http://dx.doi.org/10.18623/rvd.v23.5632

THE TRANSITION FROM AUTOMATION TO AUGMENTATION: A CONSTRAINT-AWARE FRAMEWORK FOR AGENTIC AI ADOPTION IN CORPORATE FINANCE FUNCTIONS

A TRANSIÇÃO DA AUTOMAÇÃO PARA O AUMENTO DE CAPACIDADES: UMA ESTRUTURA QUE LEVA EM CONTA AS RESTRIÇÕES PARA A ADOÇÃO DE IA AGENTE NAS FUNÇÕES DE FINANÇAS CORPORATIVAS

Article received on: 12/8/2025
Article accepted on: 3/6/2026

Khaleel Ahmed Jalaluddin*
*The New College, Chennai, Tamil Nadu, India
Orcid: https://orcid.org/0009-0003-5163-4249
laxmi@westernglobaluniverisity.us

The authors declare that there is no conflict of interest

Abstract
Corporate finance functions are moving through a new digital inflection point. Earlier waves of automation in finance centered on robotic process automation, rules engines, workflow tools, and predictive analytics that standardized high-volume transactions and reduced cycle times. The current wave is different because large language models, generative AI, and agentic systems can interpret unstructured data, coordinate across tools, and recommend or execute multi-step decisions. Yet the finance function cannot adopt autonomy in the same way as less regulated business areas. Financial reporting, planning, controllership, treasury, tax, internal audit, and investor-facing processes are constrained by internal controls, segregation of duties, auditability, data lineage, policy compliance, and accountability to boards, regulators, and external assurance providers. This review paper examines how the literature from 2020 to 2026 explains the shift from automation to augmentation and develops a constraint-aware framework for adopting agentic AI in corporate finance. Following a structured review and content analysis of recent academic studies, standards, and practitioner reports, the paper synthesizes four themes: the evolution of AI use cases in finance and accounting; the distinction between automation, augmentation, and bounded autonomy; the organizational and technical constraints that shape deployment; and the governance mechanisms needed for reliable value capture. The review argues that the most realistic pathway for finance is not unrestricted autonomy but graduated augmentation, in which AI agents expand analytical capacity, accelerate close and planning cycles, and improve stakeholder alignment while humans retain

Resumo
As funções de finanças corporativas estão passando por um novo ponto de inflexão digital. As ondas anteriores de automação em finanças concentraram-se na automação robótica de processos, mecanismos de regras, ferramentas de fluxo de trabalho e análises preditivas que padronizaram transações de alto volume e reduziram os tempos de ciclo. A onda atual é diferente porque grandes modelos de linguagem, IA generativa e sistemas agentivos podem interpretar dados não estruturados, coordenar entre ferramentas e recomendar ou executar decisões em várias etapas. No entanto, a função financeira não pode adotar a autonomia da mesma forma que áreas de negócios menos regulamentadas. Relatórios financeiros, planejamento, controladoria, tesouraria, tributação, auditoria interna e processos voltados para investidores são limitados por controles internos, segregação de funções, auditabilidade, linhagem de dados, conformidade com políticas e prestação de contas a conselhos, reguladores e prestadores de garantia externos. Este artigo de revisão examina como a literatura de 2020 a 2026 explica a mudança da automação para o aumento de capacidade e desenvolve uma estrutura que leva em conta as restrições para a adoção de IA agênica nas finanças corporativas. Após uma revisão estruturada e análise de conteúdo de estudos acadêmicos recentes, normas e relatórios de profissionais, o artigo sintetiza quatro temas: a evolução dos casos de uso de IA em finanças e contabilidade; a distinção entre automação, aumento e autonomia limitada; as restrições organizacionais e técnicas que moldam a implantação; e os mecanismos de governança

accountability over material judgments and irreversible actions. Building on the synthesis, the paper proposes a five-layer framework that aligns use-case selection, autonomy design, control requirements, human oversight, and performance metrics. The framework helps organizations match agentic capability to task criticality, data quality, reversibility, and regulatory exposure. The paper contributes a finance-specific conceptualization of agentic adoption, unique research objectives, an implementation sequence for CFO organizations, and a future research agenda focused on control redesign, explainability, operating models, and the changing role of finance professionals.

Keywords: Agentic AI. Corporate Finance. Augmentation. Automation. Finance Transformation. AI Governance. Internal Controls. Controllership. FP&A. Auditability.

necessários para a captura confiável de valor. A revisão argumenta que o caminho mais realista para as finanças não é a autonomia irrestrita, mas o aumento gradual, no qual os agentes de IA expandem a capacidade analítica, aceleram os ciclos de fechamento e planejamento e melhoram o alinhamento das partes interessadas, enquanto os humanos mantêm a responsabilidade sobre julgamentos materiais e ações irreversíveis. Com base nessa síntese, o artigo propõe uma estrutura de cinco camadas que alinha a seleção de casos de uso, o projeto de autonomia, os requisitos de controle, a supervisão humana e as métricas de desempenho. A estrutura ajuda as organizações a adequar a capacidade dos agentes à criticidade da tarefa, à qualidade dos dados, à reversibilidade e à exposição regulatória. O artigo contribui com uma conceituação específica para finanças da adoção de agentes, objetivos de pesquisa exclusivos, uma sequência de implementação para organizações de CFO e uma agenda de pesquisa futura focada no redesenho de controles, explicabilidade, modelos operacionais e a mudança no papel dos profissionais de finanças.

Palavras-chave: *IA Agênica. Finanças Corporativas. Aumento. Automação. Transformação Financeira. Governança de IA. Controles Internos. Controle de Gestão. FP&A. Auditabilidade.*

1 INTRODUCTION

Artificial intelligence is no longer entering finance only through narrow prediction models or robotic task automation. It is increasingly entering through language interfaces, copilots, orchestration layers, and agents that can retrieve information, reason over exceptions, interact with enterprise systems, and recommend or trigger actions across end-to-end workflows. Across the broader enterprise, the current debate has therefore shifted from whether organizations should automate individual tasks to how they should redesign work when AI can support or coordinate complex activity at scale (McKinsey & Company, 2025; World Economic Forum, 2025a). In finance, this shift is especially significant because many core processes still depend on repetitive reconciliation, period-end reviews, narrative reporting, manual policy checks, and fragmented analytical handoffs between teams. These characteristics make finance highly attractive for AI-enabled productivity gains, but they also make it highly sensitive to control failures.

The office of the chief financial officer (OCFO) sits at the intersection of operational data, executive decision-making, external reporting, and organizational accountability. As a result, AI adoption in finance has implications that extend beyond efficiency. Decisions made in controllership, treasury, tax, financial planning and analysis (FP&A), and internal audit affect management credibility, covenant compliance, assurance outcomes, capital allocation, and the confidence of boards, investors, lenders, and regulators. This is why finance cannot treat agentic AI as a simple extension of earlier automation programs. A journal-entry assistant that proposes entries, a close-management agent that chases dependencies, or a planning agent that drafts scenarios may appear similar to digital assistants in other functions, yet the acceptable error tolerance, required documentation, and responsibility structure are materially different (NIST, 2023; Financial Stability Board, 2024).

Recent literature suggests that AI is already reshaping accounting, auditing, and finance by increasing analytical speed, supporting anomaly detection, improving forecast generation, and shifting professional work from mechanical processing toward judgment-intensive activities (Bahoo *et al*., 2024; Cao *et al*., 2024; Abbas, 2025; Abdo-Salloum & Chehade, 2026). At the same time, the literature consistently warns that adoption depends on explainability, trust, control clarity, data governance, and organizational readiness rather than model capability alone (Lehner *et al*., 2022; Anh *et al*., 2024; Yeo *et al*., 2025). In practice, many finance functions remain in a pilot phase, often using generative AI for summarization and drafting while delaying deeper workflow integration because policy, systems architecture, and human accountability mechanisms are not yet mature (Deloitte, 2024; L.E.K. Consulting, 2025; KPMG, 2025a). This creates a strategic tension. If finance moves too slowly, it may miss substantial opportunities to improve close-cycle responsiveness, scenario analysis, working-capital visibility, and cross-functional decision speed. If it moves too quickly, it risks model errors, inappropriate actions, weak audit trails, and over-reliance on systems that were not designed for material finance judgments.

This paper addresses that tension by examining the transition from automation to augmentation in corporate finance. Automation is defined here as the substitution of repetitive, rules-based tasks through pre-specified logic, while augmentation refers to the expansion of human capability through AI-generated recommendations, explanations, workflow support, and bounded decision assistance. Agentic AI adds a further dimension: systems that not only generate content but also plan multi-step actions, use tools, maintain context, and coordinate execution across applications or actors (PwC, 2024; World Economic Forum & Capgemini, 2025; OECD, 2026). In finance, however, the move toward agency must be bounded by constraints, because the function is structurally accountable for precision, traceability, and policy compliance. The central argument of this review is that successful adoption depends less on maximizing autonomy and more on calibrating autonomy to task materiality, reversibility, control criticality, and data reliability.

The paper is structured as follows. Section 2 reviews the shift from traditional automation to augmented and agentic finance work. Section 3 explains the review procedure and content-analysis methodology, drawing on the format of the sample article shared by the user while adapting it to a structured review design. Section 4 synthesizes the literature around use cases, benefits, and adoption barriers. Section 5 develops a constraint-aware framework for agentic AI adoption in corporate finance. Section 6 discusses implementation implications for finance leaders and stakeholders. Section 7 identifies a future research agenda, and Section 8 concludes.

2 LITERATURE REVIEW: FROM FINANCE AUTOMATION TO AUGMENTED AND AGENTIC WORK

2.1 The automation era in finance

For much of the last decade, digital transformation in finance has been associated with standardization, shared services, enterprise resource planning (ERP) integration, process mining, and robotic process automation (RPA). These technologies were well suited to deterministic, repetitive, and transaction-heavy finance processes such as invoice routing, account reconciliations, data extraction, report assembly, and exception flagging. Their value proposition was based on throughput, accuracy, and cost reduction. Practitioner surveys continue to show that finance leaders associate early AI value with productivity, error reduction, and improved access to structured data, rather than with fully autonomous decision-making (Boston Consulting Group, 2024; KPMG International, 2024; McKinsey & Company, 2024).

The academic literature frames this period as one in which accounting and finance adopted machine learning and analytics selectively, primarily to enhance classification, forecasting, fraud detection, audit testing, and the processing of alternative data (Bahoo *et al*., 2024; Cao *et al*., 2024). In accounting and auditing, AI was initially used to support pattern recognition and anomaly identification rather than to redesign professional authority. Leocádio *et al*. (2024) argue that AI's early impact on auditing was strongest in evidence collection, risk assessment, and exception detection, where computational power could complement human review. Similarly, Abbas (2025) shows that management accounting has absorbed AI through digitalization, predictive tools, and analytical assistance, but still depends on human interpretation for business partnering, performance narrative, and strategic coordination.

This first wave of automation generated important operational gains, but it also revealed structural limits. Rule-based automation performs poorly when policies are ambiguous, documents are unstructured, exceptions are contextual, or a process spans multiple systems and informal interactions. Finance work is full of such situations. Period close is not just a sequence of tasks; it is also a coordination mechanism involving controllers, business units, external data sources, materiality thresholds, and judgment-based sign-offs. Budgeting is not just forecast computation; it involves assumptions, negotiation, interpretation, and executive challenge. These features constrained the transformative potential of classic automation because many finance processes could be accelerated without being fundamentally reconfigured.

2.2 The rise of augmentation

Generative AI widened the scope of finance automation because it can work on language, narratives, policies, contracts, emails, and commentary in addition to structured numerical data. This ability is important for finance because a large share of work inside the function is communication-intensive. Management discussion and analysis, variance commentary, controls documentation, board packs, accounting memos, policy interpretation, and auditor requests all involve large volumes of text that were not easily addressed by earlier tools. Recent reviews of generative AI in finance conclude that the technology is especially powerful where firms need to interpret heterogeneous data, generate contextual explanations, or accelerate decision support in highly dynamic environments (Lee *et al*., 2024; OECD, 2024).

The literature increasingly describes the next stage not as straight-line automation but as augmentation. Augmentation means using AI to expand the analytical and coordination capacity of finance professionals rather than to displace their accountability. In practice, this includes copilots that draft narrative reports, assistants that summarize reconciliations, tools that propose planning assumptions, and agents that surface likely policy deviations for human review. McKinsey & Company (2025) reports that organizations generate greater value from AI when they redesign workflows and management practices instead of merely layering tools on top of old processes. In finance, that redesign often means reducing low-value manual effort so that controllers, analysts, and finance business partners can spend more time on scenario analysis, challenge sessions, stakeholder communication, and risk interpretation.

Several recent studies reinforce this interpretation. Cao (2025) notes that AI is shaping accounting and finance not only by improving efficiency but also by changing transparency expectations and decision processes. Kerr (2025) similarly argues that management accounting is being influenced by AI through shifts in information production, analytical depth, and the role of accountants in organizational sensemaking. The most relevant implication for corporate finance is that AI can create a leverage effect: it can compress the time needed to produce analysis while increasing the number of scenarios, explanations, or anomaly hypotheses that finance professionals can evaluate. This is augmentation because the human role moves upward rather than disappearing.

2.3 From copilots to agents

The newest step in the literature is the discussion of agentic AI. Agentic systems are distinguished from simple chat interfaces because they can plan a sequence of actions, call tools, maintain state, apply rules, and work toward goals with partial autonomy (PwC, 2024; World Economic Forum & Capgemini, 2025; OECD, 2026). KPMG (2025a) describes this shift as the movement from question-answering tools to finance agents that can perform cross-process work such as reconciliation support, close coordination, journal preparation, narrative drafting, and variance follow-up. The promise is compelling: instead of waiting for a finance user to orchestrate every micro-step, an agent may pull data from the ERP, compare it against policy thresholds, draft an explanation, highlight unresolved items, and route exceptions to the appropriate reviewer.

Yet the same literature also emphasizes that agentic systems increase complexity. Because they can initiate or coordinate actions, they expand the risk surface related to access management, tool misuse, prompt injection, policy drift, weak supervision, and unclear responsibility (Microsoft, 2025; World Economic Forum & Capgemini, 2025; IMDA, 2026). The issue is not simply whether the model's answer is correct; it is whether the overall system behaves reliably under changing inputs, exceptions, and organizational controls. In finance, this distinction is decisive. A draft forecast narrative produced by a copilot is recoverable and reviewable. An agent that updates assumptions, submits requests, or escalates tasks can affect process timing, approvals, and evidence. Therefore, agentic finance requires a design philosophy that treats control integrity as a core architectural variable.

2.4 Why finance is different

The literature on AI in finance, accounting, and auditing converges on a key point: the adoption problem is not purely technological. It is socio-technical and institutional. Financial data is often fragmented across ERP instances, planning platforms, spreadsheets, and email-based workflows. Policy interpretation can depend on local business context. Materiality thresholds vary across processes. Many outputs must be explained to auditors, regulators, and executives after the fact. As Lehner *et al*. (2022) show, AI-based decision-making in accounting raises ethical and governance issues because the legitimacy of a decision depends on more than predictive performance. Yeo *et al*. (2025) extend this by showing that explainability in finance is not an optional design attribute; it is tied to model acceptance, validation, and responsible use. The consequence is that finance leaders cannot adopt agentic AI with the same tolerance for opacity that may be acceptable in experimental productivity tools elsewhere in the enterprise.

These observations create the foundation for a constraint-aware approach. Rather than asking whether finance should move from automation to full autonomy, the more relevant question is how finance should sequence augmentation so that value creation and control reliability improve together. That is the gap this review addresses.

3 METHODOLOGY

3.1 Review design

Following the general structure of the sample Sustainability paper shared by the user, this study adopts a content-analysis procedure, but applies it to a review-paper context instead of a questionnaire study. The methodology therefore combines a structured literature review with qualitative thematic synthesis. The purpose is not to estimate a pooled effect size, but to identify recurring constructs, constraints, and implementation logics relevant to agentic AI adoption in corporate finance.

3.2 Search strategy

A structured search was conducted across Google Scholar, ScienceDirect, Scopus-indexed sources available through publisher sites, Web of Science-accessible journals, SSRN-linked working papers, and authoritative institutional reports published by standards bodies, regulators, and major professional organizations. The search window covered January 2020 to March 2026 to reflect the user's requirement for contemporary sources and to capture the period in which generative and agentic AI entered mainstream enterprise discourse. Search terms included "artificial intelligence in corporate finance", "AI in accounting and auditing", "management accounting and artificial intelligence", "generative AI in finance", "agentic AI governance", "finance function AI adoption", "explainable AI in finance", "AI internal controls", and "AI agents enterprise governance".

3.3 Screening and eligibility

The initial search identified a broad pool of academic articles, white papers, and governance documents. Sources were included when they met at least one of three criteria: first, they addressed AI use in finance, accounting, auditing, treasury, FP&A, or closely related CFO activities; second, they provided governance, risk, or control guidance directly relevant to high-accountability business functions; or third, they discussed agentic, generative, or augmented AI in enterprise settings with clear implications for finance. Sources were excluded when they focused solely on consumer finance, algorithmic trading, computer vision applications unrelated to finance operations, or purely technical benchmarks without organizational adoption relevance. The final analytical corpus comprised 34 sources, including peer-reviewed journal articles, review papers, standards documents, and major practitioner reports.

3.4 Coding and synthesis procedure

Each source was coded against five analytical dimensions: (1) primary finance use case, (2) value mechanism, (3) constraint or risk category, (4) governance or control recommendation, and (5) implications for the role of finance professionals. During the synthesis, recurring patterns were grouped into higher-order themes. Four dominant themes emerged: the shift from rule-based automation to language-enabled augmentation; the growing relevance of bounded agency; the centrality of controls, explainability, and accountability; and the need for adoption frameworks tailored to finance rather than generic enterprise AI. The output of the content analysis is the constraint-aware framework proposed in Section 5. Figure 1 illustrates the review and framework-development process used in this paper.

3.5 Methodological limitations

Because the field is moving rapidly, some high-value evidence comes from institutional reports and practitioner surveys rather than mature longitudinal studies. This is an accepted limitation of an emergent topic. To mitigate that limitation, the paper triangulates between academic literature on accounting and finance, official governance frameworks, and enterprise adoption surveys. The result is a review that prioritizes conceptual coherence and practical relevance while remaining grounded in recent, citable sources.

Figure 1

Structured review and framework-development procedure.

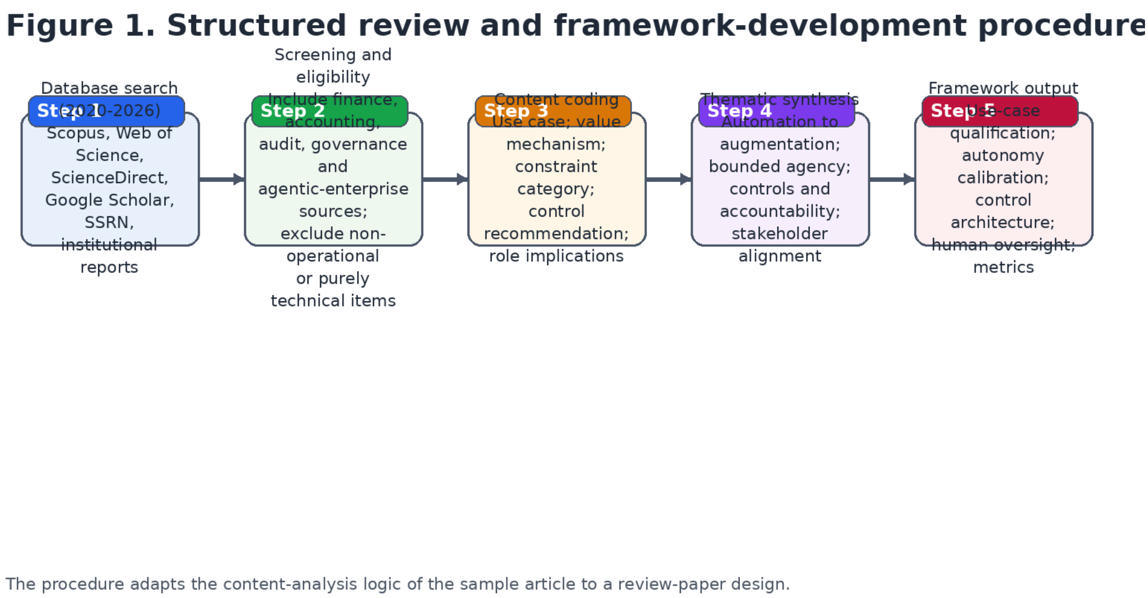

4 FINDINGS AND THEMATIC SYNTHESIS

4.1 Value is concentrated in finance workflows that combine structured data with narrative interpretation

The literature suggests that the strongest current opportunities for AI in finance lie in workflows that blend structured transaction data with recurring interpretation, coordination, and explanation. Examples include account reconciliation support, close management, variance analysis, policy compliance reviews, working-capital diagnostics, cash forecasting, management commentary, and scenario planning (Cao *et al*., 2024; Lee *et al.*, 2024; KPMG, 2025a). These are not purely manual or purely computational processes. They require pulling data from multiple systems, checking completeness, identifying unusual patterns, drafting explanations, and escalating unresolved items. In such environments, generative and agentic tools can reduce cycle times because they are better than earlier automation at handling unstructured inputs and contextual prompts.

This distinction matters because it explains why finance is moving from isolated task automation to orchestrated support. The close process illustrates the point. Traditional automation can schedule jobs, transfer files, or run reconciliation rules. Augmented AI can additionally summarize late submissions, draft issue logs, detect inconsistencies in commentary, and propose follow-up actions. A bounded agent can go one step further by monitoring dependencies, routing reminders, and preparing exception packets for controller review. In this sense, agentic capability increases the "coordination bandwidth" of the finance function. It does not eliminate the need for oversight, but it compresses the time required to coordinate across entities, accounts, business units, and approval chains.

Evidence from finance surveys is consistent with this interpretation. KPMG International (2024) finds rapid adoption of AI in financial reporting processes, while McKinsey & Company (2024) reports that finance leaders who have implemented use cases often cite employee productivity and better use of existing data as the most visible benefits. Taulia (2024) similarly reports growing use of AI for decision support within finance functions. These findings support a central theme in the literature: value is not generated only by automating keystrokes. It is generated when AI reduces the friction involved in interpreting, packaging, and routing information across stakeholders.

4.2 Augmentation is a more appropriate near-term design principle than unrestricted autonomy

A second major finding is that augmentation better matches the institutional logic of finance than unrestricted autonomy. In theory, agentic AI can execute end-to-end tasks. In practice, the literature consistently recommends proportional autonomy based on context, especially for regulated or control-intensive domains (NIST, 2023; World Economic Forum & Capgemini, 2025; IMDA, 2026). Finance is one of those domains.

Material outputs such as journal entries, impairment assumptions, transfer-pricing positions, covenant calculations, revenue-recognition judgments, or disclosures cannot be delegated without explicit accountability structures.

The review reveals at least four reasons for this. First, finance processes often contain irreversible or high-consequence actions. An automated reminder can be corrected easily; an incorrectly posted entry to a material account may cascade into reporting issues. Second, financial work is often subject to segregation-of-duties requirements, meaning the same agent cannot be allowed to both prepare and approve critical actions without violating control principles. Third, many finance decisions require explanations that are legible to auditors, regulators, and executives. Fourth, acceptable risk thresholds are lower because even small process failures can affect confidence in reported information. These conditions make "human-out-of-the-loop" designs less attractive than "human-on-the-loop" or "human-in-the-loop" arrangements for most material finance activities.

Recent scholarship reinforces this conclusion. Leocádio *et al*. (2024) show that AI in auditing is most defensible when embedded in a framework of supervision, evidence, and methodological transparency. Kokina *et al*. (2025) argue that AI adoption in auditing faces both opportunity and resistance because firms must reconcile innovation with professional standards, control visibility, and litigation concerns. Abdo-Salloum and Chehade (2026) likewise find that AI adoption in accounting and auditing is constrained by trust, cost, privacy, and unequal readiness, even where the potential benefits are substantial. For corporate finance, these findings imply that the transition should be sequenced through augmentation layers such as drafting, recommendation, monitoring, scenario generation, and exception routing before moving toward direct action in materially sensitive processes.

4.3 Finance-specific constraints shape the adoption frontier

A third theme is that finance-specific constraints define the boundary conditions of viable agentic adoption. The review identifies ten recurring constraint categories.

The first is data quality and lineage. AI outputs are only as reliable as the datasets, mappings, and metadata that feed them. In finance, fragmented charts of accounts, inconsistent entity hierarchies, spreadsheet workarounds, and undocumented data transformations undermine agent performance and user trust (OECD, 2024; EY, 2026).

The second is explainability. Finance leaders need to know not only what an AI system recommends, but why. Explainability matters for model validation, user adoption, audit defense, and regulator communication (Yeo *et al*., 2025). Black-box recommendations are especially problematic in planning, treasury, and external-reporting contexts.

Third is accountability. Governance frameworks consistently emphasize that organizations remain accountable for AI-assisted outcomes even when systems display high autonomy (NIST, 2023; AI Verify Foundation & IMDA, 2024; IMDA, 2026). In finance, this translates into named process owners, approval rights, escalation paths, and documentation obligations.

Fourth is segregation of duties and access control. Agentic systems that can retrieve data, draft transactions, or trigger workflow events must be aligned with role-based permissions and prevented from creating hidden combinations of prepare-review-approve authority. This is a distinctive challenge because agents may span tools and service accounts in ways that are not visible in traditional control matrices.

Fifth is auditability. Finance functions require durable logs showing what data was used, what prompts or rules were applied, what recommendations were produced, who approved an action, and what changed after review. This requirement increases with materiality.

Sixth is policy and regulatory compliance. Tax, revenue recognition, transfer pricing, sanctions screening, treasury controls, and reporting disclosures all operate within evolving policy environments. AI systems must therefore be linked to approved policies and refreshed governance content, not just generic language reasoning.

Seventh is cybersecurity and model misuse. Agentic systems expand the attack surface by connecting models to tools, APIs, email systems, and enterprise applications. Microsoft (2025) and the World Economic Forum and Capgemini (2025) stress the need for monitoring, safeguards, and evaluation because action-enabled systems can amplify security failures.

Eighth is integration readiness. Many finance organizations still operate through heterogeneous ERP landscapes, local spreadsheets, shared-service platforms, and manual sign-offs. In such environments, agentic orchestration may fail not because the model is weak, but because process interfaces are inconsistent.

Ninth is workforce readiness. Finance professionals must know when to trust, challenge, and override AI. Anh *et al*. (2024) show that readiness and perceived usefulness influence adoption by accountants, while Mgammal *et al*. (2024) underline the importance of skills and attitudes in shaping AI's practical effect on accounting systems.

Tenth is value measurement. Several enterprise surveys show that organizations struggle to move from experimentation to scaled value because they lack disciplined adoption, governance, and performance management mechanisms (McKinsey & Company, 2025; OpenAI, 2025). In finance, this means benefits should be measured not only through cost and time but also through reliability, review effort, issue resolution quality, close stability, forecast accuracy, and stakeholder confidence.

4.4 Stakeholder alignment is an underappreciated outcome

A fourth and especially important finding is that AI adoption in finance affects stakeholder alignment, not just process speed. Finance is a translation function. It translates operational performance into management insight, converts accounting records into external reporting, converts assumptions into budgets and forecasts, and converts risk signals into control responses. Many of the pain points in finance therefore arise from misalignment: business units and controllers interpret variances differently; regional teams follow inconsistent definitions; FP&A and operational leaders work from conflicting assumptions; auditors request evidence that internal teams cannot quickly reconstruct. AI-enabled augmentation can improve alignment by standardizing definitions, drafting comparable commentary, surfacing exceptions early, and preserving traceable reasoning across workflows (KPMG, 2025a; EY, 2025; World Economic Forum, 2025b).

This insight is particularly important for the move toward agentic systems. An agent that coordinates close tasks or scenario workflows can reduce hidden handoffs and make process status more visible across participants. An explanation engine can standardize narrative templates and reduce ambiguity in management packs. A policy-aware assistant can help business-unit finance teams interpret corporate accounting rules more consistently. These uses matter because stakeholder alignment is one of the main ways finance creates enterprise value. Faster work is useful; faster and more coherent work is transformative.

4.5 The literature gap

Although the literature on AI in finance, accounting, and auditing is expanding rapidly, a clear gap remains. Most studies focus on use cases, technological possibilities, or general barriers. Fewer offer a finance-specific framework that links autonomy design to internal controls and stakeholder accountability. Similarly, the agentic AI literature is rich in technical and governance discussions but less explicit about how highly controlled corporate functions should translate those ideas into day-to-day operating models. This paper responds to that gap by proposing a constraint-aware adoption framework.

5 A CONSTRAINT-AWARE FRAMEWORK FOR AGENTIC AI ADOPTION IN CORPORATE FINANCE

Based on the review, this paper proposes a five-layer framework for adopting agentic AI in finance. The framework is designed for controllership, FP&A, treasury, tax, shared services, and internal-audit-adjacent workflows. It assumes that the objective is not maximal autonomy but reliable value creation under finance-specific constraints.

5.1 Layer 1: Use-case qualification

The first layer asks whether a use case is structurally suitable for AI augmentation or bounded agency. Four filters should be applied. The task should have repeated patterns or reviewable precedent; the data environment should be sufficiently accessible and governed; the output should be reversible or reviewable before external consequence; and the process should have a clearly named owner. Suitable early cases include variance commentary drafting, close status monitoring, reconciliation triage, policy lookup, planning scenario assembly, board-pack summarization, and internal control evidence collation. Less suitable initial cases include autonomous posting of material entries, final approval decisions, judgment-heavy tax positions, or unreviewed external disclosures.

5.2 Layer 2: Autonomy calibration

Once a use case passes qualification, the autonomy level should be calibrated. This paper proposes four modes. Mode A is assistive automation, where AI generates summaries, classifications, or drafts but cannot trigger actions. Mode B is supervised augmentation, where AI recommends next steps, escalations, or scenario paths and a human explicitly approves movement. Mode C is bounded delegation, where the agent can perform low-risk actions inside tightly defined policy and access boundaries, such as routing reminders, collecting evidence, or creating draft workpapers. Mode D is conditional autonomy, where the agent can execute limited actions in highly standardized processes with real-time monitoring and rollback. Most finance organizations should expect the majority of material use cases to remain in Modes A and B in the near term, with selected administrative tasks moving into Mode C. Mode D should be rare and reserved for low-materiality, high-standardization contexts.

5.3. Layer 3: Constraint and control architecture

The third layer translates finance constraints into explicit design requirements. At minimum, every agentic use case should specify approved data sources, identity and access scope, policy references, logging standards, review checkpoints, exception thresholds, and rollback procedures. This layer also requires segregation-of-duties mapping. If an agent drafts a journal, a distinct reviewer must approve it. If an agent compiles planning assumptions, changes to key drivers must be attributable and reviewable. If an agent monitors close dependencies, the routing logic must be transparent and its alerts auditable. This is where governance frameworks become operational rather than abstract (NIST, 2023; AI Verify Foundation & IMDA, 2024; IMDA, 2026).

A practical design rule emerging from the literature is that agentic finance systems should be “policy-bound, identity-bound, and evidence-rich.” Policy-bound means they operate against current finance rules and thresholds. Identity-bound means every action occurs within explicit permission models rather than opaque service-level access. Evidence-rich means recommendations and actions leave a reviewable trail. These principles reduce the gap between AI functionality and finance assurance requirements.

5.4 Layer 4: Human accountability and operating model

The fourth layer defines who remains accountable. The review suggests that finance organizations should avoid diffuse ownership. Instead, every deployment should assign at least four roles: a business process owner, a control owner, a model or system steward, and a user population with clearly scoped rights. For high-impact workflows, internal audit or second-line risk functions may also need visibility into design assumptions and monitoring results. This structure aligns with emerging governance guidance that responsibility for AI outcomes cannot be delegated to the technology itself (Financial Stability Board, 2024; BIS, 2024; World Economic Forum & Capgemini, 2025).

The operating model should also redefine work. If AI reduces time spent on compilation and first-draft analysis, finance leaders should intentionally shift capacity toward review quality, challenge dialogue, scenario stress testing, and stakeholder communication. Without this redesign, organizations may deploy sophisticated tools but preserve low-value review routines, thereby limiting return on investment.

5.5 Layer 5: Value and reliability measurement

The fifth layer measures outcomes. Traditional automation metrics such as hours saved and cost reduction remain relevant, but they are insufficient. Finance adoption should also be assessed using reliability indicators. These may include close-cycle stability, percentage of AI outputs accepted after review, number of control exceptions, frequency of override, time to resolve reconciliations, planning-cycle responsiveness, audit inquiry turnaround, and user trust scores. In other words, the right question is not only whether the agent is faster, but whether the finance process becomes more dependable, explainable, and aligned.

The framework therefore defines successful adoption as the combination of five outcomes: analytical acceleration, control integrity, stakeholder alignment, workforce leverage, and scalable governance. An initiative that saves time but weakens auditability should not be treated as a success. Conversely, a deployment that modestly reduces cycle time while materially improving exception visibility and management alignment may create disproportionate strategic value.

5.6 Implementation sequence

The framework implies a staged implementation path. Stage 1 is discovery and use-case inventory. Stage 2 is data, control, and policy mapping. Stage 3 is prototyping in assistive or supervised modes. Stage 4 is control validation, security review, and evidence-design testing. Stage 5 is scaled deployment with performance dashboards and recurring policy refresh. Stage 6 is selective movement toward bounded delegation where reliability has been demonstrated. This sequence reflects the broader finding that agentic adoption should follow control maturity, not precede it.

Figure 2 summarizes the proposed framework and shows how use-case qualification, autonomy design, constraint controls, human oversight, and outcome measurement interact.

Figure 2

Constraint-aware framework for agentic AI adoption in corporate finance.

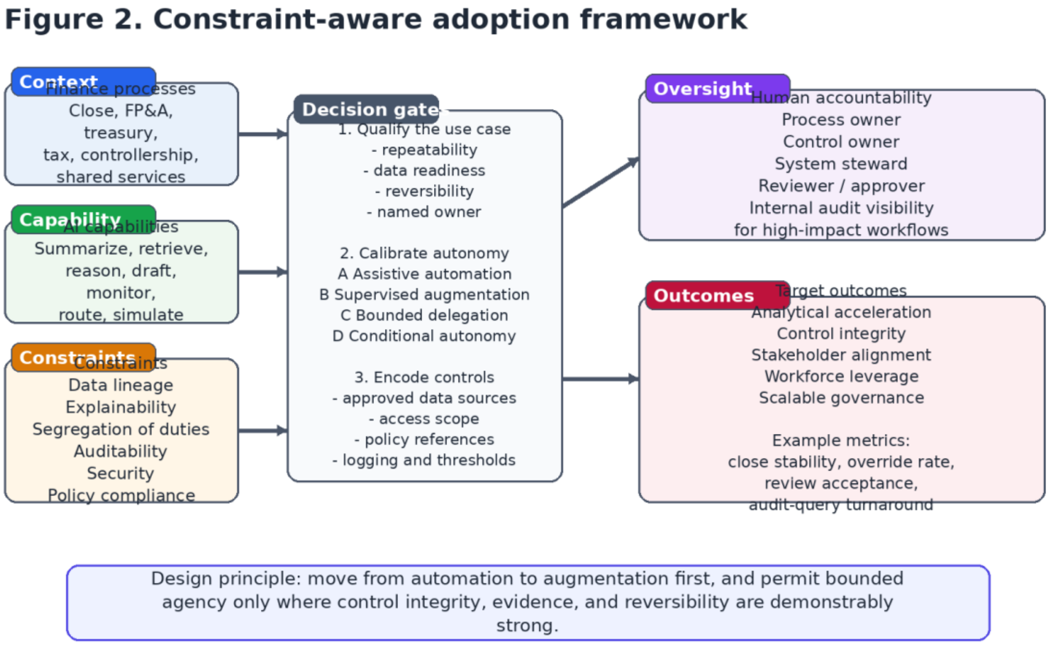

6 DISCUSSION

The review highlights that agentic AI adoption in corporate finance should be approached as a redesign of work rather than merely the purchase of new tools. For CFOs and transformation leaders, the most valuable use cases are those that improve decision speed, review quality, and issue visibility, rather than flashy but poorly integrated applications. In this context, bounded augmentation is presented as more appropriate than full autonomy, because finance operates under high accountability and control requirements. AI can add value by supporting retrieval, monitoring, drafting, synthesis, and escalation, while humans retain responsibility for material judgments and approvals.

The paper also stresses that governance and assurance models must evolve. Traditional control systems were designed for humans and deterministic software, whereas agentic AI introduces hybrid actors whose behavior depends on models, prompts, and system constraints. As a result, organizations may need new evidence trails, including model versions, retrieval sources, prompts, and records of human intervention.

For finance professionals, the findings suggest not simple job replacement but role reconfiguration. Routine manual tasks may decline, but the importance of judgment, explanation, control literacy, and cross-functional influence will increase. This means finance capability building should expand toward AI evaluation, prompt design, data governance, and model-risk awareness.

Overall, the review concludes that augmentation is the best near-term design principle for finance. Properly governed AI can improve reporting, communication, and stakeholder alignment, while preserving trust, accountability, and control integrity across the finance function.

REFERENCES

Abbas, K. (2025). Management accounting and artificial intelligence: A comprehensive literature review and recommendations for future research. *The British Accounting Review*.

Abdo-Salloum, A. M., & Chehade, S. (2026). The role of artificial intelligence in transforming accounting and auditing practices: A systematic review. *SAGE Open*. https://doi.org/10.1177/21582440251403296

AI Verify Foundation & Infocomm Media Development Authority. (2024). *Model AI governance framework for generative AI*. https://aiverifyfoundation.sg

Anh, N. T. M. (2024). The effect of technology readiness on adopting artificial intelligence by accountants. *Journal of Risk and Financial Management*, *17*(1), Article 27.

Anwar, A. (2026). Integrating artificial intelligence in audit workflow. *AI, 2*(1), Article 4.

Bahoo, S., Aschfalk-Evertz, D., Hassan, M., & Kchouri, B. (2024). Artificial intelligence in finance: A comprehensive review through bibliometric and content analysis. *SN Business & Economics*. https://doi.org/10.1007/s43546-023-00618-x

Ben Ahmed, D. (2026). The impact of artificial intelligence on accounting. *Journal of Risk and Financial Management*, *19*(1), Article 90.

Binh, N. T. T. (2025). Transforming auditing in the AI era: A comprehensive bibliometric analysis. *Information*, *16*(5), Article 400.

Boston Consulting Group. (2024). *Unlocking impact from AI in the finance function*.

Cao, S. S. (2024). Applied AI for finance and accounting: Alternative data and opportunities. *Pacific-Basin Finance Journal*.

Cao, Y. (2025). How AI is shaping accounting and finance. *The British Accounting Review*.

Deloitte. (2024, March 5). *Generative AI in finance: 2023 lookback, 2024 outlook*.

EY. (2025). *How AI is transforming FP&A.*

EY. (2026). *Unlocking financial potential*.

Financial Stability Board. (2024). *The financial stability implications of artificial intelligence*. https://www.fsb.org

Infocomm Media Development Authority. (2026). *Model AI governance framework for agentic AI*. https://www.imda.gov.sg

Kerr, D. (2025). A review of AI and its impact on management accounting and society. *Journal of Risk and Financial Management*, *18*(6), Article 340.

Kokina, J., Blum, M., Cunningham, B., & Pickerd, J. (2025). Challenges and opportunities for artificial intelligence in auditing. *International Journal of Accounting Information Systems*.

KPMG International. (2024). *Global AI in finance report*.

KPMG. (2025). *The agentic AI advantage: Finance agents that move the numbers*.

L.E.K. Consulting. (2025). *A study of AI in the OCFO*.

Lee, D. K. C., Lim, K. G., Low, R. K. Y., & Tan, Y. X. (2024). A comprehensive review of generative AI in finance. *FinTech*, *3*(3), Article 25.

Lehner, O. M., Ittonen, K., Silvola, H., Ström, E., & Wührleitner, A. (2022). Artificial intelligence-based decision-making in accounting and auditing: Ethical challenges and normative thinking. *Accounting, Auditing & Accountability Journal*, *35*(9), 109–133.

Leocádio, D., Malheiro, L., & Reis, J. (2024). Artificial intelligence in auditing: A conceptual framework for auditing practices. *Administrative Sciences*, *14*(10), Article 238.

McKinsey & Company. (2024). *McKinsey on finance, number 87*.

McKinsey & Company. (2025). *The state of AI: How organizations are rewiring to capture value*.

Microsoft. (2025). *Responsible AI transparency report*.

National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)*. U.S. Department of Commerce. https://doi.org/10.6028/NIST.AI.100-1

OpenAI. (2025). *The state of enterprise AI*.

Organisation for Economic Co-operation and Development. (2024). *Generative artificial intelligence in finance*. https://www.oecd.org

Organisation for Economic Co-operation and Development. (2026). *The agentic AI landscape and its conceptual foundations*. https://www.oecd.org

PwC. (2024). *Agentic AI: The new frontier in GenAI — An executive playbook*.

Taulia. (2024). *The rise of AI in the finance function*.

World Economic Forum & Capgemini. (2025). *AI agents in action: Foundations for evaluation and governance*.

World Economic Forum. (2025a). *AI in action: Beyond experimentation to transform industry*.

World Economic Forum. (2025b). *Artificial intelligence in financial services*.

World Economic Forum. (2025c). *Advancing responsible AI innovation: A playbook*.

Yeo, W. J., Ng, A. H. H., Lim, K. G., Poh, W. T., & Lee, D. K. C. (2025). A comprehensive review on financial explainable AI. *Artificial Intelligence Review*. https://doi.org/10.1007/s10462-024-11077-7

Authors' Contribution

All authors contributed equally to the development of this article.

Data availability

All datasets relevant to this study's findings are fully available within the article.

How to cite this article (APA)

Jalaluddin, K. A. THE TRANSITION FROM AUTOMATION TO AUGMENTATION: A CONSTRAINT-AWARE FRAMEWORK FOR AGENTIC AI ADOPTION IN CORPORATE FINANCE FUNCTIONS. Veredas Do Direito, e235632. https://doi.org/10.18623/rvd.v23.5632