The CFO’s AI Agenda: From Automation to Advantage

May 2026
By Nick Verhoeven, Marc Rodt, Alexander Roos, Michael Demyttenaere, Edoardo Palmisani, and Ichiro Kaku

Every CFO has seen the demos. An AI that drafts variance commentary in minutes. A planning agent that builds scenario analyses from a single prompt. An autonomous engine that reconciles thousands of transactions overnight without human intervention. The capabilities are impressive—and, increasingly, they are moving from prototype to product.

But it is also clear that most finance organizations still have not turned that potential into measurable impact. **BCG's AI Radar 2026** found that spending on AI is set to double as a share of revenue, and 94% of organizations plan to continue investing despite uncertain short-term returns. At the same time, the gap between expectation and impact remains wide: while nearly 90% of CEOs expect AI agents to deliver measurable return on investment, relatively few organizations can point to consistent results today.

It's not surprising that expectations are rising faster than returns. The technology is genuinely new, still maturing in important ways, and no team has years of experience to fall back on. But technology is not the only factor. BCG's research consistently shows that only about 10% of AI success can be traced to the models themselves, and another 20% to the underlying technology platform. The remaining 70% depends on organization, workforce, and skills. These are the foundations that most finance teams are still building.

As new AI capabilities arrive, finance functions must get ready to capture value by strengthening the data, processes, and capabilities that allow these tools to perform in real operating environments. AI only creates value when it operates on a trusted data foundation, close to the business processes it serves.

What's Coming

For CFOs, two developments are particularly important: the evolution from assistive tools to more autonomous, agentic systems and a rebalancing of the planning cycle.

The Evolution to Agentic Systems

Most AI in finance today works reactively. It summarizes documents, generates charts, and drafts commentary when prompted. This is useful but fundamentally passive. The human defines the task, triggers the tool, and validates the output. That is the copilot model, and it delivers a real—but modest—productivity gain.

What's emerging now is different in kind. Agentic AI systems do not wait for prompts. They monitor data continuously, detect deviations, reason through potential causes, and propose, or in some cases, execute responses. Think of an agent that detects weakening order intake in a key customer segment, traces the impact through the revenue forecast, and proposes a revised scenario — adjusting pricing, channel mix, or production plans – for the planning team to review. That is a fundamentally different way of organizing work, not simply a smarter chatbot.

The capabilities available today vary by domain. In high-volume transactional processes—invoice processing, cash application, reconciliation, and expense management—agentic capabilities are already being used. Specialized vendors report touchless automation rates above 90% for cash application, while auto-matching engines handle the bulk of reconciliation workload.

In planning and performance management, the picture is more nuanced. Vendors have introduced AI-driven model generation, conversational planning, and continuous reforecasting, but most capabilities are still at an early stage or are limited in scale. In domains such as tax and treasury, agentic capabilities remain largely aspirational.

At the same time, agentic AI is adding a complementary cycle for more operational, ad hoc decision making. While the plan–do–check–act cycle structures tactical and strategic planning, agentic AI enables a faster loop of sense–reason–act. Systems continuously sense signals across the business, reason through causes and constraints, and propose actions in near real-time. This type of short-cycle decision making was historically difficult to implement within enterprise planning environments, which were optimized for structured, periodic workflows. With stronger data foundations, it is now becoming feasible at scale.

While the ambition is real, so is the need for realism—on two fronts. First, with production readiness varying widely by domain, finance leaders need a clear view of what is deployable today versus what will emerge over the next 12 to 18 months. Second, they should be mindful that agentic AI is not a move toward unsupervised automation. Human validation remains essential, particularly in the early stages, with autonomy expanding only as systems demonstrate reliability. This is why the emphasis on people in AI transformations—the 70/20/10 split—matters: success depends less on the model than on the people who set the guardrails, review the outputs, and decide what to act on. Rather than replacing controllers, the goal is to equip them with a team of digital colleagues whose work they still own.

A Rebalanced Planning Cycle

A common misconception is that AI in planning replaces the management cycle. In reality, organizations will still plan, execute, review, and respond. What changes is where human effort is concentrated.

Planning becomes faster and more iterative. AI agents can generate planning structures from natural-language inputs—dimensions, measures, driver logic, and allocation rules—and deploy them directly into planning environments. The planner shifts from building models to validating them. At the same time, execution remains human-led, including decisions on capital allocation or headcount.

Review is where the impact is most immediate. Instead of periodic, manual comparisons of actuals against plan, agentic systems enable continuous monitoring. They sense not only internal deviations as they emerge but also external signals such as market shifts, competitor moves, or macroeconomic changes that previously required separate manual research. Deviations are flagged as they occur, not weeks later.

Response builds on this. Agents can trace variances through driver trees, identify root causes, simulate corrective actions, and propose responses grounded in the organization's data and logic. Finance and the business still decide, but faster and with better insight.

Two capabilities make this practical. Conversational interfaces allow planners to interact with systems in plain language, reducing reliance on technical model knowledge. And preconfigured domain content accelerates implementation, giving organizations a starting point rather than a blank slate.

There is a critical dependency, however. The quality of agentic output depends on the semantic richness of the data. An agent can only trace a variance if relationships are explicitly modeled, and it can only reason across units if definitions are consistent.

This does not require a multiyear transformation upfront. Data foundations can be improved alongside specific use cases, with each step increasing impact. But it does mean that data, definitions, and processes require as much attention as the AI itself.

How to Get Started

The impact of these innovations will depend less on the technology itself than on whether the ground it lands on is ready to support it. Finance organizations can start now, on several fronts.

Get your data house in order

When AI underperforms in a finance environment, the instinct is often to question the model. Yet almost always, the issue is that the data wasn't ready.

AI in finance fails not because the models lack capability but because the data lacks context. Most finance organizations operate on a patchwork of systems—aging ERPs, bolt-on planning tools, standalone spreadsheets, and manual integrations—each holding a partial view of the business. Human professionals navigate this complexity by means of institutional knowledge: they know which numbers to trust, how structures map across systems, and where adjustments are needed. AI has none of these advantages. A general ledger populated with raw account numbers and transaction codes gives a model nothing to reason about.

A related implication is that AI capabilities are most effective when embedded close to where data already resides, rather than layered on top through additional abstraction. Adding new layers can reintroduce inconsistencies in definitions and logic—the very issues that limit AI effectiveness. Keeping AI tightly coupled to core data and processes helps preserve consistency and trust and reduces the need for reconciliations.

The practical agenda starts with moving toward integrated data foundations that connect financial, operational, and planning data. This requires establishing live access to transactional data so planning models and AI agents work with current signals, not static extracts. And it requires enabling integration with external data—such as market signals, competitor data, and macroeconomic indicators—without fragile, manual pipelines.

This does not necessitate a single, monolithic data migration. That approach leads to multiyear programs, budget overruns, and organizational fatigue, and it is the wrong framing for an AI-readiness agenda. The more effective path is incremental and use case-driven. Focus on a specific planning domain, whether workforce, sales, or capital, and harmonize the data it requires. Replace manual extracts with live connections. Remove reconciliations that exist only because definitions are inconsistent. Each step measurably improves both AI readiness and finance performance.

Organizations seeing early returns from AI in finance share a common trait: they did not wait for perfect data. They picked a planning domain that mattered, harmonized the data it needed, connected the live feeds, let each step unlock the next—and then moved on to the adjacent domain. Momentum matters more than perfect data.

Harmonize your financial language, building the semantic layer for AI

Data availability is necessary but not sufficient. Even when finance data is accessible and timely, it often lacks the semantic richness AI requires to generate meaningful output. The issue is interpretability, not volume.

Consider an AI agent asked to explain a margin decline across regions. If "margin" is defined differently in Europe than in Asia-Pacific—because of differences in cost allocation, revenue recognition, or product hierarchies, for example—the agent has no stable foundation to reason from. It may produce a plausible, but unreliable, answer. Semantic fragmentation, the reality that financial terms mean different things across the same organization, is one of the most common reasons that AI in finance disappoints.

Addressing this requires work on two levels. The first is harmonizing the underlying data: establishing a unified chart of accounts, standardizing master data, and building consistent cost center, profit center, and product hierarchies. This is familiar finance transformation work. What makes it newly urgent is that AI amplifies every inconsistency. A human analyst can adjust for differing definitions. An AI agent can't, or it may do so without making the adjustment visible. Beyond preventing errors, harmonization encodes the company's operating model into the data itself—the hierarchies, relationships, and definitions that make transactions mean something. That encoded meaning is what AI actually reasons on.

The second level goes further: building a semantic layer that encodes the relationships AI needs to reason effectively. This means making driver relationships explicit: linking revenue to pricing, volume, and channel mix and linking costs to production, headcount, and supplier terms. It also means enriching transactions with context—not just account and amount, but region, product, channel, and customer segment. And it requires connecting financial and operational data, so an agent tracing a variance can follow it from the general ledger into the underlying business drivers. A large global retailer illustrates what this looks like in practice. (See "Answering 'Why' With Confidence.")

Answering "Why" With Confidence

A global retailer built a driver-based model of its P&L, linking top-line metrics such as operating income and SG&A down to operational drivers, including variable labor hours, wage rates, and cases processed per hour across the distribution network.

When an analyst asks why a category is falling short of plan, the engine first walks down the driver tree. It traces the variance through sales velocity into purchases per order and out to specific drivers—such as competitor pricing gaps and inflation assumptions—that explain the miss. This phase keeps the reasoning within the company's own data and its explicitly modeled relationships. The output is a structured retrieval, with its quality inherited directly from the underlying data.

After identifying the relevant drivers, the engine draws on its broader knowledge to hypothesize why they moved. By this point, the driver tree has narrowed the inquiry from "Explain the company's performance" to "Why did this specific driver move?" The engine reasons about a bounded problem with clear context, not the full breadth of company data. That framing significantly improves the quality of the hypotheses. The engine flags them as conjecture, separate from the walk down the tree. The analyst sees at a glance which parts of the answer are grounded in company data and which are informed speculation, and weighs them accordingly. The result: less manual work behind variance analysis without any loss in the ability to defend the conclusion.

This is where the data foundation becomes truly AI-ready. And like the broader data agenda, the work is modular—tackled for each region, entity, and use case. Every step improves both machine interpretability and human interpretability, making reports more meaningful and decisions more grounded.

Standardize processes and lead with value

In a pattern that repeats across industries, an organization launches a promising AI pilot that works in a controlled setting, leadership declares success—but then the pilot does not scale and the impact does not reach the P&L.

Research suggests that the vast majority of AI pilots follow this trajectory. The reason is rarely the technology. Most organizations try to layer intelligence onto processes that were never designed for consistency. When invoice processing follows different workflows across shared service centers, automating one does not fix the others. When the close relies on undocumented, entity-specific steps, an AI agent inherits every workaround. Automating a fragmented process scales fragmentation rather than eliminating it.

Most organizations deploying AI in finance are still engaged in the early stages of experimentation or in isolated demonstrations of capability. They have built the technology and, in some cases, added infrastructure and controls. But they have not yet tackled the harder work: redesigning processes end to end, adapting the operating model, training people, and tracking value rigorously enough to show impact on the P&L.

Getting there requires inverting the typical approach. Instead of starting with technology and searching for applications, start with value. Define specific, measurable use cases—a faster close, a more accurate forecast, or reduced reconciliation effort—and derive the data, process, and technology requirements from them. Clean, standardize, and simplify in parallel with early AI deployment, not sequentially before it. Use early gains to fund the next wave.

This value-led sequencing changes the economics of transformation. Each process standardized and each exception eliminated expands where AI can operate. The program begins to generate returns early and reinvests them, making it more sustainable and easier to defend.

The organizations extracting real value from AI in finance are not those with the most advanced models. They're the ones that made their processes consistent enough for those models to be useful.

Build AI fluency across the finance team

Of all the readiness dimensions, talent may be both the most consequential and the most underinvested. In early 2026, a Gartner survey of CFOs identified building AI talent within the finance function as their most pressing near-term challenge—not technology or budget.

The reason is structural. As AI takes on more analytical and transactional work, the human role shifts from executing tasks to navigating outcomes. Controllers who once compiled variance reports now need to validate AI-generated analyses, challenge underlying assumptions, and decide which findings warrant action. Financial planning and analysis professionals who built models manually now evaluate models generated by agents and direct them toward better outputs. The skills that made someone effective in yesterday's finance function are necessary but no longer sufficient.

This does not mean that every finance professional needs to become a data scientist. It means building role-specific AI literacy that is practical enough to be useful and structured enough to scale. At a foundational level, teams need to understand what AI can and cannot do, how to prompt it effectively, and how to identify outputs that appear plausible but are incorrect. At a more advanced level, finance leaders need frameworks to determine which tasks to delegate to AI and which to retain, considering feasibility, risk, and judgment. Not every task that AI can handle is one that it should.

It's a mistake to treat this as a training program—a one-time investment that produces a "skilled" workforce. AI fluency is built iteratively, by working alongside AI in progressively more complex workflows. Start with low-stakes applications where teams can build confidence, such as AI-drafted commentary, automated data gathering, or scenario generation. Expand into higher-stakes domains as both the technology and governance mature.

At the leadership level, this starts with personal practice. CFOs who use AI tools daily—for drafting, analysis, scenario testing—develop an intuition for the technology's strengths and limits that no training program can replicate.

The organizations that will struggle most are not those with the weakest technology. They are those whose teams never learned to work with it.

Treat governance as architecture, not compliance

When most organizations think about AI governance, they think about policies: acceptable use guidelines, ethical principles, or review committees. These matter, but they are not enough for finance.

Finance operates under constraints that most functions do not face. Precision is defined by regulatory standards, such as SOX controls, GAAP, and IFRS, where even small errors carry material consequences. Explainability is essential: when a CFO presents a variance analysis to the board, "the AI said so" is not an acceptable basis. Audit trails must link outputs back to source transactions. And segregation of duties becomes more complex when autonomous agents operate across process boundaries.

These are not abstract concerns. They are why governance in finance must be designed into how AI operates, not added afterward. That means building data lineage so outputs can be traced to inputs, defining escalation rules for when agents act or defer, and designing approval workflows that preserve segregation of duties even as agents take on more steps. It also means ensuring that when an agent generates an output—such as a planning model, a driver calculation, or an allocation rule—the result can be inspected by the humans responsible for approving it. An agent-generated model expressed as readable structure, not opaque code, lets the controller or auditor verify what was built and challenge it where needed.

Governance does not need to be fully mature from the start. It can, and should, evolve alongside adoption. Early use cases may rely on clear human-in-the-loop checkpoints and simple escalation rules. As systems demonstrate reliability, autonomy can expand incrementally. The starting posture is "AI proposes, human disposes," with the boundary shifting as trust is earned.

What matters is treating governance as a design decision from day one, not a compliance exercise triggered by the first audit. Organizations that delay will find governance far more expensive—and disruptive—to retrofit than to build from the outset.

The Real Race in Finance AI is Building Foundations

The agentic capabilities described here are not speculative. In transactional finance, they are already in use. In planning and performance management, they are less than a year away. The gap between early movers and laggards is closing faster than most organizations realize.

But this is not a race to deploy AI first. It is a race to build the foundations that make AI deliver value. Each step—harmonizing data, standardizing processes, building team fluency, and establishing clear governance—raises an organization's capacity to capture that value. The work is cumulative, and the impact compounds.

Perhaps the most important investment is in people. Finance teams must work alongside AI with confidence, judgment, and accountability. Technology is only as powerful as the human expertise that guides it.

These are no-regret investments. Better data improves reporting quality. Standardized processes reduce cost and error. More capable teams make better decisions. Strong governance reinforces control. Each of these strengthens the finance function regardless of how quickly AI evolves.

If AI delivers on its promise—and the evidence suggests it will—the finance organizations that build these foundations will capture value that others are not structurally equipped to realize. And by that time, it may be too late to catch up.

About the Authors

Nick Verhoeven
Global Product Marketer SAP Enterprise Performance Management
nick.verhoeven@sap.com

Michael Demyttenaere
Managing Director and Partner at BCG
demyttenaere.michael@bcg.com

Marc Rodt
Partner and Director at BCG, Center for CFO Excellence
rodt.marc@bcg.com

Edoardo Palmisani
Managing Director and Partner at BCG
palmisani.edoardo@bcg.com

Alexander Roos
Managing Director and Senior Partner at BCG
roos.alexander@bcg.com

Ichiro Kaku
Managing Director and Senior Partner at BCG
Kaku.Ichiro@bcg.com

Boston Consulting Group
Boston Consulting Group partners with leaders in business and society to tackle their most important challenges and capture their greatest opportunities. BCG was the pioneer in business strategy when it was founded in 1963. Today, we work closely with clients to embrace a transformational approach aimed at benefiting all stakeholders—empowering organizations to grow, build sustainable competitive advantage, and drive positive societal impact.

Our diverse, global teams bring deep industry and functional expertise and a range of perspectives that question the status quo and spark change. BCG delivers solutions through leading-edge management consulting, technology and design, and corporate and digital ventures. We work in a uniquely collaborative model across the firm and throughout all levels of the client organization, fueled by the goal of helping our clients thrive and enabling them to make the world a better place.

SAP
Founded in 1972, SAP has grown from a small, five-person endeavor to a multinational enterprise headquartered in Walldorf, Germany, with more than 105,000 employees worldwide.

With the introduction of its original SAP R/2 and SAP R/3 software, SAP established the global standard for enterprise resource planning (ERP) software. Now, SAP S/4HANA takes Cloud ERP to the next level by using the power of in-memory computing to process vast amounts of data, and to support advanced technologies such as artificial intelligence (AI) and machine learning.

The company's integrated applications connect all parts of a business into an intelligent autonomous suite on a fully digital platform, thereby replacing the process-driven, legacy platform. Today, SAP has more than 230 million cloud users, more than 100 solutions covering all business functions, and the largest cloud portfolio of any provider.