eFuturesCFO Masterclass Series

AI Workflows for the Modern CFO

PART 4

The Governance Framework

Embedding Discipline Before Deployment

Hindol Datta

Chief Content and Product Officer, eFuturesCFO

A Note Before Part 4

Part 4 is the longest of the foundational parts of this masterclass and, in my judgment, the most consequential. It is the governance framework that will determine whether the five use cases of Parts 5 through 9 produce durable value or durable damage. Most executive AI education places governance at the end, as a chapter of warnings about what to be careful of after the deployment is built. This masterclass places governance at the front, as the design discipline that precedes everything else.

The reason is structural. Governance applied at the end of an AI deployment is governance bolted onto a structure that may have already made decisions incompatible with the controls. Governance applied at the front is governance embedded in the architecture from the first design decision. The cost of the first approach is rework. The cost of the second approach is time. Time is the cheaper expenditure.

John Campbell committed to drafting the governance framework within the first thirty days following acceptance of his memo to Sarah. Part 4 is that framework. It is written in the same voice as Part 3, with the same combination of internal document conventions and executive prose. The reader watches the framework take shape section by section, with explicit reasoning for each choice rather than mere declaration.

Read Part 4 carefully. It will, more than any other part of the masterclass, change the way you read the use case parts that follow. Every workflow in Parts 5 through 9 will be analyzed through the lens of the framework established here: which risk tier applies, where the human review sits, what the audit trail captures, what vendor risk is most relevant, how governance is embedded in the design. The framework is reusable, both within this masterclass and within your own organization. The closing artifact of Part 4 is John's actual governance charter, which is short, declarative, and designed to be adapted with relatively minor edits to fit any finance organization.

> *Governance applied at the end is rework. Governance applied at the front is design. The same controls cost differently depending on when in the lifecycle they are introduced.*

Hindol Datta

Contents

Each section is approximately four to six pages. Section 7 is the longest because it provides the reusable methodology that Parts 5 through 9 will apply to each of the five use cases. Section 9 is the shortest because the charter itself is deliberately concise. Both are equally important.

Section 1 · The Philosophy of AI Governance in Finance

Why finance is qualitatively different

Artificial intelligence in finance is not the same operational category as artificial intelligence in marketing, in customer service, in product design, or in engineering. The differences are not differences of degree. They are differences of kind, and they shape the entire governance framework that follows.

A hallucinated marketing email is embarrassing. The brand recovers. A hallucinated customer service response is unfortunate. The relationship survives in most cases. A hallucinated product description is misleading and may be remediable. A hallucinated piece of code is caught in code review. None of these failures is desirable. None of them is existential to the function in which they occur.

A hallucinated forecast that reaches the audit committee is a different category of event. A hallucinated revenue recognition entry that flows into the financial statements is a different category of event. A hallucinated board commentary that misrepresents the operating reality of the company to the directors who govern it is a different category of event. These are not embarrassments that can be apologized for and corrected in the next cycle. They are failures of the financial reporting and governance system itself, with consequences that range from internal control deficiencies to restatements to regulatory action to litigation.

The standard of care in finance is therefore qualitatively different from the standard of care in adjacent functions. The governance framework must reflect this difference, not merely scale it. A governance framework that treats AI in finance as a more controlled version of AI elsewhere will produce a finance function that is incrementally safer than its peers. A governance framework that treats AI in finance as a categorically different problem will produce a finance function that can defend itself against the failures that will, statistically, eventually occur.

> **The standard-of-care principle**
>
> AI in finance is not AI in marketing scaled up. It is a qualitatively different problem governed by qualitatively different rules. The framework that follows is built on that premise, not on the premise that finance is merely a more careful version of every other function.

What the CFO is uniquely positioned to do

The CFO is the right executive to own AI governance in finance, and in many organizations, the CFO is the right executive to own AI governance more broadly. The argument rests on four properties of the role.

The first property is consequence proximity. When an AI failure produces a misstatement, a control deficiency, an audit issue, or a regulatory inquiry, the consequence flows to the CFO's office. The CFO has

the standing and the accountability that match the failure modes. No other executive sits closer to the consequences that matter most.

The second property is institutional authority. The CFO is the executive whose authority over controls, over policies, and over financial reporting is established by the role itself, by SOX in public companies and by board mandate in private ones. The authority to require controls is not something the CFO has to negotiate. It comes with the seat.

The third property is cross-functional visibility. The CFO is the executive most consistently exposed to every function in the company because every function eventually reports through finance. This visibility is what makes the CFO capable of governing AI across functional boundaries, which is necessary because AI will be deployed across functional boundaries.

The fourth property is temperamental fit. The discipline of finance, properly practiced, is the discipline of skepticism about claims that have not been verified, of evidence-based reasoning, of attention to controls, of patience with process. These are precisely the dispositions that AI governance requires. The CFO who has internalized them is naturally suited to the work, in a way that an executive whose function rewards rapid action and tolerance for failure is not.

What governance is not

It is worth saying clearly what governance is not, because the term is often misused to mean things that would weaken rather than strengthen the framework.

Governance is not a checklist. A checklist provides verification that specified items have been addressed. Governance is the system that determines which items should be on the checklist, why, how they relate to each other, and what to do when the checklist is incomplete or inadequate. A checklist without governance is a procedure. Governance without a checklist is a philosophy. Both are necessary. Neither substitutes for the other.

Governance is not approval. A governance framework that consists primarily of an approval process is a framework whose principal function is to slow deployment. Slowing is not the objective. Discipline is the objective. A well-designed framework approves quickly the workflows that meet the discipline and refuses entirely the workflows that do not. Speed of approval is not a virtue. Quality of judgment is.

Governance is not separable from operations. A framework that exists in a binder rather than in the daily operation of the system is a framework that has failed. The test of governance is whether the day-to-day operation of the AI workflows is observably governed, not whether the binder is up to date.

Governance is not finished. The framework that follows is a living document. It will be revised as the regulatory environment evolves, as the technology matures, as the company grows, and as the workflows accumulate operational history. The framework is built to be revised. Sections 5 and 9 of this part explicitly address the revision cadence.

Section 2 · The Three Pillars: Data, Model, Output

The governance framework rests on three pillars. Each addresses a distinct phase in the life of an AI workflow. Each has its own controls, its own risks, and its own failure modes. The framework treats them separately not because they are independent (they are deeply interconnected) but because the controls appropriate to each are different.

Pillar One: Data governance

Data governance asks what data the AI is allowed to see, where it lives, who owns it, how its quality is assured, and how it is protected throughout the workflow.

Five questions define the data governance posture for any workflow.

What data flows into the workflow?

Every workflow must have an explicit, written inventory of its data inputs. The inventory specifies the system of origin, the data classification level, the refresh cadence, and the owner. Workflows whose data inputs cannot be fully enumerated are not approved for deployment.

What classification level applies?

Helix uses a four-tier data classification: public, internal, confidential, and restricted. Each tier carries specific handling requirements. The data classification of every input is determined at the time of workflow design and is documented in the workflow registry.

Where does the data go?

The data flow map specifies every system, every vendor, every storage location through which the data passes between origin and output. The map is updated when the workflow changes. Data that crosses jurisdictional boundaries (in or out of the European Economic Area, for instance) is flagged for additional review.

How is data quality assured?

Data quality assurance is the responsibility of the workflow owner. The workflow registry records the data quality checks applied at each stage. Workflows whose output depends critically on data quality must specify how out-of-bounds data is detected and handled.

What retention applies?

Retention periods are specified at the workflow level for every data input and every intermediate artifact. Retention beyond the periods required by audit, regulatory, and contractual obligations is the exception rather than the default. Data that is no longer needed is purged on a defined cadence.

Data governance in summary

Five questions, asked of every workflow before deployment. What flows in? What classification applies? Where does the data go? How is quality assured? What retention applies? Workflows whose answers are incomplete are not approved.

Pillar Two: Model governance

Model governance asks which models are approved, how they are versioned, how they are tested before deployment, and how vendor changes are tracked and absorbed.

Five questions define the model governance posture.

Which model is approved for which workflow?

The approved model list is maintained by the CFO's office with input from the security and legal functions. Models are approved at the level of model family and specific version. A workflow may not invoke a model that is not on the approved list. Adding a model to the list requires review of the vendor's data handling commitments, security posture, and regulatory alignment.

How are versions controlled?

Every workflow specifies the model version it uses. When a vendor releases a new model version, the workflow does not automatically migrate. The migration is treated as a change management event, with testing of the new version against the workflow's established outputs before cutover.

How are models tested before deployment?

Every workflow must demonstrate the quality of its output against a defined test set before reaching production. The test set is workflow-specific and reflects the actual distribution of inputs the workflow will encounter. Acceptable performance is defined in advance, not retrospectively.

How are deprecations handled?

Vendors regularly deprecate older model versions, with advance notice that varies from weeks to months. The workflow registry tracks deprecation dates and ensures that migration to a current version completes before the deprecation takes effect. No workflow may operate on a deprecated model in production.

How are vendor relationships diversified?

Single-vendor dependencies are minimized where feasible. Critical workflows are designed to permit model substitution to a comparable model from a different vendor with minimal rework. This protects against vendor pricing changes, vendor product changes, and vendor relationship disruption.

Pillar Three: Output governance

Output governance asks what the AI is allowed to produce, who reviews it, who approves it, and what audit trail is preserved.

Five questions define the output governance posture.

What output is the workflow producing?

Every workflow specifies its outputs explicitly. The specification includes the format, the destination system, the audience, and the consequence chain. Workflows that produce outputs with no defined consequence chain are not approved.

Who reviews the output before release?

The human review checkpoint is specified at the workflow level. The reviewer is named, not a role. Backup reviewers are specified for absences. The reviewer's authority and training are documented. Workflows that produce financial reporting outputs require finance team review. Workflows that produce external-facing outputs require additional review by the relevant function head.

Who approves the output for use?

Approval for use, where required, is a separate act from review. The approver may be the same person as the reviewer or a different person. The approval is logged with timestamp, identity, and the version of the output approved.

What audit trail is preserved?

The audit trail captures the input, the model version, the prompt or workflow definition, the raw output, the reviewed output, the human reviewer, the approver, and timestamps for each event. The audit trail is immutable and is retained for the periods required by audit, regulatory, and litigation hold obligations.

What happens when the output is wrong?

Every workflow specifies its error detection and correction process. Errors discovered after the output has been used trigger a documented remediation, including identification of who consumed the wrong output, what downstream consequences resulted, and what corrective communication is required.

The pillars in practice

In practice, the three pillars are evaluated for every workflow at the design stage. Section 7 of this framework specifies the methodology for that evaluation. The pillars are not separate processes. They are three lenses through which the same workflow is examined. A workflow that passes data governance but fails output governance is not approved. A workflow that passes output governance but fails model governance is not approved. The pillars are three necessary conditions. None of them is sufficient alone.

Section 3 · Human-in-the-Loop Architecture

Every AI workflow in finance must include a human checkpoint. The argument for this is established in Part 3 as the second of the seven architectural principles. The present section addresses the design question that principle leaves open: where in the workflow should the checkpoint sit, and what should the human actually do at the checkpoint?

Why checkpoint placement matters

A checkpoint that is placed too early in the workflow adds friction without adding control. The human is asked to review work that has not yet been done. A checkpoint that is placed too late lets errors propagate through downstream steps before they can be caught. The propagation often makes the errors more expensive to correct. The right placement is the earliest point at which the reviewer has enough information to make a meaningful judgment.

The wrong placement also undermines the human reviewer. A reviewer who is asked to approve outputs they cannot meaningfully evaluate becomes a rubber stamp. The rubber stamp is worse than no checkpoint at all, because it creates the appearance of control while delivering none of the substance. The framework must therefore distinguish between checkpoint patterns that produce real control and checkpoint patterns that produce only the appearance.

A taxonomy of checkpoint patterns

Five checkpoint patterns are recognized in this framework. Each is appropriate to different workflow types. Each has different operational implications. The workflow design specifies which pattern applies.

Pattern One: Review-before-output

The workflow runs through its analytical steps and produces a draft output. The draft is presented to the reviewer before any external release. The reviewer either approves the draft as final or returns it for revision. The pattern is appropriate for outputs that go to external audiences (board, auditor, investor) or that have material consequences.

Pattern Two: Review-before-distribution

The workflow produces and finalizes the output, which is then held in a staging area until a reviewer approves distribution. The pattern is appropriate for outputs that are mechanically produced and reviewed in batch, such as a set of routine variance commentaries that are all released together rather than individually.

Pattern Three: Review-before-action

The workflow recommends an action but does not take it. The action is taken only after a reviewer approves the recommendation. The pattern is appropriate for workflows that propose interventions in operational systems, such as a workflow that flags vendor contracts for renegotiation but does not initiate the renegotiation itself.

Pattern Four: Review-after-action with reversal capability

The workflow takes an action automatically and then presents the action to a reviewer for post-hoc evaluation. The action is reversible if the reviewer determines it was incorrect. The pattern is appropriate only for actions whose reversal is straightforward and whose immediate execution produces meaningful value. The pattern requires particularly strong audit trails because errors are caught after the fact.

Pattern Five: Sampled review

The workflow produces many outputs of a similar type. A statistically meaningful sample of the outputs is reviewed. The pattern is appropriate only after the workflow has established a track record of consistent quality, and only for outputs whose individual error consequence is small. Sampled review is the most operationally efficient pattern and the most easily abused. It requires explicit approval from the CFO's office before adoption.

> **The default pattern**
>
> The default checkpoint pattern is Review-before-output. Any deviation from this default requires explicit design rationale documented in the workflow registry. Sampled review in particular is reserved for workflows that have demonstrated operational maturity over multiple cycles.

What the reviewer actually does

Specifying the checkpoint pattern is not sufficient. The framework must also specify what the reviewer actually does at the checkpoint. Three review activities are recognized.

Verification review

The reviewer verifies that the output is consistent with the underlying data. The review focuses on factual accuracy. Verification review is appropriate for outputs that summarize or analyze data, where the question is whether the summary or analysis accurately reflects the source.

Judgment review

The reviewer applies their own judgment to the conclusion or recommendation produced by the workflow. The review focuses on whether the conclusion is appropriate given the context. Judgment review is appropriate for outputs that produce a recommendation or interpretation.

Compliance review

The reviewer evaluates the output against a defined set of compliance criteria. The review focuses on whether the output meets the specified standards for tone, format, completeness, and policy alignment. Compliance review is appropriate for outputs that follow a defined template or must meet specified communication standards.

Most workflows require a combination of two or three review activities. The combination is specified in the workflow design. The reviewer is responsible for performing all specified activities at the checkpoint. A reviewer who performs only verification when judgment is also required is performing the review incompletely.

Backup and escalation

Every named reviewer requires a named backup. Workflows that depend on a single individual whose absence interrupts the workflow are not operationally resilient. The backup is identified at the workflow design stage and is trained on the review responsibility.

Every workflow specifies its escalation path for cases the reviewer cannot resolve at the checkpoint. The escalation path is typically to the function head, then to the CFO. Workflows that produce outputs with material consequences may have shorter escalation paths.

Section 4 · Audit Trail and Observability

Every AI workflow in finance must produce an immutable audit trail. The principle is established in Part 3. The present section specifies what the audit trail captures, how it is stored, who can access it, and how it is used.

The minimum audit trail

Every workflow invocation produces a record containing, at minimum, the following fields.

Field	Description
Invocation ID	A unique identifier for this specific workflow invocation, used to correlate all related records.
Workflow name and version	The specific workflow that was invoked, including the version of its definition.
Initiating user	The identity of the user, system, or schedule that initiated the workflow. Service accounts are identified explicitly.
Timestamp (initiated)	The time at which the workflow began, recorded with timezone in UTC and local time.
Inputs	The data inputs provided to the workflow. For large inputs, a cryptographic hash plus a pointer to the stored input is acceptable.
Model and version	The specific AI model invoked, including provider, family, and exact version identifier.
Prompt or workflow definition	The prompt template, system message, or workflow definition used. Version-controlled in the prompt registry.
Raw output	The exact output produced by the model, before any human review or editing.
Human reviewer	The named individual who reviewed the output, if any, with their role and authority at the time.
Reviewed output	The output as approved after human review, with any edits applied.
Approval timestamp	The time at which the output was approved for use.
Downstream actions	A record of where the approved output was sent, used, or stored.
Error events	Any errors, retries, fallbacks, or exceptional conditions that occurred during the invocation.

The minimum trail is not optional. A workflow that does not produce all of the above fields is not approved for deployment. The audit trail is not an output of the workflow. It is part of the workflow.

Excellence in audit trails

Beyond the minimum, excellent audit trails support capabilities that the minimum trail does not. The framework encourages workflows to implement these capabilities where feasible.

First, the excellent audit trail supports reconstruction. Given an invocation ID, the trail permits full reconstruction of what happened, in what order, with what inputs, producing what outputs, reviewed by whom. The reconstruction is sufficient for the auditor to evaluate the workflow without consulting the people who built it.

Second, the excellent audit trail supports pattern analysis. The trail is queryable across invocations, permitting analysis of patterns over time. How often does this workflow produce outputs that are revised at review? Which reviewers most frequently revise outputs? Are there inputs that consistently produce poor outputs? Pattern analysis is how the framework learns from operational experience and refines itself.

Third, the excellent audit trail supports compliance reporting. The trail can produce, on demand, the evidence needed to respond to a customer audit question, a regulator inquiry, or an external auditor review. The compliance report is a derivative of the trail, not a separate artifact that must be constructed when needed.

Storage and retention

Audit trail data is classified as restricted under the data classification scheme. It contains information about workflow operations, model outputs, and human decisions that is sensitive in itself and must be protected.

Storage is in a dedicated audit trail system, separate from the operational systems whose activity it records. The separation prevents the audit trail from being modified by the same processes it observes. The storage is immutable: records can be added but not changed or deleted within the retention period.

Retention periods are workflow-specific and are set with reference to the audit, regulatory, and litigation hold obligations that apply to the workflow's outputs. For workflows that touch financial reporting, the retention period aligns with the financial records retention policy, which at Helix is seven years. For workflows that do not touch financial reporting, the retention period is specified at the workflow design stage and reviewed periodically.

Access controls

Audit trail access is controlled by role. Three access levels are defined.

Read access

Granted to the workflow owner, the CFO, the General Counsel, the Head of Security, the external auditor on request, and named investigators when an incident requires. Read access permits reconstruction and pattern analysis. Read access is logged.

Query access

Granted to a narrower set, including the workflow owner, the CFO, and named analysts. Query access permits ad-hoc analysis across the trail without requiring file downloads. Query access is logged.

Administrative access

Granted only to the Head of Security and one named administrator in the engineering organization. Administrative access is required to configure the audit trail system itself, including retention policies and access grants. Administrative access is logged with particular care, because it is the level at which the trail could potentially be compromised.

> **The principle behind audit trail access**
>
> The audit trail is the most consequential record in the AI program. Its integrity must be defended at the same level as the integrity of the financial statements themselves. Access is granted narrowly, logged completely, and reviewed periodically.

Section 5 · Risk Classification

Not every AI workflow carries the same risk. A workflow that generates internal exploratory analysis is lower risk than a workflow that produces a board report, which is lower risk than a workflow that posts to the general ledger. The governance framework must apply proportionate discipline to workflows of differing risk, or it will either under-govern the high-risk workflows or over-govern the low-risk ones. Either failure mode undermines the framework.

Four risk tiers

The framework defines four risk tiers for AI workflows. Every workflow is assigned to a tier at the design stage. The tier determines the level of review, the depth of audit trail, the frequency of monitoring, and the escalation paths.

Tier One: Internal Exploratory

Workflows that produce outputs consumed only by named individuals for analytical or exploratory purposes, with no downstream consequence outside the analytical work itself. Examples: an analyst-facing copilot that helps with ad-hoc questions against the warehouse, where the output is consumed only by the analyst and used to inform their own further work.

Tier One workflows require basic audit trail, single reviewer at the user level (the user themselves), and lightweight governance. They are approved at the workflow-design stage by the function head, with the CFO notified.

Tier Two: Internal Operational

Workflows that produce outputs consumed by internal operations, with downstream consequences in operational systems but not in financial reporting or external communications. Examples: a workflow that classifies GL transactions for human review, a workflow that monitors AWS utilization and produces engineering alerts.

Tier Two workflows require complete audit trail, named human reviewer, periodic quality monitoring, and CFO approval before deployment. They are documented in the workflow registry and reviewed quarterly.

Tier Three: Financial Reporting Adjacent

Workflows that produce outputs which influence financial reporting, board materials, or external financial communications, even if they do not directly post entries to the books. Examples: the board reporting workflow, the forecasting workflow, the pipeline intelligence workflow.

Tier Three workflows require complete audit trail, named human reviewer with finance team authority, monthly quality monitoring, CFO approval before deployment, and audit committee notification. They are reviewed at every audit committee meeting.

Tier Four: Financial Reporting Direct or High-Risk Regulatory

Workflows that produce outputs which directly enter the financial records or that fall into high-risk regulatory categories under applicable law. Examples: a workflow that posts automated journal entries to the general ledger, a workflow that makes employment decisions, a workflow that determines creditworthiness of customers.

Tier Four workflows require complete audit trail, multiple-reviewer architecture (preparer and approver, separated), continuous quality monitoring, CFO approval, audit committee approval before deployment, external auditor coordination, and legal review. Tier Four workflows are exceptional. At Helix, no workflow currently planned in the eighteen-month horizon falls into Tier Four.

The classification process

Every candidate workflow is classified at the design stage. The classification is recommended by the workflow owner and reviewed by the governance working group, which consists of the CFO, the General Counsel, and the Head of Security. The classification is documented in the workflow registry and cannot be changed without explicit governance working group approval.

Classification is conservative. Where reasonable people might disagree about which tier applies, the higher tier governs. The cost of over-classifying a workflow is incremental review work. The cost of under-classifying a workflow is the risk that the controls applied are inadequate to the actual risk.

Reclassification

Workflows are reviewed for reclassification on a defined cadence. The cadence varies by tier: annually for Tier One, semi-annually for Tier Two, quarterly for Tier Three, and continuously for Tier Four. Reclassification is triggered by material changes in the workflow, by regulatory developments, by accumulated operational experience, or by the discovery of unanticipated risks.

A workflow may be reclassified up (to a higher tier) at any time by the governance working group. A workflow may be reclassified down (to a lower tier) only after demonstrating sustained operational maturity at the higher tier, with documented evidence supporting the reclassification.

> **The tier principle**
>
> Four tiers. Conservative classification. Periodic review. Easy upgrade. Difficult downgrade. The framework treats risk asymmetrically because the consequences of under-controlling exceed the consequences of over-controlling.

Section 6 · Vendor and Model Risk

Deploying AI in finance creates vendor risk and model risk that the finance function has not historically had to manage in this form. Vendor risk arises because critical workflows depend on third-party providers whose pricing, product, and relationship terms can change. Model risk arises because the AI models themselves can change in ways that affect workflow behavior. The two risks interact and must be governed together.

Vendor risk dimensions

Six dimensions of vendor risk apply to every AI vendor relationship.

Dimension	Question
Pricing	How may the vendor change pricing, and what advance notice applies? What is our exposure if pricing increases materially?
Product	What product changes may the vendor make, and what advance notice applies? Can the vendor remove capabilities our workflows depend on?
Data handling	What commitments does the vendor make about our data? Are commitments contractual or merely policy? What happens if the policy changes?
Security posture	What is the vendor's security posture? What independent assessments exist? What incident notification commitments apply?
Subprocessor list	Which third parties does the vendor use to deliver the service? What change notification applies if the subprocessor list changes?
Termination	Under what conditions may the vendor terminate the relationship? What transition support applies? How quickly must we substitute?

Every AI vendor we engage must have written answers to these six questions before the relationship begins. The answers are reviewed by the General Counsel and the Head of Security. The relationship is documented in the AI vendor register maintained by the operations function.

Model risk dimensions

Distinct from vendor risk, model risk addresses how the AI model itself can change. Three model risk dimensions apply.

Version evolution

Vendors release new model versions on a cadence ranging from quarterly to multiple times per year. Each new version is, in principle, a different model. The behavior of a workflow on a new version may differ from its behavior on the prior version, even when the new version is presented as an improvement.

The framework requires that workflows specify the exact model version they use, that the vendor's deprecation schedule be tracked, and that migration to a new version be treated as a change management event with testing before cutover. The migration is not automatic and is not assumed to be safe merely because the vendor labels the new version as superior.

Behavioral drift

Even within a single model version, vendor-side updates such as fine-tuning adjustments, safety tuning, or system-level changes may alter the model's behavior in subtle ways. The framework requires monitoring of workflow output quality over time, with attention to drift that cannot be explained by changes in the workflow itself.

Capability changes

Vendors occasionally restrict capabilities of their models in response to safety concerns, regulatory requirements, or commercial decisions. The framework requires that workflows be tested periodically against their full requirement set, not merely against the aspects most frequently exercised.

Substitutability and the model substitution path

The fifth of the seven architectural principles requires that workflows be designed for model substitutability. The present section operationalizes that principle.

Every workflow specifies, at the design stage, what its model substitution path looks like. The specification answers three questions: which alternative models could serve the workflow if the primary model became unavailable, what rework would be required to migrate, and how long the migration would take.

The substitution path is exercised periodically. At least once per year, every Tier Three and Tier Four workflow is tested against its alternative model to confirm that the substitution path remains viable. The test is not a full production cutover; it is a verification that the workflow can be migrated within the documented time and rework budget.

Vendor concentration

Vendor concentration is the risk that too many critical workflows depend on a single vendor. At Helix, the current concentration is high. Most of our AI workflows use Anthropic's Claude models, with a smaller portion using OpenAI's GPT models. This concentration is monitored by the governance working group.

The framework does not mandate a specific level of concentration. It does require that concentration be tracked, that it be reviewed at every audit committee meeting, and that any movement toward higher concentration be deliberate rather than accidental. The Model Context Protocol, where adopted, reduces concentration risk by making the underlying model more easily substitutable.

> **The vendor and model risk principle**
>
> Vendor risk and model risk are categories of risk that finance has not historically managed. They are real, they are growing, and they require disciplines that the finance function must develop. Tracking concentration, maintaining substitutability, and monitoring behavioral drift are the three operational practices that constitute mature handling of these risks.

Section 7 · Embedding Governance into Workflow Design

This is the longest and most important section of the framework. The five preceding sections establish what governance requires. The present section specifies how governance is operationalized at the design stage for every new workflow. The methodology is reusable: every one of the five use cases in Parts 5 through 9 will apply this methodology, and so should every workflow deployed in any finance function adopting this framework.

The eight-step methodology

Every workflow design proceeds through eight steps before deployment is approved. The steps may be completed iteratively rather than strictly sequentially, but no workflow proceeds to production until all eight are complete and documented.

Step One: Articulate the business problem

Every workflow begins with a problem statement that specifies the business pain being addressed, the function that owns the pain, and the magnitude of the pain. The problem statement is one to three paragraphs and is signed off by the function head whose function owns the pain. Workflows that cannot articulate a clear business problem do not proceed.

Step Two: Specify the output

The workflow's outputs are specified in detail before any architectural work begins. The specification includes the format, the audience, the consequence chain, the frequency, and the quality criteria. Output-first design is essential because the rest of the workflow is constrained by what the output must look like. Workflows designed without a clear output target tend to drift into general-purpose tools that serve no specific purpose well.

Step Three: Classify risk

The workflow is assigned to one of the four risk tiers established in Section 5. The classification is recommended by the workflow owner, reviewed by the governance working group, and documented in the workflow registry. The risk tier determines the depth of governance work that follows, so the classification must be done early.

Step Four: Map data flow

The complete data flow map is constructed. Inputs are enumerated. Classifications are applied. Vendor subprocessors are identified. Jurisdictional movements are flagged. Quality assurance points are specified. The data flow map is the principal artifact of data governance for this workflow.

Step Five: Design the human review pattern

The checkpoint pattern from Section 3 is selected. Named reviewers and backups are identified. The review activities (verification, judgment, compliance) are specified. Escalation paths are defined. The human review design is documented and approved before any engineering work begins.

Step Six: Specify the audit trail

The workflow's audit trail is specified using the minimum from Section 4 as the baseline. Tier Three and Tier Four workflows specify additional capabilities beyond the minimum. Retention periods are set. Access controls are specified. The audit trail design is documented and approved.

Step Seven: Define the substitution path

The model substitution path is documented. Alternative models are identified. Rework estimates are produced. Migration timing is specified. The substitution path is tested in concept before deployment and tested in practice on the periodic cadence specified in Section 6.

Step Eight: Approval and registration

The complete design package is submitted to the governance working group for approval. Approval is based on completeness of the design rather than on reviewer enthusiasm. Once approved, the workflow is registered in the workflow registry and engineering work begins. Pre-approval changes to the design are permitted and require re-approval.

> **The eight-step methodology in summary**
>
> Articulate problem. Specify output. Classify risk. Map data. Design review. Specify audit trail. Define substitution. Approve and register. Eight steps, completed before deployment, documented in writing, signed by named accountable individuals.

Design artifacts

The eight-step methodology produces a defined set of design artifacts. The artifacts are the documentation of the workflow's governance. They are stored in the workflow registry alongside the operational records of the workflow itself.

Artifact	Purpose
Problem statement	One to three paragraphs articulating the business problem the workflow addresses.
Output specification	Detailed description of what the workflow produces, in what format, for what audience.
Risk classification record	The assigned tier with reasoning and the governance working group approval.

Artifact	Purpose
Data flow map	Visual or tabular representation of the complete data flow with classifications and subprocessors.
Review design	Specification of checkpoint pattern, reviewers, activities, and escalation paths.
Audit trail design	Specification of fields captured, storage, retention, and access controls.
Substitution path	Documentation of alternative models, rework estimates, and migration timing.
Approval record	Governance working group sign-off with named individual approvals.

Common design failures

Across the workflows reviewed by the governance working group, certain design failures recur. Naming them in the framework helps designers avoid them.

Failure One: Output specification by example

Designers sometimes specify the output by showing a single example rather than by articulating the structural requirements. A single example is useful but insufficient. The specification must capture what makes any output acceptable, not merely what one good output looks like.

Failure Two: Review checkpoint as afterthought

Designers sometimes propose the workflow architecture first and then ask where to insert the human review. This is the wrong order. The review checkpoint should be designed alongside the workflow, as a constraint that shapes the workflow rather than as a step inserted into an existing flow.

Failure Three: Audit trail as logging

Designers sometimes confuse audit trail with system logging. System logging captures technical events. Audit trail captures governance events. They overlap in practice but are distinct in purpose. A workflow whose audit trail is merely its system log has not produced an audit trail. The audit trail must be specifically designed.

Failure Four: Substitution path as principle without test

Designers sometimes claim substitution as a principle without specifying which alternative model would be used or estimating the rework required. The principle is meaningful only if the path is documented. Untested substitution paths tend not to work when needed.

Failure Five: Risk classification by self-assessment

Designers sometimes classify their own workflows optimistically, assigning a lower tier than the workflow warrants. The classification must be reviewed by the governance working group, not accepted as the owner specifies. The independent review is what makes the classification trustworthy.

How the methodology will appear in Parts 5 through 9

The remaining parts of this masterclass present the five use cases. Each part will apply the eight-step methodology to its specific workflow. The reader will see the methodology operationalized five times in five different contexts, with concrete artifacts produced for each. The pattern is intentional: by the end of the masterclass, the reader has internalized the methodology through repeated exposure to its application.

The methodology is not bureaucratic ornament. It is the discipline that distinguishes a finance function that can defend its AI deployments to its auditor, its investor, and its regulator from a finance function that has deployed AI without that defensibility. The first category of finance function is what Helix is building. The second category is what most peer companies will find themselves in by the end of 2026, and what they will then spend years remediating.

Section 8 · Regulatory and External Context

The governance framework operates within a regulatory and external environment that has evolved meaningfully in the past three years and continues to evolve. The framework must accommodate the environment without being captured by it. This section addresses the principal external considerations and how the framework responds to each.

Sarbanes-Oxley and public company controls

For public companies in the United States, the Sarbanes-Oxley Act of 2002 establishes internal control requirements over financial reporting. AI workflows that produce inputs to financial reporting fall within the scope of these requirements. The relevant question is not whether AI is permitted but whether the AI workflow operates with controls sufficient to support management's assertion about the effectiveness of internal control.

Helix is private. We are not currently subject to SOX. We are operating, however, on a path toward potential public-company status within the foreseeable horizon, and we design our controls to be SOX-ready rather than SOX-absent. The cost of designing controls properly from the start is lower than the cost of retrofitting them when a public offering is contemplated. The cost is also lower than the cost of an institutional investor in the Series C round discovering that our controls do not meet the standard they expect.

The framework therefore treats SOX-style controls as the design standard, even though the legal requirement does not currently apply. Specifically: workflows that touch financial reporting are classified Tier Three or Four, require multiple-reviewer architectures where appropriate, maintain complete audit trails, and are reviewable by an external auditor without preparation.

The General Data Protection Regulation and adjacent privacy regimes

The General Data Protection Regulation governs the processing of personal data of individuals in the European Economic Area. Helix has customers in Europe and therefore processes some personal data subject to the GDPR. AI workflows that process personal data are subject to specific requirements, including the right to explanation for automated decisions that materially affect individuals.

The framework requires that workflows processing personal data identify this in the data flow map, that the workflow be reviewed by counsel before deployment, and that the principle of data minimization be applied. Where a workflow can accomplish its purpose without processing personal data, it must do so. Where personal data is genuinely required, the workflow specifies the lawful basis for processing and complies with applicable transparency and rights requirements.

Adjacent privacy regimes, including the California Consumer Privacy Act, the Brazilian General Data Protection Law, and emerging regimes in other jurisdictions, are tracked by counsel. The framework is designed to be compatible with the strictest applicable regime rather than the most permissive.

The EU AI Act

The EU AI Act, fully applicable since 2026, classifies AI systems by risk and imposes corresponding obligations. Most enterprise productivity AI falls in the limited-risk tier, with transparency requirements but no significant operational burden. High-risk uses include employment screening, creditworthiness assessment, certain critical infrastructure applications, and others. Prohibited uses are narrower and largely irrelevant to finance.

The framework requires that every workflow be evaluated against the EU AI Act risk classification, even if no European operations are currently in scope. The reason is that the Act's reach extends to systems placed on the market or put into service in the Union. A workflow that is high-risk under the Act's definition triggers obligations that include risk management, data quality requirements, human oversight, transparency, and conformity assessment.

Workflows currently planned at Helix do not fall into the high-risk categories. Workflows that might in the future (particularly any future employment or credit application) will require additional review before approval.

US federal landscape

Federal AI regulation in the United States in mid-2026 is fragmented. The executive order issued in late 2023 was revoked. No comprehensive federal AI legislation has been enacted. Sectoral guidance exists from financial services regulators, the Federal Trade Commission, the Equal Employment Opportunity Commission, and others. States have been moving faster than the federal government, with material AI legislation in California, New York, Texas, Colorado, and others.

The framework treats the US landscape as fluid. Counsel tracks developments quarterly. Workflows that touch employment, credit, or other regulated decision categories receive enhanced legal review. The framework is designed to be adaptable to material federal action when it occurs.

External auditor coordination

Our external auditor, Marsh and Henning LLP, is briefed on the AI deployment plan as part of standard audit planning. AI workflows that touch financial reporting will be subject to audit testing. Early coordination with the auditor is more productive than late surprise.

The framework requires that Tier Three and Tier Four workflows be discussed with the auditor before deployment, that their audit trail design be reviewed for auditor usability, and that the auditor be given access to the audit trail on request during the audit period. Auditor feedback influences the framework over time; we expect the framework to mature as audit experience accumulates.

Customer contractual commitments

Enterprise customers, particularly in regulated industries, are increasingly including AI clauses in their procurement agreements with us. Typical clauses require disclosure of AI subprocessors, commitments against training on customer data, audit trails for AI-driven decisions, and notification of material changes in the AI subprocessor list.

The framework supports compliance with these clauses by maintaining the AI vendor register, the data flow maps, and the audit trails that customers may request. The framework also requires that any new customer contractual commitment regarding AI be reviewed against the framework before signature. Commitments that exceed what the framework supports are flagged before signature rather than discovered after.

Section 9 · The Governance Charter

The preceding sections of this framework establish the philosophy, the pillars, the architecture, the controls, and the methodology. The present section is the operational charter that distills the framework into the form that will govern day-to-day decisions. The charter is short. It is designed to be referenced quickly when a decision must be made.

AI GOVERNANCE CHARTER

Helix Cloud Systems, Inc.

This charter establishes the principles, accountabilities, and controls that govern the deployment and operation of artificial intelligence in the Helix Cloud Systems finance function and in any function in which finance has visibility or authority. It is adopted by the Chief Financial Officer with the endorsement of the Audit Committee Chair. It is reviewed annually and may be amended at any time by the Chief Financial Officer with notification to the Audit Committee.

Article I: Authority

The Chief Financial Officer is the executive accountable for the governance of artificial intelligence in the finance function. The CFO's authority under this charter includes the approval or refusal of any AI workflow deployment, the maintenance of the approved model and vendor list, the classification of workflows by risk tier, the establishment of review requirements, the periodic review of the governance framework, and the escalation of governance matters to the Audit Committee. The authority is exercised in consultation with the General Counsel, the Head of Security and IT, and the function heads whose functions are affected by a given workflow.

Article II: Principles

The seven architectural principles established in the CFO's foundational memo govern every deployment.

First, workflows over agents. Every AI deployment is designed as a workflow with deterministic flow control unless a specific case is made and approved for an autonomous agent.

Second, human-in-the-loop for all financial outputs. No AI-generated output that affects financial statements, board materials, customer-facing financial communications, or external reporting is released without named human review.

Third, no production AI without an audit trail. Every workflow produces an immutable audit trail containing the fields specified in this framework.

Fourth, no sensitive data leaving approved environments. Customer personally identifiable information, employee compensation data, acquisition-sensitive information, and data classified as confidential or restricted does not leave the approved AI environment without specific authorization.

Fifth, no vendor lock-in without a substitution path. Every workflow specifies its model substitution path before deployment.

Sixth, every workflow has a named human owner. Ownership cannot be a committee. Ownership must be a person.

Seventh, governance precedes deployment. No workflow is deployed without completing the eight-step methodology and receiving approval from the governance working group.

Article III: Governance Working Group

The Governance Working Group consists of the Chief Financial Officer, the General Counsel, and the Head of Security and IT. The Group meets monthly and on demand for specific workflow approvals. The Group is responsible for the approval or refusal of workflow designs, the classification of workflows by risk tier, the maintenance of the approved model and vendor list, the review of operational incidents affecting governance, and the periodic review of the governance framework.

Article IV: The Workflow Registry

The Workflow Registry is the central record of every AI workflow deployed at Helix. The Registry is maintained under the authority of the Chief Financial Officer. Every workflow in the Registry is documented through the design artifacts specified in Section 7 of the governance framework. Operational records of workflow invocation, including audit trail data, are stored in association with the Registry entry.

Article V: Risk Tiers and Review Cadence

Workflows are classified into four risk tiers: Internal Exploratory, Internal Operational, Financial Reporting Adjacent, and Financial Reporting Direct or High-Risk Regulatory. Tier assignments are reviewed on the cadence specified in Section 5 of the governance framework: annually, semi-annually, quarterly, and continuously, respectively. Reclassification upward may occur at any time. Reclassification downward requires demonstrated operational maturity.

Article VI: Audit Committee Oversight

The Audit Committee maintains a standing AI Governance agenda item at every meeting. The Chief Financial Officer provides a written governance update to the Committee at least quarterly. Material developments are reported to the Audit Committee Chair within days, not at the next scheduled meeting. Tier Four workflows require Audit Committee approval before deployment.

Article VII: External Review

At the twelve-month anniversary of this charter's adoption, the Chief Financial Officer commissions an external review of the governance framework, the deployed workflows, and the governance posture. The reviewer is an outside advisor with AI and finance expertise. The review produces a written report delivered to the Audit Committee Chair and the Chief Executive Officer. The review cadence may be adjusted in subsequent years based on the outcome of the first review.

Article VIII: Policy Compliance

All employees are subject to the AI Usage Policy, which derives from this charter and is published separately. Violations are addressed through the standard disciplinary process. Repeated or material violations are reported to the Audit Committee. The Policy is published in employee handbooks and is acknowledged by employees at onboarding and annually thereafter.

Article IX: Amendment

This charter is reviewed annually by the Chief Financial Officer in consultation with the General Counsel and the Head of Security and IT. Amendments may be proposed at any time. Material amendments are presented to the Audit Committee for endorsement before taking effect. Non-material amendments take effect upon adoption by the Chief Financial Officer with notification to the Audit Committee.

Adopted on this thirty-third day of the Chief Financial Officer's tenure.

John Campbell, Chief Financial Officer

Endorsed by Diana Reyes-Okonkwo, Audit Committee Chair

The charter as written is the artifact. Every word matters. The charter is short because it must be referenced quickly. It is declarative because it must be unambiguous. It establishes authority, principles, accountabilities, and review mechanisms. It does not specify operational details, because the operational details are in the framework sections that precede it. The framework is the explanation. The charter is the operating constitution.

A reader adopting this masterclass for their own organization should adapt the charter with minimal edits. The names change. The company changes. The auditor changes. The structure should not. If a section feels removable for a particular organization, the right question is whether that organization is making a deliberate exception that should be documented or merely cutting corners that will come back later.

End of Part 4

The Governance Framework

The governance framework is now complete. The philosophy is articulated. The three pillars are established. The human-in-the-loop architecture is specified. The audit trail requirements are codified. The risk tiers are defined. The vendor and model risk disciplines are documented. The eight-step methodology for embedding governance into workflow design is set forth. The regulatory and external context is mapped. The governance charter is adopted.

In Parts 5 through 9, the masterclass turns to the five use cases themselves. Each part will apply the eight-step methodology established in Section 7 above. Each part will demonstrate, in concrete operational detail, how the framework operates when actual workflows are designed, reviewed, approved, deployed, and monitored. The framework is not background. The framework is the operating reality against which the five use cases unfold.

Part 5 begins with the Finance Operations Copilot, the workflow John Campbell sequenced first because it builds team capacity, operates on clean data, and deploys into the function that owns the governance framework. The reader will watch the eight-step methodology applied for the first time. The framework, established in this part, becomes the lens through which Parts 5 through 9 are read.

Before proceeding, take the assessment that follows. It is longer in scenario weight than the prior assessments because governance is more naturally tested through judgment than through recall. Read each scenario carefully. The questions are calibrated to the level a CFO would actually encounter.

Appendix A · Glossary of Governance Terms

Terms specific to the governance framework. Most are introduced or refined in this part. Terms covered in Parts 1 through 3 are not repeated unless their meaning is extended here.

Audit trail
The immutable record of every AI workflow invocation, capturing input, model version, output, human review, and timestamps. Specified in Section 4. Distinct from system logging.

Behavioral drift
The phenomenon by which a model's behavior changes over time even within a single nominal version, due to vendor-side updates or system-level changes. Monitored under Section 6.

Checkpoint pattern
The architectural decision about where in a workflow the human review occurs and what the reviewer does. Five patterns are recognized in Section 3.

Concentration (vendor)
The degree to which critical workflows depend on a single vendor. Tracked by the Governance Working Group under Section 6.

Data classification
The assignment of data to one of four tiers: public, internal, confidential, restricted. Determines handling requirements under data governance.

Data flow map
The complete documentation of where workflow data originates, where it travels, where it is stored, and which subprocessors touch it. Required artifact under Section 7.

Governance Working Group
The body consisting of the CFO, the General Counsel, and the Head of Security and IT. Approves workflow designs, classifications, and the approved model and vendor list. Established in the charter.

Lawful basis (for processing personal data)
The GDPR concept that personal data processing must rest on one of six lawful bases. Specified for every workflow that processes personal data.

Output specification
The detailed description of what a workflow produces, including format, audience, consequence chain, and frequency. Required artifact under Section 7.

Pillar (governance)
One of three categories under which governance controls are organized: data, model, output. Established in Section 2.

Reviewer (named human)
The specific individual who reviews a workflow output before release. Identified by name, not role. Backed up by a named alternate. Specified under Section 3.

Risk tier
The classification of an AI workflow into one of four tiers based on the consequences of failure. Established in Section 5.

Substitution path
The documented plan for migrating a workflow to an alternative model if the primary model becomes unavailable. Required under Section 6.

Tier One: Internal Exploratory
Workflows producing outputs consumed only by named individuals for analytical or exploratory purposes. Lightest governance.

Tier Two: Internal Operational
Workflows whose outputs influence internal operations but not financial reporting. Moderate governance.

Tier Three: Financial Reporting Adjacent
Workflows whose outputs influence financial reporting, board materials, or external financial communications. Strong governance.

Tier Four: Financial Reporting Direct or High-Risk Regulatory
Workflows that directly enter the financial records or that fall into high-risk regulatory categories. Strongest governance.

Workflow Registry
The central record of every AI workflow deployed at Helix, with all required design artifacts and operational records. Established in the charter.

Appendix B · Assessment

Twenty questions on Part 4. Twelve multiple choice, five short answer, three scenario-based. The scenarios in this assessment are longer and more demanding than in prior parts because governance is most usefully tested through judgment under realistic conditions.

Part I: Multiple Choice

1. The philosophy of AI governance in finance rests on which premise?

(a) Finance AI is a more careful version of AI in other functions.
(b) Finance AI is qualitatively different from AI in other functions and requires qualitatively different rules.
(c) Finance AI should be governed by the same framework that governs traditional financial controls.
(d) Finance AI requires less governance than AI in customer-facing functions because it is internal.

2. The three pillars of AI governance in this framework are:

(a) People, process, technology.
(b) Strategy, execution, measurement.
(c) Data, model, output.
(d) Input, processing, storage.

3. The default human review checkpoint pattern in this framework is:

(a) Review-before-distribution.
(b) Review-before-output.
(c) Sampled review.
(d) Review-after-action with reversal capability.

4. The minimum audit trail includes thirteen fields. Which of the following is NOT one of them?

(a) Invocation ID.
(b) Initiating user.
(c) Model and version.
(d) The reviewer's emotional state at the time of review.

5. Risk Tier Three workflows are:

(a) Workflows that produce outputs consumed only by named individuals for analytical purposes.

(b) Workflows whose outputs influence financial reporting, board materials, or external financial communications.

(c) Workflows that directly enter the financial records.

(d) Workflows that fall into prohibited categories under applicable law.

6. The framework treats risk classification asymmetrically. Specifically:

(a) Workflows may be reclassified up easily but down only after demonstrating sustained operational maturity.

(b) Workflows may be reclassified down easily but up only after material incidents.

(c) All reclassifications require Audit Committee approval.

(d) Reclassification is permitted only at annual review.

7. Which of the following is NOT one of the six dimensions of vendor risk specified in Section 6?

(a) Pricing.

(b) Product changes.

(c) Subprocessor list.

(d) Vendor's market capitalization.

8. The eight-step methodology for embedding governance into workflow design begins with:

(a) Designing the technical architecture.

(b) Articulating the business problem.

(c) Selecting the model.

(d) Defining the audit trail.

9. Which design failure occurs when designers propose the workflow architecture first and ask where to insert human review afterward?

(a) Failure One: Output specification by example.

(b) Failure Two: Review checkpoint as afterthought.

(c) Failure Three: Audit trail as logging.

(d) Failure Five: Risk classification by self-assessment.

10. Helix is private. The framework treats SOX-style controls as:

(a) Inapplicable, since SOX does not legally apply.

(b) The design standard, on the grounds that retrofitting controls later is more expensive.

(c) Optional, to be applied only if the company plans an IPO.

(d) A regulatory burden to be minimized.

11. The EU AI Act is most directly relevant to Helix because:

(a) Helix is headquartered in the European Union.

(b) The Act's reach includes systems placed on the market or put into service in the Union, and Helix has European customers.

(c) The Act applies extraterritorially to all SaaS companies.

(d) The Act prohibits most enterprise AI deployments.

12. The Governance Working Group consists of:

(a) The CEO, the CFO, and the CTO.

(b) The CFO, the General Counsel, and the Head of Security and IT.

(c) The CFO, the Controller, and the Senior FP&A; Manager.

(d) The Audit Committee, the CFO, and the external auditor.

Part II: Short Answer

13. In two or three sentences, explain why the framework asserts that AI in finance is qualitatively different from AI in marketing, customer service, or product. What is the practical implication of this assertion for the governance framework?

14. The framework distinguishes carefully between governance, checklists, approval processes, and operational reality. In two or three sentences, explain why the framework warns against treating governance as primarily an approval process.

15. Sampled review is recognized as a valid checkpoint pattern but is treated as exceptional and requires explicit CFO approval. In two or three sentences, explain the reasoning behind this caution and what would justify approving sampled review for a specific workflow.

16. The framework requires that risk reclassification upward be easy and reclassification downward be difficult. In two or three sentences, explain why this asymmetry is deliberate and what it accomplishes.

17. The eight-step methodology in Section 7 produces a defined set of design artifacts that must exist before deployment. In two or three sentences, explain why the framework requires the artifacts to exist in writing rather than as informal understandings among the workflow team.

Part III: Scenario-Based

18. Scenario: A workflow owner submits a Tier Three workflow design for governance working group approval. The data flow map is complete, the output specification is detailed, the audit trail design is documented, and the human review pattern is specified as Review-before-output with the workflow owner herself as the reviewer. She argues that, because she designed the workflow and understands it most deeply, she is the most qualified reviewer. As a member of the governance working group, in one paragraph of executive prose, describe how you would respond to this proposal, what governance principle is at stake, and what alternative arrangement you would propose.

19. Scenario: Twelve months into the AI program, the audit committee chair informs you that the external review (mandated by Article VII of the governance charter) has been completed by an outside advisor. The advisor's report praises the framework architecture but identifies three operational gaps: (a) the data flow maps for two of the deployed workflows have not been updated since deployment, despite material data source changes; (b) the model substitution paths for all five deployed workflows are documented but have never been tested in practice; (c) the audit committee briefings have included status summaries but no analysis of behavioral drift in the deployed models. In one paragraph, describe how you would respond to the audit committee, what remediation you would commit to, and what changes to the framework or its operation you would propose to prevent recurrence.

20. Scenario: Your company is acquiring a smaller competitor in a stock-and-cash transaction. As part of integration planning, you learn that the acquired company has deployed three AI workflows in its finance function with no governance framework, no risk classification, no audit trail, and ad-hoc vendor relationships. The acquired company's CFO is staying on for a transition period and is defensive about the workflows, noting that they have produced no incidents and the company has had clean audits. You will become responsible for the combined finance function within sixty days. In one paragraph of executive prose, describe how you would handle the AI governance integration during the transition, what posture you would take toward the acquired CFO's workflows, and what risk you would most want to mitigate before integration completes.

Appendix C · Answer Key with Explanations

The scenario answers in this assessment are particularly extensive because governance scenarios test executive judgment under realistic complexity rather than recall of specific framework provisions.

Multiple Choice Answers

Question 1: (b)

The framework rests on the premise that AI in finance is qualitatively different from AI in other functions and therefore requires qualitatively different rules. The standard of care in finance is categorically distinct because the consequences of failure are categorically distinct. See Section 1.

Question 2: (c)

The three pillars are data, model, and output. Each pillar addresses a distinct phase in the life of an AI workflow and has its own controls and failure modes. See Section 2.

Question 3: (b)

Review-before-output is the default checkpoint pattern. Any deviation requires explicit design rationale documented in the workflow registry. Sampled review in particular requires CFO approval. See Section 3.

Question 4: (d)

The reviewer's emotional state is not part of any reasonable audit trail. The minimum trail includes the invocation ID, workflow name and version, initiating user, timestamps, inputs, model and version, prompt or workflow definition, raw output, human reviewer, reviewed output, approval timestamp, downstream actions, and error events. See Section 4.

Question 5: (b)

Tier Three covers workflows whose outputs influence financial reporting, board materials, or external financial communications, even if they do not directly post entries to the books. The board reporting workflow, the forecasting workflow, and the pipeline intelligence workflow are examples. See Section 5.

Question 6: (a)

The framework deliberately makes upward reclassification easy and downward reclassification difficult. The asymmetry reflects the fact that the cost of under-controlling exceeds the cost of over-controlling. See Section 5.

Question 7: (d)

The vendor's market capitalization is not a defined dimension of vendor risk in this framework. The six dimensions are pricing, product changes, data handling, security posture, subprocessor list, and termination. See Section 6.

Question 8: (b)

The methodology begins with articulating the business problem. Workflows that cannot articulate a clear business problem do not proceed. See Section 7.

Question 9: (b)

Failure Two is the review checkpoint as afterthought. The right order is to design the review checkpoint alongside the workflow as a constraint that shapes it. See Section 7.

Question 10: (b)

Helix treats SOX-style controls as the design standard even though SOX does not currently apply. The reasoning is that retrofitting controls when a public offering is contemplated or when an institutional investor conducts diligence is more expensive than designing them properly from the start. See Section 8.

Question 11: (b)

The EU AI Act applies to systems placed on the market or put into service in the Union. Because Helix has European customers, its AI deployments may have indirect reach into the Union and therefore must be evaluated against the Act's classifications. See Section 8.

Question 12: (b)

The Governance Working Group consists of the CFO, the General Counsel, and the Head of Security and IT. The composition reflects the cross-functional nature of governance. See Article III of the charter.

Short Answer Explanations

13. Why AI in finance is qualitatively different

Finance produces outputs whose failure modes are not merely embarrassing but categorically more serious: financial misstatement, control deficiencies, audit issues, regulatory action, and litigation. A hallucinated marketing email is recoverable. A hallucinated board commentary or revenue recognition entry is not, in the same sense. The practical implication is that the governance framework cannot be a more careful version of generic AI controls. It must be designed from the premise that the standard of care is categorically different, which justifies the four-tier risk classification, the human-in-the-loop default, the immutable audit trail requirement, and the conservative reclassification asymmetry.

14. Why governance is not primarily approval

A framework whose principal function is to approve or refuse workflows reduces governance to a checkpoint and conflates speed of approval with quality of judgment. The framework warns against this because approval is downstream of the substantive work, which is the discipline of designing the workflow correctly in the first place. Governance is the system that determines what good design looks like, how it is evaluated, and what controls operate during the workflow's life. A well-designed framework approves quickly the workflows that meet the discipline and refuses entirely those that do not, but the approval itself is the smallest part of the governance work.

15. The caution against sampled review

Sampled review is operationally efficient but is the most easily abused checkpoint pattern. The risk is that the sampling rate becomes too sparse to catch errors that occur at low frequency but with high consequence, or that the sampling becomes routine in a way that erodes the reviewer's vigilance. Approval is justified only when the workflow has demonstrated consistent quality over multiple cycles, when individual error consequences are small, when the sampling methodology is statistically valid, and when the workflow has accumulated enough operational history that meaningful sampling is possible. The requirement of CFO approval is itself a discipline that prevents drift into sampled review by default.

16. The reclassification asymmetry

The asymmetry exists because the costs of under-controlling and over-controlling are themselves asymmetric. Over-controlling a workflow that turns out to be low-risk produces incremental review work that is annoying but recoverable. Under-controlling a workflow that turns out to be high-risk produces failures that may be material and may be unrecoverable. Making reclassification upward easy preserves the option to add controls quickly when new information emerges. Making reclassification downward difficult prevents the gradual erosion of controls that occurs when workflows that have been operating without incident accumulate pressure to be reclassified to lower tiers. The asymmetry is a structural defense against the natural human tendency to relax discipline over time.

17. Why design artifacts must exist in writing

Written artifacts are durable, transferable, and verifiable in ways that informal understandings are not. A workflow designed on the basis of verbal agreements among the original team becomes opaque the moment a team member leaves, becomes contested the moment a disagreement arises, and becomes indefensible the moment an external party asks for evidence of the design. Written artifacts also force the design to be explicit, which surfaces gaps that informal discussion conceals. The discipline of writing is itself the discipline of designing rigorously. A workflow whose design cannot be expressed in writing is, almost always, a workflow whose design is incomplete.

Scenario Discussions

18. The workflow owner as her own reviewer

The right response is to decline the proposal politely and firmly, while acknowledging the owner's substantive expertise. The governance principle at stake is the separation of preparer and reviewer, which is one of the most established disciplines in financial controls and which applies with at least equal force to AI workflows that produce financial outputs. The workflow owner's familiarity with the workflow is precisely what makes her unsuitable as the reviewer: she is too close to her own design to detect the systematic blind spots that an independent reviewer would catch. The alternative arrangement should identify a reviewer with sufficient subject-matter expertise to perform meaningful review but with sufficient independence from the workflow design that they bring fresh judgment. For a Tier Three workflow, the reviewer should be a finance team member with finance authority, distinct from the workflow owner, with named backup coverage. The workflow owner remains the designer and the operational accountability, but the review responsibility is separated. The governance working group should make this requirement explicit in the workflow registry record before approving the workflow for deployment.

19. The external review findings

The right response is to take the findings seriously, acknowledge the gaps fully, and treat the report as the validation that the external review process is working as intended. The first remediation is operational: bring all data flow maps current within sixty days, with named ownership for ongoing maintenance, and establish a quarterly review cadence that prevents future drift. The second remediation is also operational: schedule and complete substitution path tests for all five deployed workflows within ninety days, with a written report for each documenting the actual rework, time, and any issues encountered. The third remediation is more substantive: update the audit committee briefing template to include a quantitative behavioral drift analysis for each Tier Three workflow, with monthly comparison against the deployment baseline. The framework changes I would propose are two. First, the data flow map should be a living artifact with named owners and a mandatory review at every material data source change, not merely at annual review. Second, substitution path testing should be moved from an annual cadence to a semi-annual cadence, recognizing that the discipline of testing is what makes the substitution path real. The deeper lesson is that the framework was sound in design but operational discipline drifted in the absence of structural enforcement. The external review caught what internal review did not. That is precisely what the external review is for, and the audit committee should be thanked for insisting on it.

20. The acquisition integration

The right approach is firm but graduated, recognizing that the acquired CFO's defensiveness is partly emotional and partly substantive. In the immediate term, before the transaction closes, demand a full inventory of the three workflows with the same artifacts our framework requires: business problem, output specification, data flow map, risk classification, review design, audit trail, and substitution path. If the artifacts do not exist, they must be constructed retroactively from operational reality. The acquired CFO can lead this work during the transition with my support. The posture toward the workflows is not that they must be shut down but that they must be brought under governance; clean audits and absence of incidents are favorable evidence but not a substitute for the discipline of documentation. The risk I would most want to mitigate is the inheritance of liability without the documentation needed to defend it. Once we close the transaction, the responsibility for the workflows becomes ours, including the responsibility for any historical errors that may not yet have surfaced. Without the artifacts, we have inherited risk we cannot bound. I would set a hard requirement that all three acquired workflows be either documented to the framework standard within ninety days of close or be suspended pending documentation. I would communicate this requirement to the acquired CFO during the transition, in writing, so the expectation is clear. The integration team should also conduct a discovery process within the acquired company similar to the thirty-day conversations I conducted at Helix, to surface any shadow AI usage that the formal inventory may have missed. The principle is that we do not assume the acquired finance function meets our governance standard; we verify, document, and remediate where needed.

