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. One agent can make thousands of decisions an hour; no org chart can review that. The decision itself must become the unit of management: judgment your experts sign in advance, checked before the agent acts.

1 · THE FAILURE MODE NOBODY IS WATCHING FOR

The rule was followed perfectly. What it was for was missed completely.

A Dutch insurer writes its senior adjusters' rules into a claims agent, including one written after a fraud ring cost it seven million euros: any fourth claim from one claimant within sixty days goes to the fraud team.

WHAT THE AGENT CHECKED

- ✓ Policy active, premium current
- ✓ Fraud score low, amount under threshold
- ✓ Claims from this claimant in 60 days: two
- ✓ Approved. Every log line green.

WHAT THE FRAUD TEAM FOUND SIX WEEKS LATER

- ! The fourth claim from the same household in six months
- ! Split across two names and two policies
- ! Sixty-three days apart

Decision-basis failure: the workflow runs correctly, every 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.

2 · WHY ESTABLISHED CONTROLS LEAVE IT UNRESOLVED

Each layer answers a real question. None answers the one that decides whether an outcome can be trusted.

LAYER	WHAT IT ANSWERS	WHAT IT CANNOT ANSWER
Authentication	Who may act	**Was this decision sufficiently supported to act autonomously?** A question about basis, not execution. It must reflect this organization's judgment, be specific to the domain, respect who owns each piece of it, and run before the agent acts, every time, at machine speed.
Permissions and policies	On what	
Orchestration	With which tools	
Observability	What happened	
Evaluations	How well the model performed	

Human review cannot fill the gap. Attention is a biological budget, not a setting:

Thousands of decisions an hour from one agent deployment, around the clock

20 complex decisions an hour that one executive can review, for 3–4 focused hours a day

Meanwhile three things become hard to answer: 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 assumed decisions could be reviewed before they execute, and that oversight is a pyramid. Both assumptions have expired.

The book rebuilds management on a different unit: the decision itself. Herbert Simon put it 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, and each decision has measurable properties: blast radius, reversibility, propagation depth, required scrutiny, basis quality.

Decision quality has three layers. Only the first can be checked at machine speed.

BASIS QUALITY

Did it consider what it had to consider?

Checkable before the decision executes, every time. The layer the architecture is built on.

JUDGMENT QUALITY

Was the call sound?

Belongs to the people who own the decision. Cannot be delegated to a verifier.

OUTCOME

Did it turn out well?

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. Outcomes can stay acceptable for months while judgment, governance and memory weaken underneath.

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, unresolved decisions, 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:** what was weighed, who authorized it and why can no longer be reconstructed

Taking the decision as the unit changes the organization's shape from pyramid to network, and joins what the classical era kept apart: management and governance become two authorities acting through one machinery.

Operators run it; responsible owners endorse what it checks; boards approve the risk it encodes. The machinery is shared; the authorities are not.

4 · THE ARCHITECTURE: A REAL-TIME COUNTERWEIGHT

Not a brake or a gate: a continuous force that keeps agent speed governable.

DIRECTION

Governance Blueprint

What each class of decision must weigh, endorsed by the person who owns that judgment. Endorsements expire unless renewed.

VERIFICATION

Steward

Compares each decision's observable trajectory with the Blueprint as it unfolds, and intervenes before a conclusion hardens.

MEMORY

Governance Ledger

Every decision on record: what was proposed, which Blueprint applied, what was found, what followed.

← the record feeds back into the Blueprint

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.

The Blueprint is not a set of instructions for the agent. The operating agent decides with its own instructions, models, tools and context. The Blueprint independently represents how the organization expects consequential decisions to be made: the **decision trajectory**, 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.

The Steward answers to the Blueprint's owners, not to the team that built the agent. Its response is graded, chosen in design rather than at the moment of failure, and escalations go to the person who endorsed the relevant point:

Pass	Nudge	Throttle	Flag	Contain	Escalate	Block
proceeds, logged	added context	reduced scope	human reviews later	cascade blocked	pauses for a named human	prevented, logged

The Ledger makes logging a property of the architecture rather than a habit. When a regulator asks how a class of decisions was governed last quarter, the answer is a query, not a reconstruction.

5 · WHAT IT CANNOT DO

The book names the limits before asking anyone to build it.

1. It cannot see what nobody named.
2. A declared basis is not necessarily the real one.
3. Endorsements can become a new rubber stamp.
4. The verifier is itself a model.
5. It costs time and money, unevenly.
6. It can concentrate judgment, or push liability onto endorsers.
7. The evidence is young.
8. Cognitive, technical and social debts lie outside its scope; it must not make them worse.

6 · THE QUESTION FOR LEADERS AND BOARDS

Oversight shifts from examining decisions to overseeing the architecture under which decisions are permitted.

Boards remain responsible for enterprises whose autonomous decisions they cannot individually see. More dashboards do not close that gap.

01

Which classes of decisions may proceed autonomously, and which need independent scrutiny?

02

Who is authorized to set and change the boundaries, and when must an agent escalate?

03

Who verifies the verifier, and what evidence survives for later challenge?

THE OLD BOARD QUESTION

~~Did every agent pass its checks?~~

THE NEW BOARD QUESTION

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, from the failure mode to the architecture, its limits, practical roles and a playbook. Six chapters are open at **meaningstack.com/pages/book**; the complete manuscript on request.

About the author. Founder of MeaningStack B.V., Maastricht, which builds decision governance for AI agents; its Steward is in pilot. An international management career across technology, innovation and venture creation, with enterprise work for Siemens, Boston Scientific and HP. MPhil, University of Cambridge; MBA, Technical University of Munich. **luciana@meaningstack.com**

In practice. MeaningStack starts organizations on this path with one consequential decision: a two-hour executive session, or a three-week Agent Decision Sprint on one agent workflow. **meaningstack.com**