

Testing Promised AI Savings Before the Deal

An Executive Decision on Provider Delivery Economics

Phil Hatch | Akholi

October 23, 2026

Executives should test whether a supplier can deliver promised savings while completing the full accepted service. The commercial decision needs evidence of total functional cost, retained buyer work, transition requirements, and the conditions under which savings remain credible.

Translate the Savings Promise into a Service Claim

The executive should ask what the supplier will deliver at the proposed price and which costs remain elsewhere. A reduction in billed staff or unit processing cost is not the same as a lower cost for the completed business function. Define the comparable outcome, quality, timing, and risk conditions before evaluating the headline percentage.

KPMG's US survey reports organizations seeing AI value alongside greater attention to costs.¹ It is executive self-reporting across AI initiatives, not proof of outsourced agentic savings. The evidence argues against assuming either universal failure or automatic return for a particular deal.

Require a bridge from the current functional price-performance baseline to the proposed service. Include provider charges, technology consumption, retained buyer work, exception handling, assurance, transition, and necessary recovery capability. Separate one-time investment from recurring cost and stated assumptions from demonstrated results.

¹ KPMG, "AI's Value Story Sharpens as Organizations Gain Confidence in Governance, Accountability and Workforce Adoption," KPMG, September 24, 2026, US survey methodology and reported value and cost review findings, <https://kpmg.com/us/en/media/news/q3-ai-pulse-2026.html>.

Locate the Work Behind the Automation Claim

The study's manual intervention and delivery-process findings identify costs that a simple automation rate can obscure.² They do not quantify another supplier's cost base. The buyer should test whether the proposed price includes the human work and buyer decisions necessary to finish cases.

For a hypothetical claims-administration service, an inexpensive automated first pass may generate additional requests, reviews, and reopened cases. The commercial comparison should follow accepted final outcomes. A low initial transaction price can coexist with a higher total functional cost.

Ask the provider to identify the operating changes on which its price depends. These may include buyer data access, process redesign, faster decisions, minimum volume, narrower case scope, or standardized tools. Price those prerequisites and assign responsibility before treating the savings as committed.

Use a Savings Credibility Record

Figure 1 proposes four tests: baseline comparability, complete delivery cost, dependency readiness, and reconciled pilot results. The supplier should demonstrate representative cases and disclose manual work, exclusions, retries, and buyer effort. Reconcile case evidence to the commercial model.

Differentiate a lower price funded by efficiency from a temporary concession, a changed scope, or a future renegotiation assumption. All can be legitimate commercial choices, but they imply different continuity and renewal risks. The executive should understand which mechanism supports the offer.

McKinsey's agentic economics analysis examines costs beyond model usage, including orchestration and repeated work.³ The buyer's assessment should therefore test the assembled delivery method. A forecast based only on token prices omits important components of the proposed service.

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 Lari Hämäläinen et al., "Is That AI Agent Worth It? Agentic Economics and the Modern Operating Model," McKinsey & Company, July 13, 2026, "Why the operating expenditure bill keeps rising" and "The implications for competitive dynamics," <https://www.mckinsey.com/capabilities/quantumblack/our-insights/is-that-ai-agent-worth-it-agentic-economics-and-the-modern-operating-model>.

Figure 1. The savings credibility record

Commercial test	Evidence required	Executive decision
Baseline	Comparable function quality and timing	What is the current complete cost
Delivery	Human technology and retained buyer work	What will the proposed service cost
Dependency	Funded prerequisites and adjustment rules	Under what conditions is the price credible
Demonstration	Reconciled accepted outcomes and costs	What commitment does the evidence support

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.

Allocate Dependencies Without Erasing Accountability

Agree what happens when buyer prerequisites are late, volume differs, or the case mix changes. Use observable triggers and defined adjustment rules. A general clause excluding all outcomes affected by customer systems can make the supplier's savings promise impossible to evaluate.

Conversely, the buyer should not require a supplier to absorb unlimited work caused by uncontrolled changes outside its authority. Establish a joint route to classify the cause and approve remedial changes. Maintain the original baseline so that repricing does not erase the ability to assess the business case.

EU, UK, Swiss, US, and other regulatory, tax, privacy, conduct, and labor conditions can affect permissible delivery design and cost. Assess relevant requirements before standardizing the global price. A cheaper location or workflow is not comparable if it cannot lawfully or reliably perform the same function.

Commit in Stages as Evidence Improves

A reasonable objection is that providers cannot reveal every proprietary cost. The buyer does not always need that information. It needs verifiable commitments about accepted outcomes, total charges, dependencies, retained work, quality, and adjustment rules. Protect supplier confidentiality while preserving commercial testability.

Use a bounded pilot or staged commitment to reconcile costs and outcomes before transferring a large workload. Define the continuation decision in advance. If the supplier cannot demonstrate a critical assumption, revise the scope, price, or commitment rather than convert uncertainty into a guaranteed savings claim.

Approve the deal when the savings are supported by a credible delivery method and a complete functional comparison. The executive is buying sustainable performance, not an automation percentage. The right commercial question is whether the proposed service can deliver the required result at the promised total cost.

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 savings-credibility record connecting supplier pricing promises to complete functional cost, buyer prerequisites, and staged commercial acceptance. 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>.

Hämäläinen, Lari, Mark Patel, Sven Blumberg, Tanguy Catlin, and Wasim Lala. “Is That AI Agent Worth It? Agentic Economics and the Modern Operating Model.” McKinsey & Company. July 13, 2026. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/is-that-ai-agent-worth-it-agentic-economics-and-the-modern-operating-model>.

KPMG. “AI’s Value Story Sharpens as Organizations Gain Confidence in Governance, Accountability and Workforce Adoption.” KPMG. September 24, 2026. <https://kpmg.com/us/en/media/news/q3-ai-pulse-2026.html>.

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