Annual Methodological Archive Research Review
http://amresearchreview.com/index.php/Journal/about
Volume 4, Issue 3 (2026)

Online ISSN 3007-3197 Print ISSN 3007-3189
http://amresearchreview.com/index.php/Journal/about

Pricing the Machine: A Framework for Valuing AI Capabilities in B2B SaaS Business Models

Hasnat Saeed
Superior, University Lahore, 54000, Pakistan
Email: hasnat.saeed2026@gmail.com

ABSTRACT

The integration of artificial intelligence into business-to-business software-as-a-service (B2B SaaS) is straining the pricing logic on which the industry was built. As value decouples from the number of people logging in, the per-seat subscription long the default no longer maps cleanly onto the value a product delivers, and vendors are migrating toward usage-, work-, and outcome-based structures. Yet the field lacks a rigorous, theoretically grounded method for answering the question these shifts raise: how much is an AI capability actually worth to a customer, and how should that worth translate into price? Drawing on two industry surveys of B2B software and AI companies (N = 240 in 2025 and N = 230 in 2026; Growth Unhinged/Tremont), this paper documents the empirical landscape hybrid pricing rising from 27% to 41% of firms within twelve months, seat-based pricing falling from 21% to 15%, AI gross margins clustering near 50% against SaaS norms of 70–80%, and a wave of AI-credit and outcome-based experimentation and identifies a gap between this fast-moving practice and the academic pricing literature. To bridge it, the paper develops the AI Economic Value Contribution (AI-EVC) framework, which extends classical economic-value-to-the-customer analysis with an explicit willingness-to-pay attenuation term, a value-capture parameter, a cost-and-margin floor, and a mapping from estimated value to hybrid pricing architecture. A worked example applied to an autonomous customer-support agent demonstrates the framework end to end and cross-checks its output against observed market prices. The paper closes with implications for pricing strategy, product development, and investor valuation, and with an agenda for empirical validation.

Keywords: AI monetization; B2B SaaS; economic value to the customer; hybrid pricing; outcome-based pricing; technology valuation; value-based pricing; willingness-to-pay; usage-based pricing

INTRODUCTION

In the vast majority of the software industry's modern history the answer to “how shall we price this?” was reassuringly simple. A salesman sold access to an application and billed by the user, by the month. The number of seats was used as a crude measure of value: the more people utilized the product, the more value was being taken out of it, so the more

Annual Methodological Archive Research Review
http://amresearchreview.com/index.php/Journal/about
Volume 4, Issue 3 (2026)

Online ISSN 3007-3197 Print ISSN 3007-3189

http://amresearchreview.com/index.php/Journal/about

seats the more value. Over the last 20 years that proxy was able to work because the cost of serving a single user was near to nothing and the value of software could be loosely argued to scale with the number of users.

AI has cracked the code to the proxy. Value, when an AI can do the work, not just assist a person to do it, no longer represents “the number of people who log in.” A support agent capable of resolving tickets independently, a coding assistant that produces an ever-larger percentage of a particular engineering team's work, or a research agent that supplants hours of analyst effort are all valuable agents whose value can increase in the face of declining human workforce involvement. A customer might require less manpower and more machine if the AI has its intended effect, industry commentators have put it simply. In that world, the seat is a deceptive and a diminishing measure of value.

This decoupling isn't a specialty topic for a few startups built for the AI sector.This decoupling is not just a small issue for a few AI-native startups. Now it's making its way to the biggest players in the business. In the last two years, companies of all sizes have shifted from flat and per-seat pricing to hybrid models where the companies charge a subscription plus a metered fee, and, at the leading edge, to pay-for-work and pay-for-value pricing. The change comes from two arguments: a value argument – the price should be based on what the AI does, rather than who can log in; and a cost argument – unlike traditional software, AI has a real, variable cost of good, so unlimited flat-rate access can be financially devastating.

These forces make it a very hard problem for the industry to solve now, and they're doing it by intuition. When a vendor is going to charge for the delivery (as an AI product) of the AI, rather than access to the AI, it needs to both estimate the value of the delivery to the customer and determine what portion of that value to capture in price. The shortcuts used by the practitioners today are cost-plus reasoning (set the price based on the cost of providing the service to the AI, plus the desired margin) and competitive benchmarking (make the price close to the price of the competitor's service). Both are comprehensible, and both are extremely slender theories. Cost plus pricing sets the price based on the vendor's costs instead of the value to the customer; competitive benchmarking sets the price based on the competitors, who are also guessing. Neither "Yes" or "No" answers the question of value.

This paper is not to be taken lightly. It is designed to provide a solid, academically sound process for measuring the value of AI services in a B2B SaaS product and how to use that measurement to establish a defensible pricing framework. There are three parts to the contribution. First, the paper combines two series of industry surveys data, creating an empirical baseline on where monetization of AI actually stands, and separating out the durable trends from the hype. Second, it points to a gap: a theoretically sound and empirically testable method that connects the attributes of the AI to customer willingness-to-pay (WTP) in the SaaS setting. Third, the authors propose a new framework, the AI Economic Value Contribution (AI-EVC), based on the existing value-based pricing theory, but adapted to the economic context of artificial intelligence, and illustrate it with a worked example whose costing is compared with actual prices in the market.

The rest of the paper is structured as follows. The industry data is presented and analysed in Section 2. Section 3 reviews the literature on pricing and valuations which is the basis for the framework. In Section 4, the research gap is clearly and precisely identified. The

Annual Methodological Archive Research Review
http://amresearchreview.com/index.php/Journal/about
Volume 4, Issue 3 (2026)

Online ISSN 3007-3197 Print ISSN 3007-3189

http://amresearchreview.com/index.php/Journal/about

AI-EVC framework is built up in Section 5 and used in a real-world example in Section 6. The implications for pricing, product development, and valuation of the investor are detailed in Section 7, and Section 8 outlines the limitations of the framework and an agenda to empirically test its validity. Section 9 concludes.

THE CURRENT MONETIZATION LANDSCAPE: INDUSTRY DATA AND ANALYSIS

DATA SOURCES AND A NOTE ON PROVENANCE

This paper's empirical supports include two surveys conducted by Growth Unhinged (and executed by Tremont) in April-May 2025 (with over 240 companies participating) and April-May 2026 (with over 230 companies participating). These are practitioner surveys and three cautionary notes are included. The samples are self-selected from a group likely to be interested in monetization in the first place, which means that they may over-represent companies actively rethinking their pricing strategies. The two waves use different samples of respondents, this means that year-over-year comparisons reflect real change as well as change in samples. And all pricing-model and margin figures are self-reported. The limitations do not compromise the robustness of the data, with healthy sample sizes and strong and consistent patterns, but simply indicate that the figures should be interpreted as directional evidence of a moving market rather than population parameters. Later, the framework is formulated to be estimated using firm-specific data exactly because it is not reliant on generalizability of any one survey.

MARKET SEGMENTATION: SAAS, AI-NATIVE, AND HYBRID PRODUCTS

AI is changing the nature of the market. Half of respondents (50%) said that their product is mostly SaaS, 14% said that their product is AI-native, and 36% said that their product is a hybrid of SaaS and AI (Poyar, 2025). After a year, the share mostly-SaaS decreased to 32%, the share of AI-native rose somewhat to 15%, and a hybrid share rose to 43%, while an additional 10% was categorized as "other" that didn't fit into one of the above buckets (Poyar, 2026). Expect the direction to be clear enough as a result of the difference in samples, too: the pure-SaaS product is being replaced by the product where software and AI are combined. This is important for pricing because the hybrid product has an embedded variable-cost AI element in a product with a very low marginal cost, which has to incorporate the change in the pricing model.

PRICING-MODEL TRENDS: THE RETREAT FROM THE SEAT

The main trend in the headline is away from access pricing and towards hybrid pricing models. For hybrid pricing, the percentage of firms increasing in the last 12 months from 27% to 41% rose, while seat-based pricing dropped from 21% to 15% and flat-fee subscriptions dropped from 29% to 22%, and usage-based pricing dipped slightly from 20% to 17% and outcome-based pricing grew from 3% to 5% (Poyar, 2025). Across the total sample, the primary model with the highest percentage is hybrid (37%), which increased year-over-year from 25% (2026); flat-fee (25%), seat-based (17%), usage-based (15%) and

Annual Methodological Archive Research Review
http://amresearchreview.com/index.php/Journal/about
Volume 4, Issue 3 (2026)

Online ISSN 3007-3197 Print ISSN 3007-3189
http://amresearchreview.com/index.php/Journal/about

outcome-based (5%) are the other primary models (Poyar, 2026).

The two waves are not completely continuous - the 2025 wave results in hybrid adoption at 41%, the 2026 wave results in hybrid adoption at 37% and the discontinuity itself is instructive. It represents the various samples mentioned above and warns against viewing any one percentage as exact. The structural story is that hybrid is now the driving force, the seat is in the minority, while pure outcome-based has yet to become a common way of doing business. In particular, it is hardly the smallest companies that are now on the bleeding edge of pricing and packaging thinking about AI (Poyar, 2026), it's the biggest ones that are.

THE AI-CREDIT WAVE AND THE PERSISTENCE OF PREDICTABILITY

There's one particular model on the rise: the AI credit or the AI token, a finite amount that counts towards using AI within a subscription-based service. Adoption of credit models jumped significantly year over year, with about 29% of firms adopting the models, with a third planning to add them in the next 6-12 months, and adoption expected to reach 62% by 2027, among firms >$50M ARR, about 25% planned to add credits in 6-12 months (Poyar, 2026). Credits are popular because they allow for a seat to be used, essentially as an entitlement to a limited quantity of use, and reconcile the competing interests of vendor control of variable cost and the buyer's interest of budget predictability. This reconciliation is addressed as a design goal in Section 5 of the framework.

THE MARGIN PROBLEM: AI IS NOT ZERO-MARGINAL-COST SOFTWARE

The most impactful fact for valuation is the cost of AI – it comes with a real, variable price tag. Respondents were asked to identify the most important factor when considering the cost of AI, and roughly 54% (2026) and 56% (2025) of respondents indicated internal costs and margins.When asked what is most important in pricing AI, a majority of respondents answered internal costs and margins (roughly 54% in 2026 and 56% in 2025). The impact is reflected in target margins—around 12% of companies are trying to achieve the margins enjoyed by classical SaaS models of 70–80%+ and the median target gross margin for AI capabilities is closer to 50%. A significant portion of respondents already claimed gross margins of less than 60% for their AI products (Poyar, 2025).

It has a rather subtle but important take-home message for a valuation framework. Cost is non-trivial in AI, and thus cannot be ignored, as often happens with zero-marginal-cost software in value-based pricing. However, cost should not be the basis of price as it is with cost-plus practice. The proper definition of cost, that was developed in Section 5, is that it is a floor under a price ceiling determined by the value, not the anchor.

WHERE THE MONEY COMES FROM, AND THE DRIFT TOWARD PAYING FOR WORK

There are two other patterns that context the valuation question. First, most AI spend is eating into existing tech budgets: AI spend is coming out of customers' software budgets, with only 70% of respondents (compared with 75% of AI-native companies introducing new AI features) pointing to AI spend in their customers' services budgets or in headcount (Poyar, 2026). That is, for the time being, most AI capabilities are being judged by the value of the

Annual Methodological Archive Research Review
http://amresearchreview.com/index.php/Journal/about
Volume 4, Issue 3 (2026)

Online ISSN 3007-3197 Print ISSN 3007-3189
http://amresearchreview.com/index.php/Journal/about

software they can replace, and not by the vastly greater numbers of labour and services they could replace.This means that for now, most of the capabilities of AI are being valued by the software they replace, not by a much larger number of labour and services, which is a direct impact on estimated WTP. Second, it is a market that is moving up a value-capture ladder, from access to usage to work done to outcomes, such that each stage of the ladder provides the vendor with an increasing share of the economic value created and transfers the risk from the buyer to the vendor (Poyar, 2025). Outcome-based pricing is at the top of this rungs, and a majority consider this a target, yet it is the primary pricing model for just some 5% of enterprises, and a quarter anticipate to accomplish it in a couple of years (Poyar, 2025).

When viewed collectively, the facts paint a picture of a market that understands the seat's demise, adapted to hybrid structures as a temporary home, has embraced credits and outcomes, and is limited by the actual costs of AI – but is pricing its products primarily by cost-plus and by imitation, simply because it does not have a principled method for estimating value. This is what the rest of the paper will address.

THEORETICAL FOUNDATIONS FOR VALUING AI CAPABILITIES

There is no attempt here to create a completely new pricing theory; the framework presented in this paper is simply the application of a well-developed and established theory to a new object. This section briefly summarizes the four streams from which it gets its water.

VALUE-BASED PRICING AND ECONOMIC VALUE TO THE CUSTOMER

The principle is the product's value to the customer, not its production costs, should determine the price (Nagle, Hogan, & Zale, 2016; Monroe, 2003). Economic value to the customer (EVC) is the construct used by a workhorse, which breaks down the value of a product into a reference value, which is the value of the customer's next best alternative, and a differentiation value, which is the net economic value of the ways the product differs from the customer's next best alternative, positive and negative (Nagle et al., 2016; Anderson, Jain, & Chintagunta, 1993). In business markets, researchers have proven that value can and should be measured monetarily in relation to the buyer's business (Anderson, Narus & van Rossum, 2006) and that a convincing value proposition is built on identifying a handful of value drivers that are important to the buyer. The backbone of EVC is the location of the ceiling on the rational buyer's willingness to pay, which is done through EVC, and the disciplined approach of constructing that ceiling upwards from identifiable drivers.

MEASURING WILLINGNESS-TO-PAY

In principle, EVC defines the value of a product, and in practice, willingness-to-pay sets the price buyers are willing to pay, which is a gap that arises due to perception, uncertainty and behaviour. There are a variety of well-developed methodological approaches to estimating WTP: direct methods include the price-sensitivity meter (Van Westendorp, 1976) and Gabor–Granger sequences, and preference-based methods include conjoint and discrete-choice analysis, which derive the value of individual attributes from choices among bundles (Green & Srinivasan, 1978; Louviere, Hensher & Swait, 2000). For AI valuation these methods are important because they allow a vendor to calculate the marginal WTP for a

Annual Methodological Archive Research Review
http://amresearchreview.com/index.php/Journal/about
Volume 4, Issue 3 (2026)

Online ISSN 3007-3197 Print ISSN 3007-3189
http://amresearchreview.com/index.php/Journal/about

new autonomous resolution, for a latency guarantee, or anything else, without calculating the overall product WTP precisely what they need to know when making their valuation decision.

VALUE CAPTURE, INFORMATION GOODS, AND USAGE-BASED PRICING

It's one thing to create value it's another thing to capture it. The price that a vendor charges is somewhere between its cost floor and the value ceiling of its customer, and the price it charges is a strategic decision that depends on its competition, switching costs, and maturity of the market. The economics of information goods comes to the fore: software's near-zero marginal cost has made it subject to value, versioning, rather than cost to consumers; differential pricing and bundling have been essential to appropriating differential WTP (Shapiro & Varian, 1998). This inheritance is complicated by the advent of AI, which brings with it a very real marginal cost, making it imperative that usage based and metered structures, which have been explored in telecommunications, utilities, and more recently software, are central to the picture. It has been found that buyers are not apathetic towards tariff choice; they have systematic preferences, typically a preference for the flat tariff because of the insurance it offers and its predictability, even if the usage-based tariff is cheaper (Lambrecht & Skiera, 2006). This is one of the reasons why the market is converging on hybrid models and on credits which are based in a predictably contained credit envelope.

PRODUCTIVITY, OUTCOMES, AND UNCERTAINTY

Finally, valuing AI requires a view of what AI does economically. Emerging evidence quantifies AI's productivity effects directly: a large field study of a generative-AI assistant deployed to customer-support agents found substantial productivity gains concentrated among less-experienced workers (Brynjolfsson, Li, & Raymond, 2025). Such estimates are exactly the kind of value-driver inputs an EVC calculation needs. But AI value is also uncertain and variable outcomes differ across customers and drift as models change which brings the economics of decision-making under uncertainty into play. Real-options reasoning, developed for irreversible investments under uncertainty (Dixit & Pindyck, 1994), offers a lens for capabilities whose value is unpredictable but potentially large, and motivates the temporal, adaptive component of the framework. The recurring practitioner observation that outcome-based pricing requires consistency, attribution, measurability, and predictability of outcomes (Poyar, 2025) is, in effect, a statement about the conditions under which this uncertainty is low enough for value to be priced directly a point the framework formalizes.

THE RESEARCH GAP

The above industry data and literature from academia are slightly mismatched. Theory to guide the movement towards pricing, though good at parts, has not been put together in a framework for this particular object, but practice is moving at a quick pace along a clear path toward something that looks like it reflects what AI does. Three aspects of the gap should be made quite clear.Three are very specific aspects of the gap that should be stated.

NO INTEGRATED VALUATION FRAMEWORK FOR AI CAPABILITIES

Annual Methodological Archive Research Review
http://amresearchreview.com/index.php/Journal/about
Volume 4, Issue 3 (2026)

Online ISSN 3007-3197 Print ISSN 3007-3189
http://amresearchreview.com/index.php/Journal/about

Value-based pricing is a full-fledged literature, as is WTP measurement, value-capture strategy, and the economics of uncertainty, but each has not been synthesized into a process for putting a value on an AI capability in a SaaS product. The classical EVC construct does not necessarily fit the object that has material and variable costs, uncertain and drifting value, and a natural pricing unit (work performed, outcomes achieved) that is different from the one at which the value is easiest to measure. There is a lack of a framework that brings these features together.

PREVAILING PRACTICE IS THEORETICALLY UNGROUNDED

The survey evidence shows what fills the vacuum: cost-plus reasoning and competitive imitation. Internal cost and margin were the most-cited pricing factor in both waves (Poyar, 2025, 2026), which is a candid admission that most firms price AI from the inside out from their expenses rather than from the customer's value. Competitive benchmarking, the second most common approach, inherits whatever errors the benchmarked firms are making. Neither approach estimates customer value, and neither can tell a vendor whether it is leaving money on the table or pricing itself out of a deal. The prevalence of dissatisfaction in the data expansion revenue was the most-cited pain point, and no pricing model's adopters reported being fully satisfied (Poyar, 2026) is consistent with a market pricing without a value compass.

THE CAUSAL LINK TO WTP IS UNVALIDATED

Finally, the assumption underlying the entire migration that pricing aligned to AI-delivered value will be accepted by buyers and will expand revenue has not been empirically validated for this context. The market has strong priors (outcome pricing is the "holy grail") but thin evidence on whether, and under what conditions, buyers actually pay more when price is tied to AI value rather than access. A framework is valuable here not only as a pricing tool but as a source of testable hypotheses: it should specify the variables whose measurement would confirm or refute the value WTP link. Section 5 is built with that dual purpose in mind.

THE AI ECONOMIC VALUE CONTRIBUTION (AI-EVC) FRAMEWORK

It adds four concepts unique to AI to the classic economic-value-to-the-customer analysis: the explicit reduction of value into a realizable willingness-to-pay, the introduction of a strategic value-capture parameter, the presence of a binding cost-and-margin floor, and the mapping of the resulting price band onto a hybrid pricing architecture. It is broken into 7 steps.

STEP 1 — DEFINE THE JOB, THE VALUE UNIT, AND THE REFERENCE ALTERNATIVE

The customer's job is the starting point for valuation, and the unit of value is the value of the customer in a unit (in this case the time of the analyst, a resolved support ticket, a qualified lead, a completed legal deliverable, or an hour of analyst time saved). The reference value R is the customer's cost of being able to do that job with the best alternative that is available to him: an incumbent software tool, in-house labour, or an outside service. The

Online ISSN 3007-3197
Print ISSN 3007-3189
http://amresearchreview.com/index.php/Journal/about

selection of the reference is important. Anchoring the AI to the software it replaces gives a relatively small R, anchoring the AI to the labour/s it replaces gives a much larger R. What the survey results have shown is that for now, the biggest share of AI spend is happening on software projects (Poyar, 2026), which suggests that there's significant unpriced value for vendors that can credibly market themselves as an alternative to labor or services.

STEP 2 — ESTIMATE THE DIFFERENTIATION VALUE

The differentiation value ΔV is the total economic value added by the AI capability compared to the reference, based on a limited number of measurable drivers, minus the cost incurred by the customer to adopt the AI capability:

$$\Delta V = (D\_labor + D\_revenue + D\_risk + D\_speed + D\_quality) - (D\_switch + D\_oversight)$$

Here D_labor is the value of labor substituted (tasks automated × loaded labor cost × automation rate); D_revenue is the value of revenue uplift due to the capability (for example, higher conversion or retention); D_risk is the value of the errors, fraud or compliance that can be saved by the capability; D_speed is the value of the cycle-time compression (priced at the customer's opportunity cost of time), and D_quality captures experience or capability improvements that cannot be counted, typically estimated by WTP methods rather than accounting figures. These are contrasted by D_switch, the cost of integration and change management, and D_oversight, the cost of human oversight, quality assurance, and the cost of the remaining risk of AI outputs going wrong. The contribution of each positive driver should be estimated conservatively, and be described in terms which are relevant to the customer's own business.

STEP 3 — COMPUTE ECONOMIC VALUE AND ATTENUATE IT TO REALIZABLE WTP

The reference value plus the differentiation value, is the reference value to the customer; it is the highest price a completely rational, completely informed buyer will pay:

$$EVC = R + \Delta V$$

Real buyers pay less than EVC, and by an amount that is not random. The framework introduces an attenuation factor $\alpha \in (0,1]$ that discounts EVC into realizable willingness-to-pay W:

$$W = \alpha \cdot EVC$$

The factor α captures how much of the theoretical value the buyer will credit to the vendor and pay for, and it is governed by the same four conditions practitioners have identified as prerequisites for outcome pricing consistency, attribution, measurability, and predictability (Poyar, 2025). When outcomes are consistent across customers, clearly attributable to the vendor, measurable in near real time, and predictable in advance, α approaches one and value can be priced directly; when they are idiosyncratic, contested, opaque, or volatile, α falls and the vendor must retreat toward pricing on inputs it can measure. Modeling α explicitly turns a vague sense that "outcome pricing is hard" into an estimable parameter, and identifies the levers better measurement instrumentation, stronger

Annual Methodological Archive Research Review
http://amresearchreview.com/index.php/Journal/about
Volume 4, Issue 3 (2026)

Online ISSN 3007-3197 Print ISSN 3007-3189
http://amresearchreview.com/index.php/Journal/about

attribution evidence that raise it.

STEP 4 — SET THE VALUE-CAPTURE RATIO

Realizable WTP is what the vendor could charge; how much of it the vendor should charge is a separate, strategic decision represented by the value-capture ratio $\gamma \in (0,1)$:

$$P\_target = \gamma \cdot W = \gamma \cdot \alpha \cdot (R + \Delta V)$$

A low γ leaves value with the customer to accelerate adoption and expansion—appropriate in competitive or land-and-expand contexts while a high γ maximizes near-term capture, appropriate where differentiation is strong and defensible. The empirical value-capture ladder (Poyar, 2025) indicates that the higher the price, the more the vendor should get paid for the outcomes achieved, and the more likely that the pricing basis will be outcome-based rather than access-based. γ is not just a number; it is the combination of a pricing basis and the aggressiveness of the pricing.

STEP 5 — IMPOSE THE COST-AND-MARGIN FLOOR

AI has to overcome a cost threshold that is different from zero-marginal-cost software. Suppose that the variable cost of providing the capability at the relevant volume is C, and the desired gross margin is m. The smallest value of the price is:

$$P\_floor = C / (1 - m)$$

The possible price range is then [P_floor, P_target]. The survey data suggesting that median target AI margins are near 50% (Poyar, 2026) suggests that P_floor is significantly elevated, relative to the cost, compared to classical SaaS and that the band can be narrow or even empty, which is exactly when a vendor should differentiate to increase ΔV or withdraw from the capability. Importantly, cost is not a determinant of price, it is a constraint on price – it is still value-based. This is a turnaround from the cost-plus model that is generally accepted in the industry, and still adheres to the economic truth that AI is not free to serve.

STEP 6 — MAP THE PRICE BAND ONTO A HYBRID ARCHITECTURE

One number is NOT a pricing model. The last construction step serves to transform the feasible band into a hybrid model that simultaneously serves the floor, captures value and provides the buyer with the predictability that the tariff-choice literature indicates that he is asking for (Lambrecht & Skiera, 2006). A generic hybrid price is of the type:

$$P = F + u \cdot q + s \cdot O$$

where F is a fee per platform or subscription, u is a per-unit usage or credit amount for usage quantity q, and s is a success or outcome rate for the realized amount of outcome O. The composition is determined by the structure of ΔV and the value of α. Lock-in and predictability are provided by a platform fee F, to recover the reference value and the fixed costs. A usage or credit term u·q matches fixed variable revenue to variable variable cost,

Annual Methodological Archive Research Review
http://amresearchreview.com/index.php/Journal/about
Volume 4, Issue 3 (2026)

Online ISSN 3007-3197 Print ISSN 3007-3189
http://amresearchreview.com/index.php/Journal/about

ensuring that margin is safeguarded by the heavy users that AI credits are now becoming better at doing. An outcome term s·O is added only to the degree α is sufficiently large to allow it; this is a portion of ΔV, for which outcomes are measurable and can be attributed. That's why the market has settled on hybrids and credits, not usage or outcome pricing the hybrid is the structure that enables a vendor to cover the real cost floor, capture the value, and sell the predictability.

STEP 7 — ADD A TEMPORAL, ADAPTIVE LAYER

As value and cost of both drift models improve, the cost of tokens decreases and the outcomes change, it is a fixed price that decays. The framework thus sets out a re-estimation schedule and adaptive mechanisms: credit balances that get rolled over to smooth consumption; adaptive tiers that reset at renewal depending on actual consumption; and tracking outcome variance that feeds into α. When the value of the capability is unknown but could be substantial, the real-options perspective (Dixit & Pindyck, 1994) advises retaining flexibility in pricing to gain experience and to implement capped pilots with terms that only become binding if they are successful.

Table 1

Core variables of the AI-EVC framework and their estimation sources.

Term	Meaning	How it is estimated
R	Reference value of best alternative	Price of incumbent tool, loaded labor cost, or service fee for the same job
ΔV	Net differentiation value	Sum of quantified value drivers minus switching and oversight costs (Step 2)
EVC	Economic value to the customer	R + ΔV; the rational-buyer ceiling
α	WTP attenuation factor	Function of outcome consistency, attribution, measurability, predictability
W	Realizable willingness-to-pay	α · EVC; cross-checked with conjoint / Van Westendorp estimates
γ	Value-capture ratio	Strategic choice given competition, maturity, and land-and-expand goals
C, m	Variable cost and target margin	Token/compute cost per unit; target gross margin (survey median ≈ 50%)
F, u, s	Platform fee, usage price, success rate	Chosen to span [P_floor, P_target] within a hybrid architecture

ILLUSTRATIVE APPLICATION: AN AUTONOMOUS CUSTOMER-SUPPORT AGENT

To show the framework end to end, consider a vendor pricing an AI agent that autonomously resolves customer-support tickets one of the most active arenas for outcome-based experimentation, where observed market prices provide a reality check. The figures below are illustrative but chosen to be plausible.

Annual Methodological Archive Research Review
http://amresearchreview.com/index.php/Journal/about
Volume 4, Issue 3 (2026)

Online ISSN 3007-3197 Print ISSN 3007-3189
http://amresearchreview.com/index.php/Journal/about

Step 1 (reference). The customer currently resolves tickets with human agents at a loaded cost of roughly $5 per ticket. Because the agent replaces labor, the reference is framed against labor rather than software, giving a per-unit reference value of about $5.

Step 2 (differentiation). Suppose the agent autonomously resolves 50% of a customer's 100,000 monthly tickets at quality parity, i.e., 50,000 resolutions. The labor-substitution driver is 50,000 × $5 = $250,000 per month. Against this, oversight and quality assurance cost roughly 10% of that value (−$25,000), and amortized integration adds a smaller deduction (−$10,000). Setting aside harder-to-quantify retention effects for conservatism, net $\Delta V \approx$ $215,000 per month.

Step 3: EVC and attenuation. The software's reference value is insignificant compared to the value of the labour that it replaces, which is around $215,000 per month for EVC. Resolution of tickets is measurable and much of the cause; consistency and predictability of customers is only moderate, with $\alpha \approx 0.6$ so $W \approx$ $129,000 per month is realizable.

Step 4 (capture). In a competitive market with a "land and expand" approach, the vendor will have to set $\gamma \approx 0.5$, which gives a target of approximately $64,500 per month, or approximately $1.29 per ticket resolved.

Step 5 (floor). Token and compute cost is about $0.30 per resolution, so $C \approx$ $15,000 per month; at a 50% target margin, P_floor = $15,000 / (1 − 0.5) = $30,000 per month, or $0.60 per resolution. Given these parameters, the possible range of resolution per dollar is therefore approximately $0.60 to $1.29 per resolution.

Reality check. Market prices for similar products that are autonomous resolution fall in the range of $0.50 to $0.99 per resolution, which is in line with market rivalry that drove per-resolution prices downwards (Poyar, 2026). The frame's value floor and bracket this range above, whereas the cost-and-margin floor is just within the range. The right reading is not that the model is incorrect, but that competitive pressure is currently pushing γ to the down side, and that the cost bottomline is the limiting factor for “laggards” when they go ahead and differentiate. If a vendor is forced to charge more than about $0.90 per resolution, they will be pushed out, and this is exactly what the margin data forecast.

Step 6 (architecture). The vendor charges a platform fee F for integration and a base platform fee, a per-resolution fee u that mimics the variable cost, and, if a customer's outcome is easily measurable in monetary terms, an option for success fee $s \cdot O$—already widespread in this class of fees.

PRACTICAL IMPLICATIONS

FOR PRICING STRATEGY

It moves the focus in the middle pricing question from "what does it cost us plus a margin?" to "what is the viable range between our cost floor and the customer's attenuated value, and where in the range should we be?" It assigns to cost its rightful place as a bottom, it allows to estimate the value of ceiling, and it shows γ as a deliberate lever, and not a chance of mimicry. It also describes, rather than just explains, the convergence of hybrids and credits in the market, because these are the only ways to sell predictability buyers want, within a real

Annual Methodological Archive Research Review
http://amresearchreview.com/index.php/Journal/about
Volume 4, Issue 3 (2026)

Online ISSN 3007-3197 Print ISSN 3007-3189
http://amresearchreview.com/index.php/Journal/about

cost floor and value ceiling.

FOR PRODUCT DEVELOPMENT

As it's constructed from named drivers, the framework can also serve as a guideline for what to build, and what to instrument. Capability that moves the biggest drivers of labor substitution, revenue uplift, risk reduction are the ones worth monetizing, and it is explicit in the business case for investment in instrumentation and attribution infrastructure that the more you can instrument the more you can raise α, which raises the realizable WTP. That is, the ability to prove value is a value driver, and product teams that embed outcome measurement in the product are not just increasing the number of telemetry points they're growing the surface area for value.

FOR INVESTOR VALUATION

There's a direct read-through of how to value AI businesses in the framework. The AI margin reality (a median close to 50% as opposed to the 70–80% norms of classical SaaS) suggests that the revenue may be inferior to SaaS revenue and the γ and α terms suggest where a specific firm is likely to be on the value-capture ladder and how long it will be on the ladder. A firm price on outcomes with high α and defensible γ has a structurally superior economics profile when compared to a reseller selling "commoditized" tokens at a razor-thin margin, at the same headliner ARR. Applying the same SaaS multiples to these companies will lead to mispricing; the framework provides some words for the diligence that separates them.

LIMITATIONS AND AN AGENDA FOR EMPIRICAL VALIDATION

There were three restrictions on the contribution. First, the empirical anchors are surveys by a self-selected sample of practitioners and self-reported figures, and the two waves are not strictly comparable, the framework reduces but doesn't eliminate this dependence, as it is estimable using firm-specific data. Secondly, it is a conceptual framework that has not yet been validated as an instrument: the theoretical derivation of the parameters α and γ in particular has not been calibrated by field data. Third, the numbers in the worked example are merely illustrative—they do not imply claims of actual data.

The limitations set the boundaries for a research agenda. The test of whether α behaves as postulated could be done in experiments that operationalize and measure outcome consistency, attribution, measurability and predictability, such as controlled conjoint and discrete choice experiments, which estimate how much actual WTP rises with improvements of these characteristics. In collaboration with vendors during mid-transition, field studies can identify if increases in price realization due to access-based to value-based pricing increase revenue and retention, as predicted by the value-capture ladder, and estimate γ from price realisations. As costs and model quality fluctuate over time, the longitudinal data can be used to test the temporal layer to determine if adaptive and credit mechanisms can be used to beat static pricing. The driver decomposition of the framework can also be tested with independent productivity evidence like in the form of evidence that is now becoming available (Brynjolfsson et al., 2025). The point is that each of these makes a part of the framework a

Annual Methodological Archive Research Review
http://amresearchreview.com/index.php/Journal/about
Volume 4, Issue 3 (2026)

Online ISSN 3007-3197 Print ISSN 3007-3189
http://amresearchreview.com/index.php/Journal/about

falsifiable claim; the framework does not provide a definitive answer, but rather a series of hypotheses about the relationship between AI value and price.

CONCLUSION

The seat is failing as a unit of value because AI has decoupled value from the number of people using software. The industry has responded with a rapid, well-documented migration toward hybrid, usage-, and outcome-based pricing, but it is navigating that migration largely by cost-plus reasoning and imitation, because it lacks a principled way to answer the question the migration raises: what is an AI capability worth to a customer, and how should that worth become a price?

This paper has argued that the pieces of an answer already exist in the value-based pricing, WTP-measurement, value-capture, and uncertainty literatures, and that what has been missing is their assembly into a framework fitted to the economics of AI. The AI Economic Value Contribution framework provides that assembly: it builds value up from named drivers, attenuates it into realizable willingness-to-pay through an explicit and estimable factor, separates the strategic question of value capture from the analytic question of value, respects AI's real cost floor without letting cost become the basis of price, and maps the result onto the hybrid architectures the market has already discovered by trial and error. Put in a real context, it provides an upper and lower bound on observed prices, and makes sense of the competition forces that are squeezing prices.

The importance of the framework may be greater with respect to use as a research tool. It establishes a series of propositions to test the relationship between AI delivered value and willingness-to-pay, turning a rapidly evolving but little-thematized practice into one that is testable. That's the work that needs to be done – and it's work worth doing – to turn those propositions into calibrated, validated estimates, where the promise of AI is ultimately realised as a viable business.

DECLARATIONS

Acknowledgement: We appreciate the generous support from all the contributor to the research and their different affiliations.
Funding: No funding body in the public, private, or nonprofit sectors provided a particular grant for this research.
Availability of data and material: In the approach, the data sources for the variables are stated.
Authors' contributions: Each author participated equally in the creation of this work.
Conflicts of Interest: The author declare no conflict of interest.
Consent to Participate: Yes
Consent for publication and Ethical approval: Because this study does not include human or animal data, ethical approval is not required for publication. All authors have given their consent.

Annual Methodological Archive Research Review
http://amresearchreview.com/index.php/Journal/about
Volume 4, Issue 3 (2026)

Online ISSN 3007-3197 Print ISSN 3007-3189

http://amresearchreview.com/index.php/Journal/about

REFERENCES

Zhang, L., Xuan, Y., Liu, Z., Du, Z., Wang, S. and Wang, J., 2025. A hybrid ensemble model to detect Bitcoin fraudulent transactions. Engineering Applications of Artificial Intelligence, 141, p.109810.

Anderson, J. C., Jain, D. C., & Chintagunta, P. K. (1993). Customer value assessment in business markets: A state-of-practice study. Journal of Business-to-Business Marketing, 1(1), 3–29.

Anderson, J. C., Narus, J. A., & van Rossum, W. (2006). Customer value propositions in business markets. Harvard Business Review, 84(3), 90–99.

Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. The Quarterly Journal of Economics, 140(2), 889–942. https://doi.org/10.1093/qje/qjae044

Dixit, A. K., & Pindyck, R. S. (1994). Investment under uncertainty. Princeton University Press.

Green, P. E., & Srinivasan, V. (1978). Conjoint analysis in consumer research: Issues and outlook. Journal of Consumer Research, 5(2), 103–123. https://doi.org/10.1086/208721

Lambrecht, A., & Skiera, B. (2006). Paying too much and being happy about it: Existence, causes, and consequences of tariff-choice biases. Journal of Marketing Research, 43(2), 212–223. https://doi.org/10.1509/jmkr.43.2.212

Louviere, J. J., Hensher, D. A., & Swait, J. D. (2000). Stated choice methods: Analysis and applications. Cambridge University Press.

Monroe, K. B. (2003). Pricing: Making profitable decisions (3rd ed.). McGraw-Hill/Irwin.

Nagle, T. T., Hogan, J. E., & Zale, J. (2016). The strategy and tactics of pricing: A guide to growing more profitably (6th ed.). Routledge.

Poyar, K. (2025). The state of B2B monetization in 2025. Growth Unhinged.

Shapiro, C., & Varian, H. R. (1998). Information rules: A strategic guide to the network economy. Harvard Business School Press.

Van Westendorp, P. H. (1976). NSS price sensitivity meter (PSM): A new approach to study consumer perception of prices. In Proceedings of the ESOMAR Congress (pp. 139–167). ESOMAR.