Open Future Forum

OPEN FUTURE FORUM

CFO AI Leverage Report

NOVEMBER 2026

November 2026 | Edition 5 | Research report | Murray Newlands | Open Future Forum

Prepared 2 November 2026. Evidence cutoff: 31 October 2026, 23:59 America/Los_Angeles. Registration exports dated 1 November. Finance buying measures are carried-forward observations; the November contribution is benefit verification and funding discipline. Publication review status is recorded separately in the collection audit.

About this publication: Open Future Forum publishes research for executive and investor audiences. Murray Newlands is its founder. This report separates OFF survey observations, external research, author analysis and illustrative scenarios. Vendor-authored guidance is identified as such and does not constitute an independent product assessment. The collection and supporting methods are available through the OFF research index: https://openfutureforum.com/research/

4,733 substantive body words | 6,200 total words
Tables, FAQs, methods and references are excluded from the body count.

In this report

01 The decision in brief
02 Executive summary
03 What November adds
04 Reading the evidence
05 1. Start with the budget decision the evidence must support
06 2. Separate a faster task from a financial benefit
07 3. Treat payback as a testable commitment
08 4. Require a cost boundary that follows the workflow
09 5. Make benefit acceptance reversible
10 6. Align purchasing authority with operating accountability
11 External research: what it changes in the finance case
12 Putting the process to work
13 The next finance review
14 Questions CFOs should be able to answer
15 Citable facts
16 Methods and limitations
17 References
18 Related research and citation
19 Editable working tool

The decision in brief

Approve expansion when the business can show what changed, who captured the benefit, and which costs remain. Faster work is evidence of capacity; it becomes cash savings only through a documented spending change. November's priority is a benefit record that finance can challenge, reconcile and reverse, connected to explicit funding decisions.

Executive summary

The most consequential number in this report is a carried-forward finding: 211 of 389 respondents selected difficulty proving return on investment as a blocker to larger AI spending. That is 54%, using any mention in a question that accepted multiple answers. It identifies a reported obstacle in this selective executive sample. It does not establish that 54% of enterprise AI projects destroy value, or that respondents had calculated returns and found them inadequate. [BLOCK-01]

The financing picture also resists a simple story. Of the same question-level base of 389, 159 mentioned new budget, 77 existing software budget and 73 budget that would otherwise support headcount. Another 120 selected uncertainty or no dedicated AI budget. These answers can coexist. A company can fund a pilot from a software line, expect hiring avoidance later and still lack a settled expansion budget. A budget category therefore tells finance where to investigate; it does not prove how the investment ultimately paid for itself. [FUND]

This edition develops the missing connection between an operating result and a financial decision. The recommended record separates cash savings, avoided hiring, released capacity, incremental contribution and risk reduction. Each has different evidence requirements. Finance should preserve those distinctions when consolidating benefits, because combining them prematurely can turn several descriptions of the same improvement into several claimed returns.

The report adds a process for qualifying a benefit, monitoring its durability and withdrawing it when the evidence changes. It includes a worked example in which a meaningful reduction in handling time produces useful capacity but does not, on its own, recover the investment in cash. A second calculation shows the precise additional operating change required to reach cash payback. Neither example describes an OFF respondent.

Demanding financial proof too early can suppress learning and strategic options. An exploratory budget should purchase a defined learning result, with an expiry and a maximum exposure. A production expansion should purchase a repeatable business outcome. Giving both investments the same approval test obscures their purposes. The CFO's task is to make those purposes visible and require the evidence appropriate to each.

What November adds

The November exports do not provide new comparable answers to the core funding, purchase-approval, spending-blocker and payback questions used here. Their cumulative bases remain 389, 467, 389 and 389 respectively. Reprinting those measures with a November title is not evidence that financing attitudes changed during October. This report carries them forward explicitly and uses them to frame decisions that remain unresolved. The statistics registry provides the exact question wording, source files and registration-date coverage. [FUND; SIGN; BLOCK; PAY]

The first addition is a benefit qualification process. October's workflow cost-accounting article established how to assemble costs. This edition asks a different question: what evidence permits a claimed benefit to enter a management investment case, and what subsequent event requires its removal? The distinction matters when time savings survive in presentations long after implementation, demand or staffing assumptions have changed.

The second addition is a separation of three decisions often compressed into a single ROI number: whether the workflow produces an acceptable result; whether the organization can use the released resources; and whether it can realize a financial benefit within the relevant budget period. A favorable answer to the first does not answer the other two. Conversely, a workflow may justify investment through service quality or resilience without promising a payroll reduction.

The third addition is a current external comparison. Deloitte's third-quarter CFO research addresses the speed of cost control in large North American companies. Academic research on customer support and METR's account of developer-productivity measurement show why observed results depend on task design, participants and measurement. FinOps guidance helps connect resource spending to a business unit. Together, these sources inform a more discriminating funding process; they do not supply a universal AI return assumption.

The report does not calculate a new market ROI, estimate average enterprise spending from broad bands, or claim a month-on-month movement in CFO optimism. Title-based role cuts can change with classification rules. The main exhibits therefore use the full question bases; a separate methodological note explains why role labels should not be treated as a dedicated CFO panel.

Reading the evidence

OFF's first-party evidence comes from event-registration and application questions. Respondents chose to engage with these events; they were not drawn randomly from the enterprise population. Invitations without an eligible response are excluded. Other registration statuses can remain eligible, so the research population is not an attendance count. Answers describe respondents' understanding of their organization, rather than finance-system records inspected by OFF.

The analysis normalizes respondent identifiers and retains the latest eligible nonblank answer separately for each comparable question. That produces different bases for purchase authority and other finance questions. Multiple selections are preserved, including combinations that would be contradictory under a

forced single-choice interpretation. Percentages are rounded from counts, and multi-select distributions may exceed 100%.

The cost-visibility context comes from a separate operator instrument: 77 respondents answered whether they had real-time visibility into what AI costs to run. That question does not establish workflow-level attribution or the completeness of the cost boundary. It is useful context for the CFO's evidence requirements, but it is not a matched subset of the 389-person financing base. Subtracting the two percentages would not measure a finance accountability gap. [COST]

Throughout, observations are labeled with registry identifiers. Recommendations describe a proposed management practice. Illustrations use explicitly invented operational assumptions and remain outside the survey statistics. The resulting investment process is intended to make a company's own evidence inspectable; it does not certify that company, replace its accounting policies, or imply audited benefits.

1. Start with the budget decision the evidence must support

The funding question is useful because it exposes different commitments hiding behind the phrase AI investment. New budget can support experimentation without immediately displacing a cost. Reallocated software budget requires a real contract or renewal decision before it becomes a saving. A headcount-related source may reflect an avoided future hire, a changed staffing plan or simply a sponsor's intention. Those are different cash flows and should enter different lines in the investment case.

In the carried-forward distribution, new budget is the largest positive funding source at 159 of 389, or 41%. Uncertainty or no dedicated budget is selected by 120, or 31%. The remaining sources are existing software budget, 77 or 20%, and budget otherwise allocated to headcount, 73 or 19%. Because respondents could select more than one answer, the result cannot be presented as a pie chart of mutually exclusive financing routes. It also cannot identify which route is most successful. [FUND]

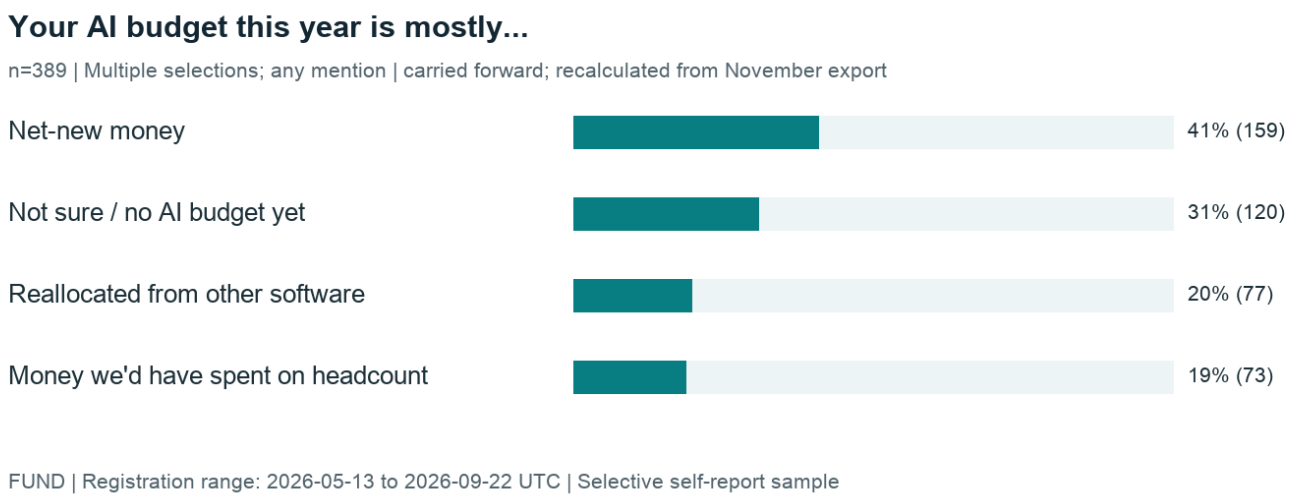
For a CFO, the useful next question is whether the proposed source can actually fund the next stage. A team may have unused license budget this quarter but face an annual commitment next quarter. A hiring plan may have room for one fewer requisition while the AI workflow requires a new specialist elsewhere. A business case should therefore specify the budget owner, decision date, contractual constraint and amount available to move. Without these details, the proposed source remains a financing intention.

An investment can still be rational when it creates no near-term savings. A company may pay to reduce waiting time, improve an important service or learn whether a new product is feasible. The discipline is to name that purchase. An experiment should have a learning question and a cost ceiling. A service-quality investment should have a quality target and an accountable service owner. A cost-reduction project should identify the expenditure that will change and the authority to change it.

This classification also changes approval timing. Finance should not wait for the annual budget cycle to discover that a successful pilot has no production owner. Before a pilot begins, the sponsor should identify who could authorize expansion and what evidence that person will require. That is not an advance commitment to expand. It is a way to avoid purchasing an experiment whose result cannot inform a real decision.

The practical test is a short funding statement: this amount buys this stage of work; this owner can move the relevant budget; this evidence determines the next release of capital; and these conditions end the commitment. The statement is more useful than an undifferentiated AI budget because it connects money to a decision that can be reviewed.

Exhibit 1. Reported AI funding sources



Selection	Count	Base	Any mention
New budget	159	389	41%
Not sure / no dedicated AI budget	120	389	31%
Existing software budget	77	389	20%
Budget otherwise used for headcount	73	389	19%

Exhibit note: FUND, carried-forward buying instrument. Selective application sample; multiple answers permitted. Counts and exact wording are in the registry. Registration dates are not verified answer-edit dates.

2. Separate a faster task from a financial benefit

The first benefit qualification question is what actually changed in the work. A measured reduction in drafting time may be real even if total completion time is unchanged. Review may take longer, more cases may need correction, or the saved minutes may fall in fragments that cannot be redeployed. Finance should ask for the end-to-end process boundary before converting a local improvement into a monetary claim.

A defensible comparison begins with the accepted output. For a support workflow, that might be a resolved case that meets the service standard and does not reopen within a defined period. For document preparation, it might be an approved document rather than a generated draft. The organization chooses the definition before examining the result, records exclusions and preserves enough evidence to reproduce the

count. The definition need not be elaborate, but it must survive the temptation to count whatever the new system produces most easily.

The next question is who can use the capacity. A team with a persistent queue may convert shorter handling time into additional completed work. A team constrained by approvals, demand or another department may not. A staffing model based on average minutes also needs to account for scheduling, peaks and specialist coverage. Reducing routine effort does not automatically remove the need for a person who handles rare but essential exceptions.

Finance should consequently maintain separate benefit categories. Cash savings require a reduced cash outflow relative to an appropriate baseline. Avoided hiring requires a credible staffing requirement and a documented decision not to fill it. Released capacity requires measured time and a named destination for that time. Revenue contribution requires evidence connecting the change to additional contribution after delivery costs and plausible alternative explanations. Risk reduction requires a defined exposure and a supportable assessment, with uncertainty retained rather than hidden inside a precise dollar figure.

These categories are not a hierarchy of human value. Better service may matter more than a small saving. They are a way to prevent an investment case from borrowing the certainty of a cash claim while relying on an untested capacity assumption. A sponsor can request approval for capacity, provided the approving executive knows what is being purchased.

The hardest cases involve more than one category. A team may use part of the released time to avoid a hire and the rest to improve service. Allocate the underlying hours once. If the same improvement is credited to both lower staffing needs and greater output without a consistent workload model, the case double counts its benefit. Finance should resolve that overlap before consolidation, when the operational owner can still explain the work.

3. Treat payback as a testable commitment

The carried-forward question about the expected timing of measurable return shows that 208 of 389 respondents selected less than six months, 101 selected six to twelve months, 37 selected twelve to twenty-four months and 53 selected uncertainty. The question permits combinations; these are reported expectations, not observed project payback periods. The wording asks about measurable return, which need not mean recovery of the entire investment; this report keeps formal payback calculations separate from that survey measure. A short stated horizon may reflect pressure for proof, small initial deployments or optimism about implementation. The data does not distinguish those explanations. [PAY]

The important financing implication is that a payback expectation must specify its starting point. A team can make the same project appear fast or slow by starting the clock at procurement, first deployment or stable production. Finance should define the investment period consistently and show preproduction spending separately. Costs incurred while cleaning data, integrating systems and training reviewers do not disappear because the calculation begins on launch day.

The numerator and denominator need equal care. Initial investment includes the relevant implementation and transition cash flows. Recurring net benefit includes recurring costs, the cash effects actually realized

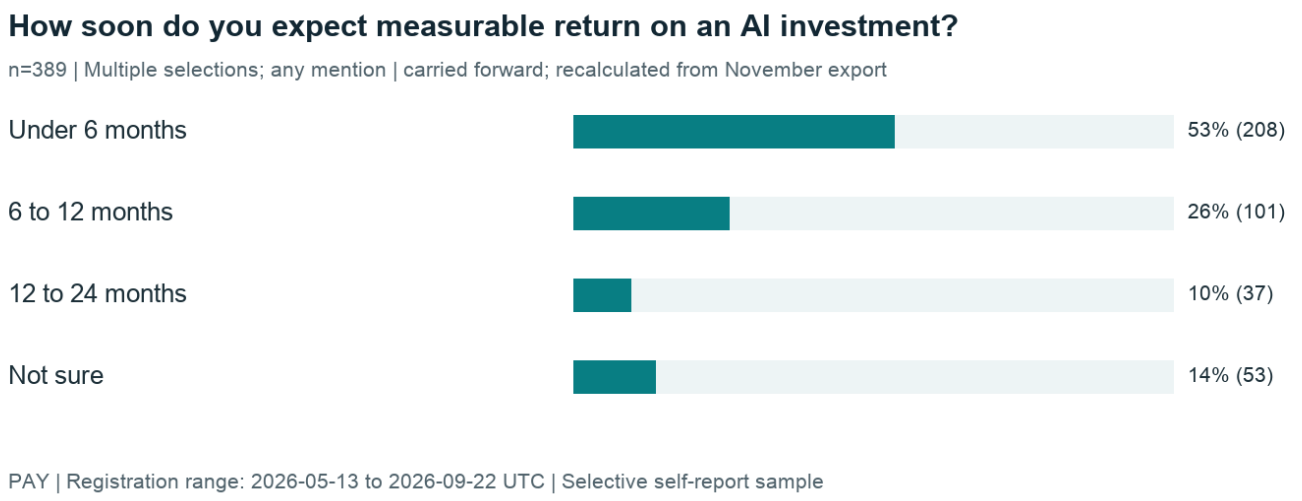
and their timing. A nominal annualized run rate is not the same as cash collected or expenditure removed during the year. A project that reaches its target in the final month cannot claim twelve months of benefit in that first year.

Some cases should not use simple payback as the only decision rule. A workflow with material operational dependence may warrant a resilience investment whose benefit is uncertain but whose failure exposure is concrete. A new commercial product may require staged learning before revenue can be estimated credibly. Finance can use scenarios and bounded commitments for those cases, while keeping the speculative components distinct from verified operating results.

The review should include a reversal condition. If the workload falls below the level needed to recover fixed costs, the previously accepted payback estimate changes. If human review expands after quality problems, net benefit changes. If a vendor introduces a new fee, or a source system becomes unavailable, the model must be reopened. The purpose of documenting these conditions is to make revision routine rather than reputationally costly.

A useful approval records a base case, a downside case and the decision that follows each. The downside is not merely a lower number in a spreadsheet. It should identify an action: narrow the workflow, renegotiate a commitment, delay hiring assumptions, restore a manual step or stop expansion. A forecast becomes governable when management knows what evidence would cause it to change course.

Exhibit 2. Expected timing of measurable return



Selection	Count	Base	Any mention
Less than six months	208	389	53%
Six to twelve months	101	389	26%
Twelve to twenty-four months	37	389	10%
Not sure	53	389	14%

Exhibit note: PAY, carried-forward buying instrument. The question asks about measurable return, not necessarily full investment recovery. Multi-select shares exceed 100%; no observed project payback distribution can be inferred.

4. Require a cost boundary that follows the workflow

The separate operator instrument provides a reason to examine cost completeness. Of 77 respondents, 34 answered yes to real-time visibility into what AI costs to run, 30 answered partly and 13 no. These counts are carried forward. They describe reported visibility, not verified workflow attribution or an audit of cost systems, and they do not establish the quality of any individual organization's calculation. [COST]

The finance consequence is straightforward: a unit-cost claim needs a stated boundary. Model charges may be easy to meter while integration support, human checking, data preparation and exception handling sit in other budgets. Omitting those costs can make an AI workflow appear cheaper by moving expenditure outside the comparison. Equally, charging all shared platform spending to the first application can make a useful workflow appear uneconomic. The allocation method must answer the decision being made.

For an expansion decision, distinguish incremental costs from costs already committed. If additional volume uses spare infrastructure, its marginal economics can differ from the economics of the platform as a whole. Management should see both. The incremental case explains whether the next unit is attractive; the fully loaded case explains whether the overall service can support its infrastructure and people over time. Using only one view can conceal either an attractive expansion or an unsustainable platform.

Human work deserves its own measurement. Review time should include normal checking, escalations and rework attributable to failures. Specialist reviewers may carry a higher opportunity cost than the people whose drafting time falls. The organization should avoid pricing both sets of hours at one convenient average if the staffing decision depends on different skills. At the same time, imputed labor value must remain distinct from an actual reduction in payroll or purchased services.

The denominator should count accepted outcomes consistently. If a workflow generates more drafts but approval rates deteriorate, cost per draft can fall while cost per approved output rises. If difficult cases are routed back to people, report that routing and include the relevant human cost. A narrow automation boundary can be legitimate; the business case becomes misleading when its exclusions are invisible.

Finally, allocation is a management convention that should remain stable enough for comparison. Record the driver, the owner and the reason for changes. A change in allocation should not be described as operating improvement. October's OFF cost-accounting article provides the detailed assembly method. November

adds the requirement that every benefit claim reconcile to the same workflow boundary, so gains cannot be counted against a conveniently incomplete cost base.

5. Make benefit acceptance reversible

Difficulty proving ROI is the largest reported spending blocker in this instrument, but the other selections matter. Integration appears in 89 of 389 answers, security in 80, data in 79 and talent in 45. These are not mutually exclusive explanations. A sponsor may struggle to demonstrate value precisely because integration work is incomplete or because the system lacks access to the information needed for reliable results. Treating ROI as a purely financial reporting problem would miss those operating dependencies. [BLOCK]

The proposed benefit record should therefore contain a state and a reason. A claimed benefit begins as proposed. It becomes observed when the workflow comparison supports the operating change. It becomes accepted for a specified management decision when the benefit owner and finance agree that the conversion into capacity, cash or contribution is supported. It becomes suspended when a material assumption fails, and retired when the workflow or claim is no longer relevant. These are recommended management states, not survey categories or accounting recognition rules.

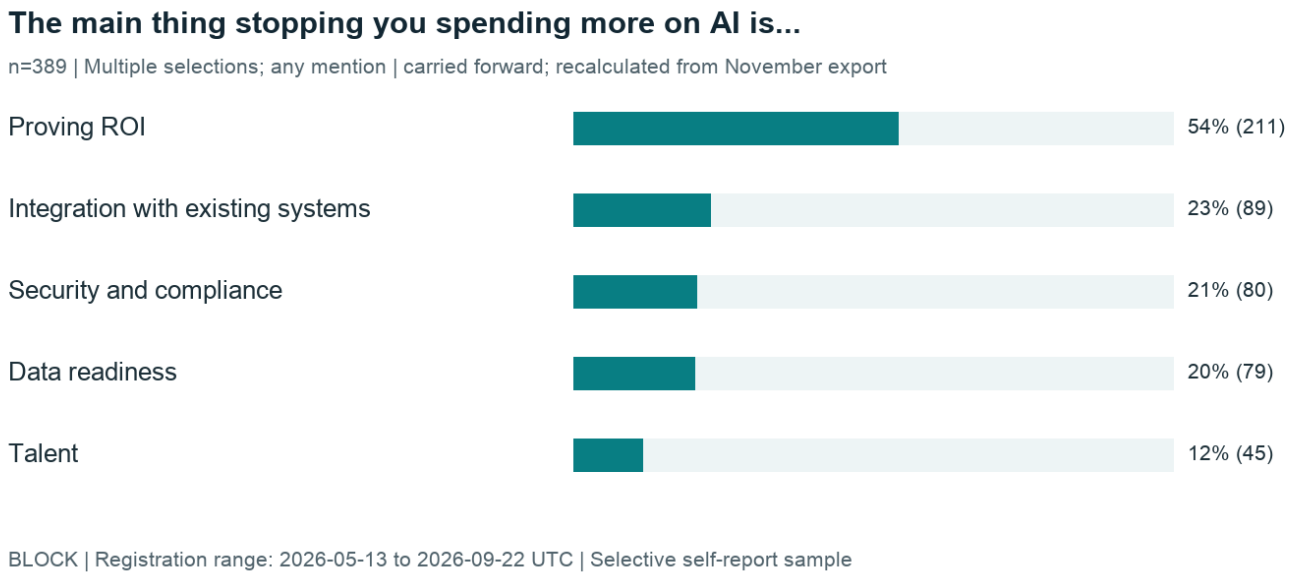
Acceptance should have a defined scope. Finance may accept a benefit for a renewal decision without accepting it for an external earnings statement. A pilot may support an expansion to one team without supporting a company-wide staffing assumption. Recording the permitted use of the evidence avoids turning a local finding into a universal claim as it moves through presentations.

The record also needs an independent challenge. The person accountable for delivering the project has valuable operational knowledge but also a reason to favor continuation. Finance should ask a different owner to check the baseline, exclusions and conversion logic. This need not become a separate bureaucracy: an existing financial planning partner, service owner or internal control reviewer may perform the role. The key is that acceptance is not simply self-certification by the sponsor.

Reversal triggers should be observable. Examples include a material rise in rework, a lost customer segment, a changed vendor contract, a staffing decision that did not occur or an interruption that changes the service boundary. The company should choose materiality thresholds in its own context. Publishing one arbitrary threshold for every workflow would conceal differences in scale, tolerance and reversibility.

A suspended benefit does not automatically mean the technology should be shut down. It means the claim is no longer supported for its stated purpose. Management may continue a bounded experiment, retain a quality improvement or revise the operating plan. This distinction makes honest reporting easier: teams can correct an overstatement without pretending that every useful part of the investment has failed.

Exhibit 3. Spending blockers reported by respondents



Selection	Count	Base	Any mention
Proving ROI	211	389	54%
Integration	89	389	23%
Security	80	389	21%
Data	79	389	20%
Talent	45	389	12%

Exhibit note: BLOCK, carried-forward buying instrument. These are stated obstacles to spending more, not observed causes of project failure. Multiple selections are retained.

6. Align purchasing authority with operating accountability

The purchase-approval question has a larger base than the other buying questions. Among 467 respondents, 227 mentioned the CEO, 144 the CFO, 79 the CTO or CIO, 45 a business-unit leader and 60 no clear owner. The instrument permits multiple selections, so these results describe reported involvement rather than a single hierarchy of final authority. The difference in denominator reflects question completion and source coverage, not a sudden change in the size of the same panel. [SIGN]

A company should translate that ambiguity into an explicit approval path. The budget owner authorizes spending; the operating owner accepts the workflow result; the security and legal owners assess risks within their mandates; and an executive resolves conflicts that exceed delegated authority. One person may hold several roles in a small company. The record should still separate the decisions, because each depends on different evidence.

The CFO should pay particular attention to commitments that outlive the operating experiment. Minimum consumption, annual licenses, data-access fees and integration commitments can make a technically reversible pilot financially difficult to stop. Procurement should identify the earliest effective exit date and the cash that remains due if usage stops. A project described as cancellable may still create substantial stranded expenditure.

Expansion rights should also be bounded. Approval for a limited workload does not necessarily authorize broader data access, a new customer population or autonomous financial actions. Those changes can alter delivery cost, review requirements and exposure. Finance should require a short change assessment when the economic or operating boundary changes, rather than treating the original purchase approval as permanent permission for every later use.

This is where the finance process connects to the CEO and AI leader reports. The CEO establishes which executive can resolve competing objectives. The AI leader supplies production evidence. The CFO determines which economic claims are accepted for a funding decision. A single shared workflow identifier can connect those records without making every function maintain the same document or repeat the same review.

The desired result is a decision that can be reconstructed. An executive who was not present at the approval meeting should be able to see the amount authorized, the evidence considered, the assumptions still open, the next review date and the person empowered to act if conditions change. That level of clarity helps a company move faster because it reduces the need to renegotiate basic authority whenever a forecast or model changes.

External research: what it changes in the finance case

The published customer-support study by Brynjolfsson, Li and Raymond reports a 15% productivity increase across 5,172 agents in its setting. Its measure is issues resolved per hour, rather than company profit. [E-QJE] This distinction is useful for finance because it locates the evidence before the conversion into a financial result. A CFO can accept a well-measured operating improvement while still requiring evidence about demand, staffing, service quality and expenditure. The study supports investigating benefits; it does not provide a percentage to insert into unrelated workflows.

METR's February 2026 update explains that selection and time-measurement problems prevent a reliable current estimate from its newer developer study. [E-METR] The implication for an internal investment review is methodological: do not assume that an apparently larger dataset automatically answers the intended question. If people decline the comparison condition, or run multiple tasks in parallel without usable time records, the observed participants and recorded durations may no longer support the comparison. Finance should record those limitations before attaching a financial conversion to the result.

Deloitte's third-quarter 2026 CFO survey concerns 200 North American CFOs at organizations with at least $1 billion in annual revenue. It reports that 74% would take several business days to several weeks to pause affected AI use after cost-control failure. [E-DELOITTE3] OFF's buying questions do not measure this response time. The sources therefore address different parts of the investment process: one concerns

purchasing expectations and obstacles, the other the ability to intervene after spending begins. A strong approval process can still be undermined by a slow response to actual consumption.

FinOps guidance distinguishes resource-level measures from business-unit economics. [E-FINOPS] For this report, the practical extension is to choose a denominator that the operating owner recognizes and finance can reconcile. Token counts may explain part of an invoice; accepted customer outcomes explain what the business purchased. Both are necessary when usage changes, because a lower technical unit price can coexist with a higher cost per useful result.

The sources also impose boundaries on interpretation. A large-company CFO survey should not be pooled with OFF's selective event sample. A customer-support field study is not a test of every agentic workflow. Developer measurement difficulties do not prove that AI lacks value. A management framework is guidance, not a measured return. The common lesson is to preserve the causal and accounting steps between technical activity and financial consequence, while allowing different investment purposes to use different tests.

Source	Relationship to OFF evidence	Finance implication	Limit
QJE customer-support study, 2025	Supplies a measured operating outcome; OFF records perceptions and expectations	Verify the conversion from throughput to a budget result	One organizational and task setting
METR update, February 2026	Complicates simple productivity extrapolation	Inspect selection, task boundaries and time records	No reliable current uplift estimate established in the update
Deloitte CFO Signals, Q3 2026	Addresses control response after approval	Include intervention speed in the funding case	Different sample and question
FinOps unit economics guidance	Provides a cost-measurement framework	Reconcile technical consumption to accepted business outcomes	Guidance rather than an empirical ROI benchmark

Putting the process to work

Begin with one material workflow whose sponsor can supply a baseline. The first review should agree on the unit of accepted output, the observation period, the relevant demand and the costs inside the boundary. Finance should then ask the sponsor to classify each claimed benefit before assigning money to it. This order matters: a monetary estimate can create false precision before the underlying operating result has been defined.

Use a benefit ledger with a stable identifier for the workflow and a separate line for each non-overlapping benefit. The operating owner supplies the observation and explains exclusions. The finance partner checks the conversion and budget consequence. The approver records the decision the claim may support. The

next review date should follow the time needed to observe the benefit and the company's exposure, rather than a universal monthly cadence.

Working tool: workflow cost and benefit ledger

Field	Required entry	Owner / evidence
Workflow boundary	Start, accepted output, exclusions, systems and users	Operating owner; process record
Baseline	Period, volume, quality, time and cost; known differences	Operating owner with finance challenge
Cost	Implementation, recurring supplier charges, human review, support, allocated shared cost	Finance; invoices and operating records
Benefit category	Cash, avoided hiring, capacity, contribution or risk reduction	Sponsor; one category per claim line
Conversion	How the operating result changes the claimed benefit	Finance; staffing, demand or commercial evidence
State and permitted use	Proposed, observed, accepted, suspended or retired; decision supported	Approver; dated decision
Review and reversal	Date, trigger, owner, required response	Operating and finance owners

Consider an illustrative support workflow processing 10,000 accepted cases per month. Baseline handling requires 15 minutes per case. The proposed workflow requires 11 minutes including routine review, with no assumed deterioration in quality. The four-minute difference releases about 667 hours monthly. At an illustrative loaded labor rate of $45 per hour, that is $30,000 of capacity value. These assumptions describe a worked example, not an OFF observation.

Suppose recurring AI and support costs are $14,000 per month and implementation costs $80,000 once. Calling the difference between $30,000 and $14,000 a $16,000 monthly cash saving would be incorrect if payroll and purchased services remain unchanged. The project may provide useful capacity, but its direct incremental cash flow remains negative until a real financial benefit is captured. The operating owner must identify additional demand, an avoided staffing requirement or another documented use of the released hours.

Now add a separate, explicit assumption: the company can eliminate $25,000 per month of purchased overflow services while maintaining the same accepted output and quality. That creates $11,000 monthly net cash benefit after the $14,000 recurring cost. Simple payback on the $80,000 implementation outlay is about 7.3 months after that steady benefit begins. If the benefit starts later or ramps gradually, calendar payback is longer. The $30,000 capacity estimate must not also be added to the $25,000 purchased-services saving for the same work.

Exhibit 4. Illustrative financial bridge, not survey data

Item	Calculation	Treatment
Released time	10,000 × 4 minutes ÷ 60 = 666.7 hours/month	Capacity observation under stated assumptions
Imputed capacity value	666.7 × $45 = $30,000/month	Not automatically cash
Recurring incremental cost	$14,000/month	Cash outflow assumption
Verified overflow reduction scenario	$25,000/month	Separate cash-benefit assumption
Net recurring cash benefit	$25,000 − $14,000 = $11,000/month	Conditional on the overflow reduction
Simple steady-state payback	$80,000 ÷ $11,000 = 7.27 months	Excludes ramp and timing effects

The downside review should change operating assumptions, not just apply an arbitrary percentage discount to benefits. If accepted volume falls, calculate whether the overflow reduction still holds. If additional human review becomes necessary, include its cost and assess whether the quality standard is still met. If the purchased-service contract cannot be reduced until renewal, move the cash benefit to the appropriate month. These changes identify the actual dependencies behind the result.

Before releasing expansion funding, require an approved output definition, a reproducible baseline, a complete cost boundary, a non-duplicated benefit calculation, a named budget action and an effective intervention plan. Record unresolved assumptions as conditions of approval. A bounded learning investment can proceed with more uncertainty, but its funding cap and expiry should be visible. A production investment claiming repeatable savings needs stronger evidence that those savings can recur.

At the next review, compare actual evidence with the accepted claim. Keep the original version so that management can distinguish forecast error, changed conditions and measurement corrections. Suspend a benefit when its premise fails; do not quietly replace the baseline. This creates an investment history that a new CFO, auditor or transaction team can understand without reconstructing informal discussions.

The next finance review

The CFO should bring three decisions to the next review. First, decide which investments are buying learning, which are buying operational capacity and which are expected to produce a specific financial return. Require each sponsor to use the corresponding evidence standard. The purpose is to preserve worthwhile experimentation while preventing speculative benefits from entering a production savings commitment.

Second, select the material benefit claims that need reconciliation. The finance partner and operating owner should agree on the workflow boundary, compare accepted output and cost, and identify the budget action that captures value. Claims that depend on future demand or a contract change should remain conditional. A change in classification is a correction to the investment case, not a reason to erase the original record.

Third, establish who can limit consumption or stop expansion when assumptions fail. The approval should name an effective decision-maker, the evidence that triggers action and the financial commitments that survive a pause. This makes the downside case operationally meaningful.

The result should be a short set of defensible decisions supported by detailed working records. Finance does not need one universal ROI threshold for every use of AI. It needs to know what each commitment is intended to achieve, which evidence supports that claim and what will happen when the evidence changes. That is the difference between a recurring AI budget discussion and an investment process that can learn.

Questions CFOs should be able to answer

Is time saved a cash saving?

Not by itself. Time saved establishes a potential capacity benefit. A cash claim additionally requires a documented expenditure change. The illustrative bridge separates these steps.

Did finance sentiment improve in October?

This report does not establish that. The buying instruments used here have no new comparable October answers in the November export set. Their results are carried forward, not presented as a new monthly panel.

Why do the funding shares exceed 100%?

The source field permits combinations. Each option is counted when mentioned by an eligible respondent. A respondent may identify more than one funding source. [FUND]

Can the company use an AI provider's reported productivity gain in its investment case?

It can inform a hypothesis, but the company should verify the relevant workflow, quality, costs and ability to capture the benefit. An external result does not establish the internal cash conversion.

Does partial cost visibility make every expansion unjustifiable?

No. It identifies uncertainty that should be bounded. A small reversible experiment may proceed with a cost cap; a large recurring savings commitment requires a more complete cost boundary. [COST; recommendation]

What changes when a benefit claim is suspended?

Its permitted use in the investment case stops until evidence supports reinstatement or revision. Management can separately decide whether the underlying workflow should continue for another purpose.

Should risk reduction be assigned a precise financial return?

Only when the exposure, assumptions and uncertainty support that estimate. Otherwise present the risk decision explicitly rather than adding an unsupported avoided-loss figure to cash benefits.

Citable facts

- 211 of 389 eligible respondents mentioned proving ROI as a blocker to more AI spending: 54%, carried forward, multi-select. [BLOCK-01]

- 208 of 389 selected measurable return within six months: 53%, an expectation rather than observed benefit or break-even. [PAY-01]

- 34 of 77 operator respondents answered yes to real-time AI running-cost visibility; 30 answered partly and 13 no. [COST]

Methods and limitations

Canonical calculation: November calculation-v1, executed from the authorized 1 November 2026 CSV exports. Byte-identical copies are excluded. Invited-only rows are excluded; other statuses are not treated as evidence of attendance. The cutoff is 31 October 2026 at 23:59:59 America/Los_Angeles. Latest eligible nonblank answers are selected by normalized email separately per question, using registration timestamp, then stable filename/row tie ordering. Registration timestamps do not establish when an answer was edited.

FUND, BLOCK and PAY use n=389; SIGN uses n=467. These are distinct question bases. COST uses a separate operator instrument with n=77. The registry supplies exact source wording, options, dates, counts, percentages, deduplication decisions and source identifiers. Multi-select results use any mention. Whole percentages use half-up rounding. No probability-sample margin of error is claimed.

The November title classification is not substituted for an unchanged October role cut. No CEO/CFO subgroup movement is asserted here. Respondent-level observations are not company-level audited outcomes. The research does not estimate market size, causal effects, realized payback or average enterprise spending. Illustrative amounts are excluded from the first-party statistical registry and tracked separately as model assumptions.

References

[OFF-DATA] Open Future Forum. Executive AI Statistics, November 2026; accompanying aggregate CSV/JSON and calculation notes.

[E-QJE] Brynjolfsson, E., Li, D., and Raymond, L. R. Generative AI at Work. Quarterly Journal of Economics 140(2), 2025, pp. 889-942. Published version. https://academic.oup.com/qje/article/140/2/889/7990658

[E-METR] METR. We are Changing our Developer Productivity Experiment Design. 24 February 2026. https://metr.org/blog/2026-02-24-uplift-update/

[E-DELOITTE3] Hardy, E., and Murphy, T. North American CFOs cite AI cost-governance gaps in third-quarter 2026 survey. Deloitte Insights, 21 October 2026. https://www.deloitte.com/us/en/insights/topics/business-strategy-growth/3q-2026-cfo-signals-survey.html

[E-FINOPS] FinOps Foundation. Unit Economics capability. Current guidance consulted 2 November 2026; not treated as a newly fielded survey. https://www.finops.org/framework/capabilities/unit-economics/

[OFF-COST] Open Future Forum. Workflow-Level AI Cost Accounting. October 2026. https://openfutureforum.com/blog/workflow-level-ai-cost-accounting

Related research and citation

Read the November CEO report for executive decision rights, the AI Leaders report for production evidence, and the Investment Banking report for revenue substantiation. These companion reports address different decisions and should not be treated as independent corroborating surveys.

Suggested citation: Newlands, Murray. CFO AI Leverage Report, November 2026. Open Future Forum. Evidence cutoff 31 October 2026.

Open Future Forum convenes executives to discuss the operating questions behind enterprise AI adoption. Research and methods: https://openfutureforum.com/research/

Editable working tool

Benefit qualification and reversal: download the editable CSV worksheet. The worksheet implements the report's recommendation; it is not a validated score, mandatory standard or completed assessment.