Deloitte.

Together makes progress

Taming the beast

What Financial Controllers need to know about AI

The creature stirs

There's a lot we can learn from Mary Shelley.

In the summer of 1818, a 19-year-old author published a story about a brilliant scientist who created something extraordinary—and then lost control of it. Victor Frankenstein didn't fail because his creation was inherently evil. He failed because he had no plan for what came next. No governance. No oversight. No guardrails. He built the thing, stood back in awe, and then ran away when it started moving.

Sound familiar?

Controllers and Chief Accounting Officers aren't in the business of reanimating corpses, but many are watching a different kind of creature come to life inside their organizations. Artificial intelligence (AI), in its many forms (machine learning, Generative AI, agentic systems), is moving into finance functions faster than many anticipated. According to Deloitte's State of AI in the Enterprise survey, 74% of surveyed leaders expect to deploy agentic AI within two years. However, only 21% of this same group report having a mature model for governance of autonomous agents.[1] Adoption is outpacing the development of the safeguards needed to manage it responsibly.

The technology is here. It's powerful. And for many controllership functions, it's already in motion.

The question isn't whether AI will transform financial reporting and controllership. The question is whether you'll be the one steering it—or the one chasing it down the hallway.

This article offers a practical framework—seven themes—for controllers who want to harness AI's potential while continuing to do what they've always done: safeguard the assets, report accurately, and maximize the value of the controllership function to the business.

Because unlike Victor Frankenstein, controllers don't have the luxury of running away. The 10-K still needs to be filed.

1. Set your objectives (before someone sets them for you)

One of the challenges facing controllership today is that AI adoption often begins with the technology rather than the problem. Enthusiasm at the board and C-suite level—driven by the pace of innovation and fear of falling behind—can translate into sweeping mandates to deploy AI before specific objectives have been defined.

In this environment, controllership teams can find themselves searching for a use case to justify tools, rather than identifying tools that address a genuine need.

Controllers should resist this approach and instead be clear-eyed about what they want AI to do for the function. That answer is broader than many professionals might expect.

When executives hear "AI in finance," the default assumption is efficiency—often interpreted as cost reduction or headcount savings. Efficiency is real, but a more complete view of AI value in controllership spans across three distinct categories:

Risk reduction and increased capability

AI can process massive data sets in ways that human reviewers simply cannot. Its ability to analyze entire populations of data, rather than samples, is a sea change in assurance.

Consider anomaly detection across journal entries, real-time monitoring of transactions for outliers, and pattern recognition in revenue data—these aren't theoretical applications.

Take a Sarbanes-Oxley (SOX) control that has historically relied on manual sampling. With the introduction of agentic AI—autonomous systems capable of ingesting documents, extracting data, applying business rules, and reporting outputs—the model can shift from limited sampling to full-population testing.

Moving from partial coverage to 100% coverage is not about doing the same work faster. It is about doing something fundamentally different and better. An AI system that flags unusual entries across every transaction provides a level of visibility that sampling approaches, by definition, cannot achieve. It is a completely different control paradigm.

For controllers responsible for Sarbanes-Oxley compliance, the risk reduction and capability expansion can enhance the confidence and transparency in financial reporting and can be a compelling business case.

Efficiency and speed

Efficiency remains an important part of the AI value proposition—but it should be framed as more than simple cost reduction. Yes, AI can automate repetitive tasks and reduce manual effort. But an equally compelling version of efficiency is accelerating the close. By automating flux analysis, generating initial variance explanations, handling reconciliations, and organizing close activities, AI can enable controllers to compress the timeline of a quarterly close that traditionally dominates the function. That's time returned to the business for analysis, insight, and judgment—activities that create value rather than just protect it.

Productivity and strategic uplift

Deloitte's recent *State of AI in the Enterprise*[2] survey demonstrated that efficiency and productivity gains are the most widely realized benefit across finance functions, with two-thirds of surveyed organizations reporting they have already achieved expected benefits in efficiency and output. But perhaps more telling is where expectations are headed.

Organizations are increasingly hoping for similar levels of success in AI in strategic and value-added areas such as enhanced decision-making, data-driven insights, and increasing revenue.

The vision is a controllership function that spends less time on historical reporting and more time providing forward-looking analysis, scenario modeling, and strategic insight to the business. AI can help make that shift possible at scale—freeing your people from the mechanics of the close and creating space for finance professionals to operate as genuine business partners.

The starting point for any AI initiative is to define clear objectives before committing to specific tools or platforms. What problem is being solved? What does measurable success look like in 12 months?

Without this clarity at the outset, AI initiatives risk becoming expensive and ineffective experiments—delivering tools but failing to deliver tangible value to the business.

2. Adopt a governance framework

If setting objectives is about ambition, governance is about discipline. And discipline is what separates a well-run AI program from a well-intentioned one.

The governance landscape for AI in financial reporting is still evolving and it's still remarkably unsettled. Standard-setters and regulators in accounting and auditing have not yet issued a mandated AI governance framework for the financial reporting processes. But "unsettled" is not the same as "unimportant." The absence of formal regulation should not be mistaken for an absence of importance. If anything, it heightens it.

Controllers, accountable for the integrity of the financial reporting process, need to ensure the governance framework provides clear answers to a series of fundamental questions: How is the AI model governed? What controls exist over the inputs and outputs? How do you know it's working correctly? Where is the supporting documentation?

These are not hypothetical questions. They are the questions that any well-designed governance framework should address—and the absence of a regulatory mandate makes it more important for controllers to build that framework proactively.

The good news is that you don't have to build from scratch.

Several credible frameworks already exist. The most prominent is the NIST AI Risk Management Framework (RMF), published by the National Institute of Standards and Technology.[3] The NIST framework is structured around four core functions:

- **Govern:** Establish organizational structures, policies, and accountability for AI risk management.
- **Map:** Identify and contextualize AI risks, including those related to bias, security, privacy, and reliability.
- **Measure**: Assess and track identified risks using quantitative and qualitative methods.
- **Manage:** Prioritize and act on risks based on their projected impact.

While the NIST AI RMF was designed for enterprise-wide AI governance—not specifically for accounting and financial reporting—it provides structured principles that controllers can interpret and apply to the unique requirements of their function.

Other relevant frameworks include COSO's *Realize the full potential of AI*[4] guidance, which applies the Enterprise Risk Management Framework specifically to AI, and ISO/IEC 42001, which provides requirements for an Artificial Intelligence Management System.

The key risks that controllers should consider when applying the framework include inaccuracies and hallucinations in AI-generated outputs, intellectual property and confidentiality concerns, bias in algorithmic decision-making, cybersecurity vulnerabilities, and a lack of transparency or explainability.

Importantly, AI governance isn't a parallel universe. AI governance should be embedded into the control environment you already have—your enterprise risk management processes, internal controls framework, and technology governance.

It is also worth noting that external auditors are developing their own methodologies to assess AI within financial reporting. The Center for Audit Quality (CAQ) has published considerations for how auditors should evaluate a company's use of Generative AI in internal control over financial reporting, with a focus on the level of human involvement, the processes for reviewing AI outputs, and whether AI controls are supported by appropriate entity-level and IT general controls.[5] Controllers who proactively address these areas can position themselves ahead of the audit conversation, rather than reacting to it after the fact.

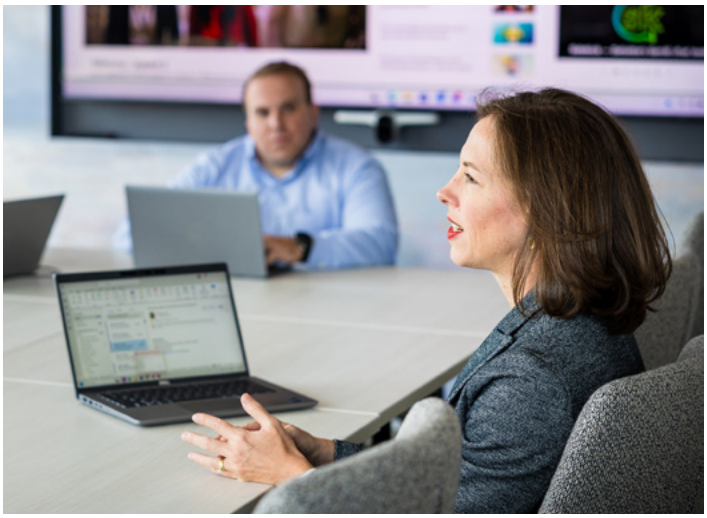

3. Test, validate, and keep testing

AI models behave differently than traditional systems. You don't install them, run a user acceptance testing cycle, and move on. They are not static. They learn from data, they evolve, and their performance can degrade over time as the underlying data changes—a phenomenon known as "model drift."

This means that **testing AI isn't a one-time gate. It is an ongoing obligation.**

For AI models to be effective in financial reporting processes, the testing challenge has a specific wrinkle: You may need to test with production data. Synthetic or sanitized data sets can be useful for initial development, but if the AI is meant to detect anomalies in your actual journal entries or generate variance commentary on your real P&L, testing with dummy data only tells so much. The model needs to see the patterns, the quirks, and the noise that exist in the actual financial environment.

This introduces a critical tension. Using **production data** improves model accuracy but raises questions about data security, access controls, and privacy. But using **only synthetic data** raises questions about whether the model will actually work in practice.

Controllers need to think carefully about this trade-off and work closely with IT, legal, and their auditors to find an approach that balances rigor with reality. Beyond initial testing, ongoing monitoring is essential.

The IIA's AI Auditing Framework offers useful guidance.[6] While written primarily for internal auditors, its recommendations apply broadly and highlight several leading practices such as establishing testing that validates AI both before and after deployment, defining success metrics and key performance indicators, reporting regularly to leadership and the board, and integrating AI into the enterprise risk management process. In practice, this means treating AI models as living components of the control environment—subject to performance benchmarks, continuous evaluation, and refinement.

4. Get your data right (it doesn't have to be perfect)

If there's one theme that appears in virtually every paper, survey, and framework on AI in finance, it's this: Data quality matters. The AI is only as good as the data it's trained on.

And it's true. But here's the part that doesn't get said enough: *Waiting for perfect data is a recipe for never starting.*

Data in most large enterprises is messy. It's fragmented across multiple systems, inconsistently formatted, governed unevenly, and not always as accurate as we'd like it to be. If you wait for data perfection before deploying AI, you'll be waiting a very long time. A pragmatic approach is to start with a data readiness assessment.

Identify where data gaps, latency, or inaccuracies exist, and then prioritize improvements based on their impact on your AI use cases.

Not all data sources carry equal weight. Financial data that directly supports financial reporting and controls should take priority; secondary sources used for commentary or analysis can tolerate a wider margin. Close those critical gaps through technology enablement, governance, and centralized data management.

Several practical steps can help improve AI readiness:

- Understand the data lineage for any AI application in your finance processes;
- Know where the data comes from, how it's transformed, and what controls exist over its integrity;
- Implement data versioning so you can trace outputs back to the specific data set that produced them; and
- Establish ownership and involvement of your data governance team early—not as an afterthought.

The bottom line: **Your data doesn't need to be perfect to get started. Know where the imperfections are, prioritize which ones matter most, and have a plan to manage informed control over data limitations.**

5. Keep a human in the loop (but define what the human actually does)

"Human in the loop" has become a defining mantra of responsible AI. It's the reassuring phrase that likely appears in every governance document, every framework, every vendor pitch: *Don't worry—there's a human in the loop.*

But what does human oversight actually mean?

The Center for Audit Quality breaks human-in-the-loop involvement into three components:

- Reviewing the accuracy and completeness of inputs;
- Understanding the explainability and interpretability of the AI's outputs; and
- Reviewing those outputs for quality, reliability, and appropriateness.

That's a useful framework, and it aligns with how controllers already think about internal controls—verifying inputs, understanding the process, and reviewing outputs.

But the presence of a human does not automatically equate to effective control.

The problem arises when human in the loop becomes a rubber stamp. If a finance professional is asked to review output from an AI system that processes thousands of transactions in minutes, what exactly is the review? Scrolling through pages of AI-generated commentary and clicking "approve" is not a control.

Meaningful and effective human involvement requires specificity. It means defining specifically what the human is expected to evaluate, what level of expertise is required, and who has the authority to override the AI's conclusions. As organizations begin to adopt agentic AI—systems that can initiate actions rather than simply analyze data—the need for oversight increases. An AI agent that extracts contract terms, drafts process narratives, and updates risk matrices is qualitatively different from a chatbot that answers questions. The level of human oversight needs to match the level of autonomy.

Deloitte's research on trust in agentic AI makes the point: Adopters should begin with low-stakes functions, experience the results for a time, and then expand to higher-stakes decisions and actions as trust builds through consistency[7]

For controllers, a practical question is: *Where in the process does the human make a decision, and what decision are they making?*

If you can't answer that with specificity, the human-in-the-loop control is a checkbox, not a safeguard.

Where in the process does the ***human*** make a decision, and what decision are they making?

6. Measure what matters

Defining objectives was the first step. Now—did you meet them?

It's a deceptively simple question, yet for many organizations, measuring the impact and value of AI remains one of the most common challenges. However, demonstrating value is what sustains momentum. Measurement should be concrete and tied directly to the objectives defined at the outset. If the goal was to accelerate the close, track the days needed to close and timeliness of reporting. If it is risk reduction, track anomalies detected, false positive rates, and resolution times. If it is increased capability, track coverage rates and previously undetected exceptions.

Return on investment deserves the same rigor. Many prior attempts at automating controllership tasks failed because the investment was disproportionate to the benefit. Generative AI and agentic systems can overcome some of these barriers through their ability to adapt to changing documentation and data—but the economics still need to be tracked. **Technology that costs more than it saves is not a success, regardless of how innovative it is.**

7. Invest in your people, not just your technology

Perhaps the most counterintuitive but biggest risk to your AI program is not the technology. It's the people.

Not because they're not capable. They just haven't been adequately prepared. Organizations are investing heavily in AI technology—and that figure is rising rapidly. But investment in the human side of the equation is often lagging behind. The Institute of Management Accountants (IMA®) has highlighted this imbalance, noting that across organizations broadly, the staff who perform the day-to-day work receive a disproportionately small share of AI training investment relative to the complexity of the tools they are expected to use.[8]

This creates a dangerous asymmetry. The technology is outrunning the talent. Organizations are deploying sophisticated AI systems while asking people who haven't been trained to serve as the human in the loop. They're implementing Generative AI in financial reporting processes while their team's understanding of how Generative AI models work—and how they can fail—remains surface-level.

Controllers need to make the case that AI investment without people investment is a recipe for failure. Addressing this requires a deliberate approach to upskilling. Training cannot be limited to a one-off awareness session. It needs to be structured, continuous, and role-specific. The IMA recommends distinct strategies for accounting staff, technical professionals, and senior leadership—each requiring a different depth of AI literacy. Finance professionals need to understand how AI models are trained, where hallucinations and biases can arise, and what constitutes meaningful oversight, including hands-on experience with the tools they're expected to supervise. Equally important is ensuring leadership has sufficient literacy to make informed decisions about AI strategy and risk.

There's a broader talent landscape that also adds urgency. The CPA and accounting professional pipeline is shrinking. The number of new candidates for the CPA exam fell from a peak of almost 50,000 in 2010 to a 17-year low of approximately 28,000 in 2024.[9] Controllership functions are competing for talent in a constrained market. AI can help address that talent gap by automating lower-value tasks, but only if the remaining team members have the skills to work effectively alongside the technology.

The organizations that get this right will be the ones that treat AI not as a replacement for people, but as a catalyst for elevating human capability. That requires investment: in training, in career development, in making controllership roles attractive to a generation that expects to work *with* cutting-edge technology, not around it.

Controllers need to make the case that AI investment without people investment is a recipe for failure.

Don’t be Frankenstein

Mary Shelley imagined a scientist who could create life but had no idea what to do next. Controllers don’t have to make that mistake. AI is powerful, it’s real, and it’s coming to your function whether you’re ready or not. But unlike Frankenstein’s creature, it isn’t inherently dangerous. It’s a tool—one that can help accelerate your close, strengthen your controls, expand your capabilities, and free your people to do their best work.

The key is intentionality. Define what you want AI to achieve. Put governance in place before you scale. Test rigorously and continuously. Get your data into the best shape you can instead of waiting for perfection. Design human oversight that’s meaningful, not cosmetic. Measure results against the objectives you set. And invest in your people at least as aggressively as you invest in the technology.

The controllership function has always been about stewardship—safeguarding assets, ensuring accuracy, managing risk. AI doesn’t change the mission. It amplifies it. The controllers who thrive in the years ahead will likely be the ones who tamed the beast rather than ran from it.

Mary Shelley wrote a horror story that didn't end well. How will you write your AI story—before it writes itself?

About the *Center for Controllership*™

CfC—a research, resource, and collaboration center—provides chief accounting officers (CAOs), controllers, and their teams direct access to valuable insights that can help organizations and leaders become more proficient, effective, and insightful— three keys to achieving world-class controllership.

Authors

Charles Carrington
Partner
Deloitte & Touche LLP
chcarrington@deloitte.com

Katie Glynn
Partner
Deloitte & Touche LLP
kaglynn@deloitte.com

Matthew Hurley
Partner
Deloitte & Touche LLP
mhurley@deloitte.com

Endnotes

1. Deloitte, State of AI in the Enterprise: The untapped edge, January 2026. Survey of 3,235 director- to C-suite-level respondents 24 countries.
2. Ibid.
3. National Institute of Standards and Technology (NIST), *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, January 2023.
4. Keri Calagna, Brian Cassidy, and Amy Park, *Realize the full potential of artificial intelligence*, Committee of Sponsoring Organizations of the Treadway Commission (COSO) and Deloitte, September 2021.
5. Center for Audit Quality (CAQ), *Auditing in the age of Generative AI*, April 2024.
6. The Institute of Internal Auditors (IIA), *Artificial intelligence auditing framework*, 2024.
7. Deloitte, "Trust emerges as main barrier to agentic AI adoption in finance and accounting, despite optimism around the tech," press release, July 19, 2025.
8. Kristine M. Brands and Mark Holtzblatt, "AI best practices for management accountants," *Strategic Finance*™, May 13, 2024.
9. American Institute of Certified Public Accountants (AICPA), *2025 Trends: A report on accounting education, the CPA exam and public accounting firms' hiring of recent graduates*, 2025.

Deloitte.