MAR 25, 2025

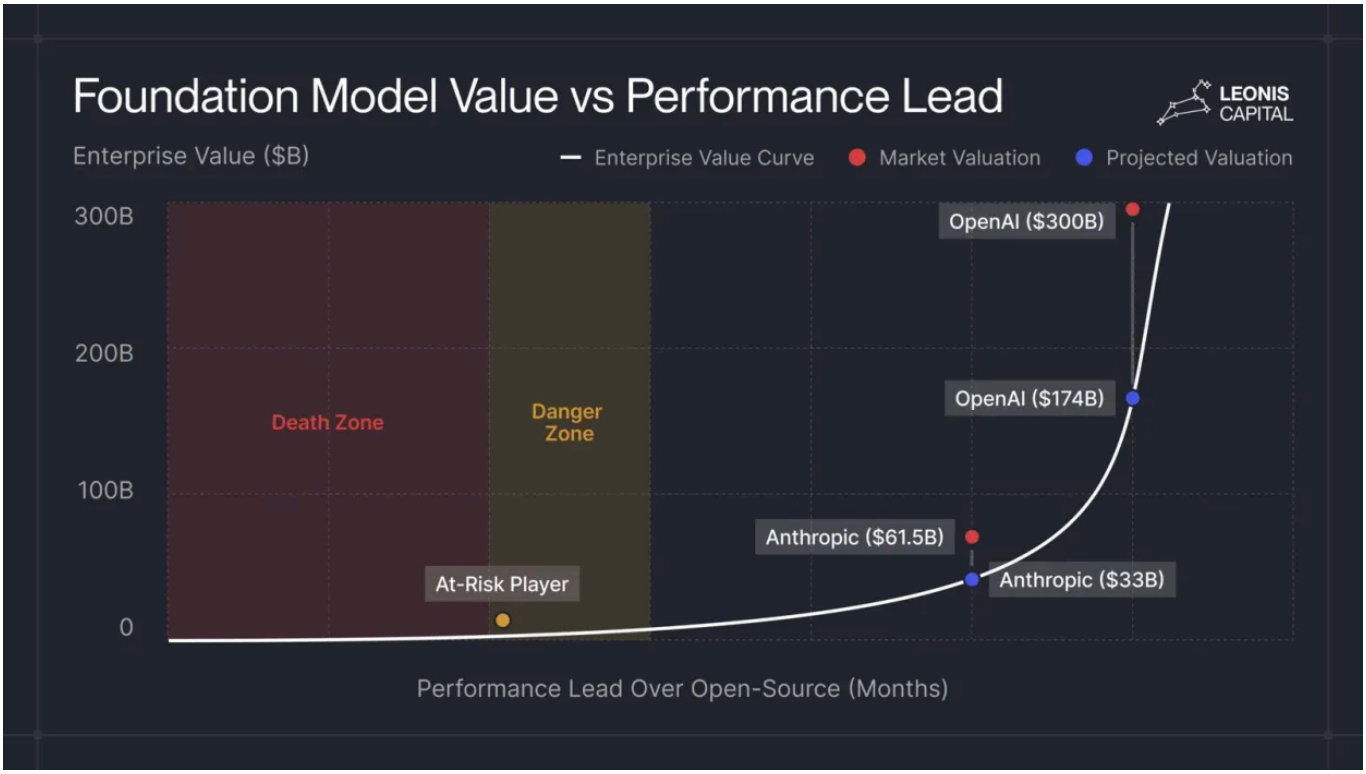
Zero or Hero: A Technical Framework for Valuing AI Companies (Part I: Foundation Models)

Why OpenAI Isn't Worth $300B (And Most AI Companies Are Worth Zero)



By Jenny Xiao and Jay Zhao

Leonis Capital · leoniscap.com/research/zero-or-hero-a-technical-framework-for-valuing-ai-companies-(part-i-foundation-models)/



SoftBank thinks OpenAI is worth $300 billion. Microsoft's investment implied a $157 billion valuation. Elon Musk wants to acquire it for $97 billion. They can't all be right—or perhaps they're all wrong.

This valuation disparity isn't just about OpenAI. It reveals a fundamental problem in AI: We don't have a systematic framework for valuing AI companies. VCs complain about sky-high valuations but offer no alternative way to value this new type of company. Traditional metrics like ARR multiples don't capture AI's unique dynamics—the exponential pace of model improvements that can create what we call **a "zero value threshold," where yesterday's cutting-edge technology**

becomes worthless as new capabilities emerge. Cloud computing analogies fall apart because AI isn't sticky—switching costs between API providers are minimal. Both foundation model providers and application startups face the harsh reality of tough competition and fleeting moats.

To gain clarity in times of chaos, we came up with quantitative frameworks for valuing both foundation model companies and AI applications. By combining technical benchmarks (model performance, training efficiency) with market metrics (penetration rates, user stickiness), we can better understand which AI companies will capture lasting value—and which might be at risk of being destroyed by the exponential improvement of AI and the impending arrival of AGI.

Our analysis will unfold as a two-part series applying a single unified valuation framework to the AI ecosystem. Part I (this article) lays out our framework and demonstrates how it applies to foundation model companies like OpenAI and Anthropic. Part II will extend the same framework to AI application companies, which face their own flavor of displacement risk. While the specific metrics and considerations differ between these categories, the core principle remains consistent: AI companies must stay ahead of a rapidly advancing zero-value threshold to maintain sustainable valuations.

Why Traditional Valuation Methods Fail for AI

Valuing AI companies using traditional VC methods isn't just wrong—it's dangerous. **Traditional metrics give investors false confidence because they assume business continuity in a market defined by discontinuity.** The standard "ARR multiple" approach that worked for SaaS creates systematic mispricing of AI companies because it fundamentally misunderstands how they compete and die.

SaaS companies typically gain or lose market share gradually. A competitor launches a better feature? You might lose 10% of your customers over the next year. But AI doesn't work this way. When a significantly better model emerges, the inferior model's value proposition deteriorates rapidly. Foundation model companies face extinction if open-source alternatives match their performance. AI applications face an equally binary threat: their entire value proposition can vanish when foundation models improve. A $10M ARR AI coding platform built on fine-tuning Claude 3.7 might become worthless when Claude 4 can handle coding natively. Single-purpose applications that had raised millions found themselves competing with foundation models that could replicate their core functionality.

This creates a fundamental problem with traditional valuation methods: T**hey don't account for displacement risk.** Even perfect SaaS metrics—100% YoY growth and 130% net revenue retention—tell you nothing about an AI company's survival. The standard formula (Value = ARR × Multiple) assumes gradual market evolution, but AI companies face existential threats that can materialize overnight.

Some might argue that displacement risk isn't unique to AI—after all, every technology faces disruption. But AI's displacement risk is fundamentally different. Foundation model companies can see advantages evaporate in months, not years, creating unprecedented speed of disruption. When models cross from "sometimes useful" to "reliably useful," adoption curves resemble step functions rather than gradual transitions. This binary outcome happens because AI is still in an early phase where each advancement creates discontinuous leaps that can quickly reset competitive landscapes.

The core metrics investors rely on become misleading because they can't capture the binary nature of displacement risk. We need a new framework that explicitly captures this risk—one that works for both foundation model companies facing open-source competition and AI applications facing obsolescence from improving foundation models.

A New Framework for Valuing AI Companies

The Zero-Value Threshold Principle

Let's start with a simple principle that applies across the AI landscape: AI companies face a zero-value threshold. Once crossed, their enterprise value effectively drops to zero, leaving only the value of talent and infrastructure. For foundation model companies, this threshold is crossing below open source performance. For AI applications, it's when foundation models can replicate their core functionality. While reality has important nuances we'll discuss, this framework helps us understand the broad dynamics that determine whether AI companies can maintain sustainable value.

This zero-value threshold exists because of fundamental technical and market dynamics that are particularly acute in AI:

For foundation model companies:

1. Minimal switching costs (standardized APIs mean customers can switch providers with far less friction than in traditional enterprise software)
2. Hard price ceiling (open source sets a price anchor that closed-source providers must justify exceeding)
3. Technical commoditization (the rapid pace at which competitors can match capabilities through distillation, architecture improvements, or talent mobility)

For application layer companies:

1. Foundation model threat (your product becomes a feature when foundation models improve—a risk more binary than traditional feature absorption)
2. Limited technical moats (fine-tuning and prompt engineering are more easily replicable than traditional software features)
3. Accelerated competition (AI-specific development tools dramatically compress the time needed for competitors to launch similar products)

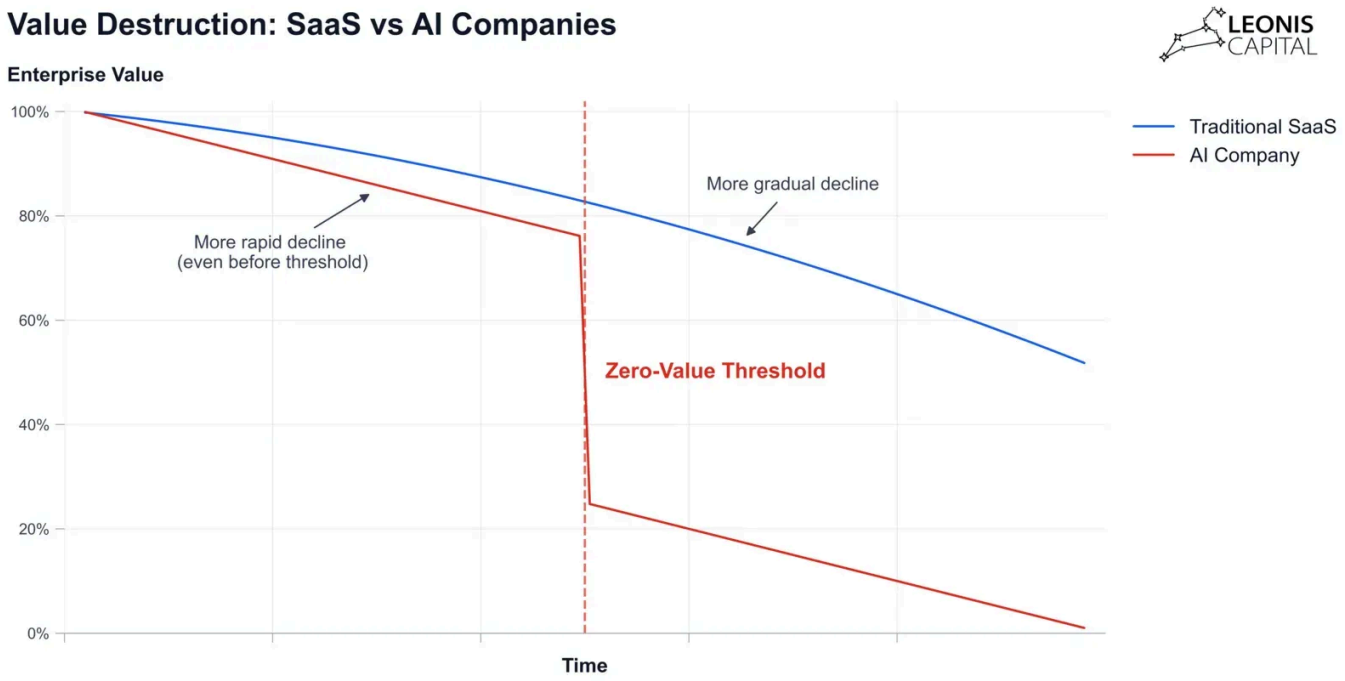


Figure 1: Enterprise value destruction curves (if companies enter the value destruction phase or get outcompeted). Traditional SaaS has a more gradual curve but accelerates over time; AI companies lose value faster and face a drastic decline at the zero-value threshold.

Figure 1 reveals the brutal truth about AI value destruction. While traditional SaaS also faces the threat of commoditization, the key difference in AI is the speed and binary nature of the transition. Traditional software bleeds out gradually—losing customers quarter by quarter as competitors chip away at their value. AI companies face a different reality: a rapid, almost vertical drop once they cross the threshold. They maintain high valuations until suddenly, they don't. When a foundation model company falls behind open source alternatives, they face a rapid pricing squeeze—they must continually justify a premium while their technical differentiation erodes. Similarly, an AI application whose core value can be replicated through foundation model prompting faces abrupt commoditization. In both cases, there's rarely a viable middle ground for prolonged periods. **You're either meaningfully ahead of the threshold or your business faces severe value destruction**, leaving investors and founders little time to course-correct once the descent begins.

A Modified ARR Framework

Having established how AI companies face zero-value thresholds, we can build a valuation framework that captures these unique dynamics. The core formula looks deceptively simple:

```
Valuation = ARR x Valuation Multiple x (1-D)
```

- ARR is Annual Recurring Revenue
- Valuation_Multiple is the standard SaaS multiple based on the company's business metrics
- D is the Displacement Risk Factor (0 to 1)

The key innovation here is the Displacement Risk Factor (D). This single number quantifies what we've been discussing: the probability that a company's core value proposition becomes commoditized. This applies differently across company types but captures the same fundamental dynamic: either you maintain clear superiority in your domain, or your value rapidly approaches zero.

For foundation model providers, D primarily captures technical obsolescence risk. Consider a company like Anthropic or OpenAI. Their D factor depends on how far ahead their models perform compared to open-source alternatives, how efficiently they can train and serve these models, and how defensible their specialized capabilities are. A foundation model company significantly outperforming open source with superior efficiency and strong enterprise features might have D = 0.1. But if open source is catching up and the company lacks unique advantages, D could surge to 0.8 or higher.

For application companies, D measures the risk of being disrupted by improving foundation models. Take a legal tech startup building document analysis tools. Their D factor depends on how far their specialized functionality exceeds what you could get from prompting the best-in-class foundation model, how deeply they're integrated into law firm workflows, and whether they've built genuine network effects through shared knowledge bases. An application deeply embedded in enterprise systems with strong network effects might have D = 0.1. But a simple application layer on top of foundation models could see D approach 0.9.

Some might question whether valuation multiples and displacement factors are redundant—after all, both seem to discount companies facing technical commoditization. But they actually measure distinct dimensions of an AI company's future. Valuation multiples primarily capture market dynamics: growth rates, margins, revenue quality, and competitive positioning. Displacement factors specifically quantify technical sustainability: the probability that a company's core technical advantage persists against rapidly improving alternatives. While correlated (companies with stronger technical moats often command higher multiples), they measure orthogonal risks. A company with excellent growth metrics and high margins might justify a 15x multiple on traditional metrics while still facing severe displacement risk. Conversely,

a company with genuine technical differentiation might have poor unit economics that warrant a lower multiple despite its sustainable advantage. By combining these factors multiplicatively, our framework captures both current business quality and the unique binary risk of AI markets.

A more simplistic approach is to just lower the multiple for AI companies to account for displacement risk, effectively creating a new "AI multiple." While this achieves similar results, the problem is that few valuation methods explicitly factor in displacement risk at all—they rely on subjective adjustments without quantifying the specific threat. Our framework makes this risk explicit and measurable.

This framework explains market behaviors that traditional ARR multiples miss. Why does an AI company's valuation sometimes collapse overnight when a new open-source model launches? Because D suddenly jumps. Why might a smaller AI company with deep technical advantages be worth more than a larger one with more revenue? **Because their D factors differ dramatically.**

The challenge, of course, is estimating D accurately. This requires combining deep technical understanding with market insight. But even rough estimates force investors and operators to confront the fundamental question: What prevents this company's value from going to zero when AI capabilities inevitably advance?

Part I: Valuing Foundation Model Companies

Technical Displacement Risk (D)

How quickly does a foundation model company's advantage evaporate once open source catches up? The data provides a sobering answer: the collapse happens faster than even skeptics expect. When open source matches a model's capabilities, the usage drop isn't gradual—it's a cliff. This creates a critical **“death zone”**—a performance range where their models either perform below open-source capabilities or maintain only a marginal lead over them.

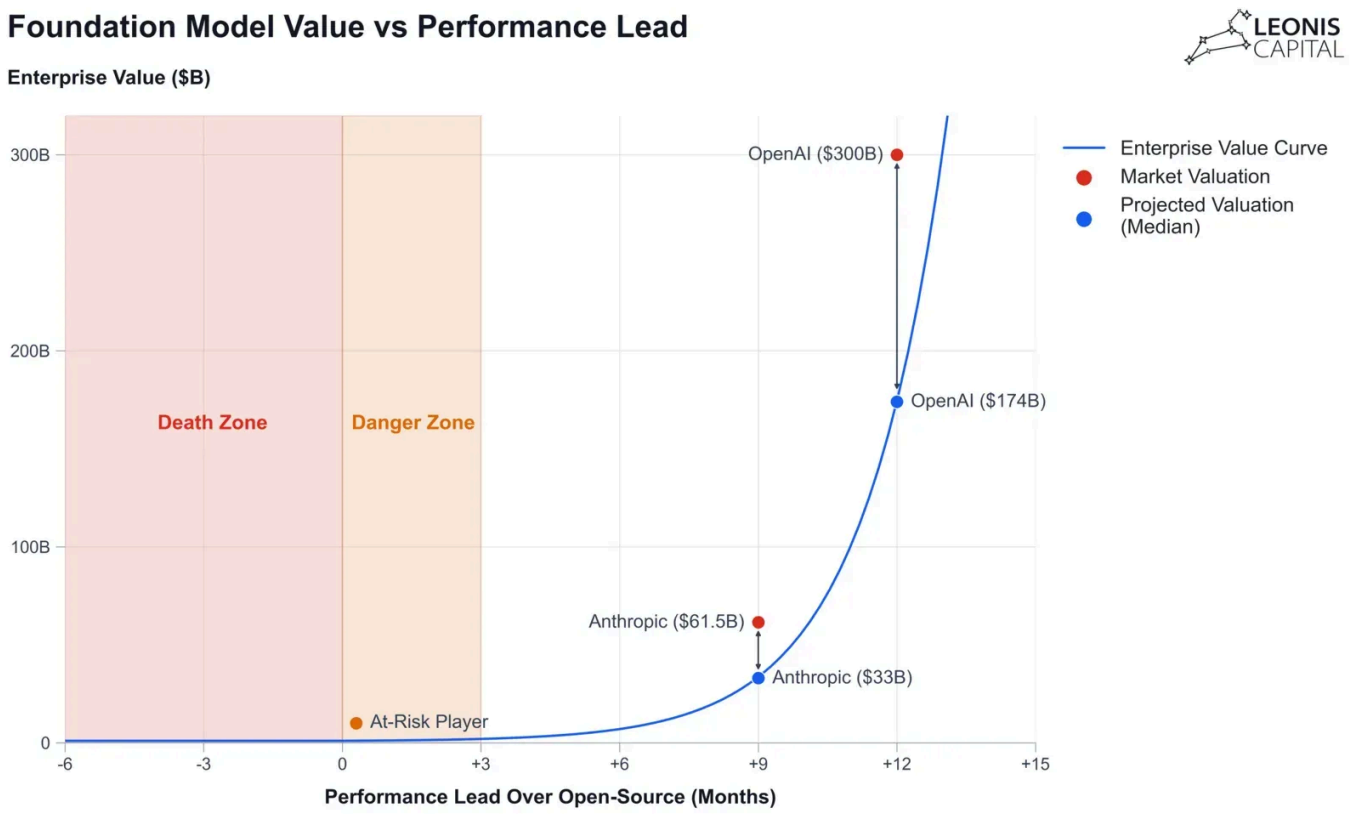


Figure 2: The exponential enterprise value curve illustrates how rapidly valuation increases with technical advantage, as the technical advantage gives companies a compounding advantage on business fronts. However, there are still significant gaps between current market valuations (red) and projected valuations (blue) based on our framework.

Recent market evidence validates this displacement pattern. Over 80% of enterprises indicate that they either use majority open-source models already or plan to switch to open-source once they match close-source performance. Few are loyal to closed-source solutions. Most startups we spoke to abandoned GPT-4o in favor of DeepSeek-V3 when the latter proved that it outperforms the closed-source model. ANZ Bank and a top-three U.S. bank have also switched from OpenAI to fine-tuned Llama models after experiencing "bill shock" from closed-source API costs and concerns over data sovereignty. Meanwhile, open-source downloads have surged, with Meta's Llama series reaching over a billion downloads and DeepSeek boasting around 47 million daily active users globally as of March 2025. It seems that the AI market is shifting from an 80-90% closed-source dominance in 2023 to an even split by 2025.

The open-source displacement cycle follows a predictable pattern. **Open source consistently matches closed-source performance with a lag of 5 to 22 months, centering around one year.** [1] Take GPQA, a Ph.D.-level reasoning benchmark: open models closed the gap in just 5 months. On MMLU, testing general knowledge, it took 16-25 months. But here's the killer: **the lag is shrinking**. The newest benchmarks show the shortest gaps. The death zone is expanding as

open-source models rapidly advance. Most notably, DeepSeek and Llama are setting new open-source performance benchmarks monthly, pushing more closed-source companies toward displacement.

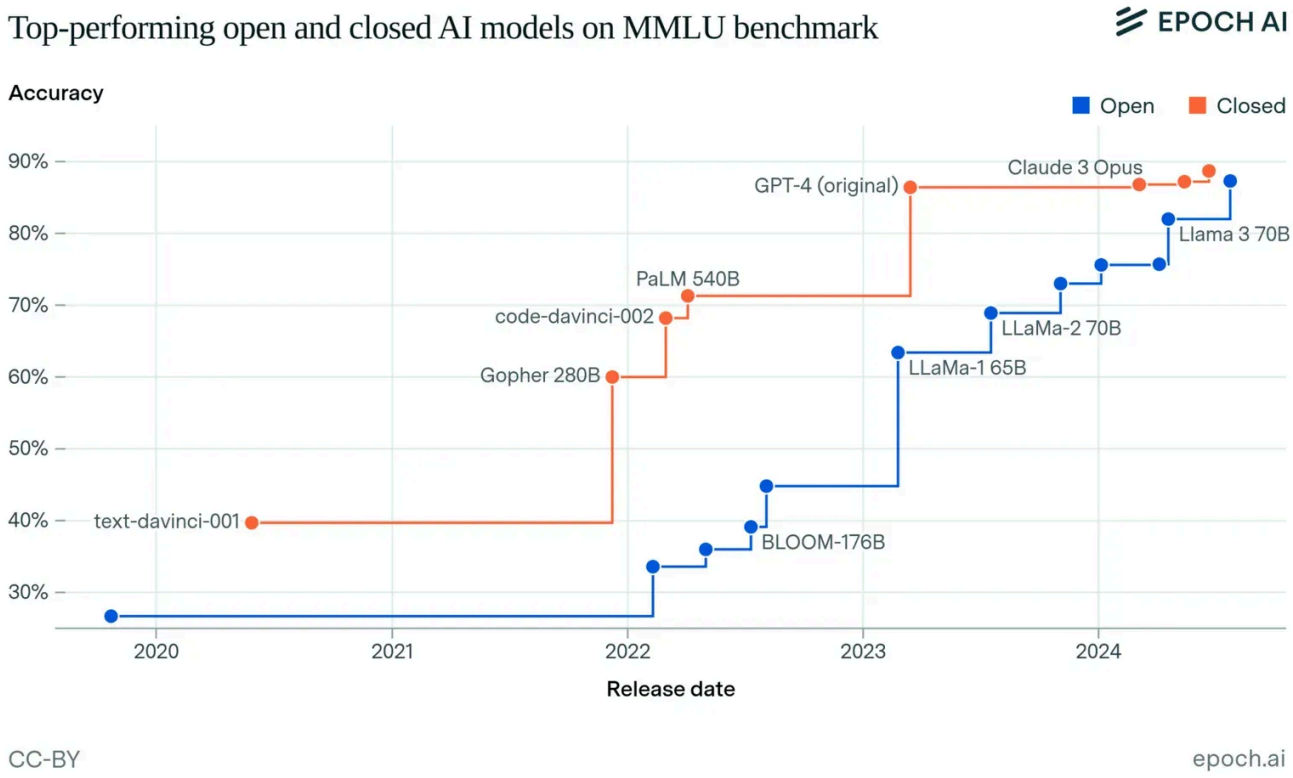


Figure 3: The gap between the top open-source and closed-source models has been closing over time.

More tellingly, **open-source models have maintained the same scaling pace as closed models in training compute—4.6x annually—for the past five years**. This isn't just correlation. When open source reaches a similar compute scale, it achieves comparable performance with almost mathematical certainty. Meta's Llama 3.1 405B proved this by matching GPT-4's scale just 16 months after its release. And the pace is accelerating: Meta has announced plans for Llama 4 to use 10x more compute than Llama 3, potentially eliminating the compute gap with closed models entirely by mid-2025. This isn't just another increment—it would mark the first time an open-source project has caught up to the absolute frontier of AI compute scale. For foundation model companies betting on maintaining a compute advantage, this should be terrifying.

Efficiency gains accelerate displacement. DeepSeek's achievement of matching PaLM 2's performance with 7x less compute is a stark warning. But this is part of a broader pattern: newer open-source models are systematically achieving more with less. Google's own Gemma 2 9B, an open-source model, matched PaLM 540B's performance while using dramatically less compute.

DeepSeek-R1 made headlines because it matched OpenAI's o1 with allegedly only $5.5M in training compute and inferior H800 chips.

Perhaps most concerning for foundation model companies, **the catch-up happens fastest in specialized domains**. The evidence is clear in benchmark data: while general knowledge tests like MMLU show open-source models lagging closed-source by 15-25 months, specialized benchmarks show dramatically shorter gaps. On GPQA (PhD-level reasoning), open-source models closed the gap in just 5 months. Even more striking, on GSM1k (grade-school math), the lag dropped to just 3-10 months. For closed source models, the very specialization that was meant to provide defense could become a liability.

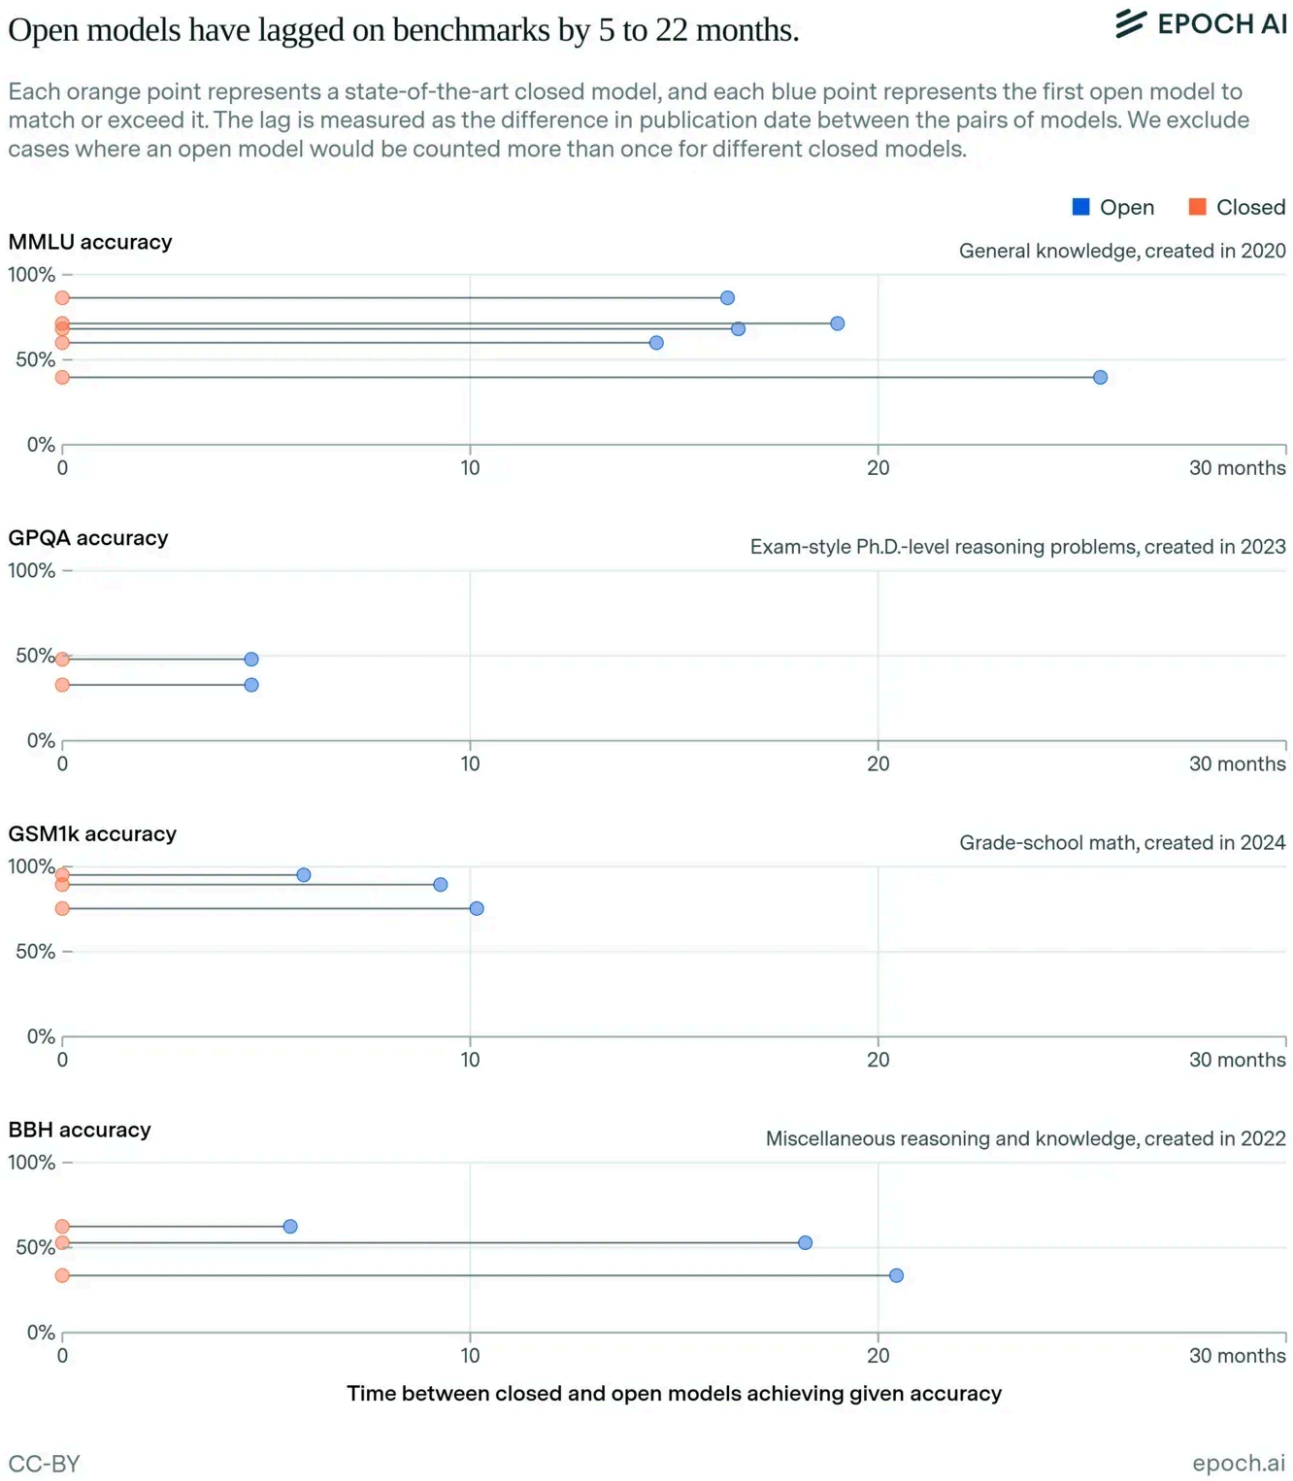


Figure 4: Open source models have closed the gap faster on newer benchmarks like GPQA and GSM1k.

Technical moats offer little protection. AI innovations spread across the industry within months through a dual infection vector: **personnel movement and published research**. Even companies

religiously guarding their secrets can't prevent knowledge transfer when their top researchers take new positions, present at conferences, or publish papers to maintain academic credibility. Even revolutionary training techniques or architectural breakthroughs buy you quarters, not years, of advantage.

The risk for closed-source companies extends beyond just catching up—fundamental architectural shifts can render existing leads irrelevant overnight. A closed-source model might be comfortably ahead of open source in transformer-based architectures, but a breakthrough in alternative architectures could reset the game entirely. This architectural displacement risk exists independently of the current performance lead.

Based on these findings and insights, here are some rough estimations of D:

Technical Lead over Open-Source	D Factor	Typical Characteristics
Significant Lead (>6 months)	0.1-0.2	• Clear technical differentiation • Novel architectures or training approaches • Proprietary advantages in data or compute • E.g., OpenAI, Anthropic
Moderate Lead (3-6 months)	0.3-0.5	• Measurable but shrinking performance gap • Some unique capabilities • Vulnerable to rapid open source progress • E.g., Google Gemini
Limited Lead (<3 months)	0.6-0.8	• Narrow performance advantage • Limited differentiation • At risk of immediate displacement • E.g., Mistral
At/Below Open Source	0.9-1.0	• No clear performance advantage • Generic features • Already technically displaced • E.g., most other foundation model companies

Table 1: Displacement risk factors (D) based on technical lead over open-source models.

The wider D factor ranges for the middle categories reflect greater uncertainty about these companies' futures. While companies with significant leads (D = 0.1-0.2) have clearer paths to sustainable value and those at/below open source (D = 0.9-1.0) face near-certain displacement, companies in the moderate and limited lead categories exist in a more ambiguous zone—their survival depends on factors beyond just technical performance, including execution speed, strategic pivots, and market positioning.

Companies entering the death zone typically attempt three escape routes:

1. **Domain Specialization:** Betting that industry-specific expertise can outrun general-purpose open source (e.g., Cohere).
2. **Infrastructure Pivot:** Building deployment and management tools rather than competing on raw performance—essentially admitting defeat on model superiority (e.g., Stability AI).
3. **Application Layer Retreat:** Embedding AI into vertical-specific software solutions—trading massive TAM for survival (e.g., Baichuan AI).

The hard truth? These strategies might keep you alive, but they're massive downgrades from the original foundation model vision. When your $20B valuation is built on being the leading foundation model provider, retreating to a $2B infrastructure play is still a 90% destruction of value.

This expanding death zone forces a brutal question: Can any foundation model company maintain enough of a technical lead to justify frontier valuations? The data suggests most can't. For investors and operators, the implications are clear: **unless you have a concrete plan to stay 6+ months ahead of open source—with the massive capital requirements that implies—you're likely already in the death zone.** In other words, most of the foundation layer will become more and more commoditized.

Mitigating the Death Zone: Can Distribution Save You?

While technical performance remains the primary driver of displacement risk, some foundation model companies have found ways to buy themselves time. The most successful players leverage scale and distribution moats that can temporarily slow their slide into the death zone—even as open source narrows the performance gap.

OpenAI's tight integration with Microsoft creates genuine switching costs that pure API providers can't match. When your model is embedded in Office, Copilot Studio, and GitHub, customers don't just switch because LLama 4 scores two points higher on MMLU. **These deep platform integrations provide a critical buffer against technical commoditization.**

Scale creates its own defense mechanisms. Companies serving billions of requests gain unique cost advantages through custom silicon, optimized infrastructure, and volume discounts. This lets them compete on price-performance even when their raw capabilities are matched. OpenAI's rumored custom chips and xAI's 100k Colossus Cluster aren't just cost-saving measures—they're displacement-delaying tactics.

The smartest players don't just compete with open source—they selectively embrace it. Meta's strategy of releasing Llama while maintaining proprietary developments shows how companies can harness open innovation while preserving key differentiators. This hybrid approach creates a talent and mindshare moat that pure closed-source companies struggle to match.

These mitigating factors matter, but they don't fundamentally change the game. They might reduce a company's D factor by 0.1-0.2, buying precious quarters of runway. But they can't eliminate the binary nature of AI displacement. Even OpenAI faces the harsh reality: fall too far behind the technical frontier, and no amount of distribution advantage will save you.

Base Multiple Considerations

Traditional SaaS multiples provide a starting point for valuing foundation model companies, but the unique economics of AI require substantial adjustments. For SaaS companies, we can start with these two simple formulas:

```
Base Multiple = Growth Rate x NRR x 10
```

```
Valuation Multiple = Base Multiple x Gross Margin Adjustment
```

Where ARR (Annual Recurring Revenue), Growth Rate, NRR (Net Revenue Retention), and Gross Margins are easy to calculate. Let's walk through some examples to understand typical SaaS multiples:

Category	Growth Rate	NRR	Base Multiple Calculation	Gross Margin	Final Multiple Range
Premium	>100% (1.0)	>120% (1.2)	10 × 1.0 × 1.2 = 12x	>80%	**20-25x**
High Growth	75-100% (0.85)	110-120% (1.15)	10 × 0.85 × 1.15 = 9.8x	>75%	**12-18x**
Moderate Growth	50-75% (0.6)	105-110% (1.08)	10 × 0.6 × 1.08 = 6.5x	>70%	**8-12x**
Low Growth	<50% (0.4)	<105% (1.05)	10 × 0.4 × 1.05 = 4.2x	>65%	**4-8x**

Table 2: Traditional SaaS revenue multiple framework.

A high-growth SaaS company with strong retention might command a 15-25x revenue multiple. A foundation model company with similar growth metrics often deserves a significantly lower multiple. Here's why.

Start with capital intensity. While SaaS companies typically spend their capital on sales and marketing, foundation model companies burn hundreds of millions on training runs before seeing a dollar of revenue. GPT-4 scale models likely cost $100M+ in compute alone. This fundamentally changes the meaning of growth metrics. A SaaS company growing 100% year-over-year typically

means efficient sales execution. A foundation model company showing the same growth might just be burning more compute.

Gross margins create a structural disadvantage for model companies. Unlike SaaS companies with relatively stable 75-85% gross margins, foundation model companies face a shifting landscape of training and serving costs, especially as training costs increase exponentially with model size. These companies are essentially charging a premium on top of underlying cloud infrastructure costs. As open-source alternatives proliferate and give away models for free, commercial providers must continuously justify this markup through performance advantages, specialized features, or enterprise support.

Switching costs and product stickiness is another massive weakness for model providers. SaaS companies with 120%+ net revenue retention often get multiple premiums of 12-20x because this stickiness compounds. Some SaaS products also have strong network effects—think of Slack, where each new user makes the platform more valuable for their colleagues, or Figma, where designers benefit from a growing community of plugin developers and shared component libraries.

Foundation model companies face minimal switching costs in both their API and subscription products. The market share volatility proves this point—from 2023 to 2024, we've seen OpenAI lose 16% market share while Anthropic gained 12%. The ecosystem increasingly facilitates this fluidity, with businesses like our portfolio company Unify.ai specifically helping developers switch between models. Even consumer-facing applications like Perplexity let users choose their preferred model. While foundation model companies can build limited moats through compliance features and customization capabilities, none have demonstrated switching costs approaching those of cloud computing platforms or standard SaaS products.

Another important consideration for both product stickiness and gross margins is the **model portfolio**. A customer using multiple specialized models across their stack (e.g., a fast model for real-time tasks, a code-specialized model for development, and domain-specific models for vertical applications) faces much higher switching costs. This approach also enables sophisticated price discrimination—customers can choose models that match their price-performance requirements—and create opportunities for upselling as customers expand their AI capabilities. This has profound implications for multiples—companies with diverse, specialized model portfolios may deserve higher multiples than those relying on a single flagship model, even if the flagship has superior capabilities.

Factor	Upward Adjustment	Downward Adjustment
Gross Margins	• High-efficiency models with lower inference costs • Portfolio of specialized or powerful models with premium pricing power	• Limited ability to optimize serving costs • Heavy reliance on general-purpose models with commodity pricing
Switching Costs	• Enterprise-grade features (security, compliance, monitoring) • Deep integration with customer workflows • Customized models and fine-tuning capabilities	• Limited enterprise features and integrations • Generic API interfaces easily replaced • Lack of customer-specific model optimization
Capital Intensiveness	• Superior training efficiency, optimized training infrastructure • Proprietary techniques for model compression and distillation	• Inefficient training processes requiring frequent retraining • Limited ability to optimize model size/ performance tradeoffs

Table 3: Foundation model company valuation adjustment framework.

We can modify the base multiple for SaaS companies to come up with a base multiple for foundation model companies, adjusting for the three AI factors (capital intensiveness, gross margins, and switching costs). Given that we are in a rapid adoption phase for AI, we have adjusted the growth rate for AI model providers to be higher than traditional SaaS.

```
Base Multiple = Base Multiple (SaaS) = Growth Rate x NRR x 10
```

```
Valuation Multiple = Base Multiple x AI Adjustment
```

Category	Growth Rate	NRR	Base Multiple Calculation	AI Adjustment	Final Multiple Range
Premium	>200% (2.0)	>120% (1.2)	10 × 1.0 × 1.2 = 24x	• Gross margins >50% • Comprehensive enterprise features and model portfolio • Superior training and inference efficiency (>30% better than baseline)	**15-20x**
High Growth	100-200% (1.5)	110-120% (1.15)	10 × 1.5 × 1.15 = 17x	• Gross margins 40-50% • Strong enterprise features • Above average training and inference efficiency (20-30% better than baseline)	**10-14x**
Moderate Growth	50-100% (0.75)	105-110% (1.08)	10 × 0.75 × 1.08 = 8x	• Gross margins 30-40% • Basic enterprise features • Average training and inference efficiency	**4-6x**
Low Growth	<50% (0.4)	<105% (1.05)	10 × 0.4 × 1.05 = 4x	• Gross margins <30% • Minimal enterprise features • Below average training and inference efficiency	**1-3x**

Table 4: AI foundation model company valuation framework with AI adjustment factors.

The classic "Rule of 40"—where a company's growth rate plus profit margin should exceed 40%—breaks down when applied to AI. A SaaS company growing at 25% with 20% margins might look attractive under this rule, but a foundation model company with the same metrics could be heading for displacement. While a SaaS company's 20% profit margins might reflect stable infrastructure costs and strong pricing power, an AI company's margins could be under pressure from rising serving costs and commoditization of general-purpose models, not to mention that likely none of the foundation model companies has achieved positive profit (not gross) margins. Similarly, 25% growth might be sustainable for a SaaS company with high switching costs and network effects, but for an AI company with standardized APIs and limited enterprise features, this growth rate could quickly erode as customers switch to more efficient or specialized providers.

The key insight is that **foundation model companies need to demonstrate not just growth but sustainable, capital-efficient growth with improving unit economics.** The base multiple must capture both the extraordinary upside of successful AI companies and the brutal economics of getting there.

Putting It All Together

At first glance, the individual components of AI valuation might not seem revolutionary. A drop from 20x to 14x revenue multiples? Manageable. A displacement risk factor of 0.3 versus 0.4? Concerning, but not catastrophic. But when we combine these factors, they reveal valuation gaps so extreme that they fundamentally reshape how we think about an AI company's worth. The multiplication of these effects—base multiples, displacement risk, and revenue scale—creates a mathematical chasm between categories that traditional analysis fails to capture.

Consider the stark reality captured in our framework:

Category	Valuation Multiple Range	D Factor	Revenue Range	Projected Valuation (Median)
Technical Leaders	15-20x	0.1-0.2	$1-10B	$82B
Pre-Displacement	10-14x	0.3-0.5	$100-500M	$2B
Near-Displacement	4-6x	0.6-0.8	$50-100M	$113M
At/Below Open-Source	1-3x	0.9-1.0	<$50M	~$0

Table 5: D factor dramatically affects valuations across different technical categories.

These categories aren't just theoretical—they reflect the brutal mathematics of displacement risk. A Technical Leader commanding a 15-20x multiple might seem only moderately ahead of a Pre-Displacement company at 10-14x. But once you factor in their respective D factors (0.1-0.2 vs 0.3-0.5), the valuation gap explodes. For a Technical Leader generating $3B in revenue, this translates to a $45B valuation. Meanwhile, a Pre-Displacement company with $300M revenue ends up worth just $2B—a difference far more dramatic than their base multiples would suggest.

The effect becomes even more severe as you move down the hierarchy. Near-displacement companies, despite having real revenue and moderate base multiples of 4-6x, see their effective valuations collapse to 1-2x revenue after accounting for their 0.6-0.8 D factors. And for companies below open-source? The D factor of 0.9-1.0 mathematically ensures their equity value approaches zero, regardless of their current revenue or growth rate.

The implications for investors and operators are clear: you must evaluate both valuation multiples and displacement risk to understand true value. A higher base multiple with high displacement risk can be worth far less than a lower multiple with strong technical differentiation. In the AI era, valuation is no longer just about current metrics—it's about your **sustainable distance from the death zone.**

The $300 Billion Question: What is OpenAI Actually Worth?

Our valuation framework points to a sobering reality for AI companies approaching or within the death zone. But what about the current market leader? Is OpenAI worth the $300B that SoftBank is paying for? Or merely the $97B that Elon Musk proposed?

Let me calculate both bull and bear cases for OpenAI using our framework:

Bull Case:

- D Factor: 0.1 (strongest technical lead, GPT-5 coming)
- Multiple: 20x (premium for market leadership, >200% growth)
- 2025 Revenue: $11.6B
- Calculation: $11.6B × 20 × (1-0.1) = $208.8B valuation

Bear Case:

- D Factor: 0.2 (DeepSeek-V4 or Llama-4 narrows gap faster than expected)
- Multiple: 15x (competitive pressure reduces premium)
- 2025 Revenue: $11.6B
- Calculation: $3.7B × 15 × (1-0.2) = $139.2B valuation

Applying our framework to OpenAI reveals a valuation range between $139.2B to $208.8B—positioning our estimate directly between the two major valuation offers the company has received. This substantial range reflects fundamental uncertainty about displacement risk rather than typical bull-bear spread. SoftBank's eye-popping $300B offer appears significantly overvalued through our lens, exceeding even our most optimistic bull case by nearly $100B. Meanwhile, Elon Musk's recent buyout offer of $97B falls well below our bear case, suggesting a much more pessimistic view of OpenAI's technical sustainability.

Notably, even the bear case assumes that OpenAI maintains significant technical leadership. If they slip into pre-displacement territory (D factor 0.3-0.5), the valuation could collapse below $100B. **This mathematical cliff—where small changes in displacement risk create massive valuation swings—explains why investors are so focused on OpenAI's ability to maintain its technical lead.**

A Tale of Five Models: Valuing the Rest of the Field

While OpenAI commands the headlines, our framework reveals equally stark realities for other foundation model companies. Let's apply the same analysis:

Company	D Factor	Multiple	Revenue	Projected Valuation (Median)	Market Valuation	Main Considerations
Anthropic	0.1-0.2	15-20x	$2.2B (2025)	$33B	$61.5B	• Slightly behind OpenAI in cutting-edge performance • Strong enterprise features and differentiation
xAI	0.3-0.5	10-14x	N/A	N/A	$75B	• Revenue unclear, mostly a bet on founder and ecosystem synergy • Unfiltered reasoning
Mistral	0.6-0.8	4-6x	$515M (2025)	$773M	$6B	• Small models still perform the best and the cheapest • May lag behind leading open source model soon
Cohere	0.6-0.8	4-6x	$85M (2025)	$128M	$5.5B	• Focused on enterprise customization, limited visibility on model capabilities • Certain features are competitive

Table 6: Valuation analysis of selected foundation model companies based on our framework. There is a dramatic difference between category leaders and second-tier players.

Our analysis errs on the generous side for second-tier players. For Mistral and Cohere, we've assigned D factors in the 0.6-0.8 range, assuming their specialized features and enterprise traction provide some defense against displacement. Many would argue these companies should already carry a D factor closer to 0.9 or 1.0, given that their base models have been overtaken by open-source alternatives. Only their targeted optimizations for specific use cases and enterprise integration prevent complete displacement—for now.

The gap between our framework's valuations and current market prices is telling. For market leaders, the differential is notable but somewhat rational—OpenAI's $300B valuation and Anthropic's $61.5B valuation represents roughly 1.7x our projected median valuations of $174B for OpenAI and $33B for Anthropic. But **for second-tier players, the disconnect becomes extreme.** Mistral's $6.4B valuation is nearly 8.3x our calculated $773M value, while Cohere's $5.5B represents a staggering 43x premium over our $128M estimate. These gaps suggest the market is systematically underestimating displacement risk, particularly for companies without clear technical leadership.

Of course, the most heated valuation debate centers on the two leaders: Is OpenAI worth $300B, and is Anthropic's $61.5B valuation justified? Our framework provides a starting point, but several critical factors complicate this analysis.

OpenAI's first-mover advantage has created significant brand equity and distribution leverage, particularly in consumer markets. This head start is reflected in their revenue composition: approximately 73% from direct consumer subscriptions and only 27% from API access, while Anthropic shows almost the opposite mix at 85% API, 15% subscriptions. This isn't merely an accounting distinction—it represents divergent bets on where sustainable value lies. OpenAI's direct-to-consumer approach yields higher margins today, while Anthropic's enterprise-focused strategy through AWS sacrifices immediate profitability for deeper integration with mission-critical systems.

These choices create hidden trade-offs that raw valuations obscure. OpenAI's consumer-first approach generates impressive revenue growth but faces potential margin compression through revenue-sharing agreements with Microsoft. Meanwhile, Anthropic's enterprise partnerships with companies like Thomson Reuters create a slower revenue ramp but potentially stronger defensibility against both open source and competitor threats.

Both companies face immense capital requirements, with reinvestment needs that challenge unit economics. Neither business model has proven sustainable at their current valuation levels, especially considering the looming competitive threat from big tech players with a lot more resources—Amazon, Google, Meta, and Microsoft have the capacity to seriously invest in models and potentially reshape competitive dynamics. Grok, despite being a newer entrant, has also demonstrated the ability to raise substantially and compete aggressively. These well-capitalized competitors cast long shadows over pure-play AI companies' valuations.

All in all, our analysis deliberately simplifies complex business models for clarity, but investors should layer in these additional considerations when applying our framework to their own valuations.

The Brutal Truth of Foundation Model Valuations

The evidence is undeniable: maintaining 'Technical Leader' status becomes exponentially harder as the AI race accelerates. Open source alternatives are closing the gap at breathtaking speed—Meta's Llama 4 and DeepSeek-V4 threaten to eliminate the compute advantage entirely. Simultaneously, tech giants like Google and Microsoft are pouring unprecedented resources into AI, creating an infrastructure arms race that smaller players simply cannot match. The moats these companies desperately dig are filling in faster than they can deepen them.

For investors, this suggests a brutal reality: the market likely supports only 2-3 players maintaining the technical lead necessary for premium valuations. Everyone else risks sliding into the death zone—where even impressive revenue growth can't outrun the shadow of displacement risk.

Our framework exposes an uncomfortable reality: Most AI companies are worth dramatically less than their valuations suggest. The combination of displacement risk and traditional metrics creates a mathematical cliff separating sustainable businesses from those racing toward zero value. For foundation model companies, the path forward is ruthlessly narrow. Only those maintaining significant technical leads while building genuine distribution advantages have sustainable futures. For everyone else, strategic pivots aren't options—they're survival imperatives.

In Part II, we'll turn this same analytical lens toward AI application companies—a field that we invest heavily in. While they face different displacement dynamics, the zero-value threshold applies with equal force. The results may be even more sobering as we examine which categories have staying power and which are simply temporary features awaiting absorption by advancing foundation models.

This isn't pessimism but strategic clarity. In an ecosystem defined by binary outcomes, the only unforgivable mistake is misunderstanding the fundamental dynamics that determine whether your AI company ends up at zero—or among the rare heroes capturing enduring value.

Thanks especially to Liang Wu and Chris Jan for detailed feedback on earlier versions of this draft.

FOOTNOTES

1. Of course, open models are also getting better at 'cramming for the leaderboard'—optimizing specifically for benchmark performance. Nevertheless, most of their performance and efficiency gains come from genuine architectural breakthroughs.