

The AI Report: AI Pricing Through the Buyer's Lens

2026



Table of Contents

01	Executive Summary <i>The headline patterns shaping how buyers experience AI pricing today.</i>	Pages 03-05
02	Demographics <i>Who answered - 296 buyers across industries, sizes, and roles.</i>	Pages 06-07
03	Pre-Contract Evaluation <i>Why AI pricing is hard to evaluate before signing.</i>	Pages 08-19
04	AI Spend Experience <i>What happens to budgets after the contract - who overspends, why, and who owns the risk.</i>	Pages 20-33
05	AI Pricing Sentiment <i>Which models buyers love, like, tolerate, or avoid and why.</i>	Pages 34-42
06	AI versus Traditional SaaS <i>What buyers treat differently when AI is involved.</i>	Pages 43-49

01

Executive Summary

*The headline patterns shaping how buyers experience
AI pricing today.*

Pages

03-05

Executive summary

The AI software market is shifting away from seat-based subscriptions toward consumption-based pricing models - credits, tokens, usage tiers, and outcome-based fees - under the rationale that AI's variable cost structure requires a different commercial approach than traditional SaaS. This study examines how that shift is being received by the buy side.

Pricing I/O, in partnership with Benchmarkit, surveyed 296 software buyers between February and April 2026. Respondents ranged from companies under \$20M in revenue to enterprises above \$10B, and 90% held Director-level titles or above. The survey examined four areas: how buyers evaluate AI pricing pre-contract, how AI spend has performed against budget, which pricing models buyers prefer, and how AI pricing is perceived relative to traditional SaaS. The findings indicate that buyer sentiment is not aligned with the direction the market is moving.

The report is organized around three themes.

- 1. Evaluation:** Buyers struggle to forecast AI cost before signing. Seat-based is the only model with pre-contract confidence (43% call it easiest to evaluate, 14% call it hardest), and this holds across all revenue bands. Every other model becomes harder to evaluate when AI is involved: 30% find hybrid more difficult to evaluate for AI than for SaaS, rising through seat-based (33%), outcome (37%), usage (44%), and credit/token (55%). Credit/token is the most affected model - nearly double the least-affected and the difficulty is structural rather than tied to buyer sophistication, with similar sentiment across all revenue bands (54-59%). Predictability, not price, is the primary evaluation criterion: 68% rank predictable total cost in their top three priorities; lowest entry price ranks last at 19%.
- 2. Transparency:** The leading concerns are cost unpredictability, pricing/usage-to-cost translation, and lack of usage transparency. Trust-related concerns such as hidden markups and unexpected fees rank lower. In open-ended responses, 38% asked for more predictability and 35% for more transparency; only 10% asked for stronger value or outcome alignment.
- 3. Protection:** 89% of buyers have exceeded their initial AI budget - 45% significantly, 44% moderately, with only 9% on budget. Overruns are attributed primarily to forecasting failures rather than vendor behavior: AI features driving additional usage (67%), usage scaling faster than expected (63%), and difficulty forecasting usage (38%). Only 10% cited post-sale changes to vendor pricing or terms. IT is named as the primary owner of AI cost risk by 67% of buyers, compared to 17% for Finance, and IT ownership grows with company size. Buyers prefer guardrails - soft caps with approval, predictive alerts, monthly/quarterly spend limits - over hard cutoffs or pre-purchased credit pools, and the most effective guardrail varies by concern.

The pricing models with the highest buyer preference are those anchored on a fixed metric: seat-based has the highest net preference at +17. The models with the lowest preference are the two pure variable models: pure usage at -7 and pure outcome at -19. This pattern indicates that buyers will accept higher prices in exchange for predictability and will discount models without a fixed anchor, regardless of how the variable component is defined.

Key statistics from the study

The headline numbers from the survey, organized under the three patterns that shape how buyers experience AI pricing today.

01



EVALUATION

Buyers struggle to assess AI pricing before they sign.

- 43%** say seat-based is the easiest pricing model to evaluate pre-contract
- 55%** find credit and token pricing harder to evaluate for AI than for traditional SaaS
- 44%** find usage-based pricing harder to evaluate for AI than for SaaS
- 68%** rank predictable total cost as a top-3 priority when evaluating an AI vendor
- 19%** rank lowest entry price as a top-3 priority - the lowest ranked criterion in the study

02



TRANSPARENCY

Buyers want clear pricing mechanisms, not creative ones.

- 38%** rank cost unpredictability as their top concern about AI pricing
- +17** net preference for seat-based - the highest of any AI pricing model tested
- 19** net preference for pure outcome-based pricing - the lowest of any model
- 38%** of open-ended responses asked for more predictability and stable costs
- 10%** asked for stronger value or outcome alignment

03



PROTECTION

Buyers want guardrails that prevent overruns, not hard cutoffs.

- 89%** have exceeded their initial AI budget
- 67%** say AI features driving additional usage was the top contributor to overrun
- 10%** blame vendor pricing or terms changing post-sale
- 67%** name IT as the primary owner of AI cost risk
- 62%** want soft caps with alerts and approval workflows

02

Demographics

Who answered the survey - 296 buyers across industries, sizes, and roles.

Pages

06-07

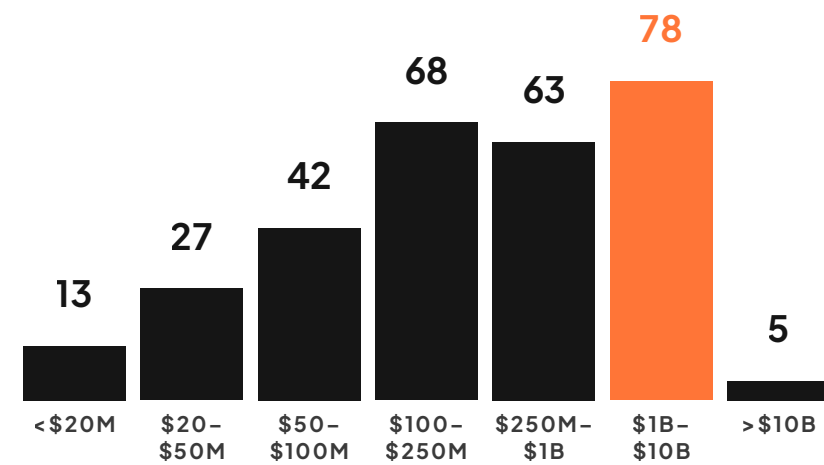
Survey Participants




296 Participants

across industries, company sizes, and seniority levels - all involved in evaluating, approving, or owning software purchases.


1. By ARR segment



2. By title



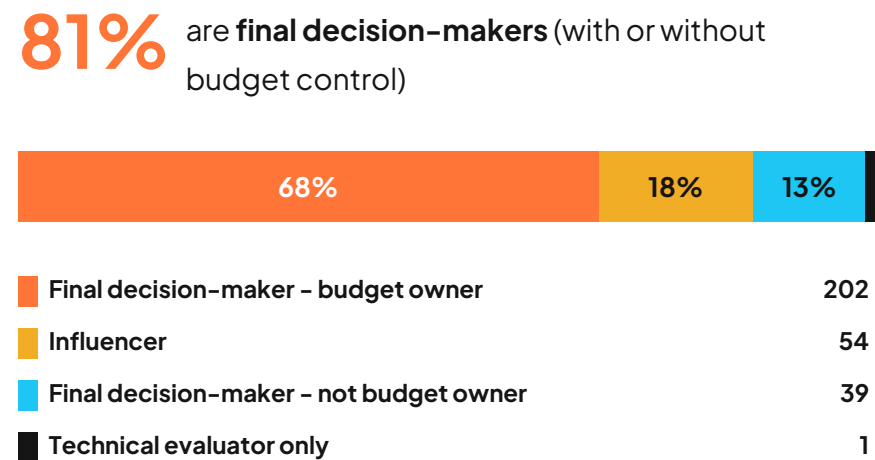
3. By industry




4. By Department




5. By purchase role



6. Software types actively bought



7. Functions approving software



A panel weighted toward buyers who own the budget. From <\$20M startups to >\$10B enterprises. 90% hold Director-level titles or above; 81% are final decision-makers.

03

Pre-Contract Evaluation

Why AI pricing is hard to evaluate before signing and what buyers measure when they try.

Pages

08-19

AI buyers were asked to provide their preferences on different pricing models

01**SEAT-BASED**

Organizations pay per user for access to the software. Spend increases based on an increase in user seats. It's a common pricing model and gives vendors predictable recurring revenue.

02**USAGE-BASED**

Organizations pay based on what they actually consume. Charges scale with increased activity like messages, storage, or data processed – similar to a utility bill. Months with lighter usage cost less.

03**CREDIT / TOKEN**

Organizations receive a pool of credits each billing cycle, and each action in the product spends some of them. Different features cost different amounts, so a simple task burns few credits and a complex one burns many.

04**OUTCOME-BASED**

Organizations pay for results, often with little or no up-front cost. The vendor only earns when a defined business outcome is achieved – such as a resolved support ticket, a recovered payment, a qualified lead.


05**HYBRID**

A blend of two or more of the above. The most common pattern is a base platform fee (seats or a subscription) plus usage or outcome charges layered on top.

SURVEY SETUP The five definitions above are the **exact framing** buyers received before expressing preferences.

Seat-based is the only AI pricing model buyers feel confident evaluating

EASIEST AND HARDEST PRICING MODELS TO EVALUATE PRE-CONTRACT · N = 292




Seat-based is generally the easiest pricing model to evaluate before signing. Other models are more polarized – buyers typically find them more difficult to evaluate. Pricing that depends on usage, outcomes, or consumption introduces variables buyers cannot reliably forecast before committing to a contract.

Seat-based is easier to evaluate at every company size; each pricing model is perceived differently across segments

NET DIFFICULTY OF EVALUATING PRICING MODELS (PRE-CONTRACT) · BY REVENUE BAND · N = 292




Credit and token pricing is uniformly difficult across every company size



Seat-based pricing is the only model that works consistently across company sizes, staying strongly positive from sub-\$50M buyers to \$1B+ enterprises. The rest are more uneven. Usage-based pricing is hardest for smaller companies, likely due to weak usage history, while hybrid pricing gets harder as companies scale. **Credit/token pricing is consistently negative across every segment,** pointing to a structural problem with the model itself. Outcome-based pricing improves with company size but still never becomes net positive.

AI buyers value predictability and clear pricing mechanisms over a lower entry price

WHAT MATTERS MOST WHEN EVALUATING AI PRICING · % RANKING TOP 3 · N = 292



Buyers will pay **more** for pricing they can predict and **less** for pricing they cannot

68% vs **19%**


rank predictable total cost as a top-3 priority

rank lowest entry price as a top-3 priority

Buyers prioritize cost predictability over cost savings. The top four priorities all describe the same underlying need: pricing buyers can forecast, explain internally, and defend to finance. Low entry price and minimal commitment rank last, indicating buyers are willing to pay more for a model they can confidently predict than less for a model they cannot.

As company size grows, the buyer's problem shifts from predicting cost to justifying it

WHAT MATTERS MOST WHEN EVALUATING AN AI PRICING MODEL · BY REVENUE BAND · N = 292



Predictable cost is a top priority across every segment, but what drives that need shifts with company size. Smaller companies need to know what they will pay for their own planning, whereas larger companies expect to spend, but need a clear value metric and a rationale they can defend to Finance. Mid-sized companies follow a similar preference for predictability, with a slight lean toward flexibility and lower entry price.


Revenue Band	predict cost	justify cost
<\$50M	77%	44%
\$50M-\$1B	68%	44%
>\$1B	63%	55%

As seniority rises, the pricing conversation shifts from proving value to defending the spend

WHAT MATTERS MOST WHEN EVALUATING AN AI PRICING MODEL · BY BUYER SENIORITY · N = 264

% of buyers ranking each as a Top-3 priority

Director: n=128 | VP / SVP / EVP: n=37 | C-Level: n=99

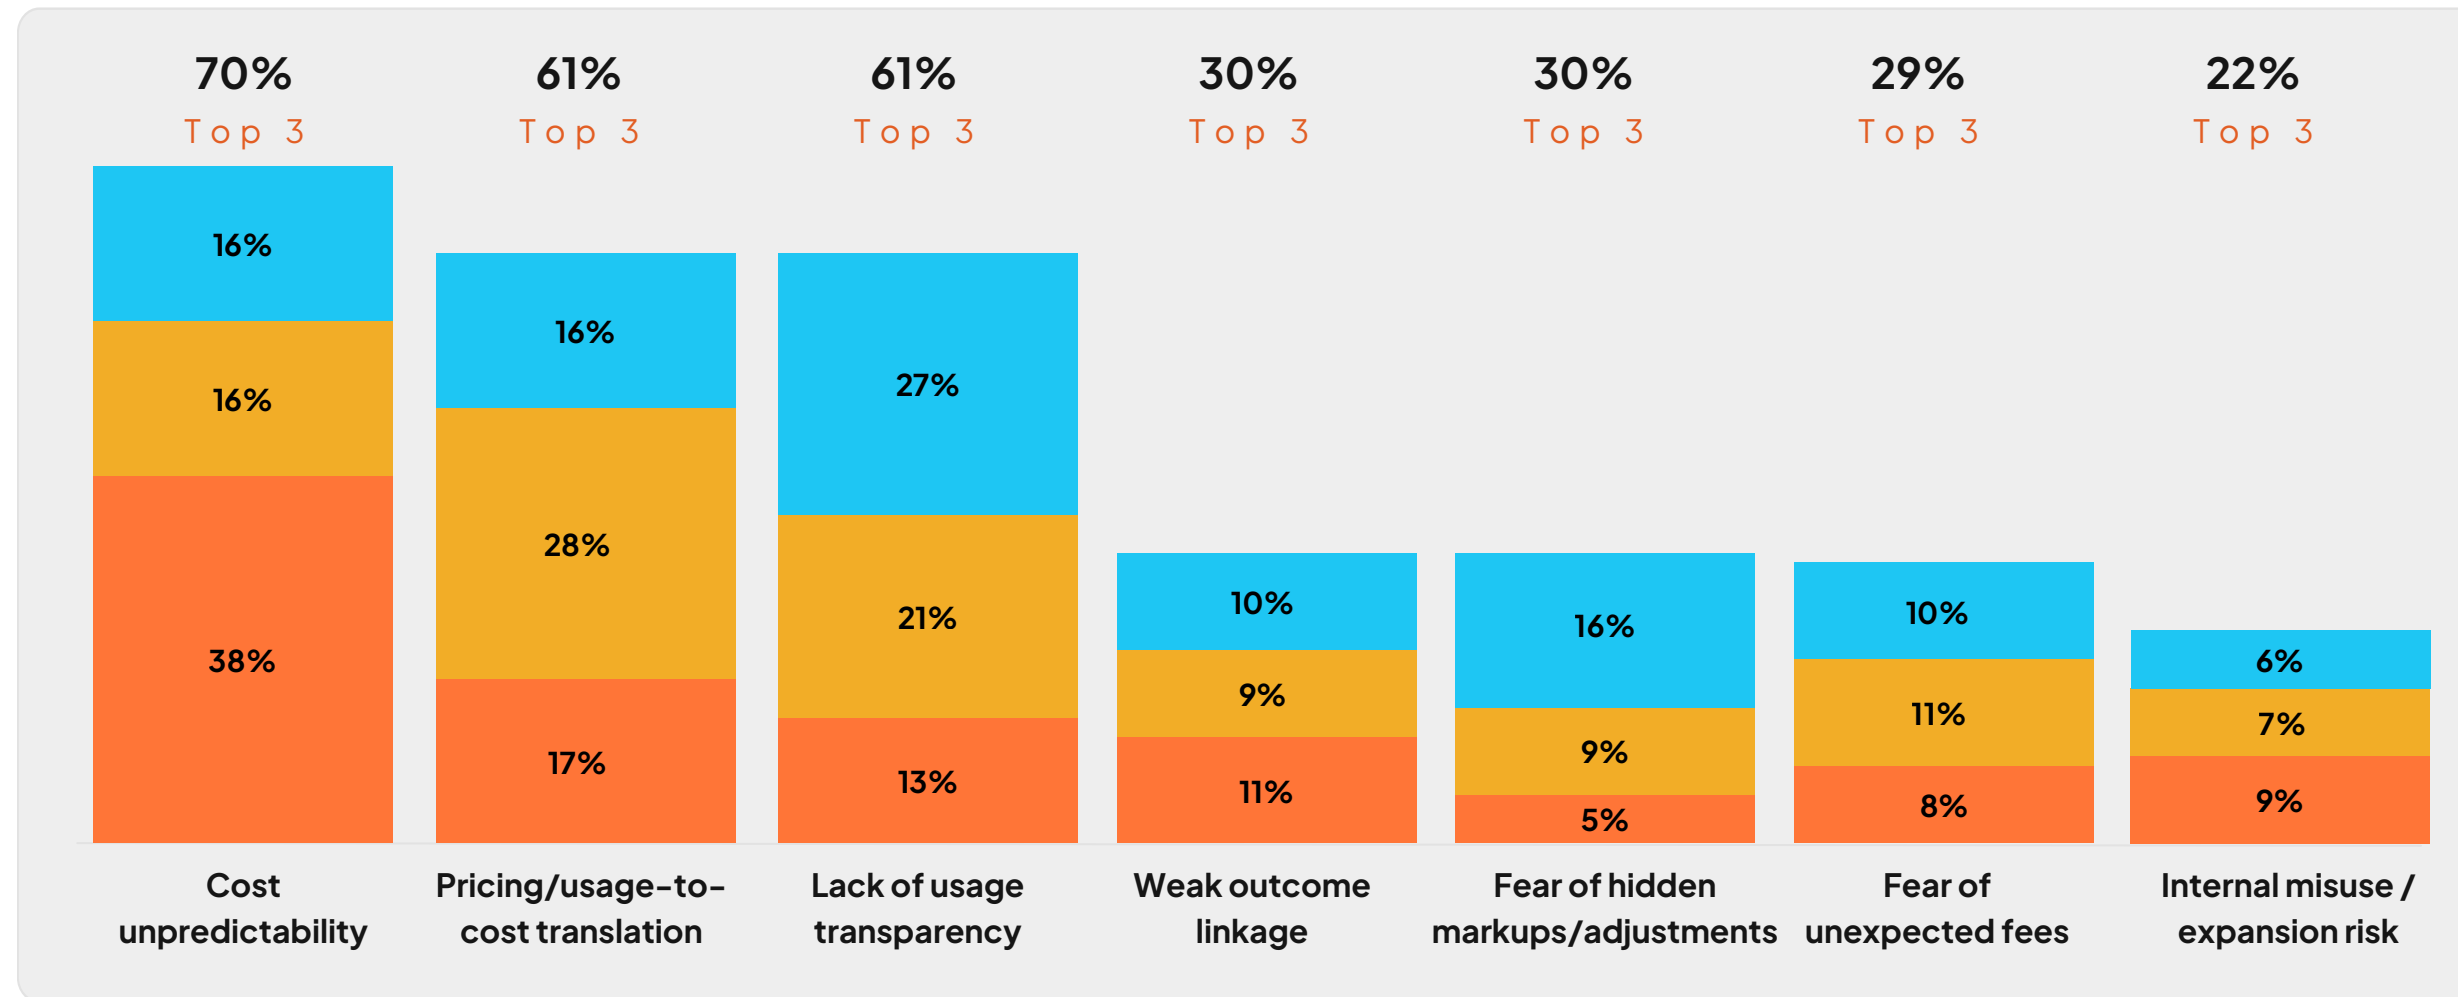


The top priorities are consistent across all seniority levels, but the pricing conversation shifts from performance to defensibility as seniority increases. **Directors are often closer to the product and place more weight on a clear ROI and price-to-value link. VPs and C-levels are less focused on proof of value and more focused on a clear value metric and justifying spend to Finance.**

Buyers' AI pricing concerns center on forecasting and transparency

TOP AI PRICING CONCERNS · % RANKING TOP 3 · N = 257

● Top Concern ● Secondary Concern ● Tertiary Concern



TOP CONCERN #1

70% Cost unpredictability

TOP CONCERN #2

61% Pricing/usage-to-cost translation










TOP CONCERN #3

61% Lack of usage transparency

The leading concerns about AI pricing all trace back to one underlying problem – buyers cannot reliably predict what the product will cost. Worries about hidden fees, markups, or weak outcome linkage rank lower, signaling that the issue is not vendor trust but a lack of visibility and forecasting confidence. Buyers are not asking vendors to be more honest; they are asking for pricing structures they can understand and plan around.

Part of the ambiguity around credits is tied to the fact that there are so many different credit models, for example:

The unit of measure varies, the cost-per-unit varies, and the definition of what triggers consumption varies. Buyers cannot evaluate what vendors cannot consistently define.

Vendor	Pricing Model	Metric	Cost Structure	What it means
 Adobe	Output-based	Generative credits	Varies across products	A credit is an action, ex: an AI image generation.
 Atlassian	Output-based	AI credits	AI credit pricing, limits not enforced	A credit is based on interactions with specific AI features; more complex interactions cost more.
 bolt.new	Cost-based	Tokens	\$0.00002 per token	Tokens are units of text that LLMs use to process information.
 Clay	Cost-based	Credits	\$0.01 to \$0.07 per credit	A credit is a data point or AI action, ex: data enrichment.
 Cursor	Cost-based	Dollars	\$0.03 per credit	Per GPT-4 request.
 Google	Cost-based	AI credits	\$0.01 per credit	A credit is used in AI generation (not a request). The amount of AI credits used for each model varies.
 HubSpot	Output-based	Credits	\$0.01 per credit	A credit is an action, ex: enrichment, a customer agent interaction.
 Lovable	Output-based	Credits	\$0.25 per credit	A credit is based on work done in response to messages; more complex actions cost more credits.
 monday	Output-based	AI credits	\$0.08 per credit	A credit is a task that's intentionally set by users and is successfully run.
 Salesforce	Output-based	Flex Credits	\$0.005 per credit	A credit is a discrete AI action, ex: updating a customer record. These translate to \$0.10 each.



A credit at
Salesforce
costs
\$0.005



A credit at
Lovable
costs
\$0.25



Buyers have
no way
to translate
one to the
other.

Five factors separate credit models buyers can plan around from those they cannot

Credit models aren't the problem. Ambiguity is.

01 Definition

02 Forecastability

03 Policy

04 Action mapping

05 Allotment design

WHAT MAKES A BAD CREDIT MODEL

⊗ Inconsistent definition

Ask three people what a credit is worth and get three answers - the black box is the product.

⊗ Buyers cannot forecast needs

Customers guess on purchase volume - then either hoard credits or hit overages and blame you.

⊗ Policies change without warning

Inconsistency over time destroys trust and customers stop buying more because they do not know what they are buying.

⊗ Wildly unequal action costs

Mixing 5-credit and 50-credit actions for the same user makes tracking feel impossible.

⊗ Per-user vs. per-team allotment is ambiguous

Nobody knows if the pool is shared or individual until someone runs out at the wrong moment.

WHAT MAKES A GOOD CREDIT MODEL

✓ Consistent definition across teams

Sales, finance, and product all describe credits the same way.

✓ Forecastable for the buyer

Customers can estimate usage before they commit - no surprise overages.

✓ Clear rollover and expiration policy

Rules are documented, enforced consistently, no exceptions for squeaky wheels.

✓ Credits map to discrete, measurable actions

One credit = one thing; no ambiguity about what "uses" a credit.

✓ Intentional definitions of allotments

Heavy users pay more because they get more; model earns trust, not resentment.

Outcome-based pricing only works when the outcome is fully owned, measured, and agreed before go-live

Three conditions must be true. **Six failure modes** emerge if any condition is missed.

WHEN OUTCOMES WORK

01

High autonomy & attribution

The AI is clearly the primary driver of the outcome, not the workflow, other factors, or the human. Attribution is unambiguous.

02

Outcomes are auditable

Results are precisely defined, trackable, and tied directly to the model's contribution. Customer has baseline metrics to quantify the delta.

03

Defined pre-purchase

Both sides agree on the outcome metric before the contract is signed. Performance is consistent and repeatable across customers.

✓ All three align

OTHERWISE

 Risk emerges

RISK OF OUTCOMES

Attribution fights

Who actually delivered the outcome - the AI, the workflow, or the human?

Buyer friction

Harder to sell when the product is new and outcomes are unproven.

Melting snowball

As automation improves, outcomes increase, but billable events collapse.

Forecast chaos

Finance can't simulate next month's bill and unpredictability becomes insurmountable.

Behavior distortion

Teams optimize for billable outcomes rather than actual customer impact.

Lumpy revenue

Outcome timing is irregular leading to revenue becoming hard to forecast.

Section Summary

Pre-Contract Evaluation

Buyers cannot reliably evaluate most AI pricing models prior to signing. Seat-based is the only model with pre-contract confidence; consumption-based, outcome-based, and hybrid pricing models are more difficult to evaluate than easy.

The friction shows up before any usage is calculated or tokens are spent, which is why predictability dominates the buyer's evaluation criteria over a lower price point.

What This Means for Buyers

- Difficulty evaluating AI pricing models is a mainly a forecasting problem.
- Increased desire to provide usage benchmarks before signing.
- Expect to pay a higher price point for more predictable usage.

What This Means for Vendors

- Leading with cost predictability, not entry price, is a clear differentiator.
- Provide a simple method (e.g., calculator) to calculate usage for prospective buyers.
- It is beneficial to opt for a model with a fixed metric (e.g., user seats) over pure variability.

43%

believe seat-based models are easiest to evaluate

Usage-based (+4%) is the only other model with a positive net difficulty score.

68%

rank predictable total cost as a top 3 priority

Low entry price ranks dead last. Buyers will pay more to be able to forecast reliably.

55%

say credit/token-based pricing is harder to evaluate for AI

The challenge in evaluating credit based models was ubiquitous across company size.

04

AI Spend Experience


What happens to budgets after the contract - who overspends, why, and who owns the risk.

Pages

20-33


For AI buyers, staying on the initial budget is now the exception

9 OF 10 BUYERS · OVER BUDGET





HAS YOUR SPEND ON AI EXCEEDED WHAT YOU INITIALLY EXPECTED?

% of buyers · n=263



Exceeding the initial AI budget is a common experience, with significant overruns nearly as common as moderate ones. This mirrors buyers' concerns around predictability: many underestimate either how much they will be charged or how much they will actually use.

Financial Services is the most burned by AI exceeding initial budget



Financial Services buyers have exceeded their initial budget the most on AI software. Tech and Software buyers are close behind, with about nine in ten exceeding what they planned to spend. Other industries are also overspending, but at lower severity. Overspending is widespread across every market, and vendors should expect buyers to come into pricing conversations already feeling burned.

AI drives more usage than buyers expect and that's the top reason budgets break

MAIN CONTRIBUTION TO HIGHER-THAN-EXPECTED AI SPEND · N = 233

01. USAGE OUTPACES PLANS

DOMINANT THEME

67% AI features / automation drove additional usage

63% Usage scaled faster than expected

37% Internal adoption expanded beyond initial users

34% Usage exceeded limits

Consumption rises faster than teams anticipate - AI products create their own demand once adopted.

02. FORECASTING IS UNCLEAR

PLANNING GAP

38% Difficulty forecasting usage

36% Pricing mechanisms / units unclear

Buyers can't model what they'll consume because credits, tokens and units don't translate to outcomes.

03. VENDOR CHANGED TERMS

MUCH SMALLER CAUSE


10% Vendor pricing / terms changed post-sale

Only one in ten buyers blame the vendor - overruns are not driven by post-sale repricing.

Budget overruns are driven by buyers underestimating how much AI their organization will actually use, not by vendors changing the rules after signing. The leading contributors all describe the same dynamic: AI products generate more activity than teams anticipate, and internal adoption spreads faster than forecasts assume.

Buyers who significantly exceeded their budget report different root causes than those who moderately exceeded

MAIN CONTRIBUTORS TO HIGHER-THAN-EXPECTED AI SPEND · BY OVERRUN SEVERITY · N = 233



SIGNIFICANT OVERRUNS · N = 117

Need consumption guardrails – not better forecasts.

Buyers who significantly exceeded their initial spend realized that the product worked, but usage compounded and the budget could not keep up. These customers do not need better forecasting; they need consumption guardrails that let them keep scaling without going back to finance every quarter.

MODERATE OVERRUNS · N = 116


Need a transparent meter – not usage caps.

Buyers who moderately exceeded their initial spend had difficulty modeling what they were buying, and the overage came from misinterpreting usage and future adoption. They do not need usage caps; they need a pricing model transparent enough to forecast against in the first place.

Each segment exceeds their budget but for meaningfully different reasons

MAIN CONTRIBUTORS TO HIGHER-THAN-EXPECTED AI SPEND · % OF BUYERS · N = 230

< \$50M revenue: n=26 | \$50M - \$1B revenue: n=138 | > \$1B revenue: n=66



Smaller companies (<\$50M) exceeding budget is a comprehension failure: buyers in this segment did not fully understand what they were buying. Mid-market (\$50M - \$1B) exceeding budget is product-driven, not buyer-driven: demand increased faster than the organization could govern it. Enterprise (>\$1B) exceeding budget is a capacity problem: this segment needs throttling, *graceful* degradation, and predictive alerts, not pricing education.

IT carries the AI budget risk – even when they don't drive the usage

PRIMARY OWNER OF AI COST (BUDGET) RISK · % OF BUYERS · N = 257

PRIMARY OWNER

67%

named IT the primary owner of AI cost risk - more than every other function combined.

EVERY OTHER OWNER

Combined: 33% of buyers

Finance		17%
Department Executive		13%
Procurement		2%
Marketing		<1%
CEO		<1%

THE TENSION

Finance watches the budget. **IT carries the risk.**

FINANCE

17%

IT

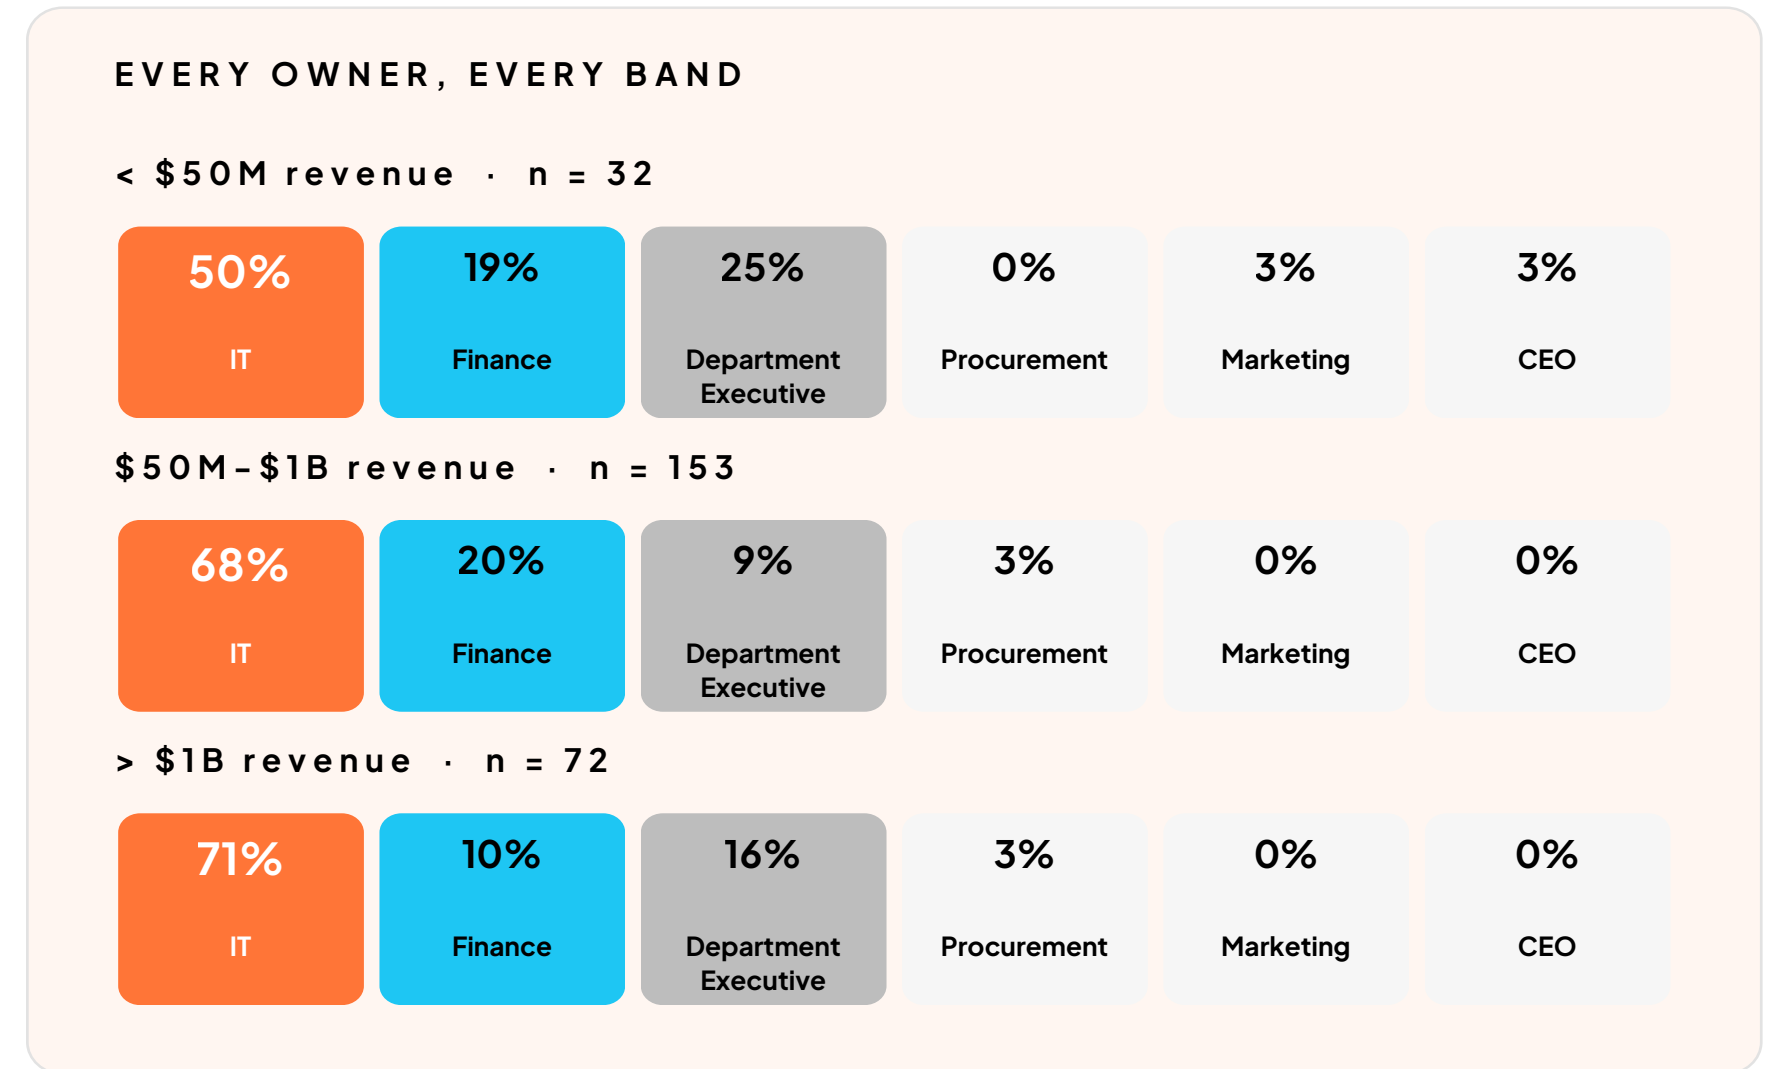

67%

That mismatch is where surprise overruns find a home.

Responsibility for AI cost risk is concentrated almost entirely within IT, with other functions playing minor roles. This creates a tension where Finance, despite watching the budget closely, is rarely named as the actual owner.

IT's ownership of AI cost risk grows with company size

PRIMARY OWNER OF AI COST (BUDGET) RISK · BY DEPARTMENT · N = 257



IT carries the largest share of AI budget risk at every company size, but its dominance grows with scale. At larger companies, IT is named as the primary owner of AI cost risk far more often than any other function. At smaller companies, Finance and Department Executives play a bigger role, likely because IT is a smaller function with less capacity to take it on.

AI buyers want visibility and control not hard cutoffs on usage

USAGE CONTROLS MOST WANTED FROM VENDORS · N = 262

THE BELL MOST WANTED



Tell me *before* the meter spikes – let me decide

Buyers want **visibility and decision points**. Soft caps, predictive alerts and explicit approvals are the controls they consistently rank highest. The product keeps working; the spend stays a choice.

62% Soft caps with alerts & approval to continue · **#1 most wanted**

55% Predictive alerts before budget is exceeded

THE WALL LEAST WANTED



Don't *stop* the work to enforce the budget.

Hard caps and prepaid credit pools rank lowest. Buyers prefer to be told "you're heading over" rather than have the product fail under load or be locked into a pool of credits they may never use.

40% Hard caps that stop usage entirely

28% Pre-purchase + draw-down credits · **least wanted**

RANK 1 · WANT

62%
Soft caps w/ alerts + approval

RANK 2 · WANT

55%
Predictive alerts before budget exceeded

RANK 3

47%
Throttle / degrade after threshold

RANK 4

45%
Pre-set monthly / quarterly limits

RANK 5 · CUTOFF

40%
Hard caps that stop usage

RANK 6 · LEAST


28%
Pre-purchase + draw-down credits

When it comes to managing AI usage, buyers prefer visibility and decision points over strict enforcement. They would rather be warned early and given a chance to approve continued usage than have the product abruptly stop (hard caps) or be locked into pre-purchased credits.

For each AI pricing concern, usage controls are ranked by how much MORE that concerned cohort wants them

N = 262

Versus overall base · orange = highest-lift control · dark grey = strongly rejected (<-3%)




Each "concern group" has a distinct top-ranked usage control, and several cohorts actively reject the controls that other cohorts prefer most. The right governance feature for your product depends on which concern your ICPs rank highest not on a market-wide default.

Buyers aren't asking for one generic fix – different reasons for exceeding budget call for different control solutions

BUYER-PREFERRED USAGE CONTROLS BY AI PRICING CONCERN · N = 262


AI PRICING CONCERN	USAGE CONTROL					
	Soft caps + approval	Predictive alerts	Throttle / degrade	Monthly / quarterly spend limits	Hard usage caps	Pre-purchase fixed credits
OVERALL	62%	55%	47%	45%	40%	28%
Cost unpredictability	60% -2% OVR	57% +2% OVR	47% no change	52% +7% OVR	46% +6% OVR	27% -1% OVR
Usage-to-cost translation	68% +6% OVR	57% +2% OVR	47% no change	51% +6% OVR	43% +3% OVR	26% -2% OVR
Lack of usage transparency	64% +2% OVR	55% no change	39% -8% OVR	47% +2% OVR	40% no change	27% -1% OVR
Weak outcome linkage	57% -5% OVR	57% +2% OVR	59% +12% OVR	38% -7% OVR	37% -3% OVR	20% -8% OVR
Hidden markups / adjustments	59% -3% OVR	47% -8% OVR	55% +8% OVR	33% -12% OVR	37% -3% OVR	30% +2% OVR



There is no universal AI-pricing fix. Each concern maps to a distinct preferred control. Vendors who match their guardrails to their buyers' actual concern profile – rather than offering a market-wide default – will be trusted with more spend.

Tech buyers demand significantly more guardrails than non-tech buyers

USAGE CONTROLS WANTED · ALL CONTROLS BY BUYER PROFILE · N = 262



Buyers are not asking for cheaper credits, instead they are asking for protection from their own forecasting failures. Every top-ranked control is a guardrail against unpredictable consumption (e.g., alerts, caps, approvals), and tech buyers want every guardrail more intensely than non-tech, suggesting the teams closest to AI usage are the ones most worried about losing control of it.

Buyers want different guardrails for different problems and predictive alerts for all of them

Soft Caps + Approval

- **Extremely effective** when buyers have difficulty forecasting usage in advance.
- **Moderately effective** when buyers' internal adoption expands beyond initial users, usage scales faster than expected, or automation drives additional usage.
- **Ineffective** when the pricing metric / units of measure is unclear.

Hard Usage Caps

- **Extremely effective** when buyers' internal adoption expands beyond initial users.
- **Moderately effective** when buyers have difficulty forecasting usage in advance or pricing metric / units of measure is unclear.
- **Less effective** when automation drives additional usage or usage scales faster than expected.

Predictive Alerts

- **Extremely effective** when buyers' internal adoption expands beyond initial users.
- **Moderately effective** when buyers have difficulty forecasting usage in advance, automation drives additional usage, usage scales faster than expected, or when the pricing metric / units of measure is unclear.

Monthly / Quarterly Spend Limits

- **Extremely effective** when buyers have difficulty forecasting usage in advance.
- **Moderately effective** when buyers' internal adoption expands beyond initial users or usage scales faster than expected.
- **Less effective** when automation drives additional usage or the pricing metric / units of measure is unclear.

Throttle / Degrade

- **Extremely effective** when the pricing metric / units of measure is unclear.
- **Less effective** when automation drives additional usage.
- **Ineffective** when buyers have difficulty forecasting usage in advance, usage scales faster than expected, or internal adoption expands beyond initial users.

Pre-Purchase Fixed Credits

- **Extremely effective** when buyers' internal adoption expands beyond initial users.
- **Moderately effective** when buyers have difficulty forecasting usage in advance.
- **Less effective** when automation drives additional usage, usage scales faster than expected, or the pricing metric / units of measure is unclear.

Section Summary

AI Spend Experience

Most AI buyers exceed their initial budget, but the cause is not simply vendor behavior. Buyers frequently underestimate their AI usage, such as through expanded internal adoption.

There is more of a governance problem than a pricing problem, and guardrails are needed to maintain buyer satisfaction in AI software purchasing.

What This Means for Buyers

- AI budgets should be built with assumptions around future spend not only initial usage.
- IT and finance need to be jointly accountable for the cost risks associated with AI spend.
- Expect that guardrails are table stakes to address common concerns such as cost unpredictability and adoption-driven increases.

What This Means for Vendors

- Usage controls - soft caps and predictive alerts - are critical if adopting a usage-based model.
- Objection handling is critical to overcome pre-existing buyer hesitancy on variable pricing.
- Anticipate that buyers may have previously exceeded budget on AI software and will arrive skeptical of variable pricing models..

89%

have exceeded their initial AI budget

45% have significantly exceeded, 44% moderately, only 9% stayed on-budget.

67%

name IT as the primary owner of AI cost risk

Finance was named only 17% of the time, despite watching the budget closely.

62%

want soft caps with alerts as protective guardrails

The majority of buyers reject hard-stop limits in favor of warnings + decision points.

05

AI Pricing Sentiment

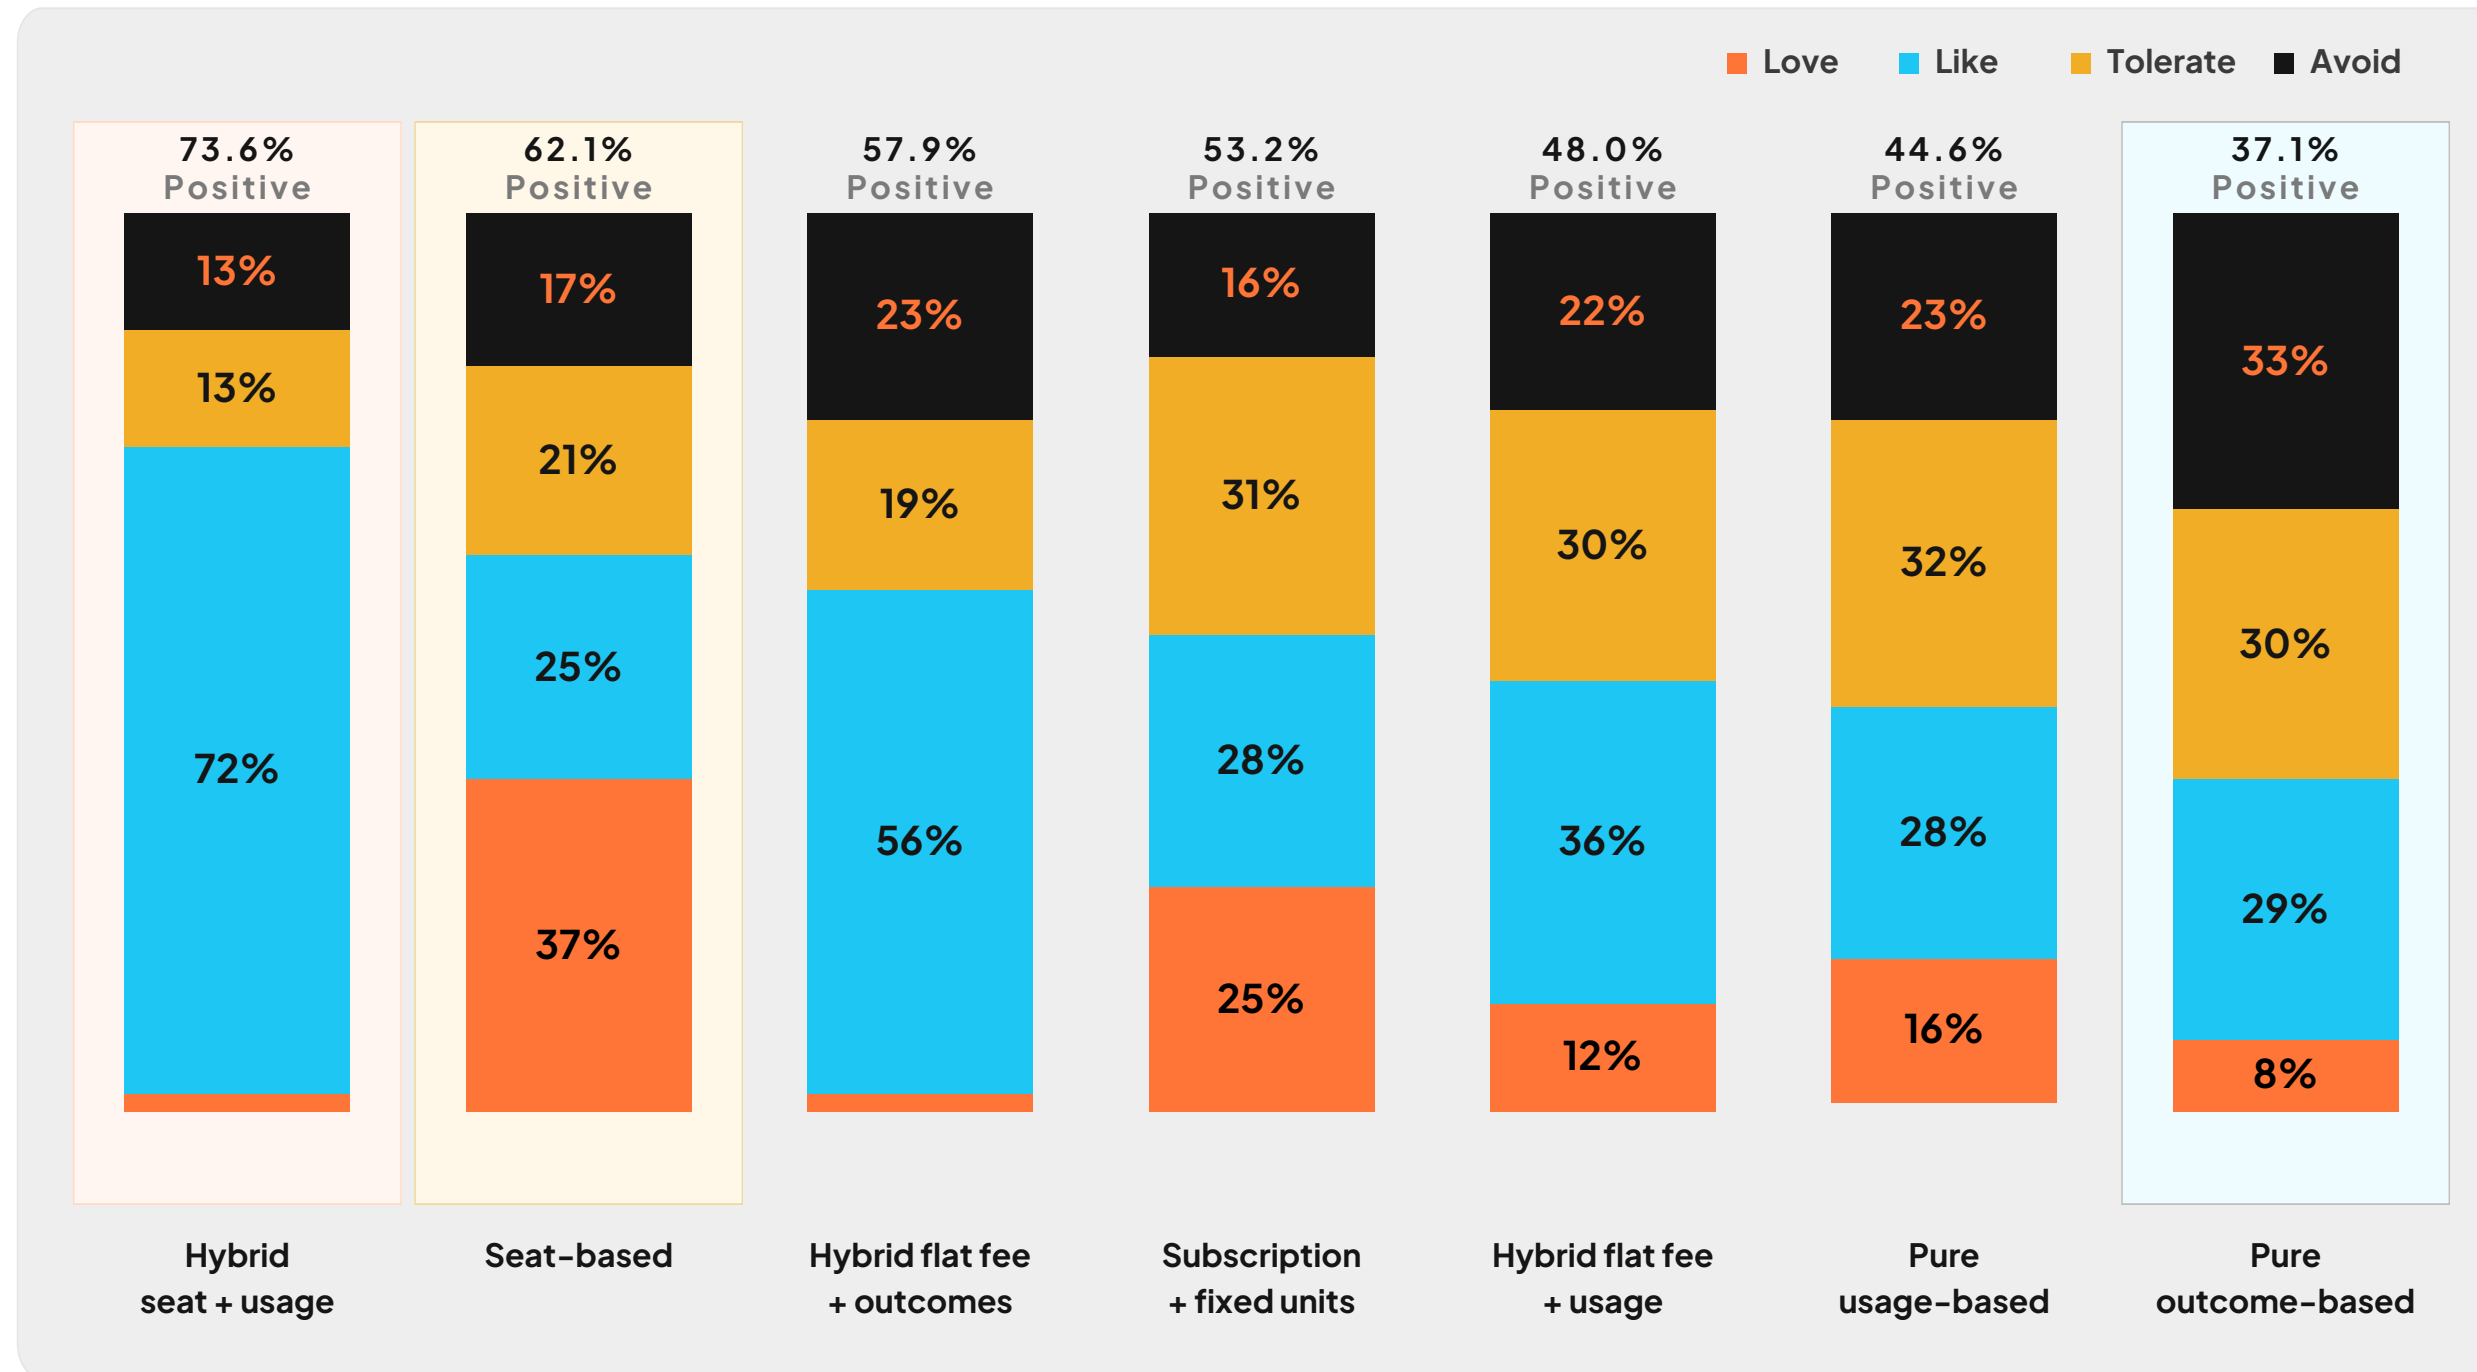
Which pricing models buyers love, like, tolerate, or avoid and why they prefer what they prefer.

Pages

34-42

AI buyers value seat-based and predictable pricing

Sentiment by AI Pricing Model · % of buyers · N = 296



01

73.6%

POSITIVE
Hybrid seat + usage

02

62.1%

POSITIVE
Seat-based





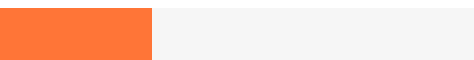
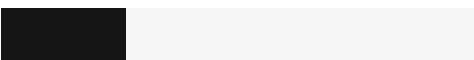
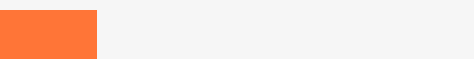
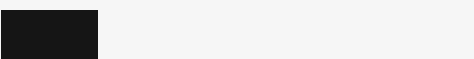
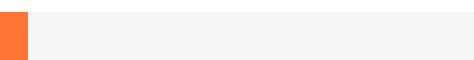
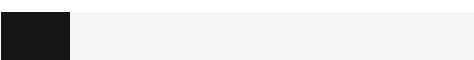
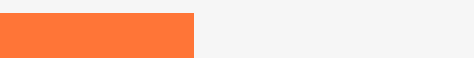

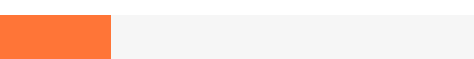

03

57.9%

POSITIVE
Hybrid flat fee + outcomes

Buyers' top preference is seat-based models – their clearest rejection is pure outcome-based

MOST AND LEAST PREFERRED PRICING MODEL · RANKED BY NET PREFERENCE · N = 296

PRICING MODEL	MOST PREFERRED	LEAST PREFERRED	NET
Seat-based	 33%	 16%	+17
Subscription + fixed units	 25%	 15%	+10
Hybrid flat fee + usage	 11%	 9%	+2
Hybrid seat + usage	 7%	 7%	+0
Hybrid flat fee + outcomes	 2%	 5%	-3
Pure usage-based	 14%	 21%	-7
Pure outcome-based	 8%	 27%	-19

Both models with the most negative net preference are the two “pure models” in the set, where pricing is tied entirely to one variable. Buyers consistently reject pricing structures that remove all fixed anchors, regardless of how the variable is defined.

For an essential AI product, predictability beats every variable alternative

PREFERRED PRICING STRUCTURE FOR AN ESSENTIAL AI PRODUCT · % OF BUYERS · N = 257

TOP CHOICE

51%

Higher fixed cost, fully predictable

Buyers will pay more for an essential product if it removes variability.
Half the market chooses certainty over savings.

OTHER OPTIONS

26%

Hybrid w/ clear guardrails

13%

Lower entry price, variable usage risk

10%


Pay only for outcomes, limited transparency

Across all four structures, roughly three-quarters of buyers choose an option with a predictable floor, while one-quarter accept variable exposure. For high-stakes AI products, packaging built on a fixed anchor is what most buyers are willing to commit to.

Finance shows directionally higher tolerance for variable usage-based pricing structure

PREFERRED PRICING STRUCTURE FOR AN ESSENTIAL AI PRODUCT · BY DEPARTMENT · N = 249

% of buyers



IT seems the most risk-averse owner of AI budget: where IT owns the spend, buyers concentrate on fully predictable pricing and reject variable or outcome-tied alternatives. Finance seems the most risk-tolerant, with noticeably higher acceptance of outcome-based and variable-usage structures. Pricing design of your product should match the function carrying budget risk, not a generic buyer.

Open-ended buyer feedback converges on a simple ask: make AI pricing clearer, more predictable, and easier to govern

WHAT BUYERS WANT CHANGED IN AI PRICING MODELS · N = 194

73% OF MENTIONS

CLARITY

38% More predictability / stable costs

35% More transparency / clearer pricing

Buyers ask first for less ambiguity - predictable costs and transparent pricing. Both top mentions describe the same thing: knowing what you'll pay.

32% OF MENTIONS

MECHANICS

10% Stronger value / outcome alignment

9% Move away from token / usage-based pricing

7% Simpler or lower pricing

6% Better usage controls / caps

Specific mechanic asks (caps, simpler tiers, value alignment, moving away from tokens) trail clarity by a wide margin.

When buyers are asked in their own words what they would change, they overwhelmingly ask for transparency, predictability, and fewer opaque usage mechanics. That reinforces the quantitative story: the market is not asking for more pricing creativity, it is asking for less ambiguity.

In buyers' own words: clarity and predictability are non-negotiable

What do software vendors **most misunderstand** about how buyers experience AI pricing?

2 themes emerged

PRICING CLARITY

- || *We need clarity and predictability. It's impossible to validate a budget otherwise.*

PRICING CLARITY

- || *We're unlikely to adopt new AI features if there is a credit/usage cap or upcharge associated with them. We're actively avoiding AI features that have this pricing mechanism.*

PRICING CLARITY

- || *Mismatch between customer outcomes, vendor revenues and vendor costs.*

VALUE PERCEPTION

- || *As a buyer I don't care about your cost model - I want a clear explanation of the benefits of your software, and transparency about how usage translates into fair pricing. I'm happy to pay for value, but don't have time to have arguments about value - help me understand up front.*

VALUE PERCEPTION

- || *The importance of lower-risk trials, even paid trials, so that you can see if the AI pricing actually lives up to the hype. I don't want to get locked into a multi-year that doesn't work or with a vendor that dies.*

VALUE PERCEPTION

- || *Too much promise of what it can do, with little evidence that it will get delivered. A bit like the move to the cloud 20 years ago. It took time and a lot of failure.*

Buyers want spend guardrails and clearer value mechanisms

If you could **change one thing about current AI pricing models**, what would this be?

2 themes emerged

THEME 01

Guardrails & Consumption Visibility

- “ I would implement clearer, more predictable pricing with built-in safeguards against unexpected spikes to make budgeting and scaling much easier.
- “ I would prioritize standardized, real-time consumption dashboards with automated soft-cap alerts. This would allow organizations to experiment with AI features without the fear of massive, unpredictable overage charges at the end of a billing cycle.
- “ More predictable pricing clear flat plans or caps would help control costs and avoid clear flat plans or caps would help control costs and avoid surprises.
- “ I'd make every vendor offer a hard usage cap with automatic shut off so we never get a surprise bill that blows our budget.
- “ I'd make pricing more predictable. Something like a clear monthly cap or flat-rate tiers with generous limits would make it way easier to plan and avoid surprise costs.

THEME 02

Standardized Value Metrics

- “ I would standardize pricing around clear and predictable units of value rather than opaque tokens or fluctuating usage.
- “ Make pricing simple, transparent, and predictable, instead of confusing usage-based billing.
- “ Fixed costs, predictability of important outcome measures, and the ability to show flexibility as needs change.
- “ The usage model seems to get a little confusing after certain milestones are hit. Seems like most charge significantly higher prices with more (usage).
- “ I would ensure there is enough flexibility to add new members as the build out of AI in our systems continues to grow. If we go with a set number of seats, it becomes challenging to do any and how it adds over time.

Section Summary

AI Pricing Sentiment

Buyer sentiment is aligned with how much fixed pricing each model contains: the two top-rated models anchor on a fixed metric (e.g., seats) with the highest rated being a hybrid of seats plus usage, and the two lowest-rated are the pure usage and pure outcome models.

When asked directly what they would change about AI pricing, buyers do not ask for more sophisticated pricing, rather one that is clear and predictable.

What This Means for Buyers

- Prioritize hybrid models that combine a fixed anchor with a variable component.
- For essential AI products, accept a higher fixed price in exchange for predictability.
- Recognize that the budget owner shapes preference: IT-owned spend skews towards fully predictable, while finance-owned spend tolerates variability.

What This Means for Vendors

- Build pricing on a fixed anchor - every top model has one; every bottom model lacks one.
- Pure outcome-based pricing is heavily product-dependent, not a market default.
- Consider seat-based or subscription + fixed-units packaging as the default for risk-averse segments.

51%

would choose a higher fixed cost / fully predictable structure

Only 10% would accept pure outcome with limited transparency.

38%

name 'more predictability / stable costs' as the one thing they would change

Only 10% of buyers asked for stronger value/outcome alignment.

37%

positive sentiment for pure outcome-based pricing

The lowest of any model tested - -19 net preference.

06

AI versus Traditional SaaS

*What buyers treat differently when AI is involved and
what they still treat the same.*

Pages

43-49


One in three buyers treats AI pricing as a genuinely different category

BUYERS' PERCEPTION OF AI vs TRADITIONAL SAAS

% of buyers · n = 257

67% see AI as similar to SaaS

33% see AI as different



CUSTOMER VERBATIMS

- " Have to consider usage in some way. Credits, tokens, etc. Have to be transparent.
- " We look for AI-specific protective language that our data won't be used by the vendor or any LLM.
- " Still trying to assess value vs cost - it's still just a tool.

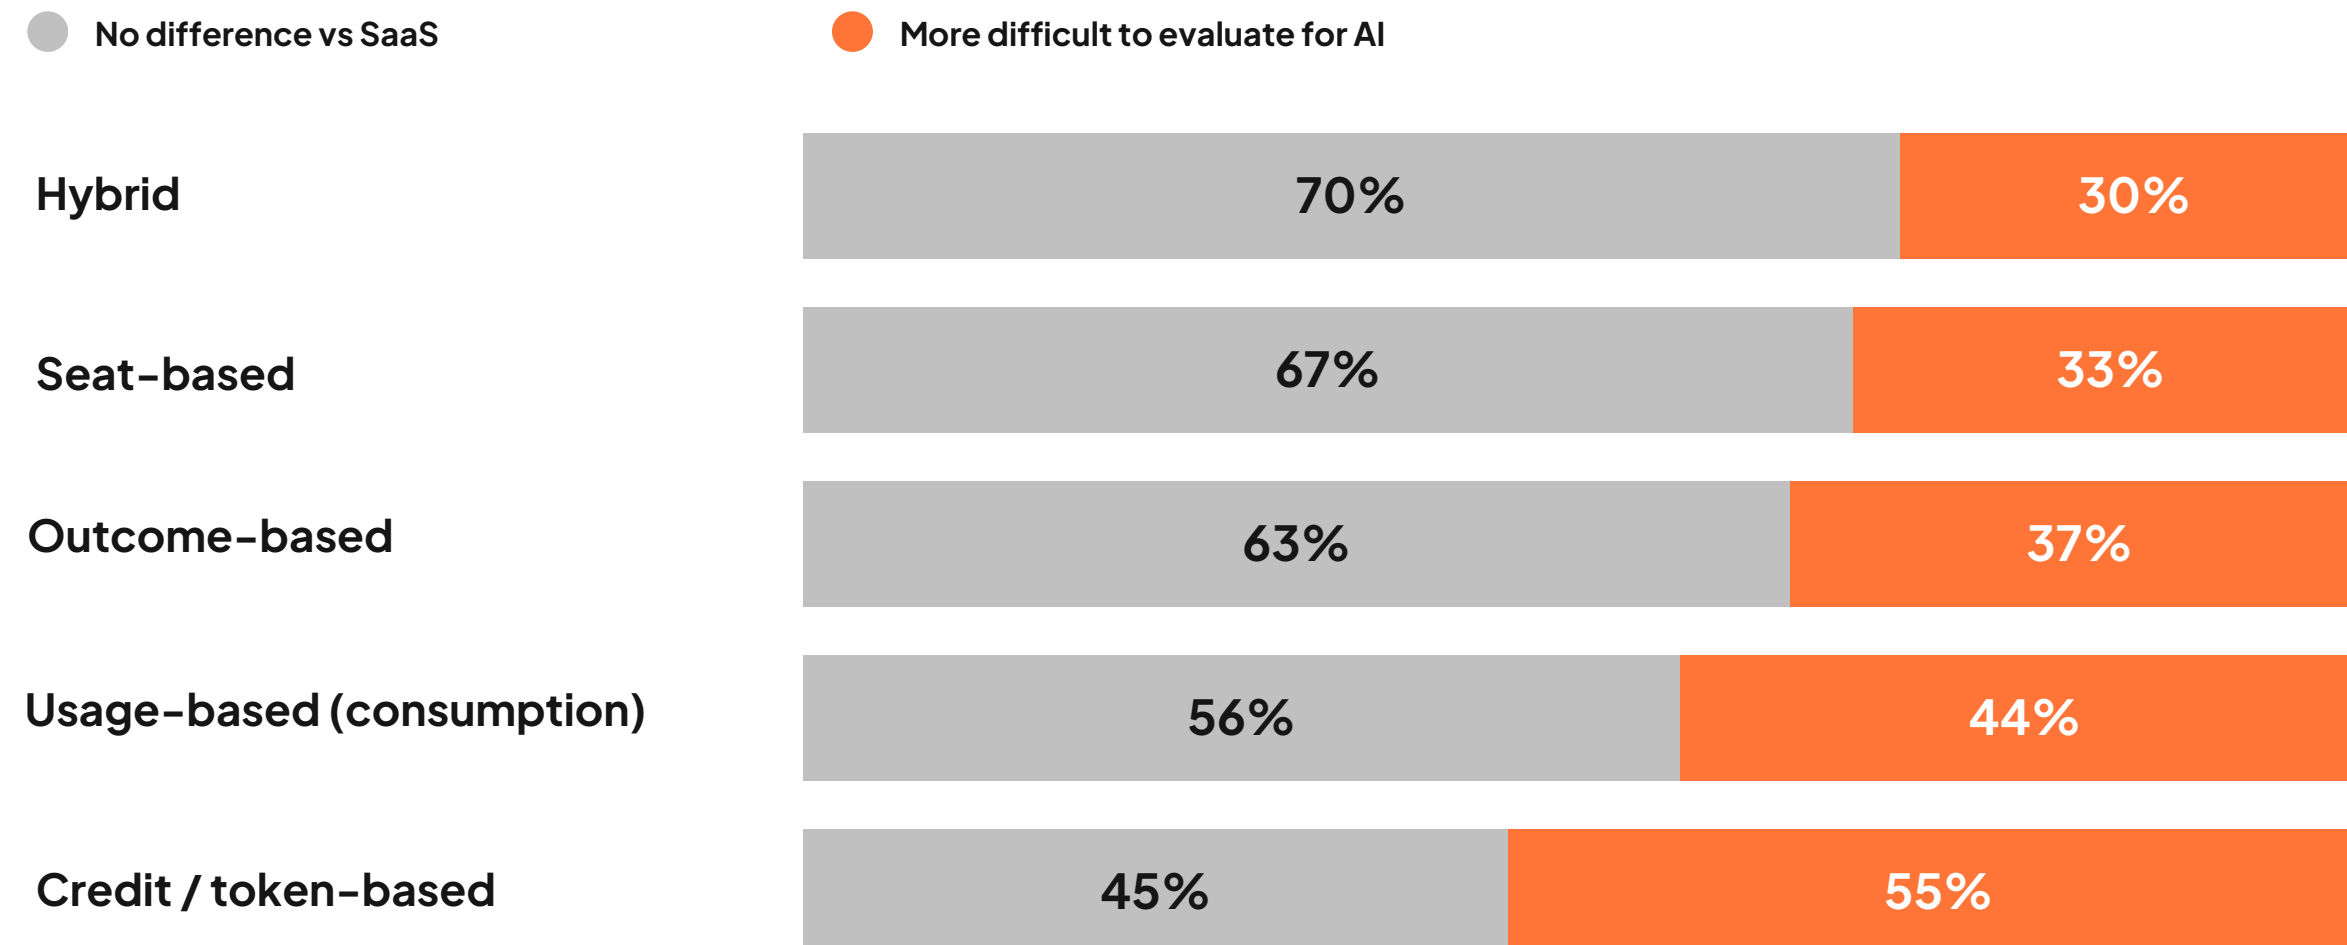
SaaS (Software as a Service): Software you access online through a subscription rather than buying and installing it on your computer. The vendor handles the hosting, updates, and maintenance - you just log in and use it. Examples: Gmail, Zoom, Slack, Salesforce.

Two-thirds of buyers approach AI pricing with similar expectations to traditional SaaS, so SaaS-style pricing models such as user-based or usage-based will feel familiar. However, this perception is not universal, and not always particularly strong, so there is room for experimentation.

More than half of buyers find credit-based pricing harder to evaluate for AI than SaaS

PRICING MODEL MORE DIFFICULT TO EVALUATE FOR AI vs SAAS · N = 292

% of buyers



HARDEST TO EVALUATE

55%

Credit / token-based


more difficult to evaluate for AI than SaaS

Buyers have no pre-AI benchmark for what a credit costs or what one buys - so they default to "too risky to forecast."

Pricing models where the unit of measure is a consumption metric (credit/tokens) can be challenging for customers when purchasing AI, despite correlation with cost-to-provide. Buyers have no pre-AI benchmark for how much they will consume, so this pricing structure becomes harder to forecast and evaluate. Models anchored on seats or a fixed component make it much easier for customers to predict spend.

Credit/token and usage-based pricing get materially harder to evaluate when AI is involved across companies of all revenue bands

PRICING MODEL MORE DIFFICULT TO EVALUATE FOR AI VS SAAS · BY REVENUE BAND · N = 292



Consumption-based pricing is considered difficult to evaluate across company sizes. These difficulties are structural to the pricing model rather than buyer size or sophistication. Vendors selling consumption-based AI are likely losing some deals to evaluation friction rather than price. Many buyers cannot forecast cost confidently before signing, so reducing this friction wherever possible is likely to improve win rates.

Tech buyers find every pricing model harder to evaluate for AI than non-tech buyers – with Seat-based being the only exception

PRICING MODEL MORE DIFFICULT TO EVALUATE FOR AI VS SAAS · BY INDUSTRY · N = 292

n = 180

TECH BUYERS

technology + software industries

Credit / token-based	60%
Usage-based (consumption)	48%
Outcome-based	42%
Seat-based	31%
Hybrid pricing models	33%

Consumption-based and credit-based pricing creates the biggest evaluation friction – tech buyers struggle with consumption models more than seat-based.

n = 112

NON-TECH BUYERS

all other industries

Credit / token-based	48%
Usage-based (consumption)	38%
Outcome-based	29%
Seat-based	37%
Hybrid pricing models	26%

Less acute friction – they may catch up as AI deployment matures across their organizations.

Tech buyers report meaningful AI-specific evaluation friction across every consumption-based and outcome-based pricing model. Non-tech buyers may eventually have stronger reactions as AI cost models and pricing maturity catch up. As non-tech AI deployment matures, their perceived evaluation difficulty may converge upward toward the tech profile.

Larger buyers reject variable pricing and treat AI as a distinct pricing category

AI PRICING FRICTION BY SEGMENT · BY REVENUE BAND · N = 257

FINDING 01

Significantly exceeded initial AI budget

● < \$50M revenue	32%
● \$50M-\$1B revenue	47%
● > \$1B revenue	44%

FINDING 02

AI pricing as different from SaaS

● < \$50M revenue	25%
● \$50M-\$1B revenue	30%
● > \$1B revenue	42%

FINDING 03

Least interest in purely variable pricing models

● < \$50M revenue	41%
● \$50M-\$1B revenue	44%
● > \$1B revenue	60%

Unexpectedly high spend affects every buyer segment, but larger buyers are more likely to have clear pricing preferences in response. They tend to reject pure variable models and see AI as its own pricing category. Vendors relying on purely variable pricing are likely to have trouble with in the largest, highest-value deals.

Section Summary

AI versus Traditional SaaS

AI buyers find credit and usage based models more difficult to evaluate than seat-based, outcome-based, or hybrid pricing models. This is true across size segments, though to varying degrees. Buyers in the tech industry especially tend to find these pricing models difficult to evaluate, relative to other industries.

These preferences, along with the sizeable fraction of AI buyers who wind up spending more than they initially expect, make price predictability a key consideration for AI products.

What This Means for Buyers

- Credit/token pricing is more difficult to comprehend than its SaaS equivalent.
- Push vendors to communicate consumption pricing in SaaS-familiar terms (per seat, per workflow, per fixed-unit subscription) wherever possible.
- Expect AI evaluation friction to grow as finance becomes less involved in the sales process.

What This Means for Vendors

- Find ways to make pricing understandable and predictable for buyers. Buyers highly value knowing how much they will pay, at least approximately.
- Be careful to keep an eye on expenses. Predictable pricing with variable costs can cause profit volatility.
- Consider who you're selling to. Different types of buyers can have different priorities.

55%

consider credit-based pricing models harder to evaluate for AI than SaaS

44% said the same about usage-based pricing generally.

10-13%

difference between tech vs. non-tech buyers in price evaluation

Tech buyers feel that usage and outcome based pricing is significantly more difficult to evaluate, compared to non-tech buyers.

60%

of buyers with > \$1B revenue selected variable pricing as least preferred

This preference was significantly stronger than it was for smaller companies.

About This Report

This report is based on a survey of 296 SaaS participants, conducted in partnership between Pricing I/O and Benchmarkit.

About Pricing I/O

Pricing I/O is a SaaS and AI pricing strategy firm helping B2B software companies design packaging, pricing, and monetization strategies. Since 2019, we've worked with 400+ clients across SaaS and AI.

www.pricingio.com

[Free Pricing Scorecard Quiz](#)



Acknowledgments

We thank the SaaS buyers and operators who contributed their time and insights to make this research possible.

About Benchmarkit

Benchmarkit is a leading source of B2B SaaS operating metrics and benchmarks, providing data-driven insights that help SaaS leaders make smarter decisions.

www.benchmarkit.ai

