IDC

IDC INSIGHTS BRIEF

Sponsored by: Sage

The Emerging Economics of Agentic AI in the Finance Function

April 2026

Written by: Kevin Permenter, Senior Research Director, IDC; Dominique Bindels, Consulting Manager, IDC

Based on the results from a global survey of finance leaders (N = 2275), North America (NA; n = 1045), Europe, the Middle East, and Africa (EMEA; n = 1230)

Why This Matters Now

Finance functions are under pressure to increase speed, reduce cost, and strengthen control simultaneously. AI (particularly more autonomous "agentic" AI) is increasingly positioned as a structural solution.

This IDC study explores two core questions:

- **Hypothesis A:** Will agentic AI fundamentally improve the economics of running a finance function?
- **Hypothesis B:** Will competitive advantage increasingly come from systems that act autonomously rather than waiting for human instruction?

The results below are final, and the patterns are consistent across regions. Interviews suggest that finance leaders are not primarily debating whether AI belongs in finance. They are debating where it belongs first, how much validation it requires, and which controls must surround it.

Conclusion

The evidence shows that:

1. AI is delivering productivity, but oversight efforts are material.
2. Approximately one quarter of productivity gains are being absorbed by explanation-related work.
3. Autonomy adoption is constrained less by capability and more by trust architecture.
4. Transparency is emerging as an economic multiplier, not a compliance feature.
5. AI adoption is clustering first around repetitive, reviewable, rule-based work such as AP, reconciliations, report comparison, and recurring close tasks.

Based on the final data, the central economic question for finance AI will not be "How intelligent is the model?" but rather "How efficiently can autonomy be trusted, verified, and scaled?"

Voice of the user: Qualitative interviews reinforce this pattern. Finance leaders are not primarily asking whether AI is powerful enough. They are asking whether outputs can be validated efficiently, whether assumptions can be surfaced clearly, and whether the system can be trusted in repeatable, bounded workflows without creating new layers of review burden.

Hypothesis A

Agentic AI will fundamentally improve the economics of running a finance function.

To test this, we examined not just automation gains but the hidden costs of making AI operationally usable in finance environments: review burden, assumption reconstruction, exception handling, and the controls needed to make outputs repeatable.

A1. The "verification tax" is already material

Finance leaders were asked how much time they spend each week checking, validating, or reconfirming outputs generated by AI or automated systems. Verification hour estimates are based on respondents' interpretation of senior finance leadership time and may reflect either individual or collective oversight effort.

A substantial share reported **15+ hours per week** dedicated to verification:

- **North America: 48.4%**
- **EMEA: 47.1%**

Midpoint analysis of reported time bands shows:

- **North America:** ~13.1 hours per week of senior finance leadership time
- **EMEA:** ~12.79 hours per week of senior finance leadership time
- **Total average:** ~12.9 hours per week of senior finance leadership time

Interpretation: AI is generating outputs, but finance teams are investing meaningful time to ensure those outputs are correct, compliant, and audit-ready. The cost of verification is measurable and economically significant. North America is reporting a similar verification burden to EMEA. This burden is set to increase if it is not addressed.

Voice of the user: Interviews suggest this burden is not limited to spot-checking outputs. In practice, finance teams are creating new validation layers around AI, including source-mapped templates, exception review workflows, repeatability checks, and debugging routines across entities. In that sense, verification is emerging as a new category of finance work rather than a temporary adoption friction.

Qualitative signal

Validation is becoming a distinct finance workflow

Interviews suggest that verification is not simply a temporary friction point in AI adoption. Finance teams are actively designing new validation layers around AI outputs, including source-mapped templates, exception reports, reconciliation checks, and repeatable review artifacts.

In practice, teams are using AI to perform the bulk of data preparation and comparison work, but they are simultaneously building structured processes to confirm that outputs align with internal controls, financial logic, and reporting standards.

This suggests that AI adoption in finance is not eliminating oversight work but reconfiguring it into a new operational discipline focused on validation at scale.

A2. Reverse-engineering efforts are absorbing productivity gains

Respondents were also asked what proportion of time "saved" through AI is spent reconstructing how outputs were generated (e.g., understanding assumptions, data lineage, or decision logic).

Midpoint analysis showed:

- Approximately **26% of AI time savings are lost to explanation recovery**, similar in NA and the EMEA region.

In practical terms, if AI saves 10 hours, roughly 2–3 hours may be consumed by clarifying how decisions were made.

Interpretation: Opacity does not eliminate work but redistributes it into reconstruction work: tracing source data, surfacing assumptions, explaining model intent, validating exceptions, and producing artifacts that can withstand executive, audit, or client scrutiny.

Voice of the user: Interviews suggest finance teams are not asking for abstract explainability but for practical review artifacts: source mapping, assumptions, exception flags, and repeatable validation templates.

Qualitative signal

Reliability issues tend to be subtle, not catastrophic

In interviews, AI failures were rarely described as dramatic breakdowns. Instead, they appeared as subtle reliability issues: scripts that do not scale across entities, outputs that pull from the wrong data context, or results that look plausible but require expert verification.

These types of errors are particularly challenging in finance environments because they may not be immediately visible and often require domain expertise to detect.

This reinforces why explainability and verification remain central adoption constraints.

Finance leaders are less concerned about whether AI can produce an answer and more concerned about how easily that answer can be validated and defended.

A3. Transparency is valued, but buyers expect it to be justified

Leaders were asked whether they would pay a premium for AI systems that provide full "glass box" reasoning traces (i.e., clear, inspectable explanation of how decisions are reached).

- Roughly **1 in 6** would pay a 20–49% premium (NA 17.1%, EMEA 16.7%).
- A smaller segment would pay significantly more (around 4%).

While this does not yet indicate a broad willingness to pay large premiums, it does suggest that a defined segment sees transparency as economically valuable.

Interpretation: Finance leaders value explainability, but may view it as an essential capability rather than a discretionary add-on. Pricing tolerance may evolve as trust-related cost savings become clearer. Further analysis of the data set is required to adequately explore the trade-offs in this area.

Voice of the user: This distinction is important. In interviews, transparency was rarely framed as a premium enhancement. Instead, it was framed as the price of admission for meaningful finance use cases, particularly where outputs affect reporting, cash movements, client recommendations, or board materials.

A4. Interest in "self-healing" agents is strong — if reporting is included

Respondents were presented with the concept of AI agents that automatically detect and correct anomalies (e.g., reconciliation issues and posting errors) while providing a full reasoning trace of what was changed and why.

- ~ 65-66%% -64% would be interested as long as there is 'Human-in-the-loop" approval step.

Interpretation: Autonomy is attractive when paired explicitly with transparency and reporting. The appeal lies in reducing manual correction effort without sacrificing control or auditability. This reinforces the idea that economic gains depend on architecture, not just intelligence.

Voice of the user: Interviews suggest that enthusiasm for self-healing agents is real but bounded by internal-control design. Finance leaders are generally more comfortable with systems that propose or pre-validate corrections than with fully autonomous correction loops that bypass human review.

Qualitative signal

The real comparison is not "AI versus no AI"

In practice, finance leaders rarely decide whether to use AI. Instead, they are deciding how to use it.

Some teams prefer embedded, vendor-built AI capabilities that have already been hardened through product development and testing. Others are beginning to experiment with building targeted workflow automations internally using general-purpose AI tools.

This creates a new decision dynamic — whether to rely on vendor-hardened AI embedded in finance software or address specific workflow bottlenecks internally using AI-enabled automation.

The survey results suggest both paths will coexist, with governance requirements and risk tolerance shaping which approach organizations prefer.

Hypothesis B

The next competitive edge may come less from broad autonomous finance operations than from systems that can operate autonomously within bounded, reviewable, high-frequency workflows.

To explore this, we assessed the appetite for autonomous transaction processing and continuous operations.

B1. Bounded autonomy has real deployment momentum

When asked about the expected deployment of transactional AI agents within two years:

- **NA: 38.9%** said they expect pilot deployment.
- **EMEA: 37.5%** said they expect pilot deployment.

Significant intent (29.1% globally) is concentrated in **"automate selected activities"** (e.g., specific workflows such as payables, reconciliations, or accruals), rather than broad, fully autonomous finance operations (9.4% globally).

Interpretation: Finance leaders are open to autonomy, but in controlled, bounded domains with defined oversight.

Voice of the user: Interviews make "bounded autonomy" highly concrete. The most credible early domains are high-volume, low-ambiguity, rule-based processes such as AP, reconciliations, report comparison, and recurring validation logic. These are attractive not because they are glamorous, but because they are repetitive, reviewable, and operationally painful.

Qualitative signal

Early autonomy is concentrated in "reviewable work"

Across interviews, the most successful early AI deployments share a common characteristic: The work is repetitive, rule-based, and easy for humans to review quickly.

Examples include reconciliations, report comparisons, AP processing logic, and recurring reporting transformations. These tasks are attractive not because they are strategically complex, but because they are operationally tedious and governed by clear rules.

Finance leaders appear comfortable allowing AI to operate autonomously in these bounded domains as long as outputs remain easily inspectable and reversible.

This reinforces the survey finding that autonomy adoption is occurring first in narrow, well-defined workflows rather than across the finance function as a whole.

B2. Continuous close becomes plausible — with audit trace

Respondents were asked how they would respond if AI enabled continuous, real-time financial updates — with a complete reasoning trace available.

- ~40-45% would adopt a **hybrid close model** (real-time insights plus a formal close for compliance).
- 25-29% would retain existing processes unchanged.

Interpretation: Structural process change (e.g., moving toward continuous close) is viable, but governance milestones remain important. Traceability appears to be a prerequisite for reconsidering long-standing processes.

Voice of the user: Interviews suggest that real-time finance is more likely to emerge first as an augmentation layer around existing close disciplines rather than a replacement for formal control points. In other words, finance leaders appear more willing to accelerate preparation, review, and issue detection than to remove approval milestones outright.

B3. Accuracy alone does not create trust

A strong majority (71.3%) indicated they would veto a 99% accurate AI system if it lacked a reasoning trace.

Interpretation: Statistical performance is insufficient in finance environments.

Without explainability, high accuracy does not overcome governance thresholds. This finding is one of the most consistent across regions and must be further explored, in conjunction with hypothesis A3 (premium for transparency).

Voice of the user: Interviews suggest that many finance leaders are implicitly evaluating AI the same way they would a junior team member — useful for speed and analysis, but expected to show work, support conclusions, and operate within explicit review boundaries.

B4. Transparency accelerates scale more than risk appetite

When asked what full transparency (100% reasoning visibility) would change:

- 58.9% said it would **accelerate the speed or volume of autonomous transactions.**
- ~24% said it would expand autonomy into higher-stakes decisions.

Interpretation: Transparency is seen primarily as a scale enabler rather than a risk escalator.

Finance leaders appear more inclined to increase throughput than to expand immediately into materially riskier domains.

Voice of the user: Interviews suggest transparency matters not only because it reduces perceived risk but also because it makes finance work repeatable. Once teams can clearly see source mappings, assumptions, and validation checks, they can institutionalize the workflow and rerun it with much greater confidence.

Qualitative signal

AI is shifting the labor mix, not simply reducing labor

Interview evidence suggests that AI is changing the composition of finance work rather than eliminating it. As manual data processing declines, more effort is shifting toward validation, exception handling, workflow design, and interpretation.

In effect, finance professionals are moving from processing transactions toward supervising and validating automated processes.

If this pattern continues, the economics of the finance function may shift less through headcount reduction and more through higher leverage of senior analytical work relative to transactional effort.

Regional observations

- Verification burden is relatively consistent between both regions, both above 15+ hours a week
- Both regions are both on par with near-term deployment intent.
- Trust thresholds and transparency expectations are closely aligned across regions.

About the Analyst

Kevin Permenter, *Senior Research Director, Financial Applications*

As a research director, Kevin M. Permenter provides insights and analysis across multiple fintech market segments, including accounting, revenue management, corporate tax, accounts payable, accounts receivable, treasury, and enterprise payment management. Kevin leads qualitative research efforts which drive a series of technology buyer-focused documents, including MarketScapes, buyer perspectives, PeerScapes, and end-user surveys. He also leads several quantitative research efforts within financial applications that feed key technology supplier-focused documents (e.g., Market Shares, Market Forecast, Market Glance, and Market Analysis Perspectives). Kevin's research includes a particular emphasis on the interplay, challenges, and trends driving financial application deployment and its role in the evolution of the complex financial technology ecosystem.

Dominique Bindels, *Consulting Manager, Custom Solutions Europe*

Dominique Bindels is a consulting manager in the IDC European Custom Solutions team, partnering with companies in the AI/ML, security, process automation, and Big Data analytics spaces. He has a background in strategic consumer market research for consumer electronics and related services and ecosystems, providing leading consumer electronics companies with insights and analysis. He is a regular speaker at industry and client events. He studied in the U.K. and Germany and has master's and bachelor's degrees in international business with finance.

IDC Custom Solutions

The content in this paper was adapted from existing IDC research published on www.idc.com.

IDC Research, Inc.
One Beacon Street
Suite 33100
Boston, MA 02108, USA
T 508.872.8200
F 508.935.4015
blogs.idc.com
www.idc.com