

■ AI & VALUATION · 2026

Earning the AI Premium

*25 Questions Public Companies Should Be Ready to
Answer*

How to connect rising AI costs to productivity, revenue, and EPS.

OPENING STATEMENT

AI spending is climbing faster than almost any other line on the corporate income statement, yet the return on it is rarely disclosed in terms investors can actually evaluate.

As the numbers grow, the market will reward not the companies that spend the most on AI, but the ones that can show that spending converts into measurable productivity, revenue, and EPS.

The 25 questions that follow — and the scorecard that closes the piece — are how to tell durable, margin-expanding AI economics from expensive experimentation.

■ THE PREMISE



CONTENTS

25 Questions Public Companies Should Be Ready to Answer

01	What is the unit economics of each major AI workload?	14	How concentrated is AI consumption among customers and use cases?
02	How much AI usage is inference versus training, fine-tuning, and experimentation?	15	What are the reliability and failure characteristics of the AI system?
03	What is the company's model-routing strategy?	16	How much technical debt is being created by rapid AI deployment?
04	How exposed is the company to model-price compression or supplier price increases?	17	How does the company prevent sensitive data from leaking through AI systems?
05	What percentage of the AI stack is proprietary?	18	How is the company managing model drift and unexpected behavioral changes?
06	Does greater usage improve the company's AI product?	19	What is the expected economic loss from AI errors?
07	How is the company measuring model quality in production?	20	Can the company demonstrate incremental causality?
08	What proportion of AI output can be delivered without human review?	21	What happens to revenue and cost when models improve by another order of magnitude?
09	Where are the bottlenecks preventing AI from scaling?	22	What is the company's build-versus-buy decision framework?
10	What is the utilization rate of committed AI infrastructure?	23	Is AI reducing headcount, avoiding future hiring, or increasing output?
11	How much latency does the application require, and what does achieving it cost?	24	What percentage of AI savings is being retained versus reinvested?
12	What is the marginal gross margin of AI-driven revenue?	25	Who owns AI risk and performance inside the company, and how is it governed?
13	Is pricing aligned with the company's AI cost structure?		

01

01 / 25

What is the unit economics of each major AI workload?

Management should disclose the cost per completed business outcome—not merely tokens consumed. Depending on the application, that might mean cost per resolved support case, software feature shipped, qualified sales lead, document processed, transaction reviewed, or customer retained.

WHAT SUPPORTS A PREMIUM VALUATION:

A credible disclosure showing that unit costs are falling quarter over quarter while quality and usage rise. Investors will pay more for evidence of an AI learning curve: declining inference cost, improving task-completion rates, and expanding gross profit per workflow.

02

02 / 25

How much AI usage is inference versus training, fine-tuning, and experimentation?

These categories have very different economics. Production inference can become recurring cost of revenue, while model training, fine-tuning, evaluations, and pilots may behave more like research and development.

WHAT SUPPORTS A PREMIUM VALUATION:

A clear breakdown showing that a growing share of expenditure is tied to scaled production workloads with identifiable revenue or savings. A company deserves a higher multiple when AI spending is migrating from open-ended experimentation into repeatable, economically productive use.

■ PREMIUM SIGNAL

A company deserves a higher multiple when AI spending is migrating from open-ended experimentation into repeatable, economically productive use.

03

03 / 25

What is the company's model-routing strategy?

Does every request go to an expensive frontier model, or does the system route tasks across frontier, smaller, specialized, open-weight, and internally developed models according to complexity?

WHAT SUPPORTS A PREMIUM VALUATION:

Evidence that the company dynamically chooses the lowest-cost model capable of meeting a defined quality threshold. Useful disclosures include the percentage of workloads handled by smaller models, cost savings from routing, and whether model substitutions can be made without rebuilding the product.

This directly addresses the episode's concern that token expenses can grow much faster than the productivity benefit.

04

04 / 25

How exposed is the company to model-price compression or supplier price increases?

AI economics can change quickly. A company may benefit when models become cheaper, but it may also be vulnerable if a dominant provider raises prices, limits access, or changes contractual terms.

WHAT SUPPORTS A PREMIUM VALUATION:

Low switching costs, multi-model compatibility, negotiated capacity, portability of prompts and workflows, and limited dependence on any single model provider. Investors should reward companies that can capture falling model costs rather than having the savings absorbed by vendors.

05

05 / 25

What percentage of the AI stack is proprietary?

Investors should distinguish between value created by a third-party model and value created by the public company's own data, workflows, retrieval systems, evaluations, user experience, and distribution.

WHAT SUPPORTS A PREMIUM VALUATION:

Evidence that the company owns a defensible layer above the foundation model—for example:

- Unique, permissioned data that improves outcomes;
- Proprietary workflow integrations;
- Domain-specific evaluations and feedback loops;
- Customer-specific memory or context;
- Exclusive distribution or embedded user behavior.

The premium should come from assets that persist even when the underlying model changes.

06

06 / 25

Does greater usage improve the company's AI product?

Traditional software does not necessarily improve when more customers use it. An AI product may develop a stronger moat when usage generates evaluation data, corrections, workflow knowledge, or outcome feedback.

WHAT SUPPORTS A PREMIUM VALUATION:

A measurable compounding loop: more usage produces better data, better data improves accuracy or automation, improved performance drives greater adoption, and adoption creates still more useful data. Management should quantify improvement in success rates, accuracy, personalization, or cost as the feedback dataset grows.

07

07 / 25

How is the company measuring model quality in production?

Benchmark scores are insufficient. Investors need to know whether the system performs reliably on the company's actual tasks and whether quality is improving or deteriorating after model or prompt changes.

WHAT SUPPORTS A PREMIUM VALUATION:

Disclosure of business-specific evaluation metrics, such as:

- Task-completion or resolution rate;
- Human override and escalation rate;
- Error or hallucination rate;
- Customer acceptance rate;
- Time saved per task;
- Accuracy on high-value edge cases;
- Performance by customer or workflow segment.

A premium is warranted when management can demonstrate disciplined, repeatable measurement rather than anecdotal success.

08

08 / 25

What proportion of AI output can be delivered without human review?

A productivity tool that drafts work but requires extensive review may save less money than headline adoption suggests. The economically important measure is the percentage of tasks completed autonomously at an acceptable quality and risk level.

WHAT SUPPORTS A PREMIUM VALUATION:

Rising "straight-through processing" or autonomous-completion rates, accompanied by stable or improving quality. Investors should look for increasing output per employee, lower exception-handling costs, and a declining need for human review.

The company should avoid presenting autonomy as a single companywide percentage; it should disclose it by workflow and risk class.

09

09 / 25

Where are the bottlenecks preventing AI from scaling?

The limiting factor may not be the model. It could be data quality, systems integration, latency, security approvals, change management, insufficient evaluations, or the inability to redesign workflows.

WHAT SUPPORTS A PREMIUM VALUATION:

Management identifies the constraints precisely, provides milestones for removing them, and demonstrates that deployment times are shortening. A particularly positive signal is a reusable internal platform that allows the second, third, and tenth AI application to be deployed faster and more cheaply than the first.

10

10 / 25

What is the utilization rate of committed AI infrastructure?

Companies may reserve GPUs, cloud capacity, data-center power, or inference commitments well ahead of actual demand. Low utilization can turn a strategic investment into stranded cost.

WHAT SUPPORTS A PREMIUM VALUATION:

High and improving utilization, flexible capacity contracts, limited take-or-pay exposure, and a close match between capacity additions and monetizable demand. Management should disclose committed capacity, actual consumption, unit cost, and the expected time to productive utilization.

For hyperscalers and infrastructure-intensive companies, this question is as important as conventional data-center utilization.

11

11 / 25

How much latency does the application require, and what does achieving it cost?

Some applications can wait minutes; others need responses in milliseconds. Low-latency systems may require more expensive models, reserved hardware, edge infrastructure, or duplicated capacity.

WHAT SUPPORTS A PREMIUM VALUATION:

Evidence that the company has optimized the cost–latency trade-off by workflow. Investors should reward businesses that meet customer requirements without overengineering every interaction to frontier-level speed or quality.

12

12 / 25

What is the marginal gross margin of AI-driven revenue?

Reported "AI revenue" may have materially different margins from conventional software revenue because of inference, retrieval, storage, human review, and third-party model costs.

WHAT SUPPORTS A PREMIUM VALUATION:

AI revenue with positive and expanding contribution margins.
Management should disclose gross margin after all directly attributable AI costs, not simply label existing subscriptions as AI-enabled revenue.
The strongest signal is that AI gross margins improve with scale rather than deteriorate as usage increases.

■ PREMIUM SIGNAL

The strongest signal is that AI gross margins improve with scale rather than deteriorate as usage increases.

13

13 / 25

Is pricing aligned with the company's AI cost structure?

Seat-based pricing can be dangerous when a small number of users generate extremely high inference costs. Unlimited usage can also cause revenue and cost to diverge.

WHAT SUPPORTS A PREMIUM VALUATION:

Pricing that captures customer value while protecting contribution margin—for example, consumption pricing, outcome-based pricing, tiered allowances, premium automation features, or contractual pass-throughs for unusually expensive workloads.

Investors will award a premium when increased product engagement predictably increases gross profit rather than merely increasing compute expense.

14

14 / 25

How concentrated is AI consumption among customers and use cases?

Average economics can conceal loss-making power users, subsidized pilots, or revenue concentrated in a few large contracts.

WHAT SUPPORTS A PREMIUM VALUATION:

Diversified adoption, profitable expansion across customer cohorts, and evidence that the economics work for ordinary customers—not only a handful of strategic design partners. Disclosures should include consumption concentration, renewal rates, expansion rates, and contribution margin by major workload class.

15

15 / 25

What are the reliability and failure characteristics of the AI system?

AI applications may fail probabilistically rather than deterministically. Investors should understand uptime, model-provider outages, fallback systems, degradation under heavy load, and the financial consequences of incorrect outputs.

WHAT SUPPORTS A PREMIUM VALUATION:

Strong service levels, automated fallbacks, multi-provider redundancy where warranted, low severity-adjusted failure rates, and rapid incident recovery. The valuation benefit comes from demonstrating that AI is dependable enough to become mission-critical infrastructure rather than an optional feature.

16

16 / 25

How much technical debt is being created by rapid AI deployment?

Companies can accumulate thousands of prompts, agents, model integrations, data pipelines, and one-off workflows that become expensive to monitor and maintain.

WHAT SUPPORTS A PREMIUM VALUATION:

A standardized AI platform with common identity controls, model gateways, evaluation tooling, observability, data permissions, and deployment processes. Reuse matters: management should quantify how much of each new application is built on shared infrastructure.

A premium is more defensible when the company is building a scalable platform rather than a collection of disconnected demonstrations.

17

17 / 25

How does the company prevent sensitive data from leaking through AI systems?

The question should cover training-data permissions, customer data, employee prompts, retrieval systems, model-provider retention policies, and cross-customer data isolation.

WHAT SUPPORTS A PREMIUM VALUATION:

Clear data provenance, contractual control of provider data use, permission-aware retrieval, strong tenant isolation, auditable access controls, and a low rate of material data incidents. These controls reduce the probability that AI growth creates future legal, customer-retention, or remediation costs.

18

18 / 25

How is the company managing model drift and unexpected behavioral changes?

Outputs may change when providers update models or when the distribution of real-world inputs changes. A system that worked last quarter may not perform identically today.

WHAT SUPPORTS A PREMIUM VALUATION:

Continuous evaluations, version control, staged deployments, regression testing, rollback capabilities, and monitoring of production outcomes. Investors should reward a company that treats models as dynamic dependencies rather than static software components.

19

19 / 25

What is the expected economic loss from AI errors?

An error in a marketing draft is different from an error in underwriting, medical workflow, industrial control, cybersecurity, or financial reporting. A raw error rate does not capture the severity of potential loss.

WHAT SUPPORTS A PREMIUM VALUATION:

Risk-adjusted metrics that combine frequency and severity, with explicit human-review thresholds for high-consequence decisions. Management should disclose material incidents, near misses, loss rates, and the controls used to keep expected losses below the economic benefit of automation.

20

20 / 25

Can the company demonstrate incremental causality?

Productivity may improve for reasons unrelated to AI: organizational restructuring, lower hiring, better demand, software upgrades, or normal process improvement. Investors need evidence that the AI system caused the claimed benefit.

WHAT SUPPORTS A PREMIUM VALUATION:

Controlled rollouts, holdout groups, before-and-after cohort analysis, or other credible attribution methods. Claims such as "developers are 20% more productive" are more valuable when they are tied to shipped output, quality, and controlled comparisons—not surveys or self-reported time savings.

21

21 / 25

What happens to revenue and cost when models improve by another order of magnitude?

This is the strategic sensitivity question. Better models may make the company's product more valuable, but they could also commoditize its features or allow customers to build substitutes internally.

WHAT SUPPORTS A PREMIUM VALUATION:

A credible explanation of why frontier-model improvement expands the company's addressable market, automation rate, pricing power, or margins. Investors should discount companies whose differentiation disappears as foundation models improve.

22

22 / 25

What is the company's build-versus-buy decision framework?

Internal development can create differentiation, but it can also produce unnecessary cost. Buying third-party tools can accelerate deployment but may reduce control and defensibility.

WHAT SUPPORTS A PREMIUM VALUATION:

Management builds only where proprietary data, workflow knowledge, scale, or customer requirements provide an advantage, while purchasing commoditized capabilities. Capital discipline—not the percentage built internally—is the premium signal.

■ PREMIUM SIGNAL

Capital discipline—not the percentage built internally—is the premium signal.

23

23 / 25

Is AI reducing headcount, avoiding future hiring, or increasing output?

These have different financial implications. Avoided hiring may not appear immediately as a restructuring benefit, while increased output may support revenue without reducing current expenses.

WHAT SUPPORTS A PREMIUM VALUATION:

A transparent bridge showing:

- Positions eliminated;
- Vacancies not backfilled;
- Future hiring avoided;
- Output per employee gained;
- Revenue capacity added;
- Restructuring or transition costs incurred.

The highest-quality outcome is generally durable operating leverage: revenue grows faster than both headcount and AI expenditure.

24

24 / 25

What percentage of AI savings is being retained versus reinvested?

A company may realize productivity gains but immediately spend them on additional models, staff, or experiments. That may be strategically sensible, but it will not yet produce margin expansion.

WHAT SUPPORTS A PREMIUM VALUATION:

Management separates gross savings, reinvestment, and net financial benefit. Investors are more likely to grant a premium when reinvestment is deliberate, time-bounded, and associated with specific future revenue or margin milestones.

25

25 / 25

Who owns AI risk and performance inside the company, and how is it governed?

As AI moves from isolated pilots into core operations, accountability can become diffuse—spread across engineering, individual business units, data teams, and risk functions with no single owner. Without clear decision rights and consistent standards, one team's model choice, prompt change, or vendor commitment can create cost, quality, or compliance consequences for the entire enterprise.

WHAT SUPPORTS A PREMIUM VALUATION:

Defined ownership of AI outcomes and risk, with decision rights, escalation paths, and standards applied consistently across business units rather than reinvented team by team. Management should be able to say who approves high-consequence deployments, how performance and incidents are reviewed, and how AI investment is prioritized against a common bar. The premium signal is disciplined governance that scales with adoption—not oversight that lags the pace of deployment.


■ PREMIUM SIGNAL

The premium signal is disciplined governance that scales with adoption—not oversight that lags the pace of deployment.

■ THE THESIS

What will drive the largest valuation premium

The market is unlikely to award the greatest premium to the company with the highest AI spending or even the largest number of AI users. The strongest premium should accrue to a company that can demonstrate five things simultaneously:



01

Falling unit cost:

AI becomes cheaper per successful outcome.

02

Improving quality and autonomy:

More work is completed correctly without human intervention.

03

Expanding contribution margin:

Revenue or savings grow faster than inference and infrastructure costs.

04

Portability and supplier leverage:

The company can change models and capture industrywide cost declines.

05


A compounding proprietary moat:

Usage produces unique data, workflow knowledge, or distribution advantages that competitors cannot easily reproduce.

EXHIBIT

AI Economics Scorecard

CATEGORY	RECOMMENDED DISCLOSURE
AI investment	Total AI OpEx, CapEx, and contractual commitments
Production maturity	Percentage spent on production versus pilots and research
Consumption	Tokens, model calls, GPU hours, or another relevant workload measure
Unit economics	Cost per completed business outcome
Business output	Revenue generated, costs removed, hours saved, or capacity added
Quality	Success, error, override, escalation, and customer-acceptance rates
Autonomy	Percentage completed without human intervention
Financial return	Gross benefit, net benefit, contribution margin, and EPS impact
Architecture	Model mix, routing strategy, switching capability, and vendor concentration
Defensibility	Proprietary data, workflow assets, feedback loops, and distribution advantage
Risk	Material incidents, expected-loss measures, and human-review requirements
Forward indicators	Twelve-month milestones for unit cost, adoption, revenue, and margin

A series of concentric circles in a light teal color, centered on the page, creating a ripple effect.

■ EARNING THE AI PREMIUM

25 Questions Public Companies Should Be Ready to Answer