Monetizing AI in SaaS

An Empirical Study of Pricing & Packaging Models for AI Features



Hitesh Sarda
Doctor of Business Administration
Under Guidance of **Dr. Latika V Shanbhag**
IIBM Institute • SIMA Switzerland
March 2026

Presentation Overview

01 **Introduction & Background**

Problem context and research motivation

02 **Research Questions & Objectives**

Primary and secondary research goals

03 **Scope & Limitations**

Study boundaries and constraints

04 **Literature Review**

SaaS economics, AI disruption, pricing models

05 **Methodology & Data Collection**

Mixed-methods sequential design

06 **Data Analysis & Key Results**

Quantitative findings from 371 companies

07 **Thematic Insights**

Executive interview qualitative themes

08 **Implications & Recommendations**

Theory and managerial frameworks

09 **Conclusion**

Summary, contributions, future research

01 | Introduction & Background

The AI-SaaS Inflection Point

The global SaaS market — valued at over USD 266B in 2024 — is undergoing a fundamental structural disruption as AI capabilities become embedded across software products.

Unlike traditional SaaS with near-zero marginal costs, AI introduces significant variable compute costs (per-token LLM inference) that scale linearly with usage. This breaks the classical per-seat pricing logic.

SaaS vendors face a critical dilemma: price for consumption to cover costs, or maintain flat pricing to satisfy enterprise budget predictability demands.

$266B

SaaS market size (2024)

$1.1T

Projected by 2032 (20% CAGR)

31.5%

AI market CAGR

$50B

India SaaS ARR target by 2030

01 | The Core Problem

"SaaS companies struggle to formulate pricing strategies for AI features that are simultaneously profitable, scalable, and aligned with customer value perceptions."

Variable Cost Shock

AI inference costs scale linearly with usage via per-token pricing — a fundamental break from the near-zero marginal cost of traditional SaaS.

Customer Anchoring Effect

Enterprise buyers anchored to $20/mo consumer AI tools (ChatGPT Plus) create a disconnect with enterprise-grade pricing needs.

Predictability Anxiety

Fear of "surprise bills" causes enterprise buyers to avoid using AI features even when free credits are offered.

Literature Gap

Existing academic literature predates AI SaaS and focuses on traditional value metrics — leaving practitioners without validated frameworks.

02 | Research Questions & Objectives

PRIMARY OBJECTIVE

Empirically investigate pricing & packaging strategies for AI features in B2B SaaS firms ($1M–$100M ARR) and develop a practical executive decision framework.

SECONDARY OBJECTIVES

1. Identify and classify primary AI pricing and packaging strategies in use
2. Determine prevalence of monetization models across the target SaaS segment
3. Analyze differences in AI monetization by revenue sub-segment ($1-10M, $10-50M, $50-100M)
4. Explore strategic considerations and competitive factors behind pricing decisions
5. Investigate how companies balance variable AI costs vs. customer demand for predictability
6. Propose an evidence-based framework supporting executive AI pricing decisions

03 | Scope & Limitations

SCOPE

- ✓ B2B SaaS companies with $1M–$100M ARR
- ✓ Companies with identifiable, publicly priced AI features
- ✓ Focus on commercial monetization strategy (not technical AI performance)
- ✓ Geographic scope: global, with US/India comparative lens
- ✓ Excludes early-stage startups and large enterprise custom contracts

LIMITATIONS

- ! Non-probability sampling limits statistical generalizability
- ! "Contact Us" pricing excluded — may skew toward smaller firms
- ! Regional nuances not visible on public pricing pages
- ! Cross-sectional design: single snapshot, no longitudinal tracking
- ! Small qualitative sample (n=9) limits interview finding generalizability

04 | Literature Review — Key Themes

1 Traditional SaaS Economics

Per-seat subscription logic; high fixed / near-zero marginal cost structure; value-based pricing; tiered and freemium models. Pre-AI cost estimation models (COCOMO, SLIM) focused solely on fixed upfront development costs.

2 AI as Structural Shift

AI transforms software from a passive tool to an active collaborator. GenAI described as 'most disruptive force in enterprise software since SaaS.' Variable per-token inference costs upend classical unit economics.

3 Monetization Challenges

Variable cost shock; customer predictability anxiety; billing infrastructure fragmentation; sales enablement gap; ethical risks (algorithmic bias, transparency). Only ~30% of firms publish quantifiable AI ROI.

4 Emerging Pricing Models

Hybrid (subscription + usage) dominant; consumption-based (token, credit, output); outcome-based pricing as strategic ideal but hard to execute; credit-based systems as psychological abstraction layer.

5 Literature Gaps

No empirical classification of mid-market ($1–$100M ARR) strategies; lack of strategic rationale documentation; no segment-specific benchmarks; cross-regional variation unexplored.

05 | Research Methodology

PHASE 1 — Quantitative

Pricing Page Analysis

- 371 SaaS company pricing pages systematically reviewed
- Predefined codebook for consistent classification
- Variables: pricing model, value metric, tier count, AI feature labeling, overages
- Revenue sub-segments: $1-10M, $10-50M, $50-100M
- Analysis: Descriptive statistics + cross-tabulations
- Tools: Excel / SPSS

PHASE 2 — Qualitative

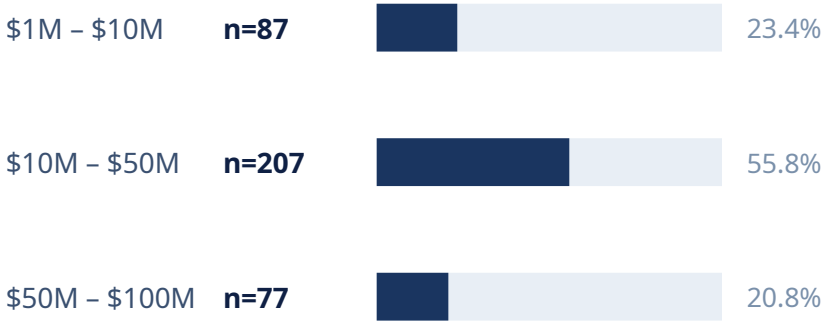
Executive Interviews

- 9 senior SaaS executives (CEO, CPO, CRO level)
- Semi-structured questionnaire format
- Verticals: ERP, real estate, analytics, fintech, retail, UX tools
- Thematic analysis: Braun & Clarke (2006) 6-phase framework
- Credibility via triangulation with quantitative data
- Confidentiality: all firms anonymized

05 | Sample Characteristics

371	215	58%	9
Companies Analyzed	Offering AI Features	AI Feature Adoption Rate	Exec Interviews Conducted

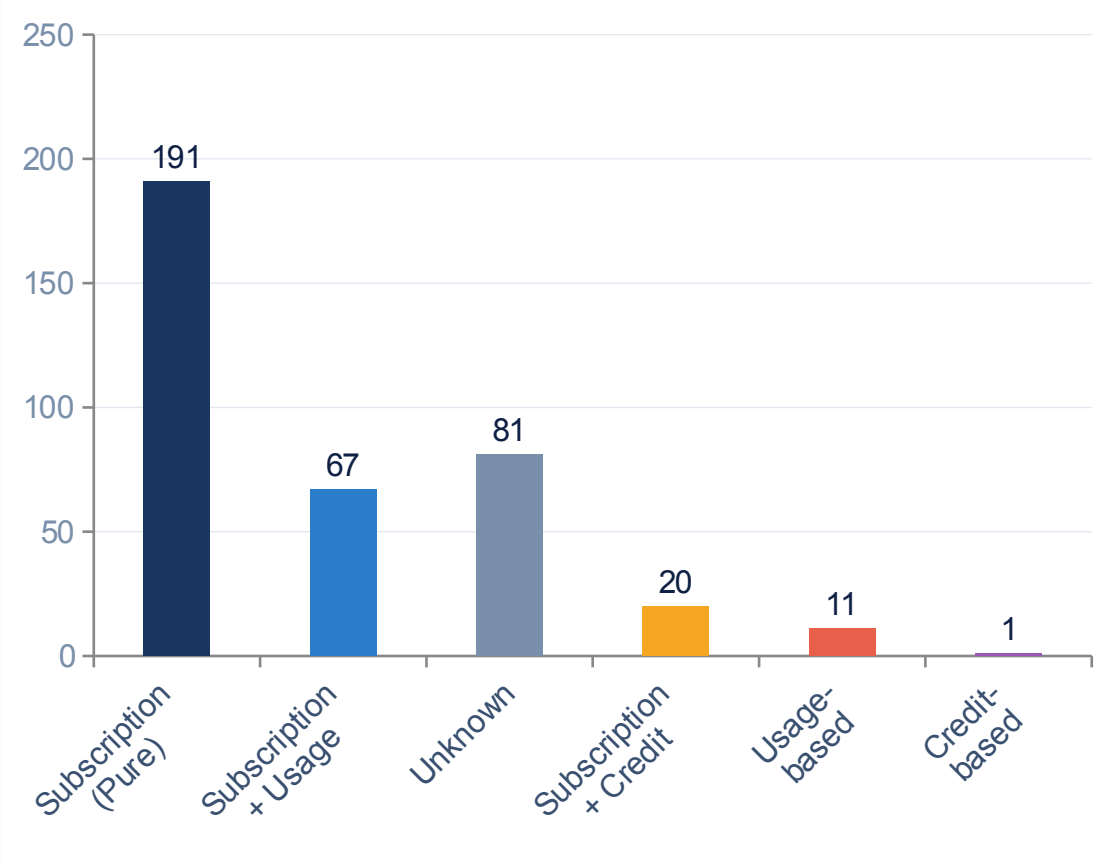
Revenue Tier Distribution



Inclusion / Exclusion Criteria

- ✓ B2B SaaS company with $1M–$100M ARR
- ✓ Publicly accessible pricing information
- ✓ Identifiable AI-powered features
- ✗ "Contact Us" as sole pricing option
- ✗ Agencies, consulting, or non-SaaS firms

06 | Pricing Model Prevalence (n=371)



51.5%

of firms use pure subscription for AI

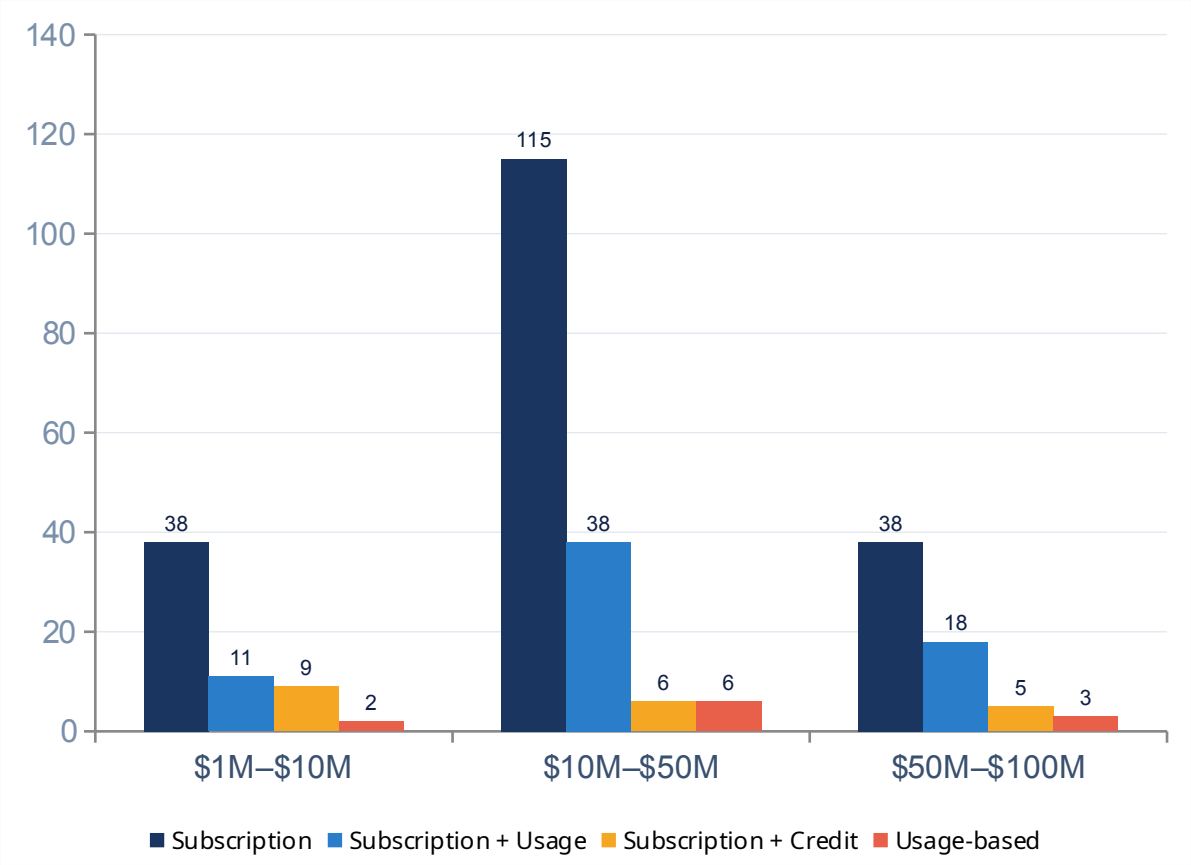
18.1%

hybrid subscription + usage-based models

5.4%

Key Finding: Pure subscription dominates despite industry discourse favoring usage-based models.

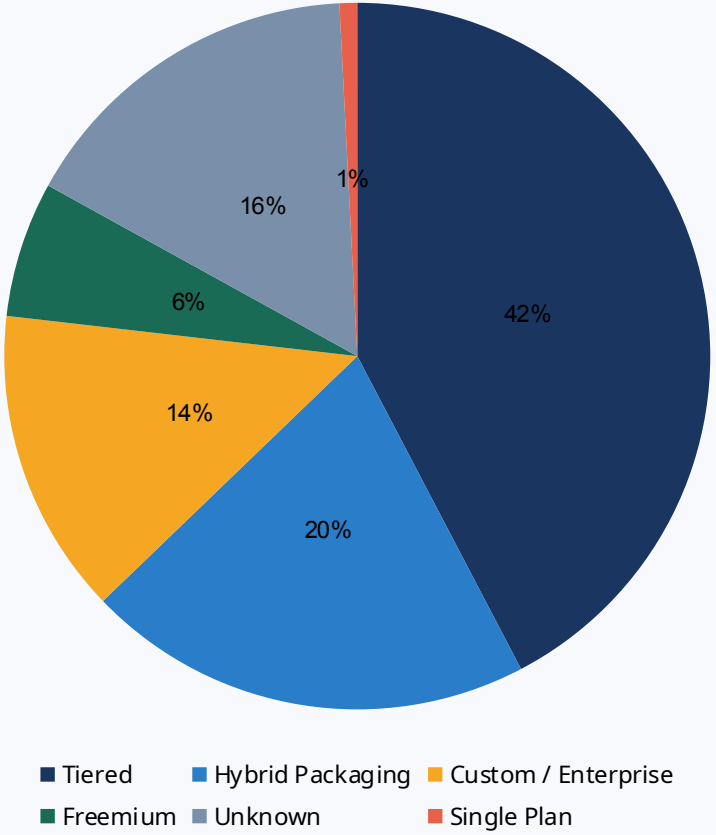
06 | Pricing Strategies by Revenue Tier



Cross-Tier Insights

- Pure subscription scales with firm size — dominant at ALL revenue tiers
- Hybrid 'Sub + Usage' grows proportionally as firms mature into higher ARR bands
- Pure usage-based remains consistently rare regardless of company size
- Higher ARR tiers show more explicit credit-based hybrid adoption

06 | AI Packaging Strategies



Tiered (42.3%)

AI capabilities mapped across plan levels (Basic→Pro→Enterprise) as a feature differentiation lever.

Hybrid (20.5%)

Basic AI in core tier; advanced AI as separate paid add-on. Bifurcated strategy capturing tiered value.

Custom/Enterprise (14.0%)

AI access requires direct sales engagement. Common for complex enterprise deals.

Freemium (6.2%)

Basic AI free up to a limit then paywalled. Useful for PLG adoption funnels.

07 | Qualitative Themes — Executive Insights

9 Senior Executives · Verticals: ERP, Real Estate, Customer Analytics, Retail, FinOps, Portfolio Mgmt, UX Validation · Braun & Clarke Thematic Analysis

Theme 1
Managing AI Cost Uncertainty

Executives treat AI infra costs as primary constraint in packaging decisions.

- Guardrailing LLM usage to customer-specific trained models only
- Absorbing costs short-term to achieve adoption at scale
- Bifurcating deterministic (chargeable) vs. probabilistic (free, guardrailed) models

Theme 2
Customer Resistance to Usage-Based Pricing

Enterprise buyers strongly prefer predictable billing; reject granular consumption metrics.

- "Avoid over-complicating usage models early on" — real estate exec
- Using credits only for highly distinct, tangible outputs (e.g., AI e-receipts)
- Pricing for job-to-be-done, not raw token consumption

Theme 3
Competitive Pressure & Feature Parity

Consumer GenAI tools anchor B2B customer expectations at near-zero cost.

- Basic AI embedded as baseline to avoid adoption friction
- Competitors set de-facto pricing ceiling for AI features
- Domain-specific hesitancy (e.g., banking/treasury) limits AI monetization

08 | 5-Stage AI Monetization Evolutionary Pathway

Stage 1	Stage 2	Stage 3	Stage 4	Stage 5
AI Feature Bundling	**Tiered Differentiation**	**Credit-Based Add-ons**	**Hybrid Metering**	**Outcome-Based Pricing**
AI embedded as table-stakes feature; zero incremental charge	Basic AI in lower tiers; advanced AI unlocked at premium tiers	Distinct AI tasks monetized via credits tied to tangible outputs	Subscription base + variable usage caps for cost recovery	Prices tied to measurable business outcomes (the strategic ideal)

Most mid-market firms ($1M–$100M ARR) currently operate between Stages 1–3. Stage 4 adoption grows with ARR maturity. Stage 5 remains aspirational.

08 | 5-Step Strategic Framework for AI Monetization

1

Diagnose Cost Structure

Map direct (per-token API) and indirect (R&D, talent, infra) AI costs. Establish unit economics at target usage volumes.

2

Assess Customer Readiness

Survey buyers on billing preferences. Identify whether consumption or outcome-based metrics resonate. Document anchoring effects.

3

Define the Value Metric

Choose billing unit tied to tangible output (credit, job, MAU). Avoid tokens/compute-minutes for enterprise segments.

4

Design Hybrid Architecture

Establish subscription base for predictability. Layer usage caps or credit overages for cost recovery. Set transparent guardrails.

5

Instrument & Iterate

Deploy real-time usage telemetry. Track margin per customer. Iterate toward outcome-based structures as ROI data accumulates.

08 | Recommendations

Margin Management

- Guardrail LLM usage to customer-specific models until cost baselines stabilize
- Absorb initial AI costs strategically for market penetration, not indefinitely
- Establish daily revenue tracking across product, finance, and GTM

Overage Policy Design

- Default to subscription + soft cap + buyable credit architecture
- Make overages transparent and predictable — avoid surprise billing
- Use credit rollover policies to reduce customer anxiety

Value Communication

- Publish quantifiable ROI metrics — only ~30% of firms currently do this
- Price AI by job-to-be-done, not raw compute consumption
- Deploy AI embedded in daily workflows, not as standalone optional tools

Future Research

- Longitudinal studies tracking AI pricing evolution over 3–5 year periods
- Large-scale performance linkage: pricing model vs. retention/ARPU/NRR
- Experimental pricing studies with A/B testing across customer segments

AI Monetization is a Business Model Transformation

DBA
2026
SIMA

- Pure subscription dominates (51.5%) despite industry enthusiasm for usage-based models
- Hybrid models are the pragmatic equilibrium — balancing cost recovery with enterprise predictability
- Credits abstraction layer is the key mechanism bridging token volatility and CFO-friendly billing
- Competitive parity forces basic AI as table stakes; advanced AI is where premium lives
- AI monetization is not a pricing tweak — it demands organizational agility and billing re-architecture

Research at a Glance — Key Findings Summary

371

SaaS companies analyzed

58%

Offer AI features

51.5%

Use pure subscription

18.1%

Hybrid usage models

42.3%

Tiered AI packaging

3

Key qualitative themes

Thank You

Questions & Discussion

Hitesh Sarda
Doctor of Business Administration
IIBM Institute · SIMA Switzerland · March 2026