Open Model Competition and US AI Valuations under Price Cost and Adoption Scenarios

A Decision Framework for Proprietary Advantage Open Weight Adoption Unit Economics and Funding

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

Open-weight models change the economics of United States artificial-intelligence companies through several channels at once. They provide customers with downloadable weights, self-hosting choices, fine-tuning control and access to a broad supplier ecosystem. They can reduce dependence on a proprietary interface, move spending from model access to infrastructure and engineering, and weaken price premiums for workloads where performance and governance are comparable. They can also enlarge the market for applications, tools, hosting, security, evaluation and specialised deployment. A valuation that treats open models only as a low-price competitor misses these second-order effects.

This paper develops a decision framework for valuing a US AI company under alternative open-model adoption paths. It distinguishes open source, open weights, structured access and proprietary services; maps customer workloads to switching rights and migration cost; reconstructs pricing, inference cost, hosting economics and research capital; and tests whether the company earns value from model scarcity, product integration, distribution, proprietary data, regulated assurance or customer workflow ownership. The framework links technical substitutability to customer retention and collected cash.

Four analytical cases are used: proprietary resilience, hybrid differentiation, open-price convergence and open-first displacement. A wholly hypothetical illustration produces annual revenue from USD 160 million to USD 320 million, contribution from USD 42 million to USD 170 million and cash before financing from negative USD 92 million to positive USD 15 million. A separate probability-weighted illustration produces enterprise value of USD 1,630 million. These figures are management assumptions used to demonstrate the framework. They are not market observations, forecasts or valuation conclusions.

The analysis finds that open-model adoption should be represented as a workload transition rather than a single market-share forecast. The investor should identify which tasks can migrate, the full cost of self-hosting, the time required to validate a replacement, the rights and governance attached to each model, and the value retained by the surrounding product. Value is strongest where the company controls customer workflow, evidence, distribution and recurring contribution while retaining the ability to route between models. Capital and transaction consideration should follow verified cohorts, measured contribution, migration evidence and a funded product and research plan.

JEL Classification: G12, G24, G32, G34, L11, L86, O31, O33

Keywords: open weight models, United States artificial intelligence, AI valuation, model pricing, inference economics, model adoption, technology investment, generative AI

Introduction

Open-weight models have become a material competitive variable for US artificial-intelligence companies. The term describes models whose weights can be downloaded or otherwise made widely

available. It does not establish that training data, source code, documentation or every commercial use is unrestricted. The US National Telecommunications and Information Administration describes a spectrum of openness and recommends analysing the marginal benefits and risks of widely available weights against closed systems and existing technologies [37-40].

The commercial effect depends on the workload. A software team can move a standard summarisation task to a self-hosted model relatively quickly when interfaces, data and evaluation are controlled. A regulated institution may require months of security, validation, legal review and operational acceptance. A consumer product may use an open model while retaining value through distribution, proprietary data and product design. The same open release can therefore reduce one company's price and increase another company's addressable market.

Provider evidence shows why a static assumption is unsafe. Meta describes Llama 4 Scout and Maverick as downloadable open-weight multimodal models with mixture-of-experts architectures [41-43]. DeepSeek reports model architecture, training and commercial-use terms for its V3 release [44,45]. Mistral and other providers continue to release models under differing licences and deployment routes [46]. Stanford's AI Index records fast declines in inference price and changing performance gaps between open and closed models [10,11].

The valuation question is specific. An investor must determine which revenue, contribution, assets and options remain defensible as open models improve and customers gain more deployment choices. The answer requires workload evidence, contract rights, model and system evaluation, full deployment cost, customer migration behaviour and the capital required to keep the product competitive.

1 Define open model competition before modelling value

Open source and open weights should not be treated as synonyms. A release may provide weights while reserving parts of the code, data or documentation. A licence may permit broad commercial use and still impose conditions on redistribution, scale, attribution or acceptable use. Structured access can provide weights to selected researchers or customers without public release. The diligence file should record the exact artefacts, licence version, permitted uses and dependencies for every relevant model [37-40].

The competitor set should be defined at workload level. A general model can appear close on a public benchmark and remain unsuitable for the customer's latency, security, language, tool-use or reliability requirement. A smaller open model can outperform a frontier model on a narrow workflow after fine-tuning. The company should maintain an evaluation matrix that compares complete systems under production conditions and records the cost of reaching an accepted outcome.

The investor should separate direct and indirect competition. Direct competition occurs when the customer replaces the proprietary model or application. Indirect competition occurs when open models reduce supplier prices, strengthen cloud negotiating leverage, accelerate new applications or shift value toward infrastructure, data and workflow providers. The same event can compress interface revenue and improve application contribution.

Market share alone is insufficient. The valuation needs the company's exposure by product, customer, workload, model, deployment route, region and contract. Each row should state whether the workload can move, who has the technical and legal right to move it, the expected validation period and the full recurring cost after migration.

2 Build a US workload and adoption ledger

A workload ledger connects open-model adoption to financial statements. Each row should identify customer, application, task, data classification, current model, alternative model, interface, hosting route, effective price, accepted output, direct cost, engineering support, security review, contract term and

collected cash. The customer and workload remain the unit of analysis because broad product revenue can hide very different substitution risks.

Adoption should be measured through observable stages. Evaluation indicates interest. A validated pilot demonstrates performance under a defined protocol. Production acceptance demonstrates that the system meets customer requirements. Sustained volume, renewal and collected cash demonstrate commercial adoption. Download counts, repository activity and public benchmarks can inform the market map but do not establish enterprise demand for the target.

The ledger should distinguish customer-controlled and vendor-controlled migrations. Customers may route traffic through a model gateway or cloud marketplace. The vendor may lower its own cost by using an open model behind the product. A migration that preserves the company's customer contract and contribution has a different valuation effect from a customer moving the entire workflow to another provider.

US customer cohorts should also be segmented by regulated status, data sensitivity, procurement cycle and deployment location. Federal, healthcare, financial and critical-infrastructure users can require evidence beyond technical quality. Developer customers may move faster when interfaces are portable. The model should apply separate adoption curves and validation cost to each cohort.

3 Test technical substitution under customer conditions

Substitution requires a controlled comparison. The protocol should fix the task population, model version, prompt and tool configuration, hardware, context, latency target, quality measure, human-review policy and failure threshold before the scored run. It should retain raw outputs and explain every exclusion. Public benchmark results provide context; company-specific evidence determines whether a workload can move.

The comparison should include the complete system. Retrieval, tools, guardrails, monitoring, support, integration and human review can determine customer outcomes. An open model that requires additional orchestration or review may have a higher accepted-task cost even when its token price is lower. A proprietary model can lose value when a customer achieves the required outcome with a smaller open model and controlled workflow.

The evaluation should measure quality, reliability, latency, throughput, availability, security, controllability and cost. It should include normal, difficult and safety-relevant cases. Benchmark contamination and adaptation to public leaderboards can weaken reported differences, so the investor should inspect test governance and independent reproduction [10,20,21].

Technical parity should not be assumed to cause immediate revenue loss. Migration can require engineering, security approval, user testing, retraining, documentation and contractual change. The valuation should model the time and cash required to move, plus the probability that the customer completes the transition.

4 Compare proprietary access and self hosted economics

The customer cost comparison should reconcile like-for-like service. Proprietary access can include the model, infrastructure, scaling, updates, safety controls, support and service levels. Self-hosting can include accelerators, cloud or colocation, orchestration, utilisation loss, storage, networking, engineering, monitoring, security, evaluation, incident response and model upgrades. A licence price of zero does not establish a zero-cost service.

The economic unit should be an accepted task under the customer's operating conditions. Cost per token or accelerator hour is an input. The full measure includes unsuccessful outputs, retries, long contexts,

tools, retrieval, human review and downtime. The workload file should show marginal cost, fully allocated cost and the customer's switching investment.

Utilisation materially affects the result. A high and stable workload can support reserved infrastructure or efficient batching. A volatile workload can leave capacity idle. Smaller models, quantisation, mixture-of-experts designs and improved kernels can reduce required compute, while larger contexts and reasoning can increase it [31,32,44,45].

The investor should model where savings are captured. A customer may retain them through self-hosting. A hosting provider may capture them through managed infrastructure. An application company may keep the price stable and expand contribution. Competitive markets may pass most savings to customers. The case must state the evidence for the assumed pass-through.

5 Estimate adoption and price pass through

Open-model adoption can affect volume and price in different directions. Lower access cost can increase experiments and production workloads. It can also reduce proprietary revenue per unit. The company should preserve historical price, usage and migration evidence so that the valuation can separate demand expansion from discounting and product-mix change.

Adoption curves should be built by workload cohort. A developer tool with portable interfaces can move rapidly. A mission-critical deployment can have a long validation period. A company with proprietary data, embedded workflow and accountable service may retain the customer while changing the underlying model. The model should use separate curves for technical evaluation, accepted production, scaled volume and renewal.

Price pass-through should be linked to competition and customer bargaining. Transparent interface prices and close substitutes can accelerate pass-through. Long contracts, differentiated system performance, regulated assurance or workflow switching cost can delay it. A partner can influence economics through marketplace fees, cloud credits, minimum commitments or bundling.

The investor should test a matrix of effective price, retained volume and contribution rather than a single market-share number. Management should identify the operating actions that change the result, including model routing, product packaging, contract design, validation support and infrastructure negotiation.

6 Reconstruct inference unit economics

Inference cost should be measured from workload telemetry and supplier invoices. The bridge includes tokens or other compute units, model mix, accelerator time, utilisation, memory, storage, network, cache writes, cache reads, batch processing, failed requests, safety calls, retrieval, observability and human review. Shared commitments should be allocated using a documented rule.

Capacity can be bought on demand, reserved, provisioned, committed or indirectly through a strategic partner. The economic cost includes consumed and unused capacity. A company can report attractive cost per active request while carrying a large minimum spend that remains outside customer contribution. The model should show both marginal and fully allocated cost.

Efficiency initiatives should be evidenced through production results. Quantisation, distillation, routing, caching, batching, prompt reduction and improved kernels can lower cost. A larger context window, richer output, more reasoning or redundant safety checks can raise it. Provider documentation confirms that cache and batch prices can differ materially from standard calls [12-19]. Realised savings depend on workload eligibility and behaviour.

The unit should connect to customer value. Tokens per dollar are an operating measure. Accepted tasks per dollar, resolved cases per dollar or retained revenue per compute dollar connect the technical system to the commercial model.

7 Treat training as a capital cycle

Frontier training is a sequence of research, data, experiments, pre-training, post-training, evaluation, safety work, deployment and continuing improvement. The valuation should separate exploratory research, platform development, major training runs and product maintenance. Each category has different success probability and useful life.

Scaling-law research shows systematic relationships between compute, model size, data and performance, while compute-optimal training work demonstrates that allocation across model and data can materially affect results [31,32]. These findings support disciplined planning. They do not turn a training budget into a predictable asset.

The capital schedule should include accelerators, cloud commitments, data acquisition, preparation, specialist labour, failed runs, evaluation, security, red-teaming, launch capacity and contingency. It should show when cash is paid, when a model becomes commercially available, which predecessor it replaces and what incremental revenue or contribution it supports.

Management should preserve forecast vintages and experiment records. A valuation can then compare expected and realised cost, schedule, capability and adoption. Persistent overruns or weak commercial conversion should affect the probability, timing and capital charge applied to future training.

8 Estimate model economic life and replacement need

A model does not need to stop functioning for its economic value to decline. Competitors can improve, prices can fall, customers can migrate, a platform can retire the interface, regulation can change or the company's own successor can cannibalise demand. Anthropic's public deprecation record illustrates why model lifecycle belongs in product and valuation analysis [16].

The company should track revenue, usage, contribution and customer count by model version. It should record migration time, incentives, breakage, support cost and churn when a model is replaced. This evidence supports an economic-life estimate for the deployed model and its customer relationships.

The valuation should distinguish the current model asset from the organisation's ability to produce successors. Weights can have a short commercial life. Research systems, data rights, evaluation infrastructure, distribution and a proven team may retain value across generations. The required reinvestment to preserve that capability belongs in cash flow.

An investor can use an obsolescence reserve or explicit replacement schedule. The model should avoid capitalising every research dollar and avoid assuming that all research disappears immediately. Evidence determines which expenditure produces controlled and transferable capability.

9 Test compute access and infrastructure commitments

Foundation-model development and deployment require access to accelerators, networking, storage, power and specialised software. Public filings from Microsoft, Alphabet, Meta and Oracle describe substantial AI and data-centre investment and the effect of infrastructure on cost and capital expenditure [26-29]. These disclosures demonstrate industry capital intensity. They do not disclose the economics of a private target.

The target's compute register should list provider, region, accelerator, capacity, term, price, minimum spend, renewal, portability, security, data location, termination, assignment, change of control and service-

level rights. Strategic investments and credits should be reconciled to legal agreements and usable capacity.

Concentration can create both advantage and risk. Preferential access may support model delivery. A tied cloud commitment can reduce switching flexibility and expose the company to one supplier's roadmap. The FTC and CMA reports provide a framework for examining these relationships [7-9].

The downside case should test delayed capacity, lower utilisation, stranded commitments, loss of credits, price changes, migration cost and inability to assign an agreement after a transaction. It should also test the capital required to move from experimental demand to reliable production.

10 Diligence data rights and provenance

Training and evaluation depend on data whose rights can differ by source and use. The register should cover licensed content, public data, customer data, synthetic data, employee contributions, purchased datasets, scraped material and generated outputs. For each source, the company should record owner, licence, permitted use, territory, term, attribution, confidentiality, privacy, deletion and model-output implications.

Legal analysis is jurisdiction- and fact-specific. The valuation consequence comes from operational dependence. If a material dataset cannot be renewed or used for a successor model, management should estimate replacement time, cost and performance effect. If customer data improves a product, the contract should support that use.

Data quality also matters. Provenance, deduplication, contamination, language coverage, safety filtering and refresh affect model performance and claims. The diligence team should trace key evaluation and training datasets into the reproducible pipeline.

Data option value should remain separate from current cash flow. A collection can support future products when rights, usability, uniqueness and demonstrated demand exist. Volume alone does not establish value.

11 Value talent and the research operating system

Foundation-model capability can depend on a small number of researchers, systems engineers, product leaders and safety specialists. An investor should identify critical teams through code, experiments, incidents, model releases, customer escalations and technical decisions. Job titles and publication counts provide context, while operating evidence shows dependency.

The research operating system includes experiment design, data pipelines, training orchestration, evaluation, review, documentation, security, launch and post-release learning. A repeatable system can retain value when a particular model ages. A founder-dependent process can lose value quickly after a transaction.

Retention analysis should include vesting, liquidity, compensation, immigration, location, non-compete enforceability, academic roles and mission alignment. Transaction structure can protect continuity through rollover equity, retention arrangements, governance and a funded research roadmap. Current legal and tax advice is required.

Knowledge transfer should be tested before attributing platform value. Another authorised team should reproduce a build, run an evaluation, diagnose a failure and execute a controlled release using documented systems.

12 Assess distribution power and model portability

Distribution determines whether open models expand or displace a US AI company. Cloud marketplaces, developer platforms, consumer applications, enterprise suites, system integrators and device channels each

affect discovery, pricing, customer identity, data and support. An open release available across many hosts can widen distribution while making model access less scarce.

The company should document who controls the customer relationship and the technical routing decision. An application provider that can switch among models may protect contribution as model prices change. A model provider that depends on one cloud route can face concentration, settlement and placement risk. A customer that controls a gateway can move volume without renegotiating the application.

Portability should be tested. The diligence team should inspect interface compatibility, prompt and tool dependence, fine-tunes, evaluation evidence, security controls, data location, support, capacity and exit rights. A claimed multi-model architecture has limited value until a representative workload has migrated without material loss of quality or service.

Distribution value belongs in the base case when contracts, observed conversion, settlement and retained contribution support it. Announced availability or a large developer community can support an option case. The valuation should keep that option separate until commercial evidence develops.

13 Reconcile governance safety and liability to economics

NIST's AI Risk Management Framework and generative-AI profile organise risk work through governance, mapping, measurement and management [22,23]. Secure-development guidance extends the evidence needed across software and model lifecycles [24]. The EU AI framework can create obligations based on role, system and use [25]. These sources provide diligence questions; the company's duties require current legal analysis.

Governance cost should appear in the operating model. Evaluation, red-teaming, incident response, abuse prevention, customer assurance, model documentation, security and regulatory work require people and systems. Weak controls can create delayed launches, customer loss, remediation and liability. Effective controls can support access to regulated or high-trust customers.

The target should maintain a claim register linking marketing, benchmark, safety, privacy, copyright, security and performance statements to evidence and an accountable owner. Customer contracts should align with product capability and upstream terms.

The valuation downside should include defined events: model withdrawal, material incident, regulatory restriction, data claim, security breach, customer remediation and service credits. Insurance, contractual limits and reserves should be evaluated against the exposure.

14 Construct four open model adoption scenarios

The proprietary-resilience case assumes that customers retain proprietary systems for differentiated quality, regulated assurance, product integration and reliable capacity. Open models expand the overall market but displace little of the company's high-value workload. Effective price declines slowly and contribution supports much of the research plan.

The hybrid-differentiation case assumes that customers and vendors use a portfolio of proprietary and open models. Standard workloads move to lower-cost models, while specialised and high-consequence tasks remain in controlled services. The company protects value through routing, workflow ownership, customer data and assurance. Price declines and volume grows.

The open-price-convergence case assumes capable open models become acceptable for a broad share of workloads. Proprietary interface prices fall, contract discounts increase and customers gain negotiating leverage. Usage expands but contribution weakens. The company needs external capital to maintain a frontier research programme and must reduce cost or narrow its product scope.

The open-first-displacement case assumes that customers default to open models and buy proprietary access only for defined exceptions. Broad interface volume declines. The company retains selected regulated or performance-sensitive cohorts, reduces frontier expenditure and concentrates on defensible applications and tools. These are analytical states rather than forecasts.

15 Build an integrated US financial model

The financial model should begin with workload cohorts and adoption stages. Revenue equals accepted billable units multiplied by effective price, adjusted for commitments, credits, service levels and partner economics. Direct cost follows each workload through model access or hosting, infrastructure, tools, support, safety and human review.

Migration should appear explicitly. The model should show the share of workload technically eligible to move, validated to move, approved to move and actually moved. It should include one-time engineering and validation cost, contract timing and any period of parallel operation. A workload does not disappear from proprietary revenue on the date an open model is released.

Operating expenditure should separate product delivery, research, sales, governance and major training programmes. Cash flow should include capacity commitments, customer collections, capital expenditure, financing and taxes. The company may benefit from lower inference cost while still requiring substantial capital for research and distribution.

The board should receive price, volume, contribution, research capital, cash before financing and minimum funding for every case. Correlated sensitivities should combine model performance, open adoption, price pass-through, hosting cost and customer retention. Each controllable response should have an owner and evidence deadline.

16 Apply valuation methods to the competitive states

Discounted cash flow can value retained customer cohorts and the capital required to serve them. The forecast should represent open-model adoption, price pass-through, migration timing, research cycles and infrastructure commitments directly. A higher discount rate cannot replace those cash-flow mechanics.

Comparable-company analysis requires careful peer selection. Model developers, cloud platforms, application software companies and infrastructure providers earn value from different layers. Revenue multiples should be adjusted for contribution, capital intensity, customer ownership, growth, concentration and open-model exposure. Reported growth without the cost of customer service and research can overstate comparability.

Replacement cost can inform the value of reproducible code, data, evaluation and operating systems. Historical research spending includes failed and obsolete work. The relevant question is the cost and time required to recreate controlled capability today, including people, rights, infrastructure and customer acceptance.

A sum-of-the-parts approach can separate retained proprietary cohorts, model-neutral applications, hosting or tooling revenue, data assets and development options. Shared costs should be allocated once. Option value should depend on a future decision, required investment and evidence gate. IFRS 13 and IVS provide useful valuation principles, while the transaction conclusion remains purpose-, date- and fact-specific [1-3].

17 Illustrative open model scenario economics

Consider a wholly hypothetical US AI company serving enterprise, developer and application customers. Every figure in this section is a management assumption used to demonstrate the method. The figures do not describe a current company or investment opportunity.

The proprietary-resilience case assumes USD 320 million of annual revenue, USD 170 million of contribution, USD 130 million of training and research expenditure and USD 15 million of cash before financing. The hybrid-differentiation case assumes USD 270 million of revenue, USD 125 million of contribution, USD 145 million of research expenditure and negative USD 45 million of cash before financing.

The open-price-convergence case assumes USD 215 million of revenue, USD 68 million of contribution, USD 135 million of research expenditure and negative USD 92 million of cash before financing. The open-first-displacement case assumes USD 160 million of revenue, USD 42 million of contribution, USD 85 million of research expenditure and negative USD 70 million of cash before financing.

The cases show that open adoption can lower both revenue and direct cost while leaving a material funding requirement. The central question is whether the company retains customer contribution and can adjust research capital before liquidity becomes constrained.

18 Value the evidence states separately

The illustrative valuation separates four evidence states. Retained cohorts contain verified production customers and contribution under current contracts. A differentiated platform adds product integration, workflow evidence and portability. A scaled hybrid platform combines proprietary and open models while preserving customer ownership and contribution. A category network adds durable distribution, proprietary data and financeable strategic options.

For demonstration, retained cohorts are assigned enterprise value of USD 450 million at a probability of 20 per cent. The differentiated platform is assigned USD 1,100 million at 35 per cent. The scaled hybrid platform is assigned USD 2,100 million at 30 per cent. The category network is assigned USD 3,500 million at 15 per cent. The probability-weighted total is USD 1,630 million.

This calculation is not a valuation conclusion. It forces the investment committee to identify the evidence supporting each state. Company-specific probabilities, values and transitions should be developed for the valuation date and should reconcile to the financial model.

19 Structure capital and consideration around adoption evidence

Capital structure should reflect the evidence and cash cycle. Equity can fund uncertain research and product transition. Customer commitments can support production capacity when acceptance and payment are clear. Infrastructure finance may suit controlled equipment or contracted capacity. Debt requires credible servicing cash and downside resilience.

Transaction consideration can protect both parties when open-model adoption is uncertain. Base value can follow retained customer contribution. Deferred consideration can follow renewal, accepted migration or collected cash. Research capital can be released against defined model, product and customer evidence. Seller rollover can preserve alignment where technical and commercial continuity matter.

The documents should define the metric, source, accounting policy, model and product version, measurement period, access right, dispute process and response to changes outside management control. A broad revenue target can reward discounting or uneconomic volume. Contribution and cash measures usually require stronger cost and settlement controls.

The company should maintain enough liquidity to run the selected product and research plan under the downside case. The financing decision should not assume that future investors will fund an unresolved model transition.

20 Govern valuation after investment

Open-model competition requires continuing evidence after closing. The board should review workload migration, price, retained volume, accepted-task cost, contribution, customer concentration, research expenditure, capacity commitments and liquidity. The same definitions used in diligence should continue into reporting.

Management should maintain a model register, workload ledger and approved routing policy. A material model release should trigger controlled evaluation before production use. A customer migration should record quality, cost, security, acceptance and commercial effect. The valuation case should be refreshed when adoption, pricing or capital need moves outside the approved range.

Capital release should follow evidence. The board can stage product, hosting and research expenditure around accepted workloads, customer commitments and available liquidity. A persistent gap between technical progress and retained contribution should lead to a scope, pricing or capital decision.

Governance should also cover licences, data rights, export controls, safety, security and model provenance. Current legal advice is required for the relevant model, transaction and jurisdiction. NTIA, NIST, FTC, DOJ and BIS materials provide diligence questions; they do not replace company-specific legal analysis [7,22-24,37-40,47-50].

21 Decision tests for US investors

An investment committee should convert the scenario analysis into a short set of decision tests. The first test concerns customer control. The company should show which party holds the contract, receives product telemetry, sets price, supports the deployment and collects cash. A model provider can have strong technical performance and weak customer ownership. An application company can retain the customer while changing the underlying model. The valuation should follow the contractual and operating relationship that produces contribution.

The second test concerns substitution evidence. Management should present a representative workload sample, the current system, the proposed open alternative, the locked evaluation protocol, the full deployment cost and the migration timetable. The committee should see accepted quality, failure rates, latency, security findings, engineering effort and customer approval. A benchmark difference without workload evidence cannot support a material value adjustment.

The third test concerns economic capture. The financial model should explain who receives each reduction in model or infrastructure cost. Customers may negotiate lower prices. Hosting providers may retain part of the saving. The company may improve contribution or fund additional product features. The evidence should reconcile supplier invoices, customer pricing, product usage and contribution by cohort. A general assumption that lower compute cost expands margins is insufficient when competition can transfer the saving to customers.

The fourth test concerns capital flexibility. The company should show which research, infrastructure and product commitments are fixed, which can be staged and which can be stopped. It should maintain a funded downside plan for slower adoption, faster price convergence and delayed customer migrations. Liquidity should remain adequate through the decision period without relying on an unspecified future financing.

The fifth test concerns portability. Management should demonstrate at least one controlled migration across models or hosting routes and document the resulting quality, cost, security and customer effect. Portability can protect negotiating leverage and reduce dependence on one model provider. It can also require continuing investment in evaluation, routing and support. The value belongs to the proven operating capability rather than the architectural claim.

The final test concerns governance. Licences, model provenance, data rights, export controls, safety requirements, incidents and customer obligations should remain connected to the model register and workload ledger. An authorised owner should approve material model changes. The board should receive exceptions, financial consequences and remedial actions. These tests create an evidence path from technical competition to price, contribution, liquidity and enterprise value.

Conclusion

Open-model competition changes US AI valuation through workload substitution, price pass-through, deployment choice and market expansion. The investor should model those mechanisms directly. A single open-model market-share forecast cannot represent customer validation, self-hosting cost, contract timing, product integration or the value retained by applications and distribution.

The practical method begins with model and licence classification, then maps customer workloads, tests technical substitution, calculates full deployment economics, measures migration and retention, and funds the product and research plan under coherent scenarios. Existing customer cash flows, required reinvestment and strategic options should remain separate.

Value is more defensible when the company controls the customer workflow, can route across models, retains contribution after price changes, holds transferable rights and governs model use. Transaction consideration and future capital should follow evidence such as accepted deployments, retained cohorts, measured contribution, controlled migrations and collected cash.

Open model competition evidence register

The evidence register should include legal entities, asset ownership, model versions, training runs, data rights, evaluation records, customer contracts, usage, invoices, cash, pricing, capacity commitments, cloud credits, direct cost, distribution terms, incidents, governance controls, financing and forecast vintages. Every investment-committee claim should link to a dated source and accountable owner.

Frequently asked questions

Why can open model competition change a US AI valuation

A software multiple can ignore frontier-training expenditure, infrastructure commitments, model replacement, distribution claims and rapid price competition. Comparables remain useful when product, revenue, margin, capital intensity and risk are aligned.

How should open model adoption be represented in valuation

The model should test price, volume, customer retention, compute cost and competitive pass-through together. Lower cost can improve margin, stimulate usage or be transferred to customers. Company-specific evidence determines the outcome.

Does a free model licence remove infrastructure cost

The model should distinguish expenditure that creates controlled transferable capability from maintenance, failed work and replacement. Accounting treatment follows applicable standards. Economic valuation should reflect useful life, success probability and required reinvestment.

What evidence proves that an open model can substitute for a proprietary model

Open-weight models are a substitution and strategic-option scenario. The analysis should test licence, capability, deployment cost, security, governance, customer acceptance and switching effort for the relevant workflow.

What can protect a proprietary price premium

Strong evidence includes accepted customer outcomes, measurable performance under production conditions, reliable service, defensible rights, low switching risk and retained contribution after discounts and delivery cost.

How should model portability enter valuation

The agreements should be modelled as linked sources of capital, capacity, distribution, revenue share, commitments, control and information rights. Usable cash and contractual obligations should remain separate.

What is the principal downside sensitivity in an open adoption case

The most informative test is usually a correlated case combining price compression, weaker retention, higher research capital and reduced capacity flexibility. The relevant case depends on the company's contracts and operating model.

Which Matchpoint practice is relevant to this work

This research connects to Matchpoint Partners' strategy and execution, valuation, capital-raising, transaction diligence and M&A work for technology businesses, investors and strategic partners.

Appendix B. Workload adoption and pricing file

The pricing file should reproduce customer invoices from workload telemetry. It should identify input, cached input, output, tools, storage, batch, reserved capacity, regional delivery, discounts, credits and service adjustments. It should retain the provider price schedule and contract version used for each period.

Appendix C. Open model substitution dossier

The dossier should contain training scope, data, compute, evaluation, safety, launch criteria, capacity, customer migration, predecessor retirement, cost, schedule and realised commercial outcomes. Forecast and actual evidence should remain separate.

Appendix D. Adoption scenario control sheet

Each case should state effective price, volume, customer retention, model mix, compute cost, utilisation, distribution share, research programme, working capital, financing and probability. Management should document correlations and operating responses.

Appendix E. Transaction evidence room

The evidence room should include corporate records, intellectual-property assignments, model and data documentation, customer and partner contracts, cloud agreements, financial reconciliations, security evidence, regulatory analysis, insurance, people arrangements, financing and the signed valuation model.

Figure 1 Open model competition valuation architecture

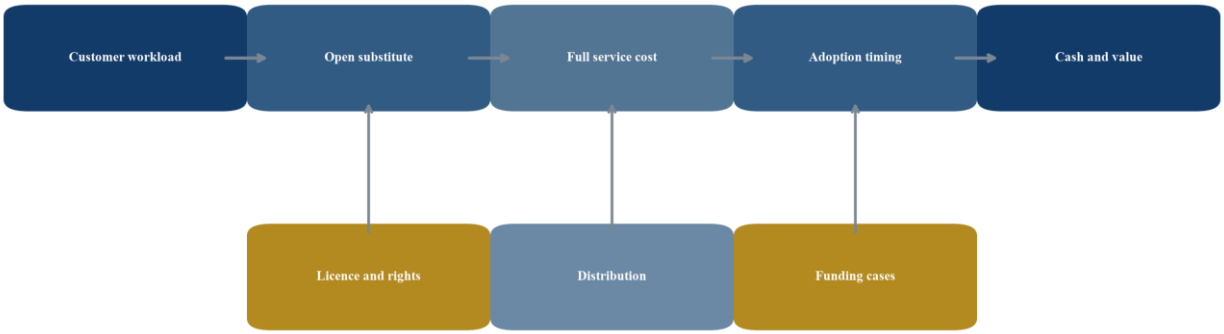

Proposed framework; every conclusion requires company-specific technical commercial and financial evidence.

Figure 2 Hypothetical annual revenue and contribution by open model adoption case

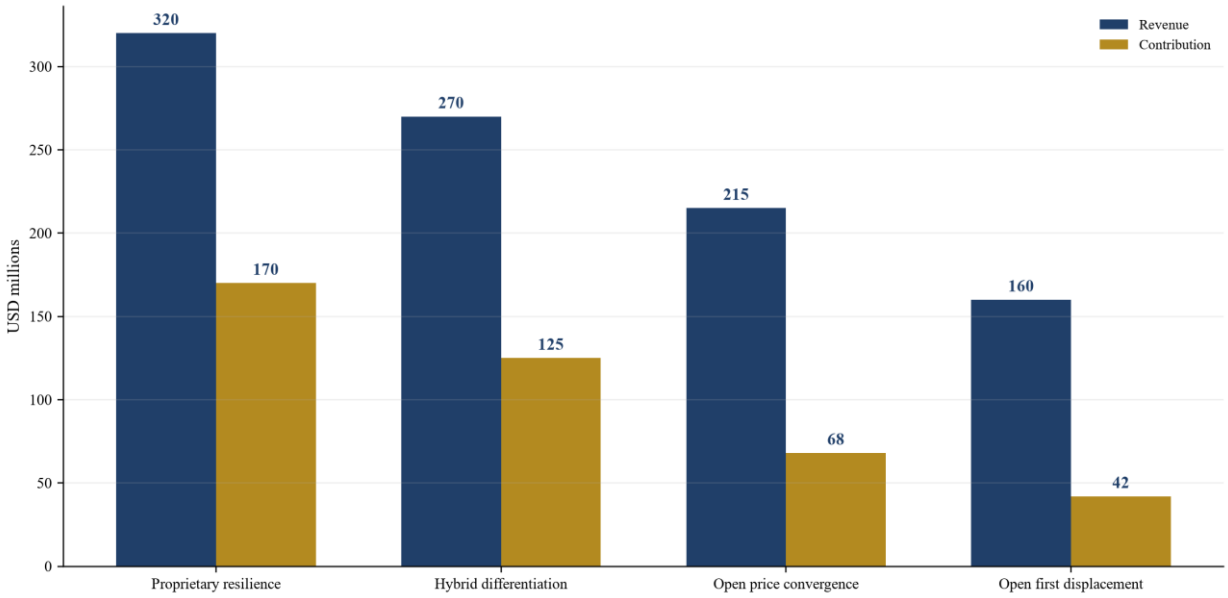

Management assumptions in USD millions; figures are not market observations forecasts or valuation conclusions.

Figure 3 Hypothetical research capital and cash before financing by adoption case

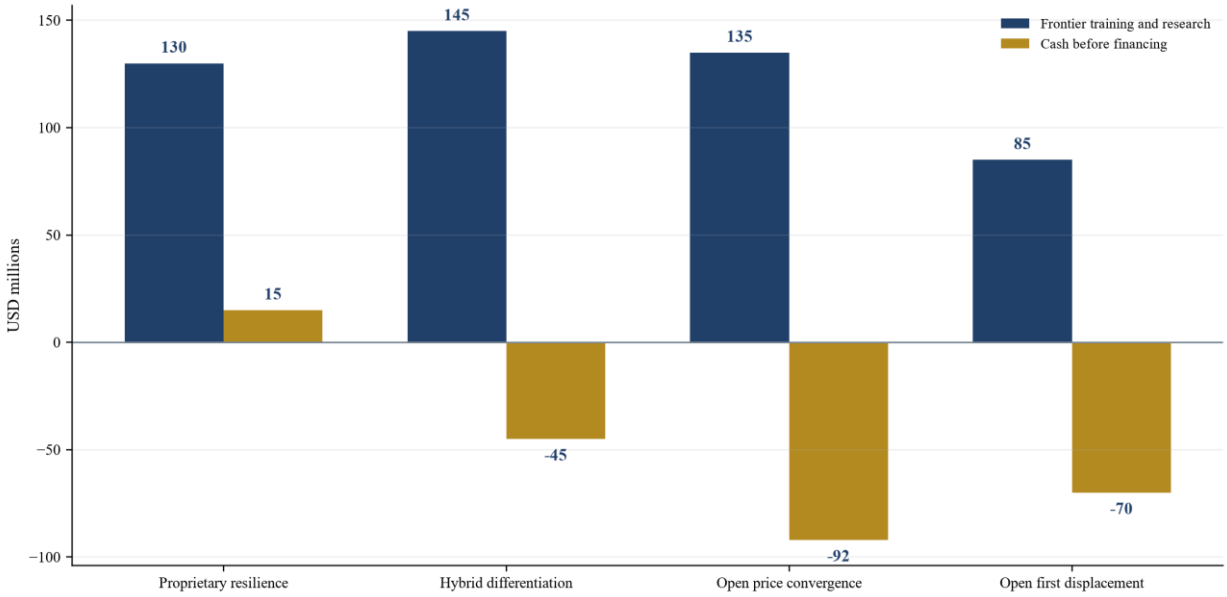

Management assumptions in USD millions; negative cash values indicate external funding need before any minimum liquidity reserve.

Figure 4 Hypothetical probability weighted value by open model evidence state

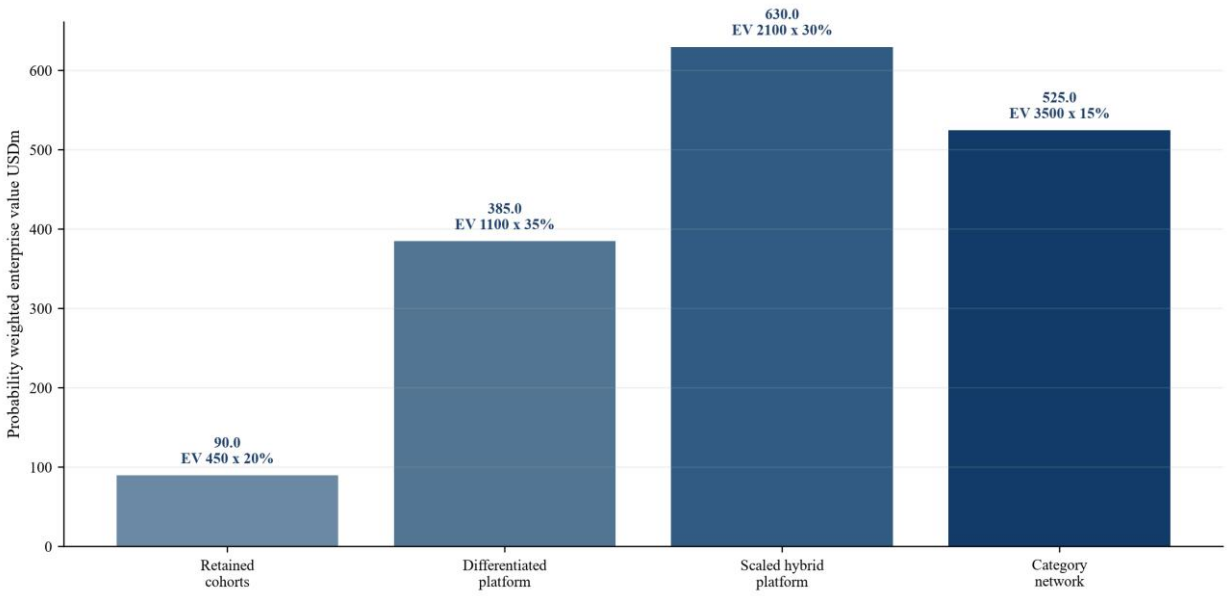

Management assumptions in USD millions; the chart is not a valuation conclusion.

Figure 5 First hundred days open model valuation control sequence

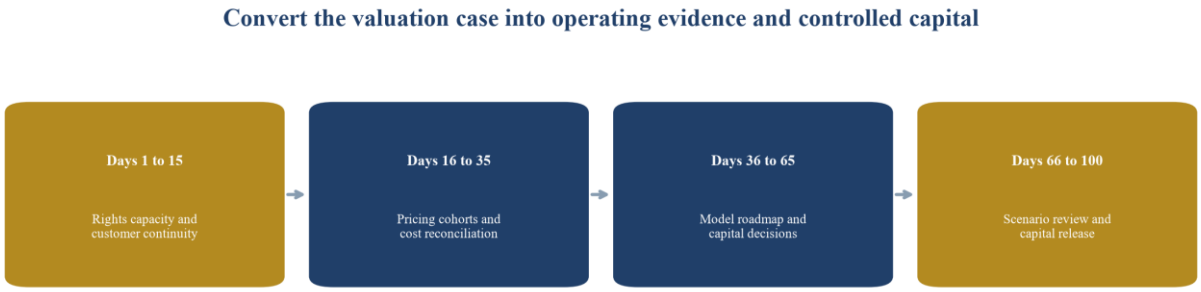

Proposed sequence; timing should follow transaction customer regulatory and financing constraints.

Table 1 Foundation model company perimeter

Component	Evidence	Valuation question	Principal risk
Models and weights	repositories releases hashes and rights	what capability is controlled and transferable	short economic life or restricted rights
Data and evaluation	provenance licences pipelines and tests	can capability be reproduced and improved	unavailable rights contamination or weak evidence
Customer products	contracts telemetry invoices and cash	which workflows create retained contribution	trials credits or bundled usage mistaken for demand
Compute and capacity	agreements utilisation invoices and commitments	can delivery and training be financed	concentration unused capacity and change restrictions
Research system	experiments people processes and forecast history	can the company produce valuable successors	founder dependence and repeated overruns
Distribution	partner agreements placement economics and customer access	who controls demand and price	bundling termination and limited customer ownership

Proposed diligence structure; evidence requirements depend on the company and transaction.

Table 2 Pricing and competition scenario design

Scenario	Effective price	Volume	Customer retention	Research posture
Premium controlled access	high and resilient	focused growth	strong in differentiated workflows	selective frontier investment funded partly from operations
Competitive central	declining with normal discounting	broad growth	mixed by cohort	continuing frontier programme with external capital
Price compression	rapid decline and high pass through	high usage growth	vulnerable developer traffic	expensive programme under funding pressure

Scenario	Effective price	Volume	Customer retention	Research posture
Open substitute	weak broad interface price	lower proprietary volume	retained specialised cohorts	focused models and defensible workflows

Analytical states; these are not forecasts.

Table 3 Customer revenue and cost ledger

Field	Required evidence	Control	Decision use
Contract and billing	signed terms order invoice and recognition policy	reconcile product term and entity	contracted revenue and renewal
Workload and acceptance	telemetry workflow outcome and service evidence	versioned customer and model identifiers	price volume and utility
Effective price	list price discount credits channel share and adjustments	invoice reproduction	price realisation and elasticity
Direct cost	compute data tools support safety and review	invoice and allocation reconciliation	contribution and cash
Collection	receivable settlement and bank receipt	aged balance and exception owner	liquidity and customer quality

Proposed minimum evidence fields for customer-level economic analysis.

Table 4 Hypothetical annual operating cases

Case	Revenue	Contribution	Frontier training and research	Cash before financing
Proprietary resilience	320	170	130	15
Hybrid differentiation	270	125	145	-45
Open price convergence	215	68	135	-92
Open first displacement	160	42	85	-70

Management assumptions in USD millions; figures are not market observations forecasts or valuation conclusions.

Table 5 Hypothetical central contribution bridge

Item	USD millions	Evidence required
Revenue	270	customer workload contract usage invoice and cash ledger
Model access and hosting	-82	workload telemetry supplier invoice and capacity allocation
Migration support and assurance	-25	engineering validation security and customer evidence
Channel and delivery cost	-38	partner settlement implementation and customer support
Contribution before central cost	125	reconciled workload cohort schedule

Management assumptions in USD millions; excludes central research sales administration financing and tax.

Table 6 Hypothetical valuation evidence states

Evidence state	Enterprise value	Probability	Weighted value
Retained cohorts	450.00	20%	90.00

Evidence state	Enterprise value	Probability	Weighted value
Differentiated platform	1100.00	35%	385.00
Scaled hybrid platform	2100.00	30%	630.00
Category network	3500.00	15%	525.00
Total		100%	1630.00

Management assumptions in USD millions; this is not a valuation conclusion.

Table 7 Capital and transaction evidence gates

Gate	Required evidence	Capital or consideration response	Post decision measure
Rights and reproducibility	ownership data rights controlled build and evaluation	closing condition or protected tranche	reproduced release and rights exceptions
Customer utility	accepted workflows retained cohorts and collected cash	base value	retention contribution and collection
Compute economics	workload invoices commitments utilisation and portability	funding reserve or price adjustment	accepted task cost and capacity use
Model roadmap	scoped run budget capacity evaluation and launch plan	staged research capital	cost schedule capability and adoption
Strategic option	defined market right investment and evidence gate	deferred consideration or option budget	exercised right and incremental cash

Proposed framework; actual instruments require current legal tax accounting and financial advice.

References

[1] IFRS Foundation. IFRS 13 Fair Value Measurement. https://www.ifrs.org/issued-standards/list-of-standards/ifrs-13-fair-value-measurement/

[2] IFRS Foundation. Educational material on measuring unquoted equity instruments within IFRS 13. https://www.ifrs.org/content/dam/ifrs/supporting-implementation/ifrs-13/education-ifrs-13-eng.pdf

[3] International Valuation Standards Council. International Valuation Standards. https://www.ivsc.org/standards/

[4] IFRS Foundation. IFRS 15 Revenue from Contracts with Customers. https://www.ifrs.org/issued-standards/list-of-standards/ifrs-15-revenue-from-contracts-with-customers/

[5] IFRS Foundation. IAS 38 Intangible Assets. https://www.ifrs.org/issued-standards/list-of-standards/ias-38-intangible-assets/

[6] IFRS Foundation. IAS 36 Impairment of Assets. https://www.ifrs.org/issued-standards/list-of-standards/ias-36-impairment-of-assets/

[7] US Federal Trade Commission. Partnerships Between Cloud Service Providers and AI Developers. January 2025. https://www.ftc.gov/reports/ftc-staff-report-ai-partnerships-investments-6b-study

[8] UK Competition and Markets Authority. AI Foundation Models Update Paper. April 2024. https://www.gov.uk/government/publications/ai-foundation-models-update-paper

[9] UK Competition and Markets Authority. AI Foundation Models Technical Update Report. April 2024. https://assets.publishing.service.gov.uk/media/661e5a4c7469198185bd3d62/AI_Foundation_Models_technical_update_report.pdf

[10] Stanford Institute for Human-Centered Artificial Intelligence. AI Index Report 2026 Technical Performance. https://hai.stanford.edu/ai-index/2026-ai-index-report/technical-performance

[11] Stanford Institute for Human-Centered Artificial Intelligence. The 2026 AI Index Report. https://hai.stanford.edu/ai-index/2026-ai-index-report

[12] OpenAI. API model and pricing documentation. https://platform.openai.com/docs/models

[13] OpenAI. Batch API reference. https://platform.openai.com/docs/api-reference/batch/object

[14] OpenAI. Scale Tier for API customers. https://openai.com/api-scale-tier/

[15] Anthropic. Claude Platform pricing documentation. https://docs.anthropic.com/en/docs/about-claude/pricing

[16] Anthropic. Model deprecations and lifecycle documentation. https://docs.anthropic.com/en/docs/about-claude/model-deprecations

[17] Google. Gemini Developer API pricing. https://ai.google.dev/gemini-api/docs/pricing

[18] Amazon Web Services. Amazon Bedrock pricing. https://aws.amazon.com/bedrock/pricing/

[19] Amazon Web Services. Prompt caching for faster model inference. https://docs.aws.amazon.com/bedrock/latest/userguide/prompt-caching.html

[20] MLCommons. MLPerf Inference benchmarks. https://mlcommons.org/working-groups/benchmarks/inference/

[21] MLCommons. MLPerf Inference datacenter benchmark. https://mlcommons.org/benchmarks/inference-datacenter/

[22] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework. https://www.nist.gov/itl/ai-risk-management-framework

[23] National Institute of Standards and Technology. Generative Artificial Intelligence Profile NIST AI 600-1. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf

[24] National Institute of Standards and Technology. Secure Software Development Practices for Generative AI and Dual-Use Foundation Models. https://www.nist.gov/publications/secure-software-development-practices-generative-ai-and-dual-use-foundation-models-ssdf

[25] European Commission. Regulatory framework for artificial intelligence. https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

[26] Microsoft Corporation. Annual Report on Form 10-K for fiscal 2026. https://www.sec.gov/Archives/edgar/data/789019/000119312526323660/msft-20260630.htm

[27] Alphabet Inc. Annual Report on Form 10-K for 2025. https://www.sec.gov/Archives/edgar/data/1652044/000165204426000018/goog-20251231.htm

[28] Meta Platforms Inc. Annual Report on Form 10-K for 2025. https://www.sec.gov/Archives/edgar/data/1326801/000162828026025534/meta-12312025x10kars.htm

[29] Oracle Corporation. Annual Report on Form 10-K for fiscal 2026. https://www.sec.gov/Archives/edgar/data/1341439/000119312526277521/orcl-20260531.htm

[30] Financial Accounting Standards Board. Accounting Standards Codification Topic 606. https://asc.fasb.org/topic&trid=2129938

[31] Hoffmann J and others. Training Compute-Optimal Large Language Models. 2022. https://arxiv.org/abs/2203.15556

[32] Kaplan J and others. Scaling Laws for Neural Language Models. 2020. https://arxiv.org/abs/2001.08361

[33] FinOps Foundation. FinOps Framework 2025. https://www.finops.org/wp-content/uploads/2025/05/English-FinOps-Framework-2025.pdf

[34] FinOps Open Cost and Usage Specification. Calculate unit economics. https://focus.finops.org/docs/use-cases/v1-4/calculate-unit-economics/

[35] World Intellectual Property Organization. Artificial Intelligence and Intellectual Property. https://www.wipo.int/en/web/frontier-technologies/artificial-intelligence/index

[36] US Securities and Exchange Commission. AI Roundtable remarks. 27 March 2025. https://www.sec.gov/newsroom/speeches-statements/crenshaw-remarks-ai-roundtable-032725

[37] US National Telecommunications and Information Administration. Dual-Use Foundation Models with Widely Available Model Weights Report. July 2024. https://www.ntia.gov/programs-and-initiatives/artificial-intelligence/open-model-weights-report

[38] US National Telecommunications and Information Administration. Spectrum of Model Openness and Background. https://www.ntia.gov/programs-and-initiatives/artificial-intelligence/open-model-weights-report/background

[39] US National Telecommunications and Information Administration. Risks and Benefits of Widely Available Model Weights. https://www.ntia.gov/programs-and-initiatives/artificial-intelligence/open-model-weights-report/risks-benefits-of-dual-use-foundation-models-with-widely-available-model-weights

[40] US National Telecommunications and Information Administration. Policy Recommendations for Open Foundation Models. https://www.ntia.gov/programs-and-initiatives/artificial-intelligence/open-model-weights-report/policy-approaches-recommendations/recommendations

[41] Meta. The Llama 4 Herd. April 2025. https://ai.meta.com/blog/llama-4-multimodal-intelligence/

[42] Meta. Llama Developer Documentation and Resources. https://ai.meta.com/llama/get-started/

[43] Meta. Llama Responsible Use Guide. https://ai.meta.com/llama/responsible-use-guide/

[44] DeepSeek AI. DeepSeek V3 Technical Report and Repository. https://github.com/deepseek-ai/DeepSeek-V3

[45] DeepSeek AI. DeepSeek V3 Weight File Documentation. https://github.com/deepseek-ai/DeepSeek-V3/blob/main/README_WEIGHTS.md

[46] Mistral AI. Models and developer documentation. https://docs.mistral.ai/getting-started/models/models_overview/

[47] US Department of Justice and Federal Trade Commission. Joint Statement on Competition in Generative AI Foundation Models and AI Products. July 2024. https://www.justice.gov/atr/media/1361306/dl?inline=

[48] US Department of Justice Antitrust Division. AI Workshop Summary Report. 2025. https://www.justice.gov/atr/media/1385586/dl?inline=

[49] US Bureau of Industry and Security. Policy Statement on Advanced Computing Items Used to Train AI Models. May 2025. https://www.bis.gov/media/documents/ai-policy-statement-training-ai-models-may-13-2025

[50] US Bureau of Industry and Security. Export Administration Regulations Part 748. https://www.bis.gov/regulations/ear/748

About the Author

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds — bridging the boardroom view to hands-on execution and close.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners he is Managing Partner, leading the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.

https://www.linkedin.com/in/ckadya/

https://www.matchpoint-partners.com/team/ck-adya.html

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.