

AGENTIC COMMERCE FRAMEWORK®

Governing autonomy

The governance of AI agents, explained and equipped

Vincent Dorange · 2026 Edition

The problem: **autonomy without governance**

- ◆ An AI agent **decides and acts on its own** — it is not a chatbot.
- ◆ It **drifts at machine speed**: the damage is done in seconds.
- ◆ **Legal, ethical, financial, reputational** — four fronts at once.
- ◆ **Who decides when the agent decides?**

AI Act penalties up to **7% of worldwide turnover**

The metatechnology — **why now**

- ◆ AI is not just another product.
- ◆ **Metatechnology**: the technology behind the technology.
- ◆ It **redistributes the power to decide**.
- ◆ Act now, **before the decision is captured**.

Mustafa Suleyman — **The Coming Wave**, 2023

Two models — laissez-faire vs control

- ◆ **Platform (US):** speed and capital — but dependency and capture of the decision.
- ◆ **Centralised (China):** state-run AI — but surveillance and the subordinated individual.
- ◆ **Two dead ends for sovereignty.**

The European **third way**

- ◆ Sovereignty through **law and design**.
- ◆ Delegated autonomy, but **governed and reversible**.
- ◆ In the lineage of **ISN × iDFRights** — Benhamou, Cavada.
- ◆ Equip the doctrine, or it stays a wish.

An **open standard** for schools — the value stays protected

The 4 new laws of robotics

- ◆ 1 · AI **complements** professionals, it does not replace them.
- ◆ 2 · AI never counterfeits humanity.
- ◆ 3 · AI does not intensify the arms race.
- ◆ 4 · AI always identifies its creators and controllers.

Frank Pasquale — [New Laws of Robotics](#), Harvard 2020

Agentic sovereignty — the definition

- ◆ Knowing **at all times who really decides**, when agents act on your behalf.
- ◆ A capability of **the organisation — and of the State**.
- ◆ The human keeps **arbitration and correction**.
- ◆ An autonomy that is always **reversible**.

The AI Act & accountability by design

- ◆ The European AI law, **in force since 2024**.
- ◆ The **company** is fined; the **executive** personally liable (national law).
- ◆ Built in **from the design stage**, not after.
- ◆ A **principle without a tool binds no one**.

Up to **7% of worldwide turnover** — AI Act administrative penalty (company)

The 4 layers of governance

1 Governance & sovereignty

who decides, who answers, in the name of which values

2 Decision policy

what can be agentified, and under which threshold

3 Agents

they perceive, reason and propose

4 Execution & supervision

action that is traced, monitored, reversible

Every decision **flows down the stack** and passes through arbitration

Governance & sovereignty — the DDAO

- ◆ Who decides, who answers, **in the name of which values**.
- ◆ **DDAO**: an accountable role, with a reasoned right of objection.
- ◆ A senior profile — **business + tech/law**.
- ◆ The ability to **say NO to top management**.

The layer that makes **everything else enforceable**

Decision policy — the 4 levels

LEVEL 1

Advisory

The agent proposes, the human decides and executes.

LEVEL 2

Co-decision

Joint agent + human decision before any action.

LEVEL 3


Supervised autonomous

The agent acts, the human monitors and can stop it.

LEVEL 4

Autonomous

The agent acts alone within a bounded scope.

— autonomy  + autonomy

The level is set **decision by decision**, never in bulk

Agents — the agent proposes

- ◆ Agents perceive, reason, **propose**.
- ◆ **The agent proposes, the human decides.**
- ◆ Every agent **declares itself as an agent**.
- ◆ A **named human owner** for each agent.

No sensitive execution **without validation**

Execution & supervision — + kill switch

- ◆ The real action, **traced and monitored**.
- ◆ Logging, alerts, **kill switch**.
- ◆ Every decision **appended, never erased**.
- ◆ **Auditable and legally enforceable**.


An **append-only** decision registry

3

Measure to govern

Three instruments to score every decision: criticality, risk, maturity. To measure is to know how far to delegate.

Measuring the **criticality** of a decision



At equal impact, it is **irreversibility** that raises the score — and brings the human back in

The **risk** matrix — likelihood × severity



Operational risk (ACF-11) ≠ decision criticality (ACF-02) — **two scales, never merged**

The 4 levels of agent maturity

N0

Classic automation

Fixed rules, scripts, RPA. The machine executes, it does not choose.

DECIDES: THE HUMAN

N1

Assisted agents

The agent proposes and recommends. The human keeps the final decision.

DECIDES: THE HUMAN

N2

Governed agents

The agent decides and acts within a frame: mandate, thresholds, guardrails, traceability.

RECOMMENDED TARGET

N3

Supervised autonomous

Broad autonomy over an extended scope, under supervision and kill switch.

DECIDES: THE AGENT

Human control

Agent autonomy

Maturity qualifies **an agent**; the 4 autonomy levels qualify **each decision**

35

The toolkit, **35 cards** ready to use

Three families to print and fill in: decide (ACF-00 to ACF-16), organise (ACF-O1 to O3) and operate (ACF-P0 to P9). From the first diagnostic to day-to-day operations.

Step 1 — diagnose & frame

ACF-00	Sovereignty Score	measure your dependency across 4 dimensions before automating
ACF-01	Decision Map	map your decisions and their level of autonomy
ACF-02	Criticality Matrix	score each decision to know how far to delegate it

You only delegate what you have first **measured and mapped**

Step 2 — deploy an agent

ACF-03	Agent Constitution	set identity, principles and forbidden zones
ACF-04	Agent Sheet	describe a concrete agent, supervisable and enforceable
ACF-05	Supervision & Governance	monitor KPIs, drift and compliance continuously
ACF-06	Kill Switch	stop an agent immediately in case of drift
ACF-07	First Agent Dossier	gather all the framing before launch

Step 3 — supervise & audit

ACF-08	Decision Registry	log every important decision for traceability
ACF-09	Action & Improvement Plan	turn incidents and audits into corrective actions
ACF-10	30-Day Governance Audit	measure governance maturity across 8 domains
ACF-11	Risk Assessment	identify and prioritise: likelihood × severity
ACF-12	Agent Mandate	formalise the enforceable scope and delegated autonomy

For the classroom — teach & practise

ACF-13	Guided Case Study	apply the framework to a real scenario, step by step
ACF-14	Teacher Guide	run, mark and assess workshops
ACF-15	Governance Simulation	practise deciding and arbitrating under pressure
ACF-16	Accountability by Design	embed the 5 pillars of accountability from the design stage

Step 4 — organise human & agent

ACF-O1	Human-AI-Agent Matrix	for each act: who executes, who answers, which rule, who stays in the loop
ACF-O2	Hybrid Organisation Map	place humans, agents and rules domain by domain
ACF-O3	Critical Skills to Keep	keep in-house what it takes to judge and take back your agents

You delegate execution, **never accountability**

Step 5 — configure & run the agent

ACF-P0	Configuration Grid	set the agent's six dials
ACF-P1 · P2	Acceptance & Go-Live	test each setting, then scale up in stages
ACF-P3 · P4	Supervision & Roles	the monitoring rhythms and who does what
ACF-P5 · P6 · P7	Escalation · Shutdown · Incident	when to hand over, cut off, record
ACF-P8 · P9	Permissions & Vendor	review access, choose the agent before governing it

One throughline — end to end

- ◆ **Lendari**, an e-commerce SME, deploys a customer-service agent.
- ◆ **Camille Roussel** (DDAO) sets the frame: score, criticality, mandate.
- ◆ Agent "**Margaux**": refund $\leq$ €100 in **supervised autonomous** mode.
- ◆ Above €100 → **co-decision**; incident → **kill switch**.
- ◆ 31 filled examples run the case end to end.

The "throughline" examples are provided **filled in** in the manual

Start this week

No **autonomy** without supervision · no supervision without governance · no governance without **sovereignty**.

1

Diagnose

Fill in ACF-00 and ACF-01 on a real case from your organisation.

2

Frame an agent

Constitution + Agent Sheet + Mandate. Set autonomy decision by decision.

3

Supervise

Registry, kill switch, 30-day audit. Make the agent enforceable.

Framework and cards: free for public and private not-for-profit teaching. Enterprise / commercial use under contract. · HD download: acfstandard.com · Vincent Dorange · 2026 Edition.