PART 15

The AI Governance, Risk, and Control Framework

AI in finance is the most consequential control environment innovation since the Sarbanes-Oxley Act. The CFO who does not lead the governance response will face the audit, regulatory, and reputational consequences personally.

The Finance Technology Stack
A CFO's Strategic and Architectural Guide

Why AI Governance Is a CFO Accountability

The artificial intelligence governance question in finance has been treated by many enterprises as a question for the chief information officer, the chief risk officer, the chief information security officer, or the emerging chief artificial intelligence officer where that role exists. The treatment reflects the historical pattern in which technology governance was the responsibility of the information technology function and the business function consumed the technology rather than governing it. The pattern was always incomplete. In the artificial intelligence era, the pattern is structurally inadequate for the finance function. The chief financial officer holds accountabilities that the artificial intelligence governance question intersects with directly, and the chief financial officer who delegates the governance to others is delegating accountability that cannot in fact be delegated.

The accountabilities the chief financial officer holds and that artificial intelligence governance affects are several. The accountability for the integrity of the financial statements, which artificial intelligence in the close and consolidation process intersects with directly. The accountability for the internal controls over financial reporting, which artificial intelligence in any process touching the general ledger intersects with directly. The accountability for the audit relationship, which artificial intelligence in the finance function will increasingly become a topic the external auditor engages with substantively. The accountability for the disclosures that public companies make and that the artificial intelligence increasingly affects through both its operational impact and the increasing regulatory requirements for artificial intelligence disclosures. The accountability for the regulatory compliance posture of the finance function across the regulatory regimes that increasingly address artificial intelligence specifically.

These accountabilities mean that the chief financial officer who allows the artificial intelligence governance to be defined elsewhere will be accountable for the consequences without having shaped the framework that produced them. When the audit identifies a control deficiency related to artificial intelligence, the board will ask the chief financial officer. When the regulatory investigation examines the artificial intelligence in the finance function, the chief financial officer will be the executive who answers. When the disclosure controls fail because the artificial intelligence handling of financial data did not satisfy the requirements, the chief financial officer will sign the restatement. The accountability follows the role regardless of how the governance was constructed. The chief financial officer who recognizes this reality will engage with the governance question deliberately. The chief financial officer who does not will discover the accountability under conditions that the engagement could have prevented.

The position taken in this masterclass is that artificial intelligence governance in finance is the most consequential control environment innovation since the Sarbanes-Oxley Act of 2002. The comparison is deliberate. Sarbanes-Oxley imposed structural requirements on the internal

controls over financial reporting that reshaped how the finance function operated. Artificial intelligence in finance is imposing structural requirements on how the function reasons about decision quality, evidence, auditability, and accountability. The structural requirements are not optional. They are being defined by the audit profession, by the regulatory community, by the standard-setting bodies, and increasingly by case law as the specific failures and the specific successes shape the field. The chief financial officer who treats artificial intelligence governance as a defensive overhead will produce a function that satisfies the minimum requirements at the maximum cost. The chief financial officer who treats artificial intelligence governance as a strategic discipline will produce a function that captures the artificial intelligence value while satisfying the requirements at reasonable cost.

The work of this part is to lay out the artificial intelligence governance framework with the seriousness the topic deserves. The part treats the hallucination problem and how to architect against it, the model risk management discipline adapted for the foundation model era, the NIST framework as the starting point, the Sarbanes-Oxley implications, the audit profession's evolving stance, and the operating disciplines that translate the framework into operational reality. The reader will leave with the framing to lead the artificial intelligence governance agenda from a position of substantive engagement rather than from a position of defensive reaction.

The Hallucination Problem and How to Architect Against It

The hallucination problem is the most consequential technical limitation of the current generation of artificial intelligence systems in finance and is the problem that the governance framework must address most directly. Hallucination is the term the field has settled on for the phenomenon in which an artificial intelligence system produces an output that appears authoritative and that is factually wrong. The output looks plausible. The output uses the vocabulary of the domain correctly. The output may even cite sources or provide reasoning. The output is incorrect in ways that may not be visible without independent verification. The phenomenon is not occasional. It is structural in the current generation of large language models that power most of the artificial intelligence capabilities being deployed in finance.

Why Hallucination Matters in Finance

Hallucination matters in finance more than in many other domains because the consequences of wrong outputs are more severe. A wrong output that becomes the basis for a financial reporting decision can produce a misstatement that requires restatement, that triggers audit issues, that affects investor decisions, and that exposes the chief financial officer and the chief executive officer to personal liability under the Sarbanes-Oxley certifications. A wrong output that becomes the basis for an operational decision can produce financial loss, customer impact, or vendor relationship damage. A wrong output that becomes the basis for a strategic decision can produce consequences that compound over years. The asymmetry between the cost of a wrong output and the cost of catching the wrong output before it produces consequences favors substantial investment in the architectures that catch the wrong outputs.

The hallucination phenomenon is particularly difficult to address because the wrong outputs do not announce themselves. A traditional software system that produces a wrong output usually does so through a recognizable failure mode that the operations team can diagnose. The artificial intelligence system that hallucinates produces an output that is indistinguishable from a correct output without independent verification. The detection of hallucination requires architectural commitment to verification rather than detection through traditional monitoring.

Why Hallucination Persists in the Current Generation

The hallucination phenomenon is not a defect that the vendors will fix in the next release. It is a structural property of the way large language models work. The models produce outputs by predicting the most likely next sequence of tokens given the context and the model's training. The prediction is statistical rather than logical. The model does not know what is true in the way that

a human knows what is true. The model knows what is likely to follow the context based on its training. When the likely output happens to be true, the model produces a correct response. When the likely output happens to be false, the model produces a hallucination. The model itself cannot distinguish between the two.

The field has invested substantially in techniques that reduce hallucination, including retrieval augmented generation that grounds the model's outputs in retrieved documents, fine tuning that specializes the model for specific domains, prompt engineering that constrains the model's outputs, and various forms of self-checking that have the model evaluate its own outputs. The techniques have reduced hallucination substantially. They have not eliminated it. The current generation of models will continue to hallucinate for the foreseeable future. The governance framework must accommodate this reality rather than depending on the elimination of the problem.

The Architectural Response to Hallucination

The architectural response to hallucination involves several layers of defense. The first layer is the architectural choice to use artificial intelligence in contexts where verification is feasible rather than in contexts where verification is impossible. An artificial intelligence that supports a human reviewing the output is fundamentally different from an artificial intelligence that operates autonomously without human verification. The first pattern accommodates hallucination by having the human verify the output. The second pattern requires architectural mechanisms that produce the verification in another way.

The second layer is the architectural choice of retrieval augmented generation over pure generation for any use case that involves factual claims. Retrieval augmented generation grounds the model's outputs in retrieved documents that the model is instructed to use as the source of truth. The architecture reduces hallucination substantially because the model's outputs are constrained by the retrieved content rather than relying on the model's training. The architecture also produces an audit trail in which the outputs can be traced back to the source documents. The architecture is not perfect. The model can still hallucinate by misinterpreting the retrieved content or by extrapolating beyond it. The architecture is substantially better than pure generation for factual use cases.

The third layer is the architectural commitment to verification mechanisms that operate independently of the artificial intelligence. For high-stakes use cases, the architectural pattern includes secondary verification through a different artificial intelligence, through structured rule checks, through human review of a sample of outputs, or through reconciliation against an independent calculation. The verification mechanisms catch the hallucinations that the primary

artificial intelligence produces. The verification architecture is more expensive than the primary architecture alone but is justified by the consequences of undetected hallucinations.

The fourth layer is the architectural commitment to monitoring that detects degradation in artificial intelligence performance over time. The artificial intelligence systems can degrade as the input distribution changes, as the model is updated, or as the underlying integrations evolve. The monitoring detects the degradation and triggers remediation before the degradation produces material consequences. The monitoring requires investment in the metrics that measure performance, the baselines that establish what good performance looks like, and the operational disciplines that respond to degradation when it occurs.

The chief financial officer who architects the artificial intelligence deployment with these four layers of defense will produce an architecture that accommodates the hallucination reality. The chief financial officer who deploys artificial intelligence without the architectural defense will produce a deployment that exposes the function to hallucination consequences that the architecture should have caught.

Model Risk Management for the Foundation Model Era

The discipline of model risk management has existed in the financial services industry for decades. The discipline emerged in banking to address the risks created by the quantitative models that banks use for credit decisions, market risk management, capital planning, and the broader analytical work that supports the banking function. The discipline produced a structured approach to model development, validation, monitoring, and governance that the banking regulators codified through guidance including the United States Federal Reserve's SR 11-7 and the broader Basel framework. The discipline is now being extended to the artificial intelligence models that the broader enterprise community is deploying. The chief financial officer of any enterprise, banking or otherwise, must understand the discipline and adapt it for the artificial intelligence use cases.

The Traditional Model Risk Management Framework

The traditional model risk management framework rests on several foundational disciplines. The model inventory is the canonical record of the models the enterprise operates, with the metadata required to understand each model and to govern it. The model development standards specify the disciplined approach to building models including the documentation requirements, the testing requirements, and the approval requirements before models enter production. The model validation discipline is the independent review of models to confirm that they perform as intended and that the documentation accurately represents the model's behavior. The model monitoring discipline is the ongoing examination of models in production to detect degradation. The model change management discipline is the controlled handling of changes to models that may affect their behavior. The framework is well-established in banking and provides the starting point for the artificial intelligence governance framework.

How Foundation Models Disrupt the Traditional Framework

The foundation models that power the current generation of artificial intelligence in finance disrupt the traditional model risk management framework in several specific ways. The traditional framework assumes that the model is built by the enterprise and that the enterprise understands what the model does. The foundation models are built by external providers, with the enterprise consuming the models without full visibility into their training data, their architecture, or their behavior. The traditional framework assumes that the model is deterministic in the sense that the same inputs produce the same outputs. The foundation models are typically non-deterministic, with the same prompt producing different outputs across invocations. The

traditional framework assumes that the model's behavior can be characterized through testing and documentation. The foundation models have behavior that is too broad to fully characterize and that may exhibit emergent capabilities or failure modes that the testing did not surface.

The disruptions do not invalidate the traditional framework. They require adaptation of the framework to the foundation model reality. The model inventory must extend to capture the foundation models and the application-specific configurations and prompts that adapt the foundation models to the specific finance use cases. The model development standards must address the prompt engineering, the retrieval augmented generation configurations, and the integration patterns that operate alongside the foundation models. The model validation discipline must include techniques that work for non-deterministic models including statistical testing across multiple invocations, adversarial testing that probes for failure modes, and structured testing of the specific use cases the model supports. The model monitoring discipline must extend to the metrics that matter for artificial intelligence including the hallucination rate, the output quality, the response time, and the consumption costs. The model change management discipline must address the version updates that the foundation model providers deploy, which may change the model's behavior without the enterprise's request.

The Adapted Framework for Finance

The adapted model risk management framework for finance includes several specific components. The artificial intelligence model inventory captures every artificial intelligence model deployed in the finance function, with the metadata required to understand and govern each model. The metadata includes the foundation model used, the specific configuration including prompts and retrieval augmentation, the use case the model supports, the data the model has access to, the decision authority the model has, the human oversight model, and the governance accountability for the model.

The pre-deployment validation discipline applies to every model before it enters production use. The discipline includes the structured testing of the model's outputs against the use case requirements, the adversarial testing to probe for failure modes, the documentation of the model's known limitations, and the approval by the appropriate governance body before deployment. The discipline is calibrated to the criticality of the use case, with the most stringent validation reserved for use cases that affect the financial statements or that operate with substantial autonomy.

The in-flight monitoring discipline operates continuously after deployment. The discipline includes the metrics that measure the model's performance, the baselines that define acceptable performance, the alerting that surfaces performance degradation, and the operational response that addresses the degradation when it occurs. The discipline is supported by the platform

capabilities that the artificial intelligence providers are increasingly delivering and by the third-party monitoring platforms that have emerged to support the discipline.

The output review discipline applies to the artificial intelligence outputs that affect financial reporting or other high-stakes decisions. The discipline requires that the outputs be reviewed by humans before they produce consequences. The review may be of every output for the highest-stakes use cases, of a sample of outputs for high-stakes use cases, or of escalation cases for routine use cases. The discipline is calibrated to the criticality of the use case and to the reliability the model has demonstrated through monitoring.

The audit trail discipline captures the inputs, the outputs, and the reasoning that the artificial intelligence produced for every operationally significant interaction. The audit trail supports the external audit, the internal investigation of any specific issues, the regulatory inquiry that may arise, and the broader governance review. The audit trail is non-negotiable for artificial intelligence in financial reporting and is increasingly expected for artificial intelligence in any operationally significant use case.

The rollback capability allows the artificial intelligence to be disabled or reverted to a previous configuration when problems are detected. The capability is essential because the artificial intelligence will produce problems despite the other layers of defense. The capability requires architectural commitment that the deployment patterns must accommodate. The chief financial officer who insists on the rollback capability before deployment will operate an artificial intelligence portfolio that can be managed when problems occur. The chief financial officer who deploys without the rollback capability will discover during the first significant problem that the deployment commitment exceeds the operational flexibility.

The NIST AI Risk Management Framework Adapted for Finance

The National Institute of Standards and Technology issued the AI Risk Management Framework in early 2023 as a voluntary framework for organizations deploying artificial intelligence. The framework has become the de facto standard reference in the United States and has been widely adopted internationally. The framework provides a structured approach to artificial intelligence risk management that the chief financial officer can use as the starting point for the governance framework specific to the finance function. The adaptation for finance involves several specific extensions and emphases beyond the general framework.

The Four Functions of the NIST Framework

The NIST framework is organized around four functions: Govern, Map, Measure, and Manage. The Govern function establishes the organizational structures and policies that support artificial intelligence risk management. The Map function identifies and characterizes the artificial intelligence systems the organization deploys and the contexts in which they operate. The Measure function assesses the risks and impacts of the artificial intelligence systems through testing, monitoring, and evaluation. The Manage function addresses the identified risks through the controls, the operational disciplines, and the decisions that govern the deployment and operation of the artificial intelligence systems.

The four functions operate continuously and iteratively rather than as a one-time exercise. The Govern function establishes the framework. The Map function maintains the inventory. The Measure function produces the assessments. The Manage function addresses the risks. The findings from the Manage function inform the Map and Measure functions as the artificial intelligence portfolio evolves. The Govern function adjusts as the experience develops. The continuous operation is essential to the framework's effectiveness and is one of the dimensions that distinguishes the NIST framework from one-time compliance exercises.

The Adaptation for Finance

The adaptation of the NIST framework for the finance function emphasizes several specific extensions. The Govern function must include the chief financial officer as the executive accountable for the artificial intelligence in the finance function, with the supporting governance structure that includes the controller, the chief audit executive, the chief information security officer, and the external relationships including the external auditor and the regulatory community where applicable. The governance framework must specifically address the Sarbanes-

Oxley implications, the audit committee engagement, and the disclosure controls that public companies must operate.

The Map function must include the specific artificial intelligence systems in the finance function with the metadata required for finance governance, including the financial materiality of the decisions the artificial intelligence affects, the audit relevance of the artificial intelligence outputs, the regulatory implications of the artificial intelligence in specific use cases, and the broader operational implications. The mapping must be specific enough to support the risk assessment that the Measure function will conduct.

The Measure function must include the specific risk categories that matter for finance including the financial reporting accuracy risk, the audit trail integrity risk, the fraud risk that artificial intelligence may create or exacerbate, the operational continuity risk if the artificial intelligence fails, and the broader risks that the artificial intelligence introduces. The measurement must produce evidence that the external auditor and the audit committee will examine.

The Manage function must include the controls that the finance function operates over the artificial intelligence including the pre-deployment validation, the in-flight monitoring, the output review, the audit trail, and the rollback capability. The controls must satisfy the Sarbanes-Oxley requirements and must be documented and operated in forms that support the audit. The chief financial officer who establishes the Manage function with reference to the Sarbanes-Oxley control framework will produce an artificial intelligence governance that integrates with the broader control environment. The chief financial officer who establishes the Manage function separately from the Sarbanes-Oxley framework will produce a governance that operates alongside the control environment rather than as part of it, with the consequent risk of inconsistency between the two.

The Sarbanes-Oxley Implications

The Sarbanes-Oxley Act of 2002 established the framework for internal controls over financial reporting that public companies in the United States operate under and that many private companies have adopted voluntarily. The framework is the most consequential control environment in the modern finance function. The artificial intelligence that the finance function deploys intersects with this framework in several specific ways that the chief financial officer must hold deliberately.

How AI Intersects with ICFR

The internal controls over financial reporting include the controls over the transactions that produce the general ledger entries, the controls over the close and consolidation work that produces the consolidated financial statements, the controls over the disclosure preparation that produces the external reports, and the broader controls over the supporting processes. Artificial intelligence is being deployed across all of these areas. Artificial intelligence in accounts payable affects the controls over invoice processing. Artificial intelligence in revenue recognition affects the controls over revenue calculation. Artificial intelligence in close acceleration affects the controls over the close work. Artificial intelligence in disclosure management affects the controls over the disclosure preparation. Each deployment creates specific control implications that the framework must address.

The specific control implications include several categories. The control design must accommodate the artificial intelligence operationally, including the segregation of duties between the artificial intelligence and the human reviewers, the authorization protocols that govern the artificial intelligence's authority, and the change management controls that handle the artificial intelligence updates. The control testing must produce evidence that the artificial intelligence operates as intended, which requires testing approaches that work for the non-deterministic and potentially evolving nature of artificial intelligence systems. The control documentation must capture the artificial intelligence's role in the control environment in forms that the external auditor can examine and that the audit committee can review. The control deficiencies that may arise from artificial intelligence issues must be identified, assessed, and remediated through the established deficiency framework.

The Management Assertion and the Auditor's Opinion

Section 404 of the Sarbanes-Oxley Act requires public company management to assess the effectiveness of the internal controls over financial reporting and to provide an annual assertion. The external auditor must then provide an opinion on the effectiveness of the controls. The

assertion and the opinion both must address the artificial intelligence that affects the control environment. The chief financial officer who attests to the effectiveness of controls that include artificial intelligence is attesting to the artificial intelligence specifically. The attestation must be supported by the documentation and the evidence that the artificial intelligence operates as the control framework requires.

The implication is that the chief financial officer cannot delegate the artificial intelligence assertion to the chief information officer or to the broader information technology function. The assertion is the chief financial officer's responsibility. The supporting documentation and evidence must be produced and maintained by the finance function with the appropriate technical support. The chief financial officer who delegates the assertion will discover during the audit that the documentation does not exist in the form required and that the attestation cannot be supported. The chief financial officer who claims the assertion responsibility will engage with the technical work substantively and will produce the documentation that the attestation requires.

The 302 Certifications

Section 302 of the Sarbanes-Oxley Act requires the chief executive officer and the chief financial officer to certify each quarter that the financial reports do not contain material misstatements, that the disclosure controls are effective, and that the certifying officers have evaluated and disclosed any deficiencies. The certifications carry personal liability for the certifying officers. The certifications must address the artificial intelligence that affects the financial reporting and the disclosure controls.

The implication is that the chief financial officer's personal certification is on the line for the artificial intelligence in the finance function. The certification cannot be supported without confidence that the artificial intelligence is operating as intended. The confidence cannot be established without the validation, monitoring, audit trail, and rollback capabilities that the governance framework provides. The chief financial officer who deploys artificial intelligence without the supporting governance is signing certifications that depend on the artificial intelligence behaving in ways the governance has not verified. The exposure is personal and substantial.

The Audit Profession's Evolving Stance

The audit profession is engaging with artificial intelligence in finance through both the regulatory framework that governs the profession and through the operational practices that the audit firms develop. The chief financial officer must understand both dimensions to anticipate how the external audit will engage with the artificial intelligence in the finance function and to design the artificial intelligence deployment in ways that support the audit relationship rather than complicating it.

The Regulatory Framework

The Public Company Accounting Oversight Board in the United States, the International Auditing and Assurance Standards Board internationally, and the various national standard setters have begun issuing guidance and standards that address artificial intelligence in the audit context. The standards address several specific dimensions. The auditor's responsibilities for understanding the entity's use of artificial intelligence in the financial reporting process. The auditor's responsibilities for evaluating the controls that the entity operates over the artificial intelligence. The auditor's responsibilities for testing the artificial intelligence outputs that affect the financial statements. The auditor's responsibilities for using artificial intelligence in the audit work itself, which is a separate dimension that the chief financial officer should understand because the audit firm's artificial intelligence affects the audit experience.

The standards are still evolving. The chief financial officer should expect substantial development in the regulatory framework over the coming years as the profession develops the experience to translate the broad principles into specific operational guidance. The development will shape the audit experience for years to come. The chief financial officer who tracks the development will be prepared for the audit engagements as the standards mature. The chief financial officer who does not will be surprised by audit positions that the standards require but that the finance function has not anticipated.

The Audit Firms' Operational Practices

The major audit firms have all invested substantially in artificial intelligence capabilities both for their own use in conducting audits and for advisory services to their clients on artificial intelligence governance. The investments have produced operational practices that vary across the firms but that share recognizable patterns. The firms have all developed approaches to evaluating the artificial intelligence in client environments. The firms have all developed positions on the documentation they expect from clients regarding the artificial intelligence. The firms have

all developed testing approaches for the artificial intelligence outputs that affect the financial statements. The firms continue to evolve the practices as the experience develops.

The chief financial officer should engage with the external auditor on the artificial intelligence agenda early in the audit cycle rather than waiting for the auditor to raise the topic during the audit itself. The early engagement supports the audit by establishing shared expectations about the documentation, the testing, and the evidence that the audit will require. The early engagement also supports the artificial intelligence deployment by surfacing the audit requirements before they become constraints on operational decisions. The chief financial officer who engages early will produce an audit experience in which the artificial intelligence is addressed efficiently. The chief financial officer who does not will discover audit requirements during the audit that may require substantial remediation under time pressure.

The Reciprocal Question of AI in Audit

A reciprocal question that the chief financial officer should understand is the artificial intelligence that the audit firms are using in conducting the audit. The audit firms have deployed artificial intelligence across substantial portions of the audit work, including the analytical procedures, the journal entry testing, the substantive testing of specific account balances, and the broader audit analytics. The artificial intelligence in the audit is affecting how the audit is conducted, what evidence the auditor requests, and how the audit findings are produced.

The implication is that the audit experience is being shaped by the audit firm's artificial intelligence in addition to the chief financial officer's artificial intelligence. The chief financial officer should ask the external auditor about the artificial intelligence the audit firm uses, the governance the audit firm operates over the artificial intelligence, and the implications for the audit experience. The understanding supports the audit relationship by establishing shared awareness of the artificial intelligence dimensions of the audit. The understanding also supports the chief financial officer's own artificial intelligence governance by surfacing the audit firm's perspective on the issues that the chief financial officer's governance must address.

The Data Lineage Requirement

The artificial intelligence in finance consumes data from across the enterprise architecture and produces outputs that affect financial reporting, operational decisions, and strategic analysis. The connection between the inputs and the outputs is the data lineage. The lineage must be capturable, examinable, and defensible for the artificial intelligence governance framework to operate. The data lineage requirement is one of the more specific architectural commitments that artificial intelligence governance imposes on the broader finance technology architecture.

What Data Lineage Means in the AI Context

Data lineage in the artificial intelligence context includes several dimensions that extend beyond the traditional data lineage that finance functions have operated. The first dimension is the data lineage from source systems to the artificial intelligence inputs, which includes the path that data takes from the operational platforms through the data architecture to the artificial intelligence model. The second dimension is the lineage within the artificial intelligence interaction, which includes the prompts, the retrieved documents in the retrieval augmented generation pattern, and the model configuration that shaped the output. The third dimension is the lineage from the artificial intelligence output to the downstream consequences, which includes the platforms that consumed the output and the decisions that the output informed. The complete lineage spans these three dimensions and must be capturable for the artificial intelligence governance to operate.

The Operational Implications

The operational implications of the data lineage requirement are substantial. The data architecture must capture the lineage of the inputs to the artificial intelligence, which requires the architectural commitments to instrumentation that the data layer treated in Part Eleven discussed. The artificial intelligence integration must capture the interaction details including the prompts and the retrieved content, which requires the platform capabilities to support the capture. The downstream platforms must capture the artificial intelligence outputs they consumed, which requires the integration architecture to maintain the connection between the source artificial intelligence interaction and the downstream usage.

The implications place specific requirements on the platforms that the finance function selects. The platforms must support the lineage capture either natively or through integration with the broader observability infrastructure. The platforms that do not support the lineage capture will produce gaps in the governance framework that the audit and the regulatory examination will identify. The chief financial officer who treats the lineage capability as a selection criterion will

produce an architecture that supports the governance. The chief financial officer who treats the lineage capability as a secondary consideration will produce an architecture in which the governance must operate around the platform limitations rather than through the platform capabilities.

The Audit and Regulatory Dimensions

The data lineage requirement is increasingly being made explicit by the audit profession and by the regulatory community. The external audit will request the lineage documentation for any artificial intelligence that materially affects the financial statements. The regulatory examination will request the lineage documentation for any artificial intelligence that affects regulated activities. The internal investigation of any specific issue will require the lineage documentation to understand what occurred and to remediate the cause. The chief financial officer who can produce the lineage documentation efficiently will navigate the audit, the regulatory examination, and the internal investigation with the artifacts that the inquiries require. The chief financial officer who cannot will discover that the absence of the lineage documentation is itself a finding of the inquiry, separate from whatever issue triggered the inquiry initially.

Operating Implications

The artificial intelligence governance framework is operated continuously rather than implemented as a project. The operating implications include the governance structure, the operating disciplines, and the cultural posture that the function adopts.

The Governance Structure

The artificial intelligence governance structure includes several specific bodies and accountabilities. The artificial intelligence steering committee provides the executive oversight of the artificial intelligence portfolio, with the chief financial officer as the executive sponsor for the finance function's artificial intelligence and with participation from the chief information officer, the chief information security officer, the chief audit executive, the controller, and the broader executive team as appropriate. The artificial intelligence working group provides the operational oversight, with finance practitioners, technology practitioners, and risk practitioners working together on the day-to-day governance work. The artificial intelligence model inventory steward maintains the canonical inventory of the artificial intelligence systems. The artificial intelligence validation function provides the independent review of artificial intelligence systems before deployment. The artificial intelligence monitoring function provides the ongoing observation of artificial intelligence systems in production.

The structure must be calibrated to the scale of the artificial intelligence portfolio and to the criticality of the use cases. Smaller portfolios with lower-criticality use cases may operate with simpler structures. Larger portfolios with high-criticality use cases require the full structure. The chief financial officer who designs the structure deliberately at the beginning will produce a governance that scales as the portfolio grows. The chief financial officer who allows the structure to emerge from the individual deployment decisions will produce a structure that requires substantial rework as the portfolio reaches the scale that demands more sophisticated governance.

The Operating Disciplines

The operating disciplines that the framework requires include several specific practices. The pre-deployment review process examines every new artificial intelligence deployment against the framework requirements before it enters production. The ongoing monitoring process examines the performance of the artificial intelligence systems against the established baselines. The quarterly portfolio review examines the cumulative artificial intelligence portfolio against the strategic agenda. The annual framework review examines the framework itself against the evolving regulatory environment and the lessons from operational experience. The incident

response process addresses the specific issues that arise in production and that warrant attention beyond the routine monitoring.

The disciplines must be operationalized with explicit accountabilities and explicit timing. A framework that exists on paper without the operating disciplines is not a framework. The chief financial officer who establishes the disciplines and maintains them will produce an artificial intelligence governance that operates effectively. The chief financial officer who establishes the framework without the disciplines will produce a governance that satisfies the documentation requirements without producing the operational reality the framework was designed to support.

The Cultural Posture

The cultural posture on artificial intelligence governance is the chief financial officer's stance on the relationship between innovation and discipline in the artificial intelligence agenda. The posture shapes how the team approaches artificial intelligence deployment, how the team responds to incidents, how the team engages with the broader artificial intelligence community, and how the function presents itself to the external auditor, the regulatory community, and the board.

The chief financial officer who articulates the cultural posture explicitly will produce a function that operates with the appropriate balance between innovation and discipline. The chief financial officer who allows the posture to be defined by individual decisions will produce a function whose posture drifts. The articulation of the posture is one of the more important communications the chief financial officer makes to the finance team and to the broader stakeholders. The posture should make clear that innovation is welcomed within the framework that the governance establishes, that the framework is non-negotiable for the categories of use case it addresses, and that the chief financial officer holds personal accountability for both the innovation and the discipline.

Synthesis

Artificial intelligence governance in finance is the most consequential control environment innovation since the Sarbanes-Oxley Act of 2002. The comparison is deliberate. Sarbanes-Oxley imposed structural requirements that reshaped how the finance function operated. Artificial intelligence is imposing structural requirements on how the function reasons about decision quality, evidence, auditability, and accountability. The structural requirements are not optional. The chief financial officer who does not lead the governance response will face the audit, regulatory, and reputational consequences personally.

The hallucination problem is the most consequential technical limitation of the current generation of artificial intelligence and the problem the governance framework must address most directly. Hallucination is structural rather than incidental, and the architectural response must accommodate the reality rather than depend on its elimination. The four layers of defense include the architectural choice of contexts where verification is feasible, the use of retrieval augmented generation for factual use cases, the independent verification mechanisms for high-stakes use cases, and the monitoring that detects degradation.

The model risk management framework from the banking industry provides the structural foundation for artificial intelligence governance, adapted for the foundation model era. The adapted framework includes the artificial intelligence model inventory, the pre-deployment validation, the in-flight monitoring, the output review, the audit trail, and the rollback capability. The framework operates continuously and integrates with the broader Sarbanes-Oxley control environment rather than alongside it.

The NIST AI Risk Management Framework provides the starting point that should be adapted for finance specifically. The four functions of Govern, Map, Measure, and Manage operate continuously and iteratively. The finance adaptation emphasizes the chief financial officer accountability, the Sarbanes-Oxley integration, the audit relationship, and the disclosure controls.

The Sarbanes-Oxley implications include the intersection with the internal controls over financial reporting, the management assertion under Section 404 that the chief financial officer signs, and the personal certifications under Section 302 that the chief executive officer and the chief financial officer make each quarter. The implications mean that the chief financial officer cannot delegate the artificial intelligence governance and that the personal exposure is substantial.

The audit profession is engaging with artificial intelligence through the evolving regulatory framework that the standard setters are developing and through the operational practices that the audit firms are deploying. The chief financial officer should engage with the external auditor on the artificial intelligence agenda early in the audit cycle rather than waiting for the auditor to raise

the topic during the audit. The reciprocal question of the artificial intelligence that the audit firms use in conducting the audit is also relevant to the chief financial officer's understanding.

The data lineage requirement is one of the more specific architectural commitments that artificial intelligence governance imposes. The lineage spans the path from source data to artificial intelligence inputs, through the artificial intelligence interaction including the prompts and retrieved content, to the downstream consequences. The lineage capability must be a selection criterion for the platforms in the artificial intelligence architecture.

The operating implications include the governance structure with its specific bodies and accountabilities, the operating disciplines including pre-deployment review and ongoing monitoring and quarterly portfolio review and incident response, and the cultural posture that the chief financial officer articulates on the relationship between innovation and discipline. The chief financial officer who establishes the framework and the supporting operational reality will produce an artificial intelligence governance that captures the artificial intelligence value while satisfying the requirements. The chief financial officer who treats the governance as defensive overhead will produce a function that satisfies the minimum requirements at maximum cost without capturing the strategic value the artificial intelligence offers.

With Parts Fourteen and Fifteen, the masterclass has completed the treatment of artificial intelligence in finance, with the strategic framework in Part Fourteen and the governance discipline in Part Fifteen forming the integrated treatment that the topic requires. The next part treats cybersecurity, identity, and fraud technology for the finance function, the related cross-cutting concern that intersects with artificial intelligence governance in several specific ways. The architectural reasoning developed across the preceding parts carries forward into the treatment of the security domain.