckers (among other newly developed transportation options), [and] wireless telephone service proved to be an adequate substitute for wireline service.”¹¹⁹

2. Network Effects Do Not Drive AI Progress

What about network effects? Will AI systems benefit from network externalities, so that, like Facebook, the more people who use them, the more valuable they become?

We doubt it. Or, at least, we doubt that the network effects in AI will be anywhere near as important as those in social media. With social media, the platform’s value—understood as the total number of possible connections—scales exponentially.¹²⁰

118. Narechania, *supra* note 3, at 1568–69. *See also* VISCUSI ET AL., *ECONOMICS OF REGULATION* at 403.

119. Narechania, *supra* note 3, at 1568–69. *See also* VISCUSI ET AL., *supra* note 119, at 403 (“Railroads possessed significant cost advantages in the late 1800s, and these advantages were considerably eroded with the introduction of trucking in the 1920s. This example introduces a new element, namely, technological change. That is, over long periods of time it is likely that the cost function will shift as new knowledge is incorporated into the production process. Hence, permanent natural monopoly is probably a rare category. Technical change can shift cost functions so as to render competition workable. And, as we will see later, a serious deficiency of regulation seems to be that it often fails to ‘disappear’ when the natural monopoly does.”).

120. *See* Katz & Shapiro, *supra* note 43, at 438.

Each new user adds more potential connections than the last one, forever.¹²¹

AI companies are not like that. Their business model does not revolve around connecting users to one another. It is true, as Narechania points out, that AI companies may collect user data to train their models.¹²² Doing so likely does improve the AI models somewhat, in turn attracting more users whose data can be used in training.

But here, the return to training on new users' data is not exponential. It is probably not even linear. Instead, the benefits of new user data are diminishing, and probably negative beyond some margin.¹²³ Consider: OpenAI's recent problems with model sycophancy suggest that too much training on user preferences can quickly degrade, rather than improve, AI models' usefulness.¹²⁴

The bigger point is that, today, user data simply isn't a core component of the processes by which AI systems get better.¹²⁵ Gone are the days where AI improved primarily by scaling "pre-training"—teaching models to predict the next token of internet text.¹²⁶

Today, the most important paradigm for improving AI systems is "reinforcement learning."¹²⁷ Here, a traditional pre-trained LLM—something like GPT-4—is improved not by predicting more internet text, but by trying to solve hard problems.¹²⁸ For example, the AI is asked difficult questions in

121. *Id.*

122. Narechania, *supra* note 3 at 1578–79.

123. *Sycophancy in GPT-4o: What Happened and What We're Doing About It*, OPENAI (Apr. 29, 2025), <https://openai.com/index/sycophancy-in-gpt-4o> [<https://perma.cc/3S4Q-J3HB>].

124. *Id.*

125. *See generally o3-mini System Card*, OPENAI (Jan. 31, 2025), <https://cdn.openai.com/o3-mini-system-card-feb10.pdf> [<https://perma.cc/8MG3-3C9K>]; Xumeng Wen et al., *Reinforcement Learning with Verifiable Rewards Implicitly Incentivizes Correct Reasoning in Base LLMs* (June 14, 2025) (unpublished manuscript), <https://arxiv.org/abs/2506.14245> [<https://perma.cc/4ES7-H9R4>].

126. Krystal Hu & Anna Tong, *OpenAI and Other Seek New Path to Smarter AI as Current Methods Hit Limitations*, REUTERS (Nov. 15, 2024), <https://www.reuters.com/technology/artificial-intelligence/openai-rivals-seek-new-path-smarter-ai-current-methods-hit-limitations-2024-11-11> [<https://perma.cc/5R4R-BWUT>].

127. OPENAI, *supra* note 126.

128. Xumeng Wen et al., *supra* note 126.

domains, like math, where the answer is easy to check. It produces many attempts, each comprising long chains of reasoning. In those rare initial cases where the system gets the right answer, the parameters that produced the correct reasoning are upweighted. This makes the model modestly more likely to reason similarly in the future. And so on.

The cumulative result is powerful. Today's most capable AI systems, from GPT-5 to Claude Opus 4.5 to Gemini 3 are all "reasoning models" trained with this approach.¹²⁹ The next paradigm shift—to AI "agents"—will be powered by reinforcement learning, as well.¹³⁰

Success in reinforcement learning is about coming up with the right kinds of problems for AIs to solve, generating lots of those problems, building virtual environments in which the AIs progress can be tracked, and giving them the right kinds of rewards for progress or success.¹³¹ User data is, to a first approximation, irrelevant.

Possibly, user data will reenter the picture as the range of tasks subjected to reinforcement learning is broadened. Imagine, for example, a reasoning model being trained to perform the work of an electrician. Here, learning by doing might be best accomplished in a *real*-world environment, rather than a virtual one. And absent sufficiently dexterous robots, the best way for the model to undertake the "doing" might be by observing the situation by video and directing some human's actions.

Whether such human-AI hybrid learning will come to dominate AI training is hard to know. But even if it did, it would not obviously be well-described as a type of training relying on "user data." In the example above, is the human the user? Or is the AI a "user" of the human teacher? Who would we expect to be paying for whose time? Furthermore, what is the evidence that the

129. See *Introducing GPT-5*, OPENAI (Aug. 7, 2025), <https://openai.com/index/introducing-gpt-5> [<https://perma.cc/VZL5-U3YS>]; *Introducing Claude Opus 4.5*, ANTHROPIC (Nov. 24, 2025), <https://www.anthropic.com/news/claude-opus-4-5> [<https://perma.cc/LN6Z-2JCT>]; Sundar Pichai, Demis Hassabis & Koray Kavukcuoglu, *A New Era of Intelligence with Gemini 3*, GOOGLE (Nov. 18, 2025), <https://blog.google/products/gemini/gemini-3> [<https://perma.cc/ZL74-QXW8>].

130. See, e.g., Xufang Luo et al., *Agent Lightning: Train ANY AI Agents with Reinforcement Learning* (Aug. 5, 2025) (unpublished manuscript), <https://arxiv.org/abs/2508.03680> [<https://perma.cc/M7QA-35RN>].

131. STUART RUSSELL & PETER NORVIG, *ARTIFICIAL INTELLIGENCE: A MODERN APPROACH* (2022).

pure volume of data, rather than well-engineered reward mechanisms, will be the scarce factor? While we cannot definitively rule out the possibility of user data someday bottlenecking AI progress, we think today's evidence in that direction is weak.

3. AI Products Are Differentiated

One final thought on competition in the market for AI models: All of the analysis above—both in favor and against the inevitability of AI monopoly—has assumed that AI systems are basically fungible. But when consumers have a rich set of preferences over subtly different products, monopolization becomes even less likely. Even if the products have subadditive costs or benefit from market effects, no firm can serve the whole market.¹³²

Consider that each of ChatGPT, Claude, and Gemini have subtly different user experiences.¹³³ For example, some users find the 'personality' of Claude less frustrating than that of ChatGPT.¹³⁴ Perhaps ChatGPT will sometimes solve a math problem that Claude can't. But Claude may be better at writing computer code. In conditions like these, consumers may be happy to pay *both* Anthropic and OpenAI—and maybe others, too—indefinitely, even if one product is cheaper or better on average.¹³⁵

C. SIZE, PROFITS, WEALTH, AND POWER IN A COMPETITIVE AI ECONOMY

These arguments against AI monopoly do not tell us very much about the broader AI economy. They suggest that AI companies will not have *market* power. But they do not illustrate much about what AI companies, or the market, or society will

132. See Herbert Hovenkamp, *Antitrust and Platform Monopoly*, 130 YALE L.J. 1952, 2021 (2021) (describing the inability for this type of firm to serve the entire market).

133. Tech resources, *Claude AI vs ChatGPT: My Honest Experience Using Both (And Which One I'd Recommend)*, MEDIUM (June 13, 2025), https://medium.com/@Tech_resources/claude-ai-vs-chatgpt-my-honest-experience-using-both-and-which-one-id-recommend-d1b3f4273618 [https://perma.cc/B3B3-YQAE].

134. *Id.*

135. *Cf.* Narechania at 1578 (“Indeed, some markets may sustain several competitors through product differentiation: One song recommendation engine might prioritize similar musical genres while another emphasizes lyrical themes, and, given this differentiation, both might co-exist in a single market.”).

look like, broadly construed. Nor about *other* kinds of power, and where it will lie.

As always, predictions are hard, especially about the future. Doubly so when many of the important factors involve nascent and fast-changing technology.

Nonetheless, here is a scenario. We think it is a plausible one, though there are many other possibilities.

We highlight this scenario because it is surprising. In it, AI companies: (1) fully automate meaningful portions of the economy, (2) lack market power, (3) have unprecedentedly large revenues, (4) have tiny profit margins, (5) have total profits that seem massive by today's standards but small by the standards of the future, and (6) cause a net reduction in wealth inequality.

Today, global GDP is about \$110T.¹³⁶ Suppose also that, as some economists think plausible, AI progress causes economic growth to speed up. Over the next 20 years, growth sits at 10% annually—a rate observed in recent decades in the real-world economies of Korea and China.¹³⁷ Then, by 2045, global GDP will be \$672T—nearly septuple that of today.

Suppose that, at that point, the four leading AI companies have automated roughly 6% of the total economy. They rake in \$10T in annual revenue *each*—about 15x what Walmart or Amazon earns today.¹³⁸

But because of sustained fierce competition, both at the frontier and from fast followers, these companies' profit *margins* are razor-thin. They can charge consumers just enough to cover the marginal energy and compute cost of running their models, but not much more. Suppose their profits are driven down to around 1%—roughly in line with the margins of US grocery stores today.¹³⁹

136. *GDP (current US\$)*, WORLD BANK GRP., <https://data.worldbank.org/indicator/NY.GDP.MKTP.CD> [<https://perma.cc/CGG9-R2NZ>] (accessed Mar. 6, 2026).

137. Graeme Wearden, *Chinese Economic Boom Has Been 30 Years in the Making*, GUARDIAN (Aug. 16, 2010), <https://www.theguardian.com/business/2010/aug/16/chinese-economic-boom> [<https://perma.cc/Y636-DCF7>].

138. *Companies Ranked by Revenue*, COMPANIES MARKET CAP, <https://companiesmarketcap.com/largest-companies-by-revenue> [<https://perma.cc/X6JC-PUH3>].

139. Allison Van Duyne, *Average Profit Margin by Industry – Benchmarks for Small Businesses*, DOORDASH (July 25, 2023), <https://merchants.doordash.com/en-nz/blog/average-profit-margin-by-industry> [<https://perma.cc/L2VH-2ZEK>].

This results in each company generating \$100B in profits annually. By today's standards, that is a lot. Only the most successful tech firms, like Apple and Google, can match it.¹⁴⁰ But in a future economy with \$672T in annual GDP, \$100B will likely appear modest. The AI companies, impressive as they are, will be capturing profits that are smaller than Google's today, as a share of global GDP. Whether \$100B will seem large as compared to other *companies'* total profits depends on the number of people and firms to which the new \$575T in annual GDP flows, and those firms' profitability.

As a result, this world may be one with *less* wealth inequality than today's. Here, venture capitalists' massive investments in inventing capable AI is generating poor returns *to financiers*. But it is producing incredible returns *to society*. The upside of AI progress, in the form of septupled global output, is flowing away from AI companies and their shareholders.

This, of course, does not guarantee that the surplus flows to ordinary people. If, for example, a player at some other level of the AI supply chain consolidates market power, it may capture the surplus and concentrate wealth in its own hands.¹⁴¹ More

140. *Apple Reports First Quarter Results*, APPLE (Jan. 29, 2026), <https://www.apple.com/newsroom/2026/01/apple-reports-first-quarter-results> [<https://perma.cc/KY8P-6VR7>].

141. Consider, for example, ASML and TSMC's market shares in the lithography and chip foundry markets, respectively. James Fox, *Is ASML Stock Poised to Outperform the Market?*, NASDAQ (July 9, 2024), <https://www.nasdaq.com/articles/asml-stock-poised-outperform-market> [<https://perma.cc/QNL6-8BK2>]; *Global Pure Foundry Market Share (Q1 2024 - Q2 2025)*, COUNTERPOINT RSCH., <https://counterpointresearch.com/en/insights/global-semiconductor-foundry-market-share> [<https://perma.cc/LKM4-C5A2>] (accessed Jan. 2, 2025). Sitaraman and Narechania have argued that chip providers and cloud providers may be such monopolists. *Supra* note 3. We agree that, today, these markets look somewhat concentrated—substantially more so than, in our view, the market for AI models is. Nonetheless, we are agnostic about how those would-be monopolists will fare as the AI revolution continues. For example, demand for frontier chips continues to rise, and NVIDIA's margins continue to remain high, companies like AMD, Intel, and Google are currently attempting to enter the frontier AI chip market. Dominic Lombardo, *AMD Is Challenging Nvidia In The Battle For AI Infrastructure*, SEEKING ALPHA (Dec. 2, 2025), <https://seekingalpha.com/article/4849542-amd-is-challenging-nvidia-in-the-battle-for-ai-infrastructure> [<https://perma.cc/T5P5-JCZM>]; Keumars Afifi-Sabet, *Google Is Rapidly Turning into a Formidable Opponent to BFF Nvidia*, TECHRADAR (Dec. 23, 2023), <https://www.techradar.com/pro/google-is-rapidly-turning-into-a-formidable-opponent-to-bff-nvidia-the-tpu-v5p-ai-chip-powering-its-hypercomputer-is-faster-and-has-more-memory-and-bandwidth-than-ever-before-beating-even-the-mighty-h100> [<https://perma.cc/QD3G-BS9S>].

generally, even if AI companies aren't monopolies, other existing monopolies may accumulate great profits from automation. But as with all historical technological revolutions, new wealth from AI will doubtless flow to smaller businesses and ordinary people, as well.¹⁴² If, as some commenters predict, AI will substantially reduce capital and labor requirements to create successful businesses, the technology may structurally favor smaller players.¹⁴³

At any rate, we raise this scenario not because it is overwhelmingly likely, but because it is a possible outcome if AI companies do not in fact have market power. The scenario thus shows one possible, and surprising, future that may emerge if our arguments here are correct.

IV. WHAT TO DO ABOUT AI MARKET POWER

Suppose that the arguments of the previous Parts are wrong, and the AI industry does end up looking like a natural monopoly. What, then, should policymakers do about it?

Antimonopoly thinkers have proposed a range of interventions: regulate AI companies using a network, platform, and utilities (NPU) framework; finance competitors; build a public option; or impose cooperative governance requirements.¹⁴⁴

Even conditional on AI monopoly, none of these seem obviously correct to us, and a few seem clearly misguided. Different kinds of natural monopolies invite different regulatory responses. And some of the proposals would be a poor fit for an AI monopoly, if it emerged. At a minimum, the proposals appear in tension with one another, such that they cannot all be appropriate together.

Classically, natural monopolies create two kinds of social costs. First, when one firm can serve an entire market more

142. Michael Solon, *The AI Revolution Will Bring Prosperity*, HUDSON (Nov. 3, 2025), <https://www.hudson.org/technology/ai-revolution-will-bring-prosperity-michael-solon> [<https://perma.cc/AR6H-59ZT>].

143. Tyler Cowen, *Interview with Sam Altman: Trust, Persuasion, and the Future of Intelligence*, CONVERSATIONS WITH TYLER (Nov. 5, 2025), <https://conversationswithtyler.com/episodes/sam-altman-2> [<https://perma.cc/GPE2-SANP>]; OECD ROUNDTABLES ON COMPETITION POL'Y PAPERS, No. 311. ARTIFICIAL INTELLIGENCE AND COMPETITIVE DYNAMICS IN DOWNSTREAM MARKETS, (2025), https://www.oecd.org/en/publications/artificial-intelligence-and-competitive-dynamics-in-downstream-markets_ccf0624a-en/full-report/component-5.html [<https://perma.cc/WBS7-ULRW>].

144. See generally, Narechania, *supra* note 3; Narechania & Sitaraman, *supra* note 3.

cheaply than two, market competition in which two firms make big investments constitutes social waste.¹⁴⁵ Second, when wasteful competition inevitably collapses into monopoly, the monopolist will raise prices and reduce quality.¹⁴⁶

But competition between firms in a natural monopoly environment can also be benign, or even efficient. Competition produces innovation, which is an extremely valuable public good.¹⁴⁷ This upside can outweigh the waste of duplicative investment.¹⁴⁸

More subtly, some amount of monopoly power may be *vital* for incentivizing innovation. Patent and copyright lawyers are familiar with this principle.¹⁴⁹ Both kinds of intellectual property amount to limited monopolies on the use of an idea—a reward designed to compensate innovators for their investment in producing public goods.¹⁵⁰ Philippe Aghion and Peter Howitt won the 2025 Nobel Prize in economics for formalizing this insight.¹⁵¹ Their work shows that innovation is maximized in a market-power “sweet spot”—where companies face enough competition to motivate progress, but not so much that they are unable to profit from it.¹⁵²

Applying these factors to a potential AI monopoly yields three insights. First, it would be unwise to subject the industry to NPU-style regulation. In an NPU framework, the government allows—or even encourages—a single firm to dominate the

145. ALFRED E. KAHN, *THE ECONOMICS OF REGULATION: PRINCIPLES AND INSTITUTIONS* 119, 173 (1970).

146. *Id.*

147. JOSEPH A. SCHUMPETER, *CAPITALISM, SOCIALISM AND DEMOCRACY* 81–85 (Harper & Bros. 1942); Kenneth J. Arrow, *Economic Welfare and the Allocation of Resources for Invention*, in *THE RATE AND DIRECTION OF INVENTIVE ACTIVITY: ECONOMIC AND SOCIAL FACTORS* 609, 615–16 (Nat’l Bureau of Econ. Research 1962).

148. *See, e.g.*, Joseph E. Stiglitz, *Economic Foundations of Intellectual Property Rights*, 57 *DUKE L.J.* 1693, 1721–24 (2008); Michael Abramowicz, *The Uneasy Case for Patent Races Over Auctions*, 60 *STAN. L. REV.* 803, 809–10 (2007).

149. *See* U.S. Const. art. I, § 8, cl. 8; *Eldred v. Ashcroft*, 537 U.S. 186, 212–13 (2003).

150. WILLIAM M. LANDES & RICHARD A. POSNER, *THE ECONOMIC STRUCTURE OF INTELLECTUAL PROPERTY LAW* 19–23 (2003).

151. Philippe Aghion & Peter Howitt, *A Model of Growth Through Creative Destruction*, (Nat’l Bureau of Econ. Rsch., Working Paper No. 3223, 1990), <https://www.nber.org/papers/w3223> [<https://perma.cc/VWE7-USED>]; *The Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel*, NOBEL PRIZE <https://www.nobelprize.org/prizes/economic-sciences> [<https://perma.cc/RE7V-K2PS>] (accessed Feb. 23, 2026).

152. *See id.*

market but regulates that firm's price and quality.¹⁵³ This eliminates wasteful duplicative competition while avoiding the pitfalls of monopoly pricing and “enshittification.”¹⁵⁴

But NPU-style regulation is the wrong choice for a highly innovative industry. As discussed above, fierce competition in the AI industry is currently producing rapid advances in AI capabilities and cost-effectiveness.¹⁵⁵ Even if AI were a natural monopoly—in the sense that one firm could eventually serve the whole market more cheaply than two—those innovations certainly outweigh the present costs of duplicative investment.¹⁵⁶

Under an NPU regime, the government mediates innovation—either by setting quality standards directly for a state-sanctions monopolist, or by soliciting bids from firms that promise to deliver a product meeting the government's specifications.¹⁵⁷ This would be a disaster for AI. This is not necessarily because the government would be *particularly* bad at identifying promising paths for improving AI—though it probably would be.

The problem is instead that *everyone* is bad at forecasting the trajectory of AI. The science of frontier AI is advancing rapidly, with even world-leading experts disagreeing deeply about the path forward.¹⁵⁸ The only way to find out who is right is by letting many companies make big bets, most of which will turn out to be wasteful, but a few of which will be revolutionary.

153. For defense in the setting of AI, see Narechania, *supra* note 3, at 1563, 1596.

154. ALFRED E. KAHN, *THE ECONOMICS OF REGULATION: PRINCIPLES AND INSTITUTIONS* 119–25 (John Wiley & Sons 1970); Harold Demsetz, *Why Regulate Utilities?*, 11 *J.L. & ECON.* 55, 63–65 (1968); JEAN TIROLE, *THE THEORY OF INDUSTRIAL ORGANIZATION* 391–96 (MIT Press 1988).

155. *See supra* Section II.A.

156. SCHUMPETER, *supra* note 148 at 81–85; Kenneth J. Arrow, *Economic Welfare and the Allocation of Resources for Invention*, in *THE RATE AND DIRECTION OF INVENTIVE ACTIVITY: ECONOMIC AND SOCIAL FACTORS* 609, 616–17 (Nat'l Bureau of Econ. Rsch. 1962); Michael Abramowicz, *The Uneasy Case for Patent Races Over Auctions*, 60 *STAN. L. REV.* 803, 809–10 (2007).

157. Harold Demsetz, *Why Regulate Utilities?*, 11 *J.L. & ECON.* 55, 57–59 (1968); ALFRED E. KAHN, *THE ECONOMICS OF REGULATION: PRINCIPLES AND INSTITUTIONS* 11–18 (John Wiley & Sons 1970); JEAN-JACQUES LAFFONT & JEAN TIROLE, *A THEORY OF INCENTIVES IN PROCUREMENT AND REGULATION* 1–5 (MIT Press 1993).

158. Yann LeCun, *A Path Towards Autonomous Machine Intelligence*, *OPEN REV.* (2022), <https://openreview.net/pdf?id=BZ5a1r-kVsf> [<https://perma.cc/3U8M-BYBY>]; Jared Kaplan et al., *Scaling Laws for Neural Language Models*, (Jan. 23, 2020) (unpublished manuscript), <https://arxiv.org/abs/2001.08361> [<https://perma.cc/58V8-A3RY>].

Indeed, today's level of concentration in the AI market might be near-optimal, from the perspective of the Aghion–Howitt innovation model.¹⁵⁹ As discussed above, the market at the bleeding edge is somewhat concentrated, dominated by roughly four firms. But the rest of the market catches up very quickly, offering AI models that would have been state-of-the-art a few months before, at low or no cost.

Assume that the frontier firms constitute an oligopoly extracting supernormal profits. This would plausibly be an excellent outcome. Those supernormal profits would reward investment in pushing the frontier. But near-frontier firms would gobble them up quickly, incentivizing new frontier innovations, and so on.

If these arguments about innovation are correct, the best thing to do about a natural monopoly in AI might be: nothing, at least for now. Presently, innovation outweighs waste from duplicative investment. And fast-following competition both prevents enshittification and limits monopoly pricing to a limited period—plausibly increasing innovation incentives even further.

We have less to say about the value of the other proposed interventions for a naturally monopolistic AI industry. The others are designed to increase competition, rather than replacing it with managed monopoly.¹⁶⁰ To that extent, they comport with our analysis—while remaining in tension with the proposal for NPU-style regulation.

Should the government fund new entrants in the AI space, or build a public option? Maybe, but we are skeptical that the benefits would be large. Private investment in the AI industry, including for new entrants, is already eye-wateringly large, raising broad fears of a speculative bubble.¹⁶¹

With so much money and labor already flowing into the industry, we think it unlikely that the government-funded or state-run entrants would be especially likely to succeed. The entire industry is already fighting fiercely over scarce ML

159. *See supra* note 152 and accompanying text.

160. *See supra* note 3.

161. Sonya Huang, Pat Grady & GPT-3, *Generative AI: A Creative New World*, SEQUOIA CAP. (Sept. 19, 2022), <https://www.sequoiacap.com/article/generative-ai-a-creative-new-world> [https://perma.cc/3U5X-D5BU]; Rogé Karma, *Just How Bad Would an AI Bubble Be?*, ATLANTIC (Sept. 7, 2025), <https://www.theatlantic.com/economy/archive/2025/09/ai-bubble-us-economy/684128> [https://perma.cc/5CUW-WFZ8].

engineering talent and cutting-edge chips.¹⁶² Consider Meta’s rumored \$100 million pay packages to poach top engineers from OpenAI and Anthropic.¹⁶³ Unless the government-led project were willing to match or exceed such investment, it would seem unlikely to win in an all-out race for dominance at the frontier. More likely, the marginal government investment would simply “crowd out” what would otherwise be private investment.¹⁶⁴

However, if the government-run project were focused on achieving a specific goal that the market was unlikely to pursue, we would favor such a venture. We have in fact proposed one such ambitious state-run AI lab in other work.¹⁶⁵

Finally, what about Sitaraman and Narechania’s endorsement of “cooperative governance,” in which AI companies’ profits are distributed to users, rather than “distant shareholders”? We think this would probably be a mistake in the short run.¹⁶⁶ For now, frontier AI companies are plowing all of their profits—and hundreds of billions of dollars more from investors—into research and development.¹⁶⁷ To the extent that cooperative governance would not change that, there would be no present profits to distribute. To the extent that they would change it, the reduction in innovation would probably harm society as a whole.¹⁶⁸

162. Hyunjoo Jin et al., *The AI Frenzy Is Driving a Memory Chip Supply Crisis*, REUTERS (Dec. 2, 2025), <https://www.reuters.com/world/china/ai-frenzy-is-driving-new-global-supply-chain-crisis-2025-12-03> [<https://perma.cc/TN8V-QYJ6>].

163. Meghan Bobrowsky et al., *Zuckerberg Leads AI Recruitment Blitz Armed With \$100 Million Pay Packages*, WALL ST. J. (June 22, 2025), <https://www.wsj.com/tech/ai/meta-ai-recruiting-mark-zuckerberg-5c231f75> [<https://perma.cc/5ZBX-6W7H>].

164. See Douglas J. Cumming & Jeffrey G. MacIntosh, *Crowding out Private Equity: Canadian Evidence*, 21 J. BUS. VENTURING 569 (2006).

165. Simon Goldstein & Peter Salib, *How to Stop an AI Arms Race* (June 1, 2025) (unpublished manuscript), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5369439 [<https://perma.cc/WL8G-H396>].

166. *Id.*

167. Jaime Sevilla, Hannah Petrovic & Anson Ho, *Can AI Companies Become Profitable?*, EPOCH AI (Jan. 28, 2026), <https://epoch.ai/gradient-updates/can-ai-companies-become-profitable> [<https://perma.cc/ZE4R-B55Z>].

168. One argument for reducing AI investment is that highly advanced AI systems pose an existential risk to humanity, and slowing progress would give society time to mitigate those risks. See, e.g., *Statement on AI Risk*, CTR. AI SAFETY <https://aistatement.com> [<https://perma.cc/D5UZ-XZXV>] (accessed Dec. 3, 2024). We are very sympathetic to those concerns. Peter Salib & Simon Goldstein, *AI Rights for Human Safety*, 88 VA. L. REV. (forthcoming). But since they are orthogonal to the question of market power, we do not address them here. *Id.*

In the long run, we strongly favor policies that broadly distribute the gains from the AI revolution. We also favor increased public governance of AI companies' decisions today, especially to the extent that they impose societal risks—like AIs permanently replacing human jobs or harming humans at scale. But it is somewhat unclear to us why novel corporate governance structures are better tools for these goals than traditional instruments like regulation and taxation.

CONCLUSION

This essay has argued that the AI industry is not a monopoly today, will probably not become a monopoly tomorrow, and thus that anti-monopoly tools are likely to be counter-productive in regulating AI.

But this does not imply that all is well in the AI industry. Market concentration is not the only kind of market failure. AI development may come with a host of externalities, positive and negative. On the positive side, AI development may lead to new scientific discoveries, cure many diseases, accelerate economic growth, and continue to push humanity beyond conditions of material scarcity.

On the negative side, AI development creates a wide range of socialized risks. It may, for example, empower rogue nations or terrorist groups by enhancing their ability to produce bio-weapons or carry out cyberattacks.¹⁶⁹ Alternatively, powerful new AIs may be a threat in their own right, if their ability to pursue long-run goals continues to advance more quickly than our ability to control those goals.¹⁷⁰

Governments can and should regulate the AI industry to help mitigate these, and other, serious risks.¹⁷¹ But to work, the regulatory schemes must fit the problems to be solved. In our view, market power is not the problem on which governments should be focusing. And, more importantly, antimonopoly

169. Cf. S.B. 53, ch. 138, 2025 Cal. Legis. Serv.

170. STUART J. RUSSELL, HUMAN COMPATIBLE: ARTIFICIAL INTELLIGENCE AND THE PROBLEM OF CONTROL (Viking 2019); Alexander Matt Turner & Prasad Tadepalli, Parametrically Retargetable Decision-Makers Tend to Seek Power (June 27, 2020) (unpublished manuscript); Nick Bostrom, *The Superintelligent Will: Motivation and Instrumental Rationality in Advanced Artificial Agents*, 22 MINDS & MACHINES 71 (2012).

171. Goldstein & Salib, *supra* note 166.

regulatory schemes are not the right tool for addressing the most important threats.