



FINANCE FUNCTION EXCELLENCE

Vibe Coding Is Coming to Finance. CFOs Need Guardrails

By [Matthew Harris](#), [Michael Demyttenaere](#), [Edoardo Palmisani](#), [Hardik Sheth](#), [Sam Hawes](#), [Nicolas Tanaka](#), [Aissa Boudadi](#), and [Laurin Henderson Menton](#)

ARTICLE JUNE 04, 2026 8 MIN READ

As AI use spreads across corporate functions, finance has been more deliberate than most in its adoption. That caution is a feature, not a flaw, since there are no small mistakes in finance. Financial outputs flow into audited results, requiring a level of precision, traceability, and data security that early, general-use AI tools could not guarantee.

That landscape is now changing with the emergence of what many call “vibe coding”—or more precisely, AI-driven custom application development. AI models have become more capable, the

platforms for building with them no longer require an engineering team, and finance teams have learned from experience which problems are better solved with AI rather than conventional software or ERP systems. Forecasting, anomaly detection, and document review are among the tasks AI handles best.

These developments have created a credible path for finance teams to build targeted custom applications for analysis, matching, anomaly detection, and other specific tasks without waiting for long development queues. (See Exhibit 1.)

EXHIBIT 1

Where Finance Teams Can Put Coding Agents to Work

Team	Core workflows (where agents apply)	Where coding agents add value	Pre-agent baseline	Value created
FP&A	Forecasting, variance analysis, scenario modeling	Localized (by geography or sub-business unit) forecasting and insights	Spreadsheet rebuilds, manual driver updates, fragmented scenario analysis, inaccurate forecasts	Improves forecast accuracy and insights for localized FP&A analysts or teams
General accounting (controllership)	Sub-ledger anomaly detection	Generates Python scripts on the fly to scan full transaction data for duplicates, round-dollar entries, off-hours postings, and outlier amounts	Manual sampling or static rules in ERP; analysts review only flagged items or trial balance rollups	Full-population testing instead of sampling; surfaces issues invisible at the trial balance level
Order-to-cash	Revenue leakage detection	Cross-checks contracts, billings, and shipped/delivered records to flag under-billed orders, missed price escalations, or unbilled services	Spot audits or annual revenue assurance reviews; most leakage goes undetected	Recovers margin lost to billing errors; surfaces systemic pricing/contract gaps
Treasury	Short-term cash position forecasting	Pulls live AR/AP aging, payment terms, and historical payment behavior by customer/vendor to model expected inflows/outflows day-by-day procedures, generate reports and narratives	Static spreadsheet rolled forward weekly; assumes invoices pay on terms; treasury analyst spends 1–2 days per cycle rebuilding	More accurate 13-week forecasts in minutes; reduced idle cash

Source: BCG analysis.

Note: AR/AP = accounts receivable/accounts payable; ERP = enterprise resource planning; FP&A = financial planning and analysis.

For CFOs, the opportunity is not to let finance teams build whatever they want with AI. It is to identify the narrow workflows where custom applications can reduce manual effort, improve control, and produce more reliable insight than spreadsheet-heavy processes allow.

How Coding Agents Create Value in Finance

The change matters because it makes desktop-level automation more accessible. Finance teams can now build targeted solutions faster, including FP&A forecasting tools that surface localized insights and controllership applications that flag anomalies across full transaction populations. (See Exhibit 2.) Analysts still need to define the logic, constraints, and desired outputs. But the AI agent can write the code, test it, refine it, embed AI insights, and help connect it to relevant data sources.

Higher reliability does not come from the model, but from where the work runs. Reliability comes from moving calculations into rule-based code, such as Python or SQL, that runs in controlled environments and provides the same predictable answers. Agent-driven coding makes that path far easier for finance teams that do not have deep software engineering skills in-house.

AI-built custom applications do not replace traditional finance systems and databases. Core platforms for accounting, planning, and reporting remain essential. Systems of record, such as enterprise resource planning (ERP), and systems of execution, including enterprise performance management (EPM) systems, consolidation tools, and upstream procure-to-pay and order-to-cash processes, remain critical to operations. What AI-coded applications provide is an augmentation layer. They do not replace the system of record; instead, they sit on top of it and make the data and insights more accessible, usable, and easier to interpret.

“ AI-coded applications provide an augmentation layer. They sit on top of the system of record and make the data and insights more accessible, usable, and easier to interpret.”

Understanding where that boundary falls is important. Processes that generate reported financials should remain in core systems. For instance, enterprise-level planning and forecasting should stay within EPM platforms, and proven workflow tools such as BlackLine for reconciliations should not be replaced. By contrast, processes centered on interpretation, input shaping, localized insight generation, or anomaly detection are often better candidates for augmentation. This layered architecture, where AI tools augment but do not replace core systems, creates a clear governance boundary. Core financial platforms handle reporting and compliance. AI-coded applications handle speed and insight. Applied well, this model lets finance teams move faster on analysis and exception handling without putting compliance or auditability at risk.

To be sure, this accessibility introduces a new risk: AI sprawl—the uncontrolled spread of custom AI tools across a finance department that can create complexity and management challenges. AI-coded tool sprawl can be as damaging as Excel sprawl, the quiet chaos of a thousand spreadsheets nobody owns. A CFO who allows the team to build freely without a governance framework risks trading “shadow Excel” for “shadow code”—undocumented scripts and

applications that exist outside official systems and lack proper oversight. Shadow code is far harder to see until something breaks. To scale AI safely, finance leaders must set clear guardrails around version control, data access, and auditability.

The solution is not to slow adoption but to enforce control in the runtime where apps built by coding agents operate. These apps work best when employees use them to analyze, explore, or automate their own work, not when they touch financial systems, ERPs, or customer-facing flows. Those domains belong to purpose-built enterprise and orchestration agents, which are strictly designed and continuously monitored.

Within safe boundaries, leading organizations embed governance rules directly into the system, run execution in isolated environments, and maintain active oversight. They also keep watching after the app ships, with a record of what happened and checking that behavior remains correct. Governance is built into how these apps operate, not added afterward.

Consider intercompany accounting, where even teams with established processes for eliminating transactions between related entities still spend hours each close updating entries that offset across those entities, and matching by amount and trading partner. A coding agent can turn those activities into a repeatable script that applies matching logic and surfaces an exception list ranked by materiality, so analysts focus on resolving breaks rather than finding them.

With AI agent-driven development, an analyst can build an AI tool that pulls transaction data from multiple ledgers, matches intercompany payables against receivables, and flags discrepancies in minutes instead of hours. This tool does not replace the reconciliation and booking in the core general ledger but uses agentic capabilities to accelerate exception handling.

Today, coding agents such as Claude Code and OpenAI Codex can write and execute code, work with structured data, and build small applications from natural language instructions. Innovation in this space is quickly expanding the breadth of capabilities that agent-driven coding platforms provide. This allows both finance analysts and IT departments to accelerate development cycles.

Recognizing Limits

The pace of innovation, however, does not eliminate the conditions that determine whether these tools create value. Sophisticated output does not fix flawed input and actually poses a risk because the polish makes bad data more convincing. AI-coded tools are no exception. Coding agents also do not replace judgment. CFOs still need employees who can interpret results, apply policy, and make decisions in ambiguous situations.

LL Sophisticated output does not fix flawed input and actually poses a risk because the polish makes bad data more convincing.

Coding agents also do not reduce governance risk on their own. Without proper controls, the spread of desk-level applications can introduce new risks related to access, version control, and auditability. These governance questions are not secondary concerns, but the requirements for moving from isolated experiments to enterprise-wide deployment.

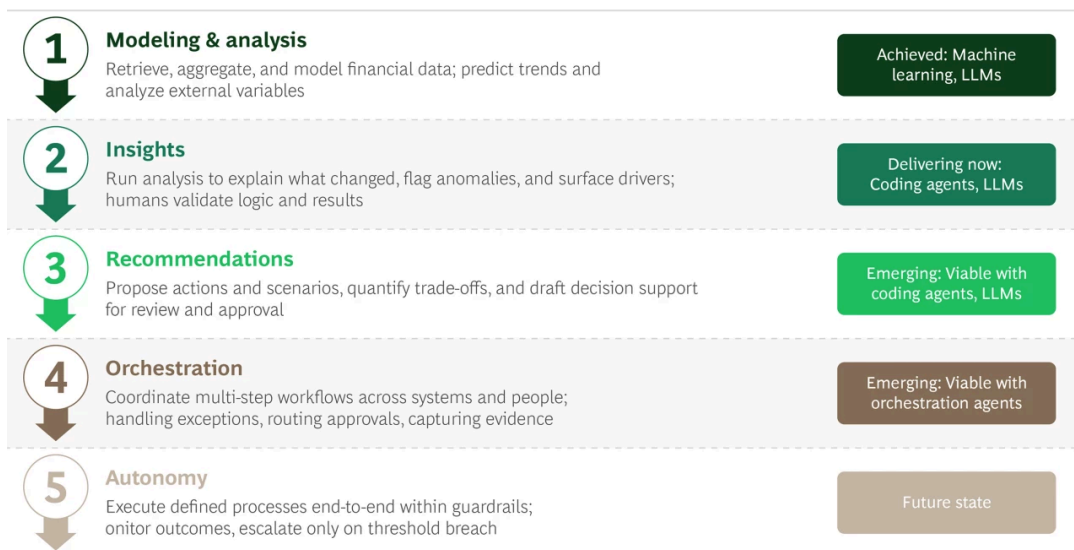
Toward Orchestration

The next phase of value creation lies in moving from point solutions to end-to-end process orchestration, alongside a more deliberate rationalization of how and where agents are used. (See Exhibit 2.) Many finance workflows follow predictable sequences but remain labor-intensive due to handoffs, approvals, and exception handling.

EXHIBIT 2

AI in Finance: From Reporting to Autonomy

Five levels of capability—Coding agents are accelerating the climb¹



Source: BCG analysis.

¹Coding agents ≠ enterprise replacement; they fill the desk-level gap.

While orchestration can streamline these workflows, scaling it requires clear decisions about when to rely on agentic capabilities versus traditional automation means. High-volume data transformations and reconciliations remain far more efficient in traditional pipelines such as SQL (Structured Query Language) and ETL (Extract, Transform, Load), with agents applied selectively for exception handling, reasoning, and coordination. The result is a hybrid model that balances automation, cost efficiency, and control.

An orchestration layer can automate multistep processes across systems. In cash application, for example, an agent can match remittance advices against open receivables, route exceptions to the appropriate analyst with supporting detail attached, and track resolution through close. AI agent-led coding is well suited to building the individual components of the workflow. Orchestration connects them into a cohesive, automated process. Orchestration can also “call up” and leverage the custom-coded tools developed in the prior use case.

A new generation of agentic frameworks is beginning to move beyond single-task automation toward more autonomous, outcome-oriented operations. Rather than executing predefined scripts, these agents can navigate applications, coordinate across systems, and adapt dynamically as conditions change, linking multiple steps into end-to-end workflows.

Over time, this enables loosely coupled networks of agents that work together toward defined objectives, such as completing a close process or resolving exceptions, with workers setting goals and guardrails rather than directing each step. Although still in its early stages, this shift nonetheless signals a move to more autonomous operations where agents pursue clear goals with greater independence.

What Finance Leaders Should Do Next

The opportunity is real, but value depends on disciplined use. Finance teams should start by building their analysts’ capabilities so they can frame the desktop-level problems best suited for AI-driven custom applications. Collaboration with IT, security, and risk functions should begin early, not after tool use has spread. Without visibility from these groups, finance teams risk building tools that create compliance gaps or security vulnerabilities. Leaders should also choose initial use cases carefully, favoring solutions where risk is manageable and the boundary between augmentation and system-of-record responsibilities is clear.

AI coding agents give finance teams a credible way to build targeted applications faster than before. Done right, AI-driven development gives finance teams a faster path to the applications

they have always needed and rarely had the resources to build. The leaders who empower their teams with these capabilities will achieve better productivity and faster insights, while maintaining financial integrity.

About the Center for CFO Excellence

The Center for CFO Excellence (CFOx) is the hub of BCG's know-how and global expertise on the CFO domain. CFOx supports CFOs, helping them deal effectively with today's most pressing finance-related challenges. Our experts work with CFOs worldwide to support their efforts in transforming the finance function's operating model to achieve higher efficiency and effectiveness. They achieve this by capitalizing on the digital disruption to unleash the power of advanced analytics and artificial intelligence, also advancing corporate performance management to the next level and enhancing the value of the finance function. They focus on improving capital allocation to ensure a productive return on investments and to outperform competitors. Aiming to orchestrate company-wide transformations, BCG helps drive the corporate value agenda, identifying the best strategy for renewing the organization's finance systems and increasing transparency and insight. If you would like to learn more about CFOx, please contact one of the BCG authors or visit cfox.bcg.com.

Authors



Matthew Harris

Associate Director
Austin



Michael Demyttenaere

Managing Director & Partner;
CFO Excellence NAMR Leader
Los Angeles



Edoardo Palmisani

Managing Director & Senior
Partner; Global Lead for CFO
Excellence
Rome



Hardik Sheth

Partner & Director, Center for
CFO Excellence
Philadelphia



Sam Hawes

Partner & Associate Director,
Artificial Intelligence
San Francisco - Bay Area



Nicolas Tanaka

Partner
San Francisco - Bay Area



Aissa Boudadi

Partner & Associate Director
Miami



Laurin Henderson Menton

Managing Director & Partner
Brooklyn



ABOUT 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.

© Boston Consulting Group 2026. All rights reserved.

For information or permission to reprint, please contact BCG at permissions@bcg.com. To find the latest BCG content and register to receive e-alerts on this topic or others, please visit bcg.com. Follow Boston Consulting Group on [Facebook](#) and [X \(formerly Twitter\)](#).