

BOARD OF DIRECTORS — RESTRICTED CIRCULATION

The Coordination Economy

AI, Coordination Debt, and the Case for Enterprise Governability

AI does not create coordination debt. It exposes the debt already embedded in the enterprise—and when deployed without governing conditions, compounds faster than the organization can absorb.

Executive Summary

Artificial intelligence is changing more than how work is performed. It is changing how work is authorized, handed off, interpreted, escalated, and accepted. In other words, AI is entering the coordination system of the enterprise.

That matters because most organizations already carry a large, largely unmeasured burden of coordination debt: accumulated ambiguity about who may decide, what counts as sufficient evidence, when exceptions must escalate, which commitments are binding, and how prior outcomes should shape future permissions. Human managers have historically absorbed this debt through meetings, judgment, workarounds, and institutional memory.

AI makes that hidden architecture visible. It can compress routine coordination dramatically. It can also accelerate unresolved ambiguity. When autonomous systems act across functions faster than decision rights, evidence standards, accountability, and memory have been made explicit, the organization does not become more governable. It becomes faster at producing drift.

The board's task is therefore larger than supervising AI risk. It is to determine whether the enterprise can continue behaving according to its intentions as intelligence, autonomy, and complexity increase. That capability is enterprise governability.

This briefing advances four propositions: coordination is an economic phenomenon; coordination debt predates AI; AI exposes and compounds that debt; and boards must establish the governing conditions under which human and machine agents may act. A constitutional model is one practical means of doing so, but the strategic objective is broader: an enterprise that can move faster without losing coherence.

SECTION 1

The Coordination Economy

The enterprise is not merely a collection of assets, functions, and people. It is a system of commitments. Someone promises to do something, to a standard, by a time, so someone else can take the next step. Revenue is the visible result. Coordination is the machinery beneath it.

Meetings, approvals, status reports, escalations, reconciliations, and handoffs are not peripheral to production. They are the means by which distributed intention becomes coherent action. The economic question is not whether coordination exists. It is how much coordination an organization must consume to produce a dependable outcome.

As enterprises add products, markets, regulations, technologies, and specialized roles, the number of dependencies grows. The organization may become more capable in each part while becoming less coherent as a whole. Exceptions multiply. Decisions travel farther. Evidence is recreated. Managers become translators between systems that no longer share the same context.

The operating cycle is the economic shadow of the commitment system.

Inventory, receivables, buffers, rework, expedited freight, management layers, and delay are often treated as separate operational or financial issues. Seen through a governability lens, many are buffers against unreliable coordination. They are the cost of uncertainty about whether commitments will be fulfilled, evidenced, accepted, and remembered in time for the next action to occur.

Coordination Debt

Coordination debt is the accumulated burden created when an organization repeatedly resolves ambiguity through effort rather than design. It appears whenever people must rediscover authority, reconstruct context, reconcile conflicting records, chase evidence, or negotiate the same exception again because the enterprise has not converted experience into governing memory.

It is not necessarily the product of poor leadership. Often it is the residue of growth. A process that worked when ten people shared context becomes fragile when hundreds of people, vendors, systems, and agents must coordinate across boundaries. The enterprise continues to function because capable people compensate. Their compensating effort hides the debt.

Where the debt becomes visible

Decision latency: Work waits because authority is unclear or distributed across too many approval points.

Reconciliation burden: Teams spend time determining which record, interpretation, or status is authoritative.

Exception recurrence: The same class of problem returns because resolution is not converted into institutional memory.

Expediting and buffering: Inventory, cash, labor, and managerial attention are held in reserve against unreliable coordination.

Outcome drift: Individually reasonable actions gradually diverge from enterprise intent.

The Hidden Cost of Coordination

In large enterprises, internal coordination — meetings, approvals, status updates, escalations, and alignment activities — consumes an estimated **30 to 50 percent of available productive capacity**. This overhead does not surface in traditional financial reporting. It is visible only in extended cycle times, elevated management-to-contributor ratios, and the accumulated cost of decisions that were delayed, diluted, or reversed for want of timely information. For a 10,000-person knowledge-work organization, even a conservative 35 percent coordination tax represents the equivalent of thousands of years of full-time roles consumed by coordination activity rather than direct production — a structural liability that compounds as organizations grow.

The coordination economy is not a metaphor. It is the operating reality of every enterprise that has scaled beyond the boundaries of direct managerial oversight. Competitive advantage, in this environment, accrues to organizations with the governance architecture and technology infrastructure to coordinate at lower cost, higher speed, and greater accuracy than their peers.

SECTION 2

AI and the Compression of Coordination Burden

Most discussion of enterprise AI begins with intelligence: better analysis, faster drafting, stronger prediction, and more automation. Those benefits are real. But the more consequential change is that AI can now participate directly in coordination.

An agent can receive an objective, access tools, interpret evidence, route work, invoke another agent, make a bounded decision, and continue until an outcome is produced. Each step is a handoff. Each handoff is a commitment. Each commitment creates the possibility of ambiguity, failure, or drift.

What AI can compress

Coordination activity	Potential compression
Status and evidence aggregation	Continuous monitoring can replace periodic collection and reporting.

Coordination activity	Potential compression
Drafting and synthesis	AI can create reviewable outputs from distributed inputs in seconds rather than days.
Routing and triage	Rules and risk thresholds can move routine cases without managerial intervention.
Cross-functional coordination	Agents can operate across systems and functions without waiting for synchronized human availability.

The compression paradox

AI can reduce coordination effort at the operating layer while increasing governability requirements at the enterprise layer. A system that makes decisions, invokes tools, and interacts with customers or regulators removes friction from execution—but it also creates more consequential actions that must remain aligned with enterprise intent.

This is the compression paradox: the easier it becomes to act, the more important it becomes to establish the conditions under which action is legitimate. Speed does not eliminate governance. It moves governance upstream.

AI does not merely automate the organization. It tests whether the organization knows what it means.

SECTION 3

AI Exposes the Debt Already There

Software engineers have long recognized the concept of technical debt: the accumulated cost of expedient architectural shortcuts that defer proper design work. When development teams prioritize speed over structure, they create a hidden liability — code that functions in the short run but becomes increasingly fragile, expensive to maintain, and prone to failure as it scales. The parallel to enterprise governability is instructive.

Coordination debt does not begin when an AI system is deployed. It begins when the enterprise relies on unwritten judgment, informal authority, fragmented evidence, and heroic intervention to remain coherent. AI reveals those dependencies because software cannot reliably operate on rules the organization has never made explicit.

A human manager can sense that a policy should not be applied literally in a particular situation. An agent needs a threshold, an exception class, an authorized override, or a route to human judgment. A veteran employee may know which data source is trusted. An agent needs an evidence hierarchy. A leader may intuit who owns a decision. An agent needs decision rights that can be invoked and audited.

When those conditions are absent, organizations often respond by adding prompts, approvals, monitoring layers, and review committees. These may contain immediate risk, but they do not retire the underlying debt. They move the burden elsewhere.

How AI compounds coordination debt

Interaction volume: Agentic systems multiply the number and speed of handoffs. Small ambiguities become repeated at machine scale.

Error propagation: One weak output becomes another agent's input, allowing uncertainty to travel farther before a human notices.

Authority diffusion: Actions occur across systems and functions faster than conventional accountability structures can assign ownership.

Memory fragmentation: Logs accumulate without becoming usable institutional memory that changes future permissions or behavior.

Oversight latency: Boards and executives receive summaries after autonomous actions have already cascaded through the operating system.

Three Early Indicators of Coordination Debt

1. Shadow AI is a demand signal, not merely a policy violation. Employees often reach for unsanctioned tools because authorized enterprise capabilities cannot meet the speed or usability their work requires. The security exposure is real, but the deeper signal is a governability gap between operational need and legitimate decision rights. The board should ask whether sanctioned systems can satisfy the needs that are driving work outside established boundaries.

2. Agentic pilots stall at the boundary between capability and legitimate authority. A pilot may demonstrate technical success yet remain unable to scale because ownership, exception handling, evidence requirements, and human override remain unresolved. What is often called pilot purgatory is more precisely a governability gap: the technology can act, but the enterprise has not yet established the conditions under which that action is legitimate. The board should ask whether stalled pilots are being treated as technical problems when the real constraint is unresolved authority, oversight, and permissions.

3. Accountability becomes visible only after failure. When several functions, vendors, models, and agents contribute to an outcome, traditional responsibility charts are insufficient. The deeper signal is a governability gap in which commitments, evidence, permissions, and consequences cannot be traced clearly enough to preserve accountability across the full chain of action. The board should ask whether the enterprise can reconstruct who or what was authorized to act, on what evidence, within which boundaries, and with what consequence before a regulatory inquiry, consequential error, or data breach forces the question.

The Regulatory and Legal Dimensions

Coordination debt becomes a board concern not only when it impairs performance, but when ambiguity about authority, oversight, and accountability acquires legal or regulatory consequence. Coordination debt is no longer a theoretical governance risk. It has acquired regulatory and legal form. The EU AI Act

is now in enforcement — its prohibited use prohibitions active and its high-risk system requirements on a phased implementation schedule toward full applicability in August 2027. Penalties reach up to 7 percent of global annual turnover. For a large multinational, that is a material financial exposure arising directly from governance gaps that a well-structured oversight framework would have identified and remediated.

Equally consequential is the board-level legal dimension. The Caremark duty — the Delaware standard holding that directors have affirmative oversight obligations for material enterprise risks — has been interpreted in recent years to extend to technology and cybersecurity failures that the board failed to monitor. AI governance gaps that produce material harm, in the presence of clear warning signs and available remediation frameworks, create conditions that plaintiff counsel will recognize. Directors cannot effectively delegate oversight of AI risk to IT or legal functions and consider the governance obligation discharged.

The analogy to financial leverage is instructive. Like debt, coordination debt accelerates performance in the short run: faster deployment, faster value realization, competitive advantage. Also like debt, it creates fragility precisely when the organization can least afford it — in a stress event, a regulatory audit, a consequential system failure. And like financial leverage, it compounds: each unreviewed deployment adds not linearly but multiplicatively to the governance exposure, as ungoverned systems interact with one another in ways that no individual deployment decision anticipated.

SECTION 4

Enterprise Governability

Governability is an enterprise's ability to continue behaving according to its intentions as complexity increases.

This definition shifts the board's attention from whether the organization has more intelligence to whether intelligence can be translated into coherent behavior. An enterprise may have excellent strategy, advanced technology, talented people, and strong controls — and still become progressively harder to steer.

Governability is not synonymous with compliance, efficiency, or centralized control. A governable enterprise may be highly distributed and adaptive. What distinguishes it is that autonomy occurs inside conditions that preserve identity, accountability, and continuity of intent.

The constitutional model

One method for producing enterprise governability is constitutional design: establishing binding conditions before work begins rather than attempting to correct every behavior after it occurs. The constitutional model is not the category. It is an architecture for achieving the category.

A useful constitutional system answers five recurring questions:

Commitment: What has an authorized agent promised to do, for whom, to what standard, and by when?

Evidence: What must be observable before the enterprise treats the commitment as fulfilled?

Permissions: What may this agent do now, given authority, readiness, risk, and prior behavior?

Consequence: What follows when a commitment is fulfilled, disputed, delayed, or broken?

Memory: What must the enterprise remember so that the same event changes future decisions?

These are not AI controls. They are the primitive conditions from which coordinated behavior emerges. They apply to people, systems, vendors, autonomous agents, and the handoffs between them. Their purpose is to make coherent behavior more likely before execution begins.

From compliance activity to durable control

When organizations feel exposed on AI governance, the instinctive response is to add more process — new committees, longer checklists, extra sign-off steps. It is understandable and almost entirely ineffective. Layering reviews and approvals on top of AI systems that were not properly set up to be overseen produces what might be called governance theater — compliance activity that looks responsible in an audit but provides no real oversight of the systems and decisions it purports to govern.

The durable solution requires a different foundation. **Constitutional governability** is the practice of setting clear rules *before* AI systems go live — analogous to how a constitution pre-commits an institution to inviolable principles that sit above day-to-day decisions and cannot easily be brushed aside when it is inconvenient. Constitutional governability does not govern AI after incidents. It establishes the operating boundaries before deployment, creating a framework within which AI can be used confidently, boldly, and at scale.

The constitutional analogy is not decorative. Constitutions work because they are difficult to circumvent and derive authority from a source above the ordinary political process. An AI governance framework that lives in an IT handbook, maintained by a working group of middle managers, has no constitutional character. One ratified at the board level, with defined consequences and active oversight, does. The governance is only as strong as the authority from which it derives.

Governance theater versus governing conditions

The distinction between governance theater and constitutional governability is not a matter of intent. Organizations that have invested in AI ethics principles, responsible AI frameworks, and review committees are genuinely trying to do the right thing. The failure is architectural: a framework that lacks board-level backing, enforcement mechanisms, and integration with the way AI is actually deployed cannot govern at the speed and complexity of modern AI. It produces documentation. It does not produce accountability.

Governance theater	Governing conditions
Principles stated after deployment	Rules established before deployment
Committees without operational authority	Decision rights that systems and people must obey
Human oversight as an aspiration	Explicit thresholds for review, override, and escalation
Logs retained for audit	Memory that changes future permissions and behavior
Compliance evidence assembled periodically	Evidence generated as part of execution

Governance theater	Governing conditions
Responsibility distributed broadly	Commitments assigned to identifiable agents

Constitutional governability is characterized by three properties that separate real governance from the appearance of it:

- **Rules set in advance.** Boundaries are established before deployment, not derived from incidents after things go wrong.
- **Oversight calibrated to risk tier.** The level of scrutiny applied to a given system is proportional to how much damage it could do, not applied as a blanket policy or aspirationally.
- **Authority must be real.** The framework is ratified at the board level so it carries genuine organizational obligation rather than residing in a policy document that operating systems and personnel are free to disregard.

The Elements of a Constitutional Governance Framework

Element	Description	Board Role
AI Decision-Rights Charter	A published document establishing which categories of decisions AI may make autonomously, which require human review, and which are prohibited regardless of capability.	Formal ratification; periodic review
Risk-Tiered Mandatory Human Judgment Thresholds	Defined, non-aspirational thresholds specifying when human oversight is mandatory based on risk classification — not general guidance subject to operational interpretation.	Approve risk tier definitions and override protocols
Standing AI Oversight Committee	A cross-functional committee with defined authority, board-level reporting lines, and a standing mandate — not a project committee that dissolves after initial deployment.	Receive quarterly reports as standing agenda item
Mandatory Deployment Inventory	A complete, maintained registry of every AI system in production, mapped to its risk classification, oversight controls, owner, and regulatory applicability.	Treat inventory completeness as a fiduciary baseline metric
Regulatory Calendar Alignment	A phased remediation roadmap aligned to EU AI Act obligations and other applicable jurisdictional requirements, with milestone accountability.	Monitor compliance milestones; approve remediation resourcing

The Competitive Advantage Case

Enterprise governability should not be framed as a brake on innovation or constraint on AI deployment. It is what allows an organization to move quickly without relying on optimism, heroics, or repeated executive intervention.

An enterprise with clear decision rights, evidence standards, exception thresholds, and institutional memory can authorize consequential AI use with greater confidence than peers operating without it. The limiting factor on AI deployment in most organizations is not technical capability or cost. It is organizational risk tolerance, shaped by the absence of a framework that would allow confident action. With that, the enterprise can distribute autonomy because it has designed the conditions that make autonomy safe. It does not need to invent governance anew for each deployment.

The less governable enterprise faces the opposite pattern. Each new capability requires more negotiation, more review, more reconciliation, and more executive attention. It may possess more intelligence yet be unable to convert that intelligence into dependable behavior.

Put simply: enterprise governability is what allows an organization to move fast without losing coherence. Constitutional design is one architecture for building that capability.

The strategic advantage will not belong only to the enterprise with the best AI. It will belong to the enterprise that can absorb more intelligence without losing coherence.

What the board should expect to become measurable

Governed deployment coverage: The proportion of consequential AI-enabled workflows operating within explicit decision rights, evidence, oversight, and memory conditions.

Exception recurrence: How often the same class of exception reappears after purported resolution.

Decision latency: The time lost between the need for a decision and authorized action.

Reconciliation burden: The labor and delay required to establish a shared account of what happened.

Human intervention rate: Where and how frequently people must rescue, reinterpret, or override autonomous workflows.

Coordination exposure: The economic value placed at risk by unresolved commitments, delayed acceptance, or fragmented accountability.

Five Principles of Constitutional Governability

1. Pre-commitment over post-incident response. Governance frameworks established before deployment govern effectively. Frameworks assembled after incidents govern defensively — and at far greater cost.

2. Proportional oversight by risk tier. Not all AI systems warrant the same scrutiny. The governance architecture should be calibrated to actual risk — rigorous for high-stakes autonomous systems, streamlined for low-risk productivity tools.

3. Full deployment inventory as a fiduciary baseline. A board cannot govern what it cannot see. The deployment inventory is not an IT asset management exercise — it is a fiduciary prerequisite. Its absence is a governance failure, not a gap.

4. Mandatory human judgment thresholds, not aspirational goals. Oversight thresholds must be defined, enforceable, and binding. "We aim to have human review for significant decisions" is not a threshold — it is an aspiration. Constitutional governability requires the former.

5. Board-level accountability, not delegation to IT. AI governance is a board-level responsibility. Delegating it entirely to technology or legal functions without board oversight, reporting structures, and ratified authority produces organizational structures that cannot bear the weight of a material AI failure.

SECTION 5

Recommended Board Actions

The board does not need to govern every model, prompt, or workflow. It must govern the conditions under which the enterprise delegates consequential action. This is a constitutional responsibility in the practical sense: determining what may not be left to improvisation.

The board should therefore not ask only, "What AI systems do we have?" but also, "What commitments are these systems authorized to make, what evidence must support their actions, how are exceptions handled, and how does the enterprise learn from the outcomes?"

The following actions are specific, sequenced, and within the board's direct authority. They are ordered by urgency and interdependency — items one through three establish the foundation upon which items four through six operate.

Building Enterprise Governability: A Board Sequence

#	Board action	Why it matters	Timing
1	Map the coordination system, not only the technology inventory. Identify consequential AI-enabled workflows and map their human and machine participants, decision rights, handoffs, evidence requirements, existing controls, and exception paths. Report the gap between deployed capability and governed capability as a standing board-level risk indicator.	A model inventory tells the board <i>what exists</i> . A coordination map reveals <i>how consequential work actually moves through the enterprise</i> , where ambiguity accumulates, and where coordination debt is likely to emerge.	Immediate

#	Board action	Why it matters	Timing
2	Ratify the enterprise's non-delegable governing conditions. Define which decisions require human judgment, which actions AI may perform autonomously, which remain prohibited regardless of capability, and the governing principles that operational urgency cannot override.	Governability begins by establishing the conditions under which legitimate action may occur. These conditions should remain stable even as technologies, organizational structures, and operating models evolve.	Immediate
3	Establish an enterprise decision-rights and permissions architecture. Specify who—or what—may act, within which boundaries, under which risk tiers, with what evidence, and through which review, escalation, and override mechanisms. Assign ongoing stewardship through a cross-functional governance body reporting regularly to the board.	Autonomous systems cannot exercise judgment responsibly unless authority, permissions, and accountability are explicit, traceable, and consistently applied across the enterprise.	Near term
4	Require evidence at the point of consequential action. Ensure that decisions, approvals, overrides, and autonomous actions generate the evidence needed for acceptance, auditability, dispute resolution, regulatory compliance, and organizational learning.	Governability depends upon observable evidence, not retrospective reconstruction. Evidence created during execution reduces ambiguity, strengthens accountability, and lowers the cost of future coordination.	Near term
5	Convert operational experience into governing memory. Require material failures, overrides, disputes, recurring exceptions, and significant interventions to change future permissions, thresholds, workflow design, or governing conditions.	Logging records history. Governing memory changes future behavior. Organizations become more governable only when experience permanently reduces future coordination burden.	Ongoing
6	Treat enterprise governability as a board-level performance dimension. Monitor governed deployment coverage, decision latency, reconciliation burden, exception recurrence, human intervention rates, coordination exposure, and the gap between	The board cannot improve what it does not observe. Governability should become as visible as liquidity, cybersecurity, operational resilience, and financial performance because it increasingly determines how effectively the enterprise converts intelligence into coordinated action.	Ongoing

#	Board action	Why it matters	Timing
	deployed capability and governed capability alongside traditional financial and operational measures.		

CLOSING

A Defining Strategic Moment

The transition underway is not simply from human work to automated work. It is from organizations in which humans informally carry the burden of coordination to organizations in which humans and machine agents must operate inside a shared governing environment.

That transition will expose weaknesses that were previously survivable: ambiguous authority, fragmented truth, undocumented exceptions, and institutional memory that resides in people rather than systems. AI will not create those weaknesses. It will reveal them — and, unless the enterprise redesigns the conditions of coordination, amplify them.

The board's responsibility is not to prescribe the technology. It is to ensure that the organization remains itself as technology changes how action occurs. That means protecting the continuity between purpose, intent, commitment, evidence, permissions, execution, acceptance, memory, and economic result.

The central question is therefore:

How much intelligence and autonomy can this enterprise absorb while continuing to behave according to its intentions?

That is the central question of enterprise governability. It belongs on the board agenda now.

DISCUSSION QUESTIONS FOR THE BOARD SESSION

1. Do we have sufficient visibility into the consequential workflows in which AI now participates — and do we know what commitments those systems may make, what evidence their actions must produce, and where human judgment remains mandatory?

This question tests inventory completeness and the practical enforceability of existing oversight thresholds. If the answer is uncertain, the deployment inventory commission is the immediate priority.

2. Have we established binding governing conditions before deployment, or are we relying on policies, committees, and retrospective review to contain ambiguity after consequential action has already occurred?

This question distinguishes genuine oversight architecture from governance theater, and surfaces whether the board's existing committee structures and reporting lines are adequate to the task.

3. What evidence would tell us that the enterprise is becoming more — or less — governable as AI adoption increases?

This question connects governance posture to strategic intent, and creates the internal accountability standard against which the board can evaluate management's deployment decisions going forward.