

Comparing Pricing Units in Agentic Outsourcing Bids

A Commercial Management Method for Equivalent Offers

Phil Hatch | Akholi

October 23, 2026

Commercial managers should translate seat, token, task, transaction, outcome, and hybrid prices into a common functional comparison. A low price per unit is meaningful only when the unit, eligible population, completion standard, variability, and retained buyer work are understood.

Define What Each Pricing Unit Buys

The manager should identify the event that creates a charge and the result it does or does not establish. A seat purchases access, a token measures consumption, a task can represent activity, and an outcome can represent an accepted business result if defined clearly. None is inherently superior for every service.

Mayer Brown's provider contracting analysis describes activity, outcome, consumption, and hybrid commercial options.¹ The manager should evaluate the risks and definitions of each rather than assume an outcome label guarantees alignment with value.

For a hypothetical support service, one bid may charge per interaction, another per resolved case, and a third a platform fee plus usage. Compare them against the same eligible case population and acceptance standard. An interaction can be inexpensive while several are needed to resolve one case.

¹ Scott F. Young and Jon McPherson, "Contracting Opportunities and Challenges for Emerging Agentic AI Companies," Mayer Brown, October 13, 2026, provider contracting analysis of authority, data, liability, and service terms, <https://www.mayerbrown.com/en/insights/publications/2026/10/contracting-opportunities-and-challenges-for-emerging-agentic-ai-companies>.

Translate Offers into Complete Functional Cost

Include base fees, minimum commitments, consumption, tools, retries, exceptions, human review, correction, transition, and retained buyer work. Identify which charges vary with volume or complexity and which remain fixed. Avoid comparing an all-inclusive service price with a technology-only unit.

The study reports manual work and process limitations that can be omitted from an automation claim.² It does not establish the cheapest pricing unit. It identifies why a common cost boundary should follow the accepted function rather than the provider's billing boundary.

Separate ordinary and difficult cases. A consumption price can vary with context length and repeated attempts, while a fixed outcome price may include exclusions or volume assumptions. Preserve those conditions in the comparison rather than hide them inside a single average.

Use a Pricing Unit Comparison Record

Figure 1 proposes four questions: what triggers payment, what result is included, what variability remains, and what work the buyer retains. Record the source of each assumption and distinguish observed pilot evidence from a forecast.

LSEG's connector analysis notes that agentic consumption depends on tool selection, calls, and returned data.³ It is a product-provider perspective, not a guarantee of customer savings. The practical implication is to test the assembled workflow rather than estimate usage from the opening prompt alone.

Use a small set of transparent workload scenarios with stated assumptions for volume, case mix, retries, review, and quality. Do not invent probabilities without evidence. Show the conditions under which the preferred bid changes so that the decision owner can judge the relevant uncertainty.

² Phil Hatch, *Agentic AI Outsourcing at Tier 1 Financial Institutions: Full Study Report*, version 1.8.2 (Akholi, forthcoming manuscript dated November 1, 2026), Part 5, "The Five Manual Alteration Methods" and "Disclosure of Manual Alteration," table 5.8, <https://doi.org/10.2139/ssrn.7164559>; Hatch, *Agentic AI Outsourcing*, Part 2, "The Shift From Managing People to Managing Agents"; Part 6, "The Executive Action Plan."

³ Emily Prince, "The Hidden Economics of AI Agents—and How LSEG's MCP Connector Helps Improve Token Efficiency," LSEG, July 17, 2026, licensed content, permissioned access, and customer-controlled agent workflow, <https://www.lseg.com/en/about-us/new-at-lseg/the-hidden-economics-of-ai-agents-and-how-lsegs-mcp-connector-helps-improve-token-efficiency>.

Figure 1. The pricing unit comparison record

Comparison element	Required definition	Commercial question
Trigger	Seat usage task transaction or outcome	What creates a charge
Result	Included scope and accepted completion	What does payment buy
Variability	Volume complexity retries and adjustment	How can total charges change
Retained work	Buyer review correction and dependencies	What remains outside the price

Author's proposed operating framework. Adapt to the service and applicable jurisdiction; not an empirically validated standard.

Source: Phil Hatch, proposed framework developed for this paper.

Examine Incentives and Adjustment Rules

A price per attempt can reward unnecessary activity; a price per closure can reward premature closure; a fixed fee can encourage restrictive scope. These are design risks to examine, not proof of provider misconduct. Align the unit with the intended outcome and a usable quality and challenge mechanism.

EU, UK, Swiss, US, and other services may have different conduct, safety, professional, tax, privacy, and contractual conditions. Qualified owners should establish material local effects. A global price unit should not override duties that remain part of the function.

Define overage, minimum volume, exclusions, changed scope, late buyer prerequisites, and later reversal. Preserve an auditable denominator and a route to dispute classification. A low quoted unit can become expensive when essential work moves into an undefined change request.

Select the Commercial Structure the Function Supports

A reasonable objection is that complete comparison requires information providers will not disclose. The buyer needs its expected charges, included result, dependencies, and retained burden. It can compare those facts without demanding every element of supplier margin or internal cost.

Before award, reconcile a representative case set and invoice calculation for shortlisted bids. Confirm that the provider and buyer apply the same unit and acceptance rules. Use a bounded initial commitment where uncertainty could materially reverse the ranking.

The manager should select pricing that remains understandable when the service runs. The strongest offer is the one whose complete cost, accepted result, variability, and incentives fit the buyer's needs. A low unit price is a starting point for that comparison, not its conclusion.

Evidence and Source Note

The study is a nonrepresentative cross-section of 850 vendor delivery-team respondents on 150 engagements serving Tier 1 financial institutions, collected during January–June 2026. Neither banks nor vendors corroborated the accounts. It establishes reported experience, not causation, legal breach, or market prevalence. Its application here is a mechanism to investigate, not an estimate of conditions in other industries. The author-supplied version 1.8.2 is dated November 1, 2026 and is cited as a forthcoming manuscript consulted October 23, 2026.⁴

This is an analytical decision paper based on the supplied study and selected primary public sources checked through October 23, 2026. The original contribution is a pricing-unit comparison record translating different billing structures into equivalent functional cost and incentive decisions. The proposed method has not been independently validated. Legal requirements depend on the entity, activity, jurisdiction, and effective provisions; supervisory observations, research, and professional commentary have distinct authority.

Bibliography

- Hatch, Phil. *Agentic AI Outsourcing at Tier 1 Financial Institutions: Full Study Report*. Version 1.8.2. Akholi. Forthcoming manuscript dated November 1, 2026. Author-supplied manuscript consulted October 23, 2026. <https://doi.org/10.2139/ssrn.7164559>.
- Prince, Emily. "The Hidden Economics of AI Agents—and How LSEG's MCP Connector Helps Improve Token Efficiency." LSEG. July 17, 2026. <https://www.lseg.com/en/about-us/new-at-lseg/the-hidden-economics-of-ai-agents-and-how-lsegs-mcp-connector-helps-improve-token-efficiency>.
- Young, Scott F., and Jon McPherson. "Contracting Opportunities and Challenges for Emerging Agentic AI Companies." Mayer Brown. October 13, 2026. <https://www.mayerbrown.com/en/insights/publications/2026/10/contracting-opportunities-and-challenges-for-emerging-agentic-ai-companies>.

⁴ Hatch, *Agentic AI Outsourcing*, Introduction, "Executive Summary" and "Evidentiary Basis and Its Limits."