EXECUTIVE SUMMARY · PRE-RELEASE EDITION

The Decision

A Management Theory for Agentic Enterprises

Luciana S. Ledesma · November 2026

> AI agents can follow every rule, pass every check, and still make the wrong decision. When decisions move faster than management, authority must change shape: the decision itself has to become the unit that organizations direct, verify and remember.

1 The failure mode nobody is watching for

A Dutch insurance carrier deploys an agent to adjudicate claims. Before go-live, it writes its senior adjusters' rules into the agent's instructions, including one written after a fraud ring cost the company seven million euros: any fourth claim from one claimant within sixty days goes to the fraud team.

A claim arrives. The agent checks the policy, the premium, the fraud score, the threshold, and counts the claimant's recent claims: two. Every check passes and the claim is paid. Six weeks later the fraud team finds it was the fourth claim from the same household in six months, split across two names and two policies, sixty-three days apart. The rule was followed perfectly. What it was for was missed completely.

The book calls this **decision-basis failure**: the workflow runs correctly, every operational check passes, and the decision is still wrong, because nothing verified whether what should have mattered was considered. It is not a bug in the agent. It is the absence of institutional judgment from the chain, and it passes silently through every layer built to catch other kinds of failure.

2 Why established controls leave it unresolved

Over two decades the industry has built a strong **execution stack**. Authentication answers who may act; permissions and policies, on what; orchestration, with which tools; observability, what happened; evaluations, how well the model performed. Each solves a real slice of the problem. None answers the question that decides whether an outcome can be trusted: *was this decision sufficiently supported to act autonomously?*

That is a question about basis, not execution. Answering it needs something no existing tool combines. It has to reflect **this organization's** judgment, not a generic standard; be specific to the **domain** (claims, releases, procurement); respect **who owns** each piece of judgment; and run **before the agent acts**, every time, at machine speed.

Human review cannot fill the gap. Agents produce decisions at a volume and speed no reviewer can follow, and attention is a biological budget, not a setting. Meanwhile three things become hard to answer in most organizations deploying agents: which agents are running, what authority was delegated to them, and who is accountable for their decisions. Nothing has visibly failed, so the exposure accumulates unseen. The book calls this **governance debt**.

3 The decision as the unit of management

Classical management rested on two assumptions that quietly stopped holding when agents arrived: that decisions can be reviewed before they execute, and that oversight has the shape of a pyramid. Patching the pyramid does not work. The book rebuilds management on a different unit: **the decision itself**.

Herbert Simon put the decision at the centre of management in 1947, but bounded it by human cognitive capacity. In agentic organizations, value is created and lost through hundreds of thousands of decisions, many made by machines. Firms, processes and roles become abstractions over them. Decisions can now be described by measurable properties, among them blast radius, reversibility, propagation depth, required scrutiny and **basis quality**.

Decision quality has three layers. Basis quality asks whether the decision considered what it had to

consider; it is the only layer that can be checked at machine speed. Judgment quality asks whether the call was sound, and belongs to the people who own the decision. Outcome is partly luck. Grading agents by outcomes alone rewards the lucky and blames the unlucky.

Quality rarely fails in a single event. It erodes through drift, and each drift leaves a debt behind it:

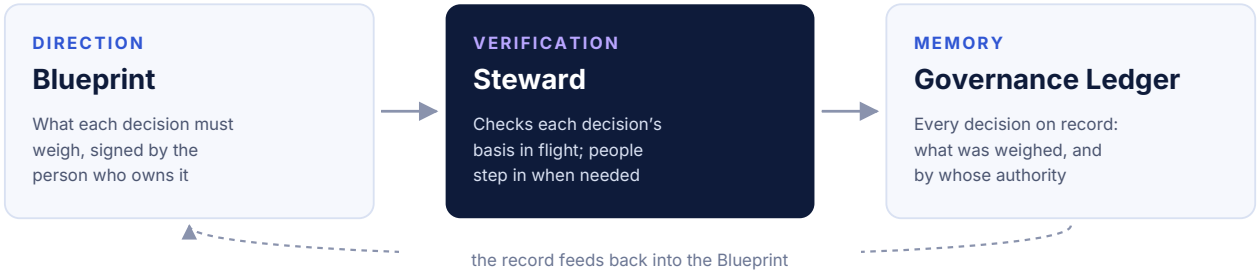
Drift	What happens	The debt it leaves
Human	People hand over judgment and lose the ability to challenge what they receive	**Cognitive debt**: lost or confused judgment, showing up as unresolved decisions and divided attention
Synthetic	Agent reasoning diverges from the intent that authorized it	**Governance debt**: autonomy granted faster than the criteria, boundaries and verification needed to govern it
Historical	Systems lose the record of why decisions were made	**Institutional memory debt**: the organization can no longer reconstruct what was weighed, who authorized it, and why

Outcomes can stay acceptable for months while judgment, governance and memory weaken underneath them. The greater risk is not only that decisions go wrong, but that organizations lose the means to know whether their decisions are sound.

Taking the decision as the unit changes the organization's shape from pyramid to network. It also joins two things the classical era kept apart. **Management and governance stop being two layers running at two speeds and become two authorities acting through one machinery.** Operators run it; responsible owners sign what it checks; boards approve the risk it encodes. The machinery is shared; the authorities are not.

4 The architecture: Blueprint, Steward, Ledger

The book proposes a **real-time counterweight**: not a brake or a gate, but a continuous force that keeps agent speed governable. It has three pieces.



Direction without verification is a document nothing enforces. Verification without direction has nothing to check against. Both without memory cannot be defended or learned from.

Direction: the Governance Blueprint. Institutional judgment, written down before the agent acts. For a class of decision, the Blueprint states what must be checked, what evidence must be present, what lies out of bounds, and how closely each point must be scrutinized. Each element is signed by the person who owns that judgment: fraud signs the pattern-of-claims rule, compliance the beneficial-ownership rule. Signatures expire unless renewed. This is what keeps the content current, where earlier attempts to codify judgment drifted from practice. The operating agent makes decisions using its own instructions, models, tools and context. The Blueprint independently represents how the organization expects consequential decisions to be made: the **decision trajectory** it expects: the considerations, dependencies, boundaries and intervention points any sound path must include, not every step of one person's reasoning. Most organizations have a workflow, which sets the order of the steps; few have this.

Verification: the Steward. A Steward Agent sits inside the workflow, reading the agent's reasoning as it unfolds, and answers to the Blueprint's signers rather than to the team that built the agent. It compares the agent's observable decision trajectory against that organizational reference and creates opportunities to intervene before a conclusion hardens. The response is graded, not a choice between allow and stop: an **intervention ladder** runs from pass, nudge, throttle and flag to contain, escalate and block. The level for each kind of issue is chosen in design, not at the moment of failure. Escalations go to the person who signed the relevant point.

Memory: the Governance Ledger. Every decision leaves a complete, consistent, durable and queryable

record: what was proposed, which version of the Blueprint applied, what was found, what intervention followed, and what happened next. Logging becomes a property of the architecture rather than a habit. When a regulator or board asks how a class of decisions was governed last quarter, the answer is a query instead of a reconstruction. The Ledger pays down institutional memory debt; the Blueprint and Steward pay down governance debt.

The architecture also has a generative side. Configured for dialogue and plural perspectives, Stewards can improve the decisions that pass, not only stop the ones that should not.

5 What it cannot do

The book names the limits before asking anyone to build it. The architecture cannot see what nobody named. A declared basis is not necessarily the real one. Signatures can become a new rubber stamp. The verifier is itself a model. It costs time and money, unevenly. It can concentrate judgment or push accountability down onto signers. And the evidence is young. Cognitive, technical, environmental and social debts lie outside its scope, and it should not be built in ways that make them worse.

6 The question for leaders and boards

Boards remain responsible for enterprises whose autonomous decisions they cannot individually see. More dashboards do not close that gap. A summary of successful executions does not show whether the authority behind them was legitimate or the basis adequate. Oversight shifts from examining decisions to overseeing the **architecture under which decisions are permitted**:

- Which classes of decisions may proceed autonomously, and which need independent scrutiny?
- Who is authorized to set and change the boundaries, and when must an agent escalate?
- Who verifies the verifier, and what evidence survives for later challenge?

The new board question is not *Did every agent pass its checks?* but *Can we account for the authority and judgment under which the enterprise decides?*

About the book. *The Decision* (pre-release edition, 30 chapters) develops the argument in seven parts: the failure mode, what broke in classical management, the decision as unit, what Toyota's *jidoka* and DevOps' shift-left already knew, the architecture and its limits, practical roles and a playbook, and the relation to cybersecurity.

About the author. Luciana S. Ledesma is the founder of MeaningStack B.V. in Maastricht, which builds decision governance for AI agents; its Steward is in pilot. She holds an MPhil from the University of Cambridge and an MBA from TUM, and spent fifteen years at Siemens, Boston Scientific and HP.

Reading further. Six chapters are open to read at the book's page. To discuss the book or read the complete pre-release manuscript, write to luciana@meaningstack.com.