

CONCORDAT

The Five Primitives of Execution Governance

A Theoretical Framework for Governing Organizational Execution

This paper establishes the theoretical foundation of Execution Governance as an organizational control discipline. It defines the five primitives — the necessary and sufficient constructs for governing whether organizational execution is progressing correctly across distributed systems, teams, processes, and actors. Each primitive is grounded in established management and organizational theory, defined as an independent governance construct, and argued as structurally necessary to a complete governance model. Together the five primitives constitute the Concordat Operating Model: the philosophical and theoretical basis upon which Concordat Beacon, the first Execution Governance platform, is built.

1. The Theoretical Basis for Execution Governance

1.1 Governance Was Always Assumed

For most of organizational history, execution governance was implicit. Management layers existed to observe work, make decisions, assign accountability, identify problems, and evaluate results. The organizational structure itself was the governance mechanism — hierarchy provided visibility, authority provided accountability, and proximity provided the situational awareness needed to detect when execution was going wrong.

This model assumed that the work being governed was visible, that the people doing it were within reach, and that the boundaries of organizational activity were knowable. Those assumptions held in a world where operational work was concentrated, largely manual, and conducted within a relatively contained organizational environment. Governance was not a designed system — it was an emergent property of how organizations were structured.

Organizational control theory formalized these implicit mechanisms. Anthony (1965) identified three distinct control layers — strategic planning, management control, and operational control — each requiring different information and different governance instruments. The management control systems literature described how organizations use budgets, performance metrics, oversight structures, and cultural norms to direct and evaluate behavior (Malmi & Brown, 2008). These frameworks assumed that execution, however complex, remained observable by the humans responsible for governing it.

1.2 The Assumption Broke

The assumption broke gradually, then completely. As organizations grew and specialized, software tools emerged to handle discrete operational tasks more efficiently. Each tool addressed a real problem — managing customer relationships, processing

payments, scheduling resources, routing workflows, storing documents. Each made the organization faster and more capable within its domain.

But each tool also made an implicit assumption: that the governance surrounding the work it handled was addressed elsewhere. The CRM assumed someone was ensuring leads were followed up correctly. The payment platform assumed someone was confirming transactions connected to the right operational context. The scheduling system assumed someone was verifying that appointments reflected actual organizational readiness. No tool was wrong to make this assumption at the individual level. The catastrophic failure was that as tools proliferated, the assumption compounded — governance got distributed across an expanding surface area of systems, teams, and processes until no human layer and no single system retained the organizational visibility required to govern execution as a whole.

The result is the condition that defines modern organizational operations: tool sprawl has created an execution environment where work happens across dozens of systems simultaneously, accountability is distributed across multiple teams and roles, decisions are made in the seams between tools, and no layer exists with the authority, visibility, or architecture to evaluate whether execution across all of it is unfolding correctly. Governance was not eliminated. It was assumed to exist in a layer that was never built.

1.3 Why Existing Categories Cannot Close the Gap

The organizational response to this condition has been to deploy more tools — workflow orchestration to automate handoffs, business intelligence to surface outcomes, process mining to reconstruct what happened, project management to track tasks. Each

category addresses a real symptom. None addresses the underlying governance problem, because none was designed to.

Every incumbent category governs from within its own execution boundary. A business process management engine governs the workflows it executes, not the work happening in systems outside its runtime. A process mining platform infers process behavior from event logs it can observe, but cannot govern work that leaves no digital trace or crosses system boundaries without a shared data model. An analytics platform measures outcomes after they have occurred, but cannot evaluate whether execution is conforming to organizational intent while work is still in progress. An orchestration platform coordinates services it invokes, but cannot govern human judgment, manual handoffs, or organizational decisions made outside its workflow.

This is not a feature limitation. It is an architectural constraint imposed by the design intent of each category. None of them takes organizational execution — across all systems, actors, and processes simultaneously — as its primary object of governance. The gap is not between tools. The gap is above them.

What is required is a governance layer that sits above the tool layer entirely — one that does not execute work, orchestrate workflows, or analyze outcomes after the fact, but continuously evaluates whether execution across the full organizational environment is conforming to a declared standard. This is the control problem that Execution Governance is designed to solve. The five primitives define what solving it requires.

2. The Five Primitives

The five primitives are derived from a first-principles question: what is the minimum set of organizational control constructs necessary and sufficient to govern execution across a distributed operational environment? The answer converges on five constructs, each grounded in an independent body of organizational theory, each addressing a distinct governance requirement, and together constituting a complete governance model.

The necessity test governs the framework: a primitive belongs here only if its absence would leave execution governance structurally incomplete — if removing it would create a governance gap that the remaining four primitives cannot compensate for. Each primitive satisfies this test. The framework is complete only when all five are present.

Concordat Beacon is the first platform to implement all five primitives as an integrated execution governance system. Its implementation is noted within each primitive as the authoritative expression of the theoretical construct in practice.

PRIMITIVE 1 DEFINE INITIATIVES

Definition: *Declare the discrete, bounded units of organizational work that require governance — the operational containers within which execution unfolds and against which all other primitives are applied.*

Theoretical Basis

Governance cannot be applied in the abstract. It requires a declared object — a bounded unit of work with a defined scope, a declared intent, and a measurable endpoint. Goal-setting theory establishes that performance requires specific, bounded objectives: without defined scope, execution cannot be evaluated against any meaningful standard (Locke & Latham, 1990). In management accounting, this construct corresponds to the responsibility center — the discrete accountability unit within which performance is measured and governed (Anthony & Govindarajan, 2007). Merchant and Van der Stede (2007) further establish that results control — one of the most fundamental forms of organizational

governance — requires that the results to be governed be explicitly declared before execution begins. The initiative is that declaration. It establishes what the organization is attempting, defines the boundary within which execution will be evaluated, and creates the operational container that all other governance constructs depend on.

Expression in Concordat Beacon

In Concordat Beacon, initiatives are instantiated as process instances within the Execution Playbook — discrete operational objects such as a membership onboarding, a client engagement, or a product launch cycle. Each initiative progresses through defined states from inception to completion. It is the unit against which decisions are captured, ownership is assigned, risks are evaluated, and outcomes are measured. The Control Tower visualizes the live population of active initiatives, their state distribution, and their aggregate governance exposure.

Why It Cannot Be Omitted: *Without defined initiatives, execution governance has no object. Every other primitive — decisions, ownership, risk, outcomes — requires an initiative as its operational anchor. Governance applied to undeclared work is not governance. It is observation without accountability.*

PRIMITIVE 2 CAPTURE DECISIONS

Definition: *Record the explicit organizational judgments made during execution — the attributable choices that determine how an initiative progresses, confirmed by the accountable owner and preserved as part of the governance record.*

Theoretical Basis

Organizational execution does not fail only because work is not done. It fails because work moves forward without the right judgments being made — because decisions that should have been deliberate were silent, assumed, or bypassed entirely. Decision rights theory establishes that organizational performance depends on making judgment explicit and attributable: the right decisions must be made by the right people with the right authority, and the governance layer must be able to confirm that this occurred (Jensen & Meckling, 1992). Perrow (1984) identified silent progression — work advancing through organizational systems without required confirmation — as a fundamental mechanism of organizational failure, particularly in complex multi-system environments. The BPM literature recognizes decision requirements as a core component of process state completeness: a state cannot be considered correctly closed if required organizational judgment has not been confirmed (Dumas et al., 2018). Governance without decision capture can confirm that work moved forward. It cannot confirm

that the right choices were made to move it forward — and in complex operations, that distinction is where accountability breaks down.

Expression in Concordat Beacon

In Concordat Beacon, decisions are declared in the Execution Playbook as named judgment checkpoints within process states — distinct in kind from system-generated completion signals. Each decision carries a name, an accountable owner, and an optional rationale field. The Concordat OS treats an unconfirmed decision as a blocking governance condition: the state remains open, risk surfaces, the owning role is alerted, and elapsed time against the unconfirmed decision accumulates as a governance signal in itself. When confirmed, the governance record captures what decision was made, by whom, and when — creating an auditable trail of organizational judgment. Decisions are declared sparingly, at the moments where organizational judgment is genuinely required and genuinely at risk of being bypassed. The OS surfaces them at the right moment to the right person with the relevant execution context already assembled — so that decision prompts carry governance weight rather than becoming operational noise.

Why It Cannot Be Omitted: *Without captured decisions, governance cannot distinguish between work that progressed correctly and work that progressed by default. Silent advancement — execution moving forward without required judgment — is among the most consequential and least visible forms of organizational failure. Decision capture is the primitive that makes organizational judgment legible, attributable, and auditable.*

PRIMITIVE 3 ASSIGN OWNERSHIP

Definition: *Establish explicit, named accountability for each stage of initiative execution — identifying the role or individual responsible for ensuring that each process state is satisfied correctly and on time.*

Theoretical Basis

Accountability without attribution is not accountability. Principal-agent theory establishes that organizational performance depends on explicit assignment of responsibility, with corresponding authority and visibility into execution state (Jensen & Meckling, 1976). When accountability is diffuse — spread across teams, assumed by default, or left unassigned at the boundary between systems — intervention cannot be directed, performance cannot be evaluated, and governance signals have no organizational recipient. Responsibility accounting theory operationalizes this principle by assigning performance targets and governance obligations to specific organizational actors, making accountability a structural property of the organization rather than an informal expectation (Anthony, 1965). Research on organizational failure consistently identifies the absence of ownership clarity —

particularly at the seams between systems and teams — as a primary cause of execution breakdown in complex operational environments (Weick & Sutcliffe, 2007). Ownership is not a label. It is the mechanism through which governance signals become directed organizational action.

Expression in Concordat Beacon

In Concordat Beacon, ownership is encoded in the Governance Schema at the state level — each process state carries an explicit assignment of the role, team, or individual accountable for ensuring that state is satisfied. The Concordat OS attributes risk conditions, SLA violations, stalled execution, and unconfirmed decisions to the owning role — enabling directed intervention rather than diffuse organizational alerting. The Control Tower surfaces ownership in the live execution view, presenting risk and revenue exposure in the context of the accountable party so that governance outputs produce immediate, directed accountability.

Why It Cannot Be Omitted: *Without assigned ownership, governance can identify that execution has stalled but cannot direct accountability. Risk without attribution is noise. Ownership is the primitive that converts governance signals into actionable organizational information and closes the loop between evaluation and intervention.*

PRIMITIVE 4 SURFACE RISKS

Definition: *Continuously evaluate whether execution is conforming to the declared operating model — detecting deviation, stalled progression, missing requirements, and exposure in real time, while intervention remains possible.*

Theoretical Basis

This primitive is the operational expression of what distinguishes Execution Governance from every prior organizational control category. The distinction is philosophical before it is technical: Execution Governance is normative and prospective — organizations declare how execution should unfold and the governance layer continuously evaluates conformance as work progresses. Analytics and reporting are empirical and retrospective — they measure what happened after it occurred. The difference is not one of speed. It is one of control logic. Simons (1995) identifies real-time exception detection as the defining capability of interactive control systems — those that enable organizations to detect emerging deviation and intervene before outcomes become irreversible. Operational risk management defines risk as the effect of uncertainty on objectives (ISO 31000, 2018) — and in execution governance, risk is not a classification applied after deviation has been confirmed. It is a continuous evaluation function that surfaces deviation while execution is still in motion. The management control systems literature distinguishes between ex ante controls, which define expected behavior, and feedback controls, which

detect deviation from that expectation (Malmi & Brown, 2008). The risk primitive unifies both: it encodes expected behavior in the governance schema and evaluates conformance against it continuously. This is the mechanism that makes governance a control system rather than a record system.

Expression in Concordat Beacon

In Concordat Beacon, the Concordat OS implements the risk primitive as the core evaluation logic of the platform. Every incoming signal is assessed against the Governance Schema to determine whether each active initiative is conformant, stalled, or in violation. Risk conditions evaluated continuously include SLA breaches, missing artifacts, unconfirmed decisions blocking state progression, ownership gaps, and execution patterns indicating systemic risk across the initiative population. Risk is surfaced at two levels: the instance level, identifying specific initiatives at risk, and the aggregate level, expressing cumulative exposure — including revenue at risk — across all active initiatives.

Why It Cannot Be Omitted: *Without risk surfacing, execution governance is reduced to a reporting system. The capability that makes governance prospective rather than retrospective — and that separates Execution Governance from analytics, process mining, and business intelligence — is the continuous, real-time evaluation of whether execution is conforming to organizational intent while intervention is still possible. Risk is the primitive that makes this possible.*

PRIMITIVE 5 MEASURE OUTCOMES

Definition: *Evaluate the measurable results that execution is intended to produce — assessing whether the operating model is achieving organizational intent and providing the empirical basis for improving it over time.*

Theoretical Basis

Governance without outcome measurement produces process conformance without organizational effectiveness. An organization can execute every initiative correctly, capture every decision, maintain perfect ownership, and surface every risk — and still fail to achieve what the work was intended to produce. The Balanced Scorecard framework established that organizational performance measurement must connect operational activity to strategic outcomes, and that the absence of outcome measurement severs the link between execution and purpose (Kaplan & Norton, 1992). Argyris and Schön (1978) distinguish between single-loop learning — detecting and correcting deviation from a declared standard — and double-loop learning — questioning whether the standard itself is producing desired outcomes. Both forms of organizational learning require outcome data. Without it, organizations can optimize execution conformance while remaining blind to whether conformant execution is achieving anything of value. Outcome measurement is also the mechanism through which

a governance model remains alive: it provides the empirical basis for refining process definitions, adjusting governance thresholds, and improving the operating model over time. A governance framework without an improvement loop is a static instrument in a dynamic environment.

Expression in Concordat Beacon

In Concordat Beacon, outcomes are encoded in the Governance Schema as KPIs and SLA benchmarks tied to execution quality — average initiative completion time, revenue realized per process instance, state completion rates, and conversion metrics across the initiative population. The Concordat OS evaluates active initiatives against declared outcome targets, not just process conformance. The Control Tower surfaces outcome metrics in aggregate, enabling leaders to evaluate whether the operating model as a whole is performing as intended. Outcome data flows back into the Execution Playbook review cycle, providing the empirical basis for improving process definitions and refining the governance model over time.

Why It Cannot Be Omitted: *Without outcome measurement, governance has no terminal condition and no improvement mechanism. The operating model becomes a static document rather than a living governance instrument. Outcomes close the governance loop and connect the mechanics of execution to the purpose of the organization.*

3. The Primitives as an Integrated Governance Model

3.1 Structural Interdependence

The five primitives are not a checklist. They are an integrated governance model in which each primitive is independently grounded in organizational theory and interdependent in governance function. The structural logic is as follows:

Define Initiatives establishes the unit of governance. Without a declared initiative, there is no container within which decisions can be captured, ownership can be assigned, risk can be evaluated, or outcomes can be measured. Every other primitive depends on the initiative as its operational anchor.

Capture Decisions establishes the governance record of organizational judgment within each initiative. Without captured decisions, execution can advance through states while bypassing the judgment checkpoints that ensure correct and accountable progression.

Assign Ownership establishes accountability for each stage of initiative execution. Without ownership, risk surfaces without attribution and intervention cannot be directed to the right organizational actor.

Surface Risks establishes the continuous evaluation function that makes governance prospective. Without risk surfacing, the framework can describe execution state but cannot govern it — the model becomes a record system rather than a control system.

Measure Outcomes establishes the terminal condition and the improvement loop. Without outcome measurement, governance confirms conformance but cannot evaluate whether conformance is producing value or identify how the operating model should evolve.

Remove any single primitive and the governance model has a structural gap that the remaining four cannot compensate for. An initiative without ownership cannot be intervened upon. Decisions without risk evaluation cannot be monitored for timeliness or completeness. Ownership without outcome measurement cannot be held accountable to results. Outcomes without initiatives cannot be attributed to specific units of organizational work. The model is complete only when all five primitives are present and operational.

3.2 The Governance Cycle

The five primitives describe a continuous governance cycle, not a static structure. Initiatives are defined and launched. Decisions are required, captured, and recorded. Ownership is assigned and accountability is maintained. Risks are evaluated and surfaced in real time. Outcomes are measured and fed back into the operating model.

This cycle is the execution-layer implementation of what organizational control theory has long identified as necessary but never previously operationalized at the cross-system level: a closed-loop control system that connects organizational intent to operational reality continuously and in real time (Anthony, 1965; Simons, 1995). The five primitives define the minimum elements of that loop. An organization that operates all five is governing its execution. An organization that operates fewer has governance gaps — identifiable, specific, and addressable by introducing the missing primitive.

3.3 The Control Plane Principle

The governance model defined by the five primitives operates on a principle analogous to the control plane architectures that govern distributed computing systems. A control plane does not execute the work handled by the systems it governs. It declares the desired state, observes the actual state, evaluates the difference, and surfaces deviation. The systems beneath it — containers, services, agents — execute work independently and without awareness of one another. The control plane governs whether that execution is collectively conforming to the declared intent.

Execution Governance applies this principle to organizational operations. Tools, operators, processes, and agents execute work independently, each within its own domain, each without full awareness of the organizational context surrounding it. The

governance layer — defined by the five primitives — sits above them, declaring the operating model, ingesting signals from across the execution environment, evaluating conformance continuously, and surfacing deviation before it produces irreversible organizational consequences.

The governance layer does not replace the systems beneath it. It does not route workflows, automate tasks, or manage records. It governs whether the work being done across all of those systems is unfolding according to the organization's declared intent. That is the function the five primitives define. That is the function no prior organizational control category has addressed.

4. Research Agenda

The Execution Governance framework generates a clear empirical research agenda. The theoretical case for the five primitives is grounded in established organizational control literature. What remains to be demonstrated through deployment is the empirical validation of three foundational propositions:

Necessity and sufficiency: Are these five primitives necessary and sufficient in real organizational deployments? Implementation across diverse operational environments will either confirm the framework's completeness or reveal additional primitives or redundant constructs. This is the foundational empirical question — one that only deployment evidence can answer.

Governance effectiveness: Does implementing all five primitives produce measurably better execution outcomes than partial implementation or no governance layer? Controlled deployments comparing governed and ungoverned operational environments will provide the outcome evidence required to establish the framework's empirical validity beyond its theoretical grounding.

Marginal governance value: Which primitive produces the highest marginal governance value when introduced into an organization that currently lacks a formal execution governance model? Does this vary by organizational size, operational complexity, or industry context? Understanding the governance value sequence will inform deployment prioritization and establish which primitives are most critical to organizational execution health.

These questions are answerable through controlled pilot deployments and represent the research path from theoretical framework to empirically validated governance model. The framework is designed to generate this evidence — each deployment producing the data required to validate, refine, and extend the primitives over time.

References

Anthony, R. N. (1965). *Planning and control systems: A framework for analysis*. Harvard Business School Press.

Anthony, R. N., & Govindarajan, V. (2007). *Management control systems* (12th ed.). McGraw-Hill.

Argyris, C., & Schön, D. A. (1978). *Organizational learning: A theory of action perspective*. Addison-Wesley.

Concordat. (2024). *The Concordat doctrine: Execution governance*. Concordat.

Dumas, M., La Rosa, M., Mendling, J., & Reijers, H. A. (2018). *Fundamentals of business process management* (2nd ed.). Springer.

ISO. (2018). *ISO 31000:2018 — Risk management: Guidelines*. International Organization for Standardization.

Jensen, M. C., & Meckling, W. H. (1976). Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3(4), 305–360.

Jensen, M. C., & Meckling, W. H. (1992). Specific and general knowledge, and organizational structure. In L. Werin & H. Wijkander (Eds.), *Contract economics* (pp. 251–274). Blackwell.

Kaplan, R. S., & Norton, D. P. (1992). The balanced scorecard: Measures that drive performance. *Harvard Business Review*, 70(1), 71–79.

Locke, E. A., & Latham, G. P. (1990). *A theory of goal setting and task performance*. Prentice Hall.

Malmi, T., & Brown, D. A. (2008). Management control systems as a package — Opportunities, challenges and research directions. *Management Accounting Research*, 19(4), 287–300.

Merchant, K. A., & Van der Stede, W. A. (2007). *Management control systems: Performance measurement, evaluation and incentives* (2nd ed.). Pearson.

Perrow, C. (1984). *Normal accidents: Living with high-risk technologies*. Basic Books.

Simons, R. (1995). *Levers of control: How managers use innovative control systems to drive strategic renewal*. Harvard Business School Press.

Weick, K. E., & Sutcliffe, K. M. (2007). *Managing the unexpected: Resilient performance in an age of uncertainty* (2nd ed.). Jossey-Bass.