

The Full Cost of an Outsourced Agentic Service

A Joint Cost Model for Sourcing Managers

Phil Hatch | Akholi

October 23, 2026

Sourcing finance managers should build a joint delivery cost model and a separate buyer expenditure view. That distinction exposes hidden work without counting provider costs twice or treating released staff time as cash savings.

Set Two Views of Cost

The sourcing finance owner should distinguish what the buyer pays from what the two firms consume to deliver the service. The buyer view supports investment and sourcing decisions. The joint delivery view explains effort, dependencies, and improvement opportunities. Combining them without a reconciliation can double-count labor already included in provider fees.

Define the business function, workload, acceptance standard, period, currency, and comparison before collecting cost. Include normal, failed, corrected, and pending cases. A model based only on successful agent runs can omit the effort needed to produce the actual service and overstate the economics of scale.

McKinsey's agentic economics analysis recommends considering fixed and variable costs, including oversight. It provides a useful cost classification, not a universal rate card. Managers should obtain local transaction and effort evidence for the model rather than transplant an external percentage.¹

¹ Chandana Asif et al., "Where AI Agents Pay Off: A Practical Guide to the Economics of Agentic Workflows," McKinsey & Company, August 24, 2026, "What drives the unit economics of AI-based workflows?," <https://www.mckinsey.com/capabilities/quantumblack/our-insights/where-ai-agents-pay-off-a-practical-guide-to-the-economics-of-agentic-workflows>.

Create a Cost Line with an Owner and Evidence

Each line should record the activity, responsible firm, cost behavior, allocation basis, source, and uncertainty. Buyer review, provider exceptions, integration maintenance, assurance, retries, recovery readiness, and change work should appear where relevant. Identify whether a cost is recurring, transitional, shared, or already included elsewhere.

The study's manual-alteration accounts illustrate work that can sit outside the intended process. Its return figures were reported by delivery teams and were not independently verified financial results. Managers should use the evidence to investigate omitted effort, not use the study as a substitute for their own cost ledger.²

Measure effort with an appropriate method and protect staff and customer information. Time records, case sampling, invoices, system events, and interviews can contribute different evidence. State where estimates replace observation and avoid precision that the collection method cannot support.

Reconcile the Buyer View to Delivery Economics

Figure 1 proposes a reconciliation. In an explicitly hypothetical service, annual buyer cost includes \$600,000 in provider fees, \$120,000 in retained review, and \$80,000 in attributable platform and assurance cost. The buyer total is \$800,000. If provider internal delivery cost is \$450,000, do not add it to that total because the buyer already pays the fee.

A separate joint resource view can replace the fee with the \$450,000 provider cost, producing \$650,000 in identified delivery resources. That comparison explains a different question and may require confidential commercial evidence. It is not the buyer's expenditure and does not establish the provider's net profit because other costs and allocations may remain.

If the period contains 80,000 accepted outcomes, buyer unit cost is \$10. The model should separately show failed attempts, pending work, and capacity. These figures are hypothetical arithmetic, not reported study results. Their purpose is to make the boundary and potential double counting visible.

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, "Five Manual Alteration Methods" and "Disclosure of Manual Alteration," table 5.8, <https://doi.org/10.2139/ssrn.7164559>; Hatch, *Agentic AI Outsourcing*, Introduction, "Executive Summary"; Part 1, "Customer Satisfaction, Delivery Performance, and Measured Return."

Figure 1. Keep the buyer and joint delivery cost views separate

Cost item	Buyer expenditure view	Joint resource view
Provider delivery	Count provider fee once	Count underlying resources instead
Buyer retained work	Include attributable buyer cost	Include same identified resources
Shared and transition work	Allocate once under stated method	Record ownership and consumption
Accepted outcomes	Use agreed acceptance denominator	Explain failed pending and repaired work

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.

Compare Alternatives on the Same Functional Basis

Evaluate human delivery, deterministic automation, redesigned workflows, and bounded agents against the same accepted outcome and service requirements. An agent can change the process, so old phase costs may no longer align. Compare the function delivered rather than force every new activity into the old process labels.

Run explicit uncertainty cases for acceptance, retained review, retry rates, model consumption, and workload. Explain the conditions behind each case instead of presenting an unsupported probability distribution. A service can remain viable across a range of assumptions or depend on one fragile estimate; managers should make that distinction visible.

BCG's 2026 analysis emphasizes tracking outcomes and building operating capability alongside AI strategy. It supplies market context for disciplined measurement. Its survey associations should not be treated as causal proof or as a forecast for this sourcing decision.³

Track the Benefit Through to Financial Capture

Record released hours, redeployed capacity, avoided future cost, and realized cash savings separately. A team that finishes work faster may absorb more demand without reducing

³ Jessica Apotheker et al., "The Formula for Agentic AI Value," Boston Consulting Group, September 30, 2026, "The Formula for Success," <https://www.bcg.com/publications/2026/the-formula-for-agentic-ai-value>.

expense. That can be valuable, but the business case should state the value mechanism and evidence rather than convert every hour into a cash saving.

When data are incomplete, show a bounded range and identify the evidence that would narrow it. For example, unrecorded buyer review effort may dominate the difference between alternatives. A short, properly designed time-and-case study can then be more useful than refining token-price assumptions. The sourcing owner should prioritize measurement according to its effect on the decision, with unresolved assumptions visible in the approval record.

For cross-border services, retain local cost and control differences. EU, UK, Swiss, and US operations may have different acceptance requirements and retained duties. State exchange-rate and inflation assumptions, and show material regional differences before aggregation. The cheapest blended average may conceal a location with unacceptable performance.

The commercial review should connect cost changes to authority and dependencies. More autonomous execution can increase assurance or recovery needs. A lower model charge can coincide with higher total consumption. The manager should explain the net change in the full service, not report a single input-price movement as savings.

Maintain the Model as an Operating Record

Assign cost lines to owners and reconcile them at an agreed service-review interval. Investigate variance through cases and process evidence. Repeated human repair should trigger a delivery redesign decision, while necessary controls should remain funded rather than be removed to improve the apparent business case.

The completed model should let a sourcing decision-maker see what the buyer pays, what delivery consumes, which outcomes are accepted, and how uncertainty affects the comparison. Keeping those views separate produces a more useful negotiation and a more defensible investment decision.

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 reconciled buyer expenditure and joint resource cost model, with independently calculable hypothetical examples. 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

Apotheker, Jessica, Vinciane Beauchene, Nicolas de Bellefonds, et al.. “The Formula for Agentic AI Value.” Boston Consulting Group. September 30, 2026. <https://www.bcg.com/publications/2026/the-formula-for-agentic-ai-value>.

Asif, Chandana, Dieter Kiewell, Lari Hämäläinen, Tom Kolaja, and Tunde Olanrewaju. “Where AI Agents Pay Off: A Practical Guide to the Economics of Agentic Workflows.” McKinsey & Company. August 24, 2026. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/where-ai-agents-pay-off-a-practical-guide-to-the-economics-of-agentic-workflows>.

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>.

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