

**FSN**<sup>®</sup>  
The Modern Finance Forum



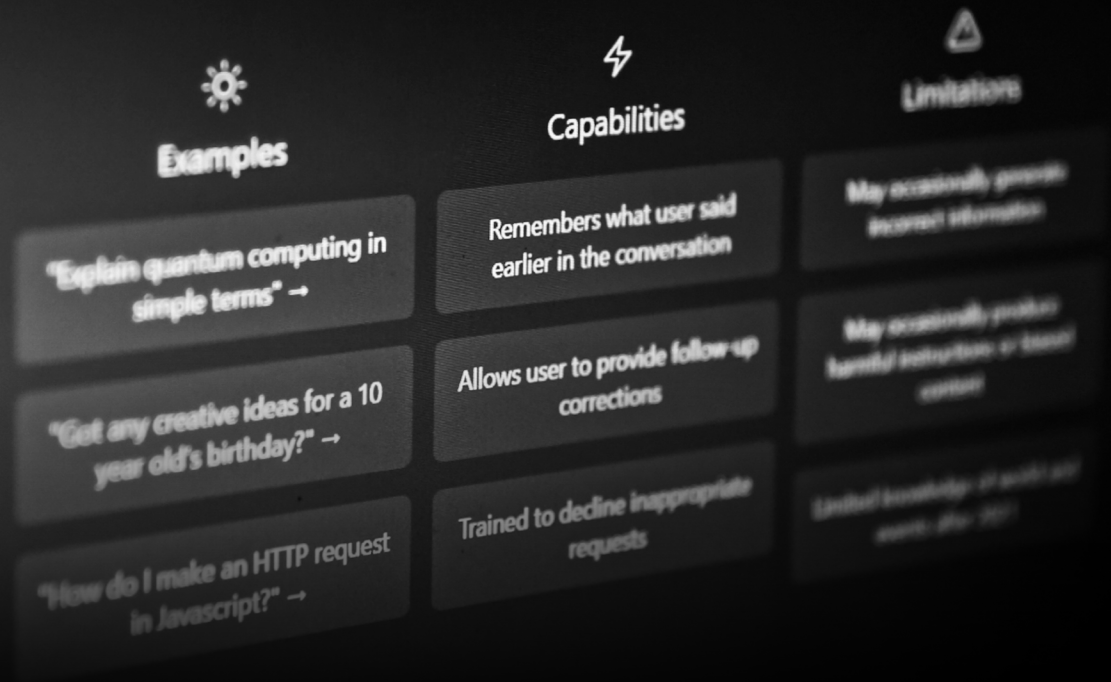
Agentic AI in Finance

# Capturing Value Without losing Control

 Prophix<sup>®</sup>

# Introduction

## ChatGPT



The advent of ChatGPT ushered in a level of excitement, curiosity and innovation that has not been seen in the finance function since the introduction of the spreadsheet. However, early experimentation, largely driven by 'fear of missing out' and short-term cost reduction, yielded mixed results.

Fortune.com, quoting an MIT report in August 2025<sup>1</sup> pointed out that how companies adopt AI is crucial. For example, purchasing AI tools from

specialised vendors and building partnerships succeed about 67% of the time, while internal builds succeed only one-third as often. MIT also found the biggest ROI is in back-office automation.

Even in this experimental phase, the direction of travel was clear: AI can deliver material value to the CFO and, with the emergence of agentic capability, will likely reshape how finance work gets done.

# From pilots to execution: the shift finance can't ignore



A 2025 Salesforce research report<sup>2</sup> found that

**74%**

of CFOs believe AI agents will not only cut costs but also drive revenue growth of almost

**20%**


There is now a growing sense that experimentation has run its course, and attention is shifting from AI that informs decisions to AI agents that act on them. In other words, agents that execute multi-step tasks within defined constraints, rather than simply generating content and recommendations. These agents are designed to behave like virtual finance teammates, automating workflows across planning, analysis, reporting and the close.

The implications are transformational on a scale that would have been hard to imagine two or three years ago. A 2025 Salesforce research report<sup>2</sup> found that 74% of CFOs believe AI agents will not only cut costs but also drive

revenue growth of almost 20%. A 2025 Deloitte survey<sup>3</sup> reports that more than half of finance chiefs (59%) have become more optimistic over the past 12 months about AI's potential to improve organisational performance. Meanwhile, KPMG's Quarterly AI Pulse Survey<sup>4</sup> suggests that AI agent deployment accelerated through 2025, more than doubling from 11% in Q1 to 26% in Q4.

However, as excitement builds, many CFOs also recognise that agentic AI introduces new risks. Enthusiasm is increasingly tempered by concerns that ungoverned use could undermine the trustworthiness of reporting and weaken established controls and governance.

# What an AI agent is — and isn't



CFOs have invested heavily in automation over the last decade, and agentic AI is the latest iteration, building on traditional automation and copilots.

Traditional automation focuses on embedded workflows and collaboration, as seen in modern consolidation and budgeting systems. For example, the progress of monthly submissions and budget iterations is tracked, and changes are routed for completion, approval or rejection. This approach has accelerated cycle times and tightened connections between people and process.

Copilots represent a further step change in capability and can be viewed as an interactive assistant. Finance asks questions and copilots draft,

summarise, explain, reconcile and propose next steps. For example, an FP&A analyst might ask, "Why did we experience such a large margin variance?" The copilot can extract driver detail (volume, price, mix), create an explanatory narrative and prepare slide-ready bullet points — compressing the time spent finding, explaining and drafting insights, all grounded in the organisation's own data.

AI agents go further because they can act on this derived information, potentially autonomously. Nonetheless, most finance leaders agree that agents require guardrails: permissions, controls and, in most cases, a human in the loop who grants or denies further action.

# Where agents fit in the finance operating model

It is still early days, but a view is emerging on where agents are most effective.

Traditional automation lends itself to accounting transactions: rules-based, repeatable, high volume and with limited contextual variability. At the other end of the spectrum, work that depends on professional judgement and multidisciplinary consultation, such as interpreting accounting policy in novel circumstances or making commercial decisions that require deep industry context — is less suited to AI automation.

Between these extremes lies the sweet spot for agentic AI. In this middle ground, AI can drive deeper analysis, surface insights rapidly, reduce time to decision and, where appropriate, take controlled action subject to human approval.

#### Examples might include:

- A “reporting agent” that investigates quarter end margin movements, highlights the main drivers of under or over-performance, identifies causes, isolates exceptional variances and offers different viewpoints, while showing its thinking every step of the way.

- A “budgeting agent” that supports revised headcount decisions or revisions to marketing spend, accepting natural language instructions and making changes to the underlying budget model based on business rules and meta data (employee grade profiles or seasonality), without manual intervention or manipulating data-entry spreadsheets.
- An “intercompany elimination agent” that identifies mismatches during consolidation, pinpoints root causes, assesses materiality and proposes balancing journals for human review – taking pressure off the close.
- A “challenger agent” that stress-tests forecast scenarios, challenges assumptions, identifies gaps and proposes alternative outcomes for review and acceptance.

Every finance use case is different and judgement is required to determine where value truly lies. But the middle ground between transactional processing and judgement-heavy deliberation appears to be the most fertile territory for agentic AI, even if autonomy is often partial rather than absolute.

# From hype to Responsible AI

As agentic capability expands, “Responsible AI” becomes the differentiator. Responsible AI is the discipline of capturing AI’s potential without compromising the trust, standing and credibility of the finance function. It ensures that every AI outcome is explainable and rests on four pillars: sound architectural choices, mastery of the data environment, transparency in AI-supported processes, and governance.

## Architectural choices

The AI marketplace is awash with different architectural choices for the finance function, for example,



standalone AI tools outside of the finance control environment



spreadsheet based or AI add-ons, layered on top of existing ways of working



custom internal builds, often assembled by IT or data teams

While these approaches may look like an easier entry point to AI, they come with familiar trade-offs: fragmented data integration, inconsistent user experiences, limited explainability, and weaker control and governance. By contrast, AI natively embedded within finance applications is typically more joined up and controllable. The finance function’s long history with spreadsheets illustrates how quickly separate, ad-hoc solutions undermine the control environment.

Custom internal builds add a different set of risks. They

often struggle with true ‘data understanding’ (definitions, lineage, context) and they create a long-term support and maintenance burden, especially when key people move on.

On the other hand, pre-built agentic capability embedded in the core of finance owned applications appears to offer a structurally better way forward. It offers faster, lower-friction access in a familiar environment with maximum data understanding, trusted data which remains private to the enterprise, and established controls with clearer accountability.

## Data mastery

In the agentic era, finance functions need to be obsessive about data quality, security permissions and protecting the data universe from unintended exposure. Finance data sets extend well beyond 'pure' financials into operational and non-financial measures which is often where much of the competitive value lies. If AI outputs are to be relied upon, the same standards of quality, lineage and control must apply across both.

A crucial advantage of AI embedded in core finance applications is that the underlying data structures,

financial intelligence and metadata (hierarchies, dimensions, cost centres, products and so on) are accessible and understood by the agent. Metadata becomes even more significant in the agentic era because it describes what data contains, who owns it and how reliable it is. The better the metadata, the more effective the AI agent can be.

But access cannot be unfettered. Users and agents must be restricted to authorised data, enforced through role-based permissions and monitored usage.

## From "black box" to "glass box" (transparency)

Early implementations of AI, particularly in forecasting, often suffered from "black box" syndrome. The modelling logic used by AI algorithms was opaque, leaving FP&A teams dependent on vendors and data scientists to explain behaviour and make changes. This lack of transparency undermined confidence in the outputs.

Conversational AI has shifted expectations. Well-designed agents are able to explain their reasoning; show their workings and cite the inputs they relied upon. The goal is explainability that allows finance to validate, challenge and approve outcomes — a level of transparency that restores control to the function.

## Governance

CFOs ignore AI governance at their peril. AI-induced errors can quickly translate into regulatory and audit exposure, unwelcome board and investor scrutiny, and, in some cases, personal accountability. It's so important that 57% of global enterprises have formal AI oversight roles, such as Chief AI Officers and Ethics Boards.<sup>5</sup>

When autonomous or semi-autonomous agents can access data, recommend decisions and initiate actions, guardrails are essential. Trust by design, governance-as-code and a human in the loop become the fabric of the new control environment. Actions cannot always be unwound, so governance must be embedded from the start.

### In practice this means



**Keep a human in the loop.** Define where an agent's authority starts and ends. Set approval rules by risk and materiality and require supporting evidence before anything is actioned. Be explicit on data permissions too: what the agent can access, share, and what must remain ring-fenced.



**Manage agents like part of the team.** Treat them like inexperienced junior analysts. Give them tight mandates, challenge the logic, review exceptions and confirm outcomes.



**Segregation of duties.** Let agents prepare, reconcile and recommend — but keep approvals, postings and external submissions under human control.



**Build auditability in by design.** Ensure a full trail: who initiated, what data was used, what ran, what controls were triggered, what changed, and who signed off.



**Governance-as-code.** Encode executable rules (data access, escalation thresholds, approval steps and permitted actions) directly into the agent's runtime, so compliant behaviour is consistent and auditable.

# Role sharing with AI agents — a new paradigm

Agentic AI is set to have profound implications for how work is managed in the finance function. This is not organisational change in the traditional sense, for example, a restructuring of responsibilities, moving staff around, or simply reallocating tasks between team members. Instead, it points to a re-imagining of roles, built around a different way of thinking about how work is shared between professionals and agents, i.e., between human and digital labour.

Rather than asking, “Which tasks do we move between people?”, the better question becomes: “Which components of existing roles can be ceded to AI agents (under supervision) so that humans can focus on judgement and accountability?”

Take an FP&A manager who oversees scenario planning. Typically, that remit includes shaping strategic assumptions, scanning market conditions, identifying variables, and building forecasts across multiple scenarios and probabilities. With an AI agent in the workflow, the role can be split into three sub-roles: assumption generation,

evaluation, and strategic alignment.

An agent could propose new assumptions based on historic performance, current market signals and the strategic parameters set by the FP&A manager. It could then test forecasts for reasonableness and internal consistency and assess alignment with stated strategic goals. The FP&A manager, as the human in the loop, remains responsible for the final judgement, for instance, whether the forecast is credible, actionable and fit for decision-making.

A similar pattern applies to the group financial controller. An agent could oversee submissions of monthly actuals from reporting entities, validate completeness and accuracy, chase missing returns, and trigger management and statutory consolidations. With submissions, validations and the initial consolidation largely in the hands of an AI agent, the controller is freed to focus on value-added complexities such as consolidation adjustments, technical accounting judgement, disclosures and performance narrative.

# The AI dividend



In the above examples, AI agents relieve highly qualified finance professionals from manually intensive, repetitive work, enabling focus on where professional judgement adds the most value. For instance, in the samples above;

- Cycle times for consolidation and scenario planning fall as agents handle preparation and initial checks.
- Routine interactions between reporting entities and the corporate centre, become fewer, faster and more targeted.

- Data collection becomes simpler, more automated and less reliant on chasing.

Workflows are enhanced by allowing agents to do the heavy lifting, while humans remain accountable for curating the process and ensuring completeness, accuracy and governance. Capacity that would otherwise be squandered in process mechanics is released for higher-value work. Agentic AI is not displacing headcount — it is elevating it.

# When will CFOs handover the reins to AI?

Artificial intelligence brings into sharp focus the dual responsibilities of the modern CFO. On the one hand, acting as a strategic adviser and business partner to the board; on the other, being a safe pair of hands — a financial steward and dependable guardian of governance, risk management and control.

AI therefore presents an unprecedented dilemma for finance. CFOs do not want to be seen to stifle progress, yet they are being asked to endorse systems that not only generate insights but initiate actions. The tension sharpens further with agentic AI, where recommendations may be executed autonomously or semi-autonomously, even with a human in the loop.

AI is already proving its worth in a range of scenarios, and most CFOs acknowledge its value. More than half (54%) of North American CFOs say integrating AI agents into finance will be one of their top finance transformation priorities in 2026.<sup>6</sup> But finance leaders will not readily embrace agentic capability if it introduces unacceptable risk

of error, material misstatement or reputational damage. The deeper question is therefore not whether AI is useful, but how CFOs gain sufficient confidence that risk is contained within tolerable limits.

As set out earlier, that confidence rests on sound architecture, data mastery, process transparency and governance. AI embedded within finance processes — rather than bolted on — and grounded in proprietary data that is understood, ring-fenced and secure, provides a more trustworthy platform for agentic capability, provided a human remains accountable for outcomes.

It is the prospect of AI acting autonomously and its ability to provide convincing answers to things that might be made up or wrong, that truly rattles nerves — even in a controlled environment. In finance, autonomy is earned gradually through proof, discipline and repeatable control. Finance professionals will need to see for themselves a solid track record of consistency and accuracy before handing over the reins completely to AI agents.

# Conclusion — will CFOs act as AI change agents?

CFOs are arguably the best qualified members of the C-suite to lead the charge. Their understanding of the balance sheet and what drives enterprise value coupled with their influence over strategy and technology decisions makes them natural business partners and credible agents of change. That big-picture perspective matters, because agentic AI is not only about new, differentiated business models; it is equally about building more progressive, agile and resilient business processes.

Just as importantly, the CFO's deep knowledge of governance and the control environment is invaluable as organisations take steps towards more autonomous operations — not only in finance, but across other business functions as well.

Autonomous finance may be the long-term vision, but the pace of progress will be defined by how quickly developments address concerns about the trustworthiness of agentic AI. Not all finance functions will move at the same rate, and adoption is likely to be incremental. Much depends on the readiness of the pillars described earlier: architectural choices, data mastery, transparency and governance.

**So the 'winning model' isn't autonomous finance at any cost. It's about increasing autonomy as trust is earned. This will happen where;**

agentic AI is embedded and sits inside finance-owned platforms and processes

there is a strong commitment and investment in data governance and data understanding

there is full transparency, "the glass box" and every material output is explainable

agents do the 'heavy lifting', but there is always a human in the loop.

Finance can lead the way by serving as a proving ground. It may not be the showiest use case, but it is demanding and, smart CFOs know that if agentic AI can operate quietly and dependably in a tightly governed finance setting, finance's lead will become a credible blueprint for the wider enterprise to follow.

# Bibliography

1. [Fortune.com quoting MIT NANDA Initiative, The GenAI Divide: State of AI in Business 2025](#)
2. [Salesforce Research Report, "From Caution to Core Strategy", 2025.](#)
3. [Deloitte CFO Survey \(UK\): finance leaders upbeat on tech investment and AI potential for the year ahead, January 2026.](#)
4. [KPMG, Quarterly AI Pulse Survey / Executive Pulse: GenAI strategy and investment trends, 2025.](#)
4. [KPMG, Agents of change: New organizational roles in the age of AI December 2025](#)
6. [Deloitte, North American CFOs reveal their top 6 expectations for 2026, January 2026](#)

## About the Author



Gary Simon is a fellow of the ICAEW and a fellow of the British Computer Society with more than 40 years' experience of enterprise performance management with global 2,000 companies and mid-market businesses. He

is the author of four books on financial systems and a popular independent speaker on market developments. He was formerly a senior consulting partner in Deloitte for more than 16 years.

# About FSN

[FSN®](#) is a global publisher of thought leadership, research and “must-have” content for CFOs and senior finance professionals around the world. FSN’s highly popular and active FSN Modern

Finance Forum on LinkedIn has a membership of more than 60,600 readers in more than 23 countries and across every major industry segment.

**FSN®**  
The Modern Finance Forum

---

# About Prophix

Prophix® is a global leader in financial performance management, empowering finance teams to lead with clarity, capacity, and confidence. From planning and budgeting to forecasting, reporting, reconciliation, and consolidation, Prophix brings it all together in one intelligent platform.

Prophix One™, the flagship Autonomous Finance Platform, combines AI, automation, and intuitive technology to simplify complex work and elevate finance to a more strategic role. With nearly four decades of innovation and a global footprint serving more than 3,000 customers in 100+ countries, Prophix is the trusted partner for organizations ready to transform finance into the driving force behind business growth.

 **Prophix®**

# Prophix UK contact details

---

Paddington Works  
8 Hermitage Street  
W2 1BE, London

Tel: +44(0) 208 050 3590  
[www.prophix.com/uk](http://www.prophix.com/uk)

#### Disclaimer of Warranty/Limit of Liability

Whilst every attempt has been made to ensure that the information in this document is accurate and complete some typographical errors or technical inaccuracies may exist. This report is of a general nature and not intended to be specific to a particular set of circumstances. The publisher, FSN Publishing Limited and authors make no representations or warranties with respect to the accuracy or completeness of the contents of this white paper and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives, or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. FSN Publishing Limited and the author shall not be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.