

Stage the Commitment and Fund the Learning

An Executive Decision on Agentic Investment Under Uncertainty

Phil Hatch | Akholi

October 23, 2026

Executives should match financial commitment and delegated authority to the evidence available. A staged investment should fund the tests that resolve material uncertainty, with explicit continuation, restriction, and exit decisions rather than an indefinite pilot or an unqualified full-scale promise.

Identify the Uncertainty That Can Change the Decision

The executive should ask which facts remain unknown and how they could affect the business case. These may include accepted quality, human work, buyer readiness, decision latency, data meaning, total cost, or recovery capability. A pilot is useful when it is designed to resolve such questions.

McKinsey's scaling blueprint emphasizes defined scope and a sequenced path from proof of concept to scale.¹ It is practitioner guidance, not a guarantee of successful investment. The buyer should establish evidence thresholds appropriate to its own service and exposure.

For a hypothetical outsourced service, the provider may demonstrate normal cases while the cost and quality of difficult cases remain uncertain. A bounded stage can test that population before the buyer transfers a large workload or removes its existing capability.

¹ McKinsey & Company, "Stacking the Odds: A Blueprint for Successfully Scaling Agentic AI," McKinsey & Company, September 14, 2026, workflow redesign, technical foundation, and operating-model blueprint, <https://www.mckinsey.com/capabilities/operations/our-insights/stacking-the-odds-a-blueprint-for-successfully-scaling-agentic-ai>.

Separate Learning Spend from Proven Return

The study reports manual work and process limitations that can affect a service's economics.² It does not determine the return from a new investment. The executive should include those mechanisms among the assumptions to test rather than treat the study as proof the proposed project will fail.

Define the functional baseline, intended benefit, cost boundary, and acceptance rule. Where historical evidence is missing, reconstruct it transparently and state uncertainty. A learning stage can be justified without claiming that its forecast is already a measured return.

Include the people, access, tools, control work, and provider support needed for a realistic test. A cheap demonstration that omits the conditions of live delivery can produce weak evidence. Learning spend should buy decision-relevant information rather than a polished showcase.

Use a Staged Commitment Record

Figure 1 proposes four decisions: question, test, exposure limit, and continuation rule. Define what evidence would support expansion, redesign, restriction, or withdrawal. Agree the rule before results arrive so that sunk effort does not become the main reason to continue.

BCG's agentic-value analysis emphasizes linking capability to business value.³ The staged approach preserves that connection by releasing commitment against evidence of the function, not merely the number of agents built. A technically successful stage can still reveal an uneconomic delivery model.

Limit spending, volume, action authority, data exposure, and dependency as appropriate. Identify the incumbent capability that must remain until replacement is demonstrated. A stage is not truly bounded if failure leaves the buyer unable to resume the required function.

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 Jessica Apotheke 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>.

Figure 1. The staged commitment record

Stage element	Required definition	Executive decision
Question	Material uncertainty and business consequence	What must be learned
Test	Representative evidence and acceptance rule	How will the question be resolved
Limit	Spend authority volume and dependency	What exposure is permitted while learning
Gate	Expand redesign restrict or stop criteria	What follows from the evidence

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.

Agree Commercial and Operating Conditions Together

The provider should know what it must demonstrate, which buyer prerequisites are supplied, how work is priced, and what happens at the decision point. Protect both firms from an open-ended pilot whose scope expands while acceptance remains undefined.

EU, UK, Swiss, US, and other legal, sector, privacy, procurement, and professional requirements can apply during testing as well as production. Qualified owners should establish the relevant conditions. Calling work a pilot does not automatically remove duties or permit consequential action.

Preserve evidence rights and transition assets from the start. The buyer should be able to use the learning and required records if the relationship does not proceed, subject to agreed lawful rights. Otherwise, the stage may create dependency without a usable continuation or exit decision.

Make the Decision When the Evidence Arrives

A reasonable objection is that staging can slow a valuable opportunity. Set a focused question, proportionate test, decision date, and authorized owner. The aim is to resolve important uncertainty quickly while preserving the ability to act on the result.

At the gate, compare findings with the agreed threshold and document limitations. Expand where justified, redesign where the mechanism is promising but incomplete, and stop where the evidence defeats the case. Do not relabel a failed criterion simply to preserve momentum.

The executive should fund learning that changes commitment. Staging is valuable when it makes uncertainty visible, limits exposure, and produces a clear decision. It becomes another form of drift when the organization keeps investing without establishing what the next stage must prove.

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 staged-commitment record tying learning spend and delegated authority to pre-agreed evidence and continuation 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

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>.

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>.

McKinsey & Company. “Stacking the Odds: A Blueprint for Successfully Scaling Agentic AI.” McKinsey & Company. September 14, 2026. <https://www.mckinsey.com/capabilities/operations/our-insights/stacking-the-odds-a-blueprint-for-successfully-scaling-agentic-ai>.

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