The Ontological Collision of the Product Regime and the System Regime[1]

How European regulation of medical AI externalizes its incoherence to manufacturers

> *«Ontological» here does not designate a philosophical abstraction, but the fact that each regime legally constructs a different object.*

Statement of facts

Six objective elements structure this dossier. All reading levels refer to them.

- **19 November 2025.** The European Commission presents the *Digital Omnibus on AI*, which amends the AI Act itself: postponement of the applicability of high-risk obligations, broadening of regulatory sandboxes, clarification of the interplay with MDR / IVDR.
- **16 December 2025.** The Commission tables proposal COM(2025) 1023 final, which amends the MDR and the IVDR. Among the simplification measures: removal of medical AI devices from the HRAIS perimeter, by displacement from Annex I Chapter §A to Chapter §B of the AI Act (Chapter §A calls for third-party conformity assessment; Chapter §B relies on existing sectoral regimes).
- **January 2026.** The Food and Drug Administration finalizes its Predetermined Change Control Plan (PCCP) framework. The PCCP qualifies *ex ante* the authorized modifications of an AI device on the basis of a change protocol declared at the moment of initial authorization.
- **13 March 2026.** The Council of the European Union adopts its position in favor of postponing, to August 2028, the applicability of the AI Act high-risk obligations to systems integrated within medical devices.
- **26 March 2026.** The European Parliament adopts its position in favor of moving Annex I of the AI Act from Chapter §A to Chapter §B.
- **Summer 2026 or 2027.** Final adoption envisaged by the two institutions, within the ordinary legislative procedure. The joint version resulting from the trilogues is not stabilized as of the date of this note.

Order of magnitude. By the end of 2025, 1,451 medical AI devices were authorized by the FDA, 76% of which in radiology. On the European side, 51 notified bodies were designated under the MDR in October 2025.

The decisive timing. The substantive proposals precede the calendar postponement. The Commission drafts before the Council defers. This ordering is the inverse of what one would observe if the movement proceeded from an industrial concession.

Methodological note. The texts (AI Act, consolidated MDR / IVDR, COM(2025) 1023 final, Digital Omnibus AI) are primary sources. The contextual elements (MedTech position, 51 MDR notified bodies in October 2025, projected adoption) come from secondary analyses. They illuminate; they do not prove. The principal thesis does not depend on whether MedTech specifically asked for 2029 rather than 2030. The analysis bears, moreover, on regulation at the Union level. Disparities of implementation at Member State level (differentiated designation and capacity of notified bodies, national product liability case law, modulation by national health authorities) are mobilized as tipping factors (§12) without constituting the core of the demonstration. A complete analysis would require examination of these variations, which exceeds the scope of this note.

First reading level: Thesis

The European Union is not in the process of softening the regulation of medical AI. It is recording an incompatibility that it does not know how to resolve.

Two regimes coexist:

- MDR / IVDR regulates a validated state (a product),
- the AI Act regulates a trajectory (an evolving system).

These two regimes do not bear on the same object.

Consequence:

- their combination is unstable,
- their cumulation is non-commutative,
- their separation creates blind spots.

The solution is neither stacking nor withdrawal. The solution is a change of regulated object: moving from {product} or {system} to {authorized state + qualified trajectory}.

But the difficulty is not only technical. The cumulation holds not because it is coherent, but because its incoherence is externalized to manufacturers. As long as absorption remains possible, resolution has no institutional reason to occur.

Second reading level: Operational reading

1. The real problem (not the theoretical one)

Consider a radiology AI model updated monthly on new corpora of annotated images.

Two readings:

- MDR: «Does the update change the device?»
- AI Act: «Does the update remain within the foreseen limits?»

Same event. Two logics.

2. The point of rupture

The system breaks the moment one asks a simple question: in what order should the two regulations be applied?

- MDR → AI Act: compliant
- AI Act → MDR: potentially non-compliant

This is not a legal ambiguity. It is a logical indeterminacy.

3. What manufacturers do

They do not solve the problem. They absorb it: double documentation, double validation cycle, invisible internal arbitration. The system functions because the cost of coherence is transferred to manufacturers.

4. Three possible trajectories

- Maintenance of cumulation: politically stable, technically incoherent.
- Unstacking: apparent simplification, loss of coverage.
- A holding regime: technical coherence, high institutional cost.

5. The conceptual shift

The problem does not resolve itself by adjusting procedures. It resolves itself by changing the regulated object: an AI device is neither a product nor a system, it is **a product in trajectory**.

Third reading level: Demonstration

1. State of the European dossier

Since 2025 / 2026:

- postponement of high-risk obligations to August 2028 (Council position, 13 March 2026),
- proposals to remove medical devices from the HRAIS perimeter (COM(2025) 1023 final, 16 December 2025; Parliament position, 26 March 2026),
- final adoption envisaged in summer 2026 or 2027.

Two readings are possible.

Naive reading. The Union is softening. This is a concession to industrial pressure. This reading is arithmetically difficult to sustain: MedTech Europe was asking for 2029, the result is 2028. The lobby obtained less than it asked for. It is hard to qualify as a concession what falls short of the demand.

Reading defended here. The regulation is recording an inability to make two ontologies coexist. The postponement is not a decision on the substance. It is the admission that the substance does not allow itself to be decided under current conditions.

2. Two regimes, two objects

MDR / IVDR

- Object: validated device,
- logic: certification of a state,
- modification $\rightarrow$ requalification.

AI Act (HRAIS)

- Object: evolving system,
- logic: qualification of a trajectory,
- update $\rightarrow$ normal if bounded.

They do not describe the same thing.

3. Operational frictions

Four tensions appear immediately.

Qualification of modifications

- MDR: potentially significant modification (criteria MDCG 2020-3 Rev.1; Annex IX §4.10),
- AI Act: normal update if bounded.

Liability

- MDR: notified body,
- AI Act: provider (article 16).

Temporality

- MDR: multi-year cycles (PSUR),
- AI Act: continuous monitoring (article 72).

Documentation

- Two distinct corpora, coherence not guaranteed.

To these frictions is added a territorial dimension that will not be developed here but must be named: notified bodies are designated by Member States and exhibit heterogeneous capacities; product liability case law varies across national jurisdictions; the practical transposition of European regulations passes through health authorities whose doctrine is not convergent. The MDR / AI Act cumulation therefore does not apply to a homogeneous regulatory system, but to twenty-seven administrative apparatuses whose average resultant is itself a variable object.

These frictions show a cost. Not yet an incompatibility.

4. Proof: non-commutativity

Consider a monthly update.

Sequence 1. MDR then AI Act → non-significant modification (MDCG) → trajectory adjusted *a posteriori* → compliance maintained.

Sequence 2. AI Act then MDR → overrun of the qualified trajectory → requalification of the system → modification now significant in the MDR sense → possible recertification.

Same event. Different results: **the regulatory system has no unified semantics**.

Guidance MDCG-AIB 2025-6 on the AI Act / MDR interplay sketches a conciliation effort. It does not resolve non-commutativity at the categorial level. It organizes it.

5. Passage to the categorial level

When two regimes bear on the same object: their combination is commutative, or can be rendered commutative by a priority rule.

When they bear on different objects: their combination is not defined.

We are in the second case. This is no longer a problem of procedure.

6. Why the system holds despite its incoherence

If the cumulation is non-commutative, if operators bear its cost, if the regulators themselves are beginning to acknowledge it, why does the system hold?

Because it has found its equilibrium. The categorial incompatibility produces, mechanically, two correlated effects:

- local over-regulation, where the AI Act's generic instruments have operational grip,
- blind spots, where they have no instruments suited to clinical specificity.

Operators react rationally:

- circumvention by reclassification of the product,
- displacement of the declared use toward indications outside HRAIS,
- innovation accumulated in the blind spots (technical differentiation without prudential friction).

The regulator then re-extends its perimeter, with the same generic instruments, thus reproducing the initial asymmetry on a larger scale. The loop closes.

This loop is stable because it distributes costs in an institutionally bearable way:

- regulators preserve a public posture of maximal coverage,

- manufacturers absorb the incoherence internally,
- patients do not see the structure, but can reveal its failures through punctual events that move from the clinical register to the political register: a pharmacovigilance signal, individual litigation, media alert. They are never explicit actors of regulation, but they are its implicit triggers.

The cost is not invisible. It is externalized. Externalized to manufacturers, who do not decide the regulation. **The cumulation holds because those who pay for it are not those who decide it.**

This is not a malfunction. It is the normal functioning of a system that has found how to dispense with coherence, by delegating it to those who have no choice but to absorb it.

Uncomfortable consequence for administrative rationality: what presents itself as arbitration is, in reality, an admission. The Commission did not get the object wrong. It simply ended up noticing.

6 bis. Minimal formalization

The preceding mechanism calls for a formalization, both to specify its domain of application and to make it transposable beyond the medical case.

Conditions of application. The mechanism appears under three cumulative conditions:

- plurality of non-aligned prudential regimes bearing on the same artifact,
- institutional impossibility of unifying them within a short political horizon,
- existence of operators structurally capable of absorbing the cost of incoherence.

Outside these three conditions, the mechanism does not apply. Under these three conditions, it can be described as follows:

- *Actors:* regulator (R), manufacturer (F), patient (P).
- *Variables:* political cost of coherence (C_pol), operational cost of incoherence (C_op), residual clinical risk (R_clin).
- *Current equilibrium:* R minimizes C_pol under the constraint that F absorbs C_op and that R_clin remains politically invisible.
- *Tipping condition:* R_clin moves from the clinical register to the political register, requalifying C_pol and imposing a decision the institution had until then avoided.

Three stabilities to distinguish. The equilibrium is not monolithic. It rests on three logically distinct stabilities, the confusion of which is the first analytical error:

- *Institutional stability:* R remains capable of deferring.
- *Economic stability:* F remains capable of absorbing without critical loss.
- *Clinical stability:* R_clin remains below the threshold of political visibility.

These three stabilities are not independent, but they do not fail simultaneously. Current European regulation holds on the first two; the third is, by construction, non-observable until the moment it ceases to hold. It is precisely this lag of visibility that explains why a system can be institutionally stable while being clinically unstable, and why the second instability becomes regulatory only when it becomes political.

Domain of validity. This framework is not specific to medicine. It describes a class of prudential equilibria under constraint. European regulation of medical AI is a remarkable particular case because the externality there falls on subjects, not on assets; but the same structures are encountered in financial certification, in critical industrial safety, or in the regulation of foundation

models themselves.

7. The false reasoning of cumulation

The implicit hypothesis (more regulation = more coverage) is valid only if regulations bear on compatible objects. When the objects differ, regulations do not add: they overlap badly and leave holes. Cumulation then increases cost without guaranteeing coverage. This is precisely the configuration produced by MDR + AI Act.

8. The holding regime

Change of object: {validated state} + {qualified trajectory}.

Five properties:

- the update is integrated from the outset,
- variability is bounded *ex ante*,
- compliance is dynamic,
- proof bears on behavior,
- documentation describes a trajectory.

Industrial parallel. For a CTO, the holding regime is not a novelty. It is what software engineering has been doing for ten years: feature flags + canary deployments + observability contracts. There, compliance is a property of a governed pipeline, not a state of a frozen version. The software industry has, over a decade, developed robust mechanisms for the operational governance of evolving systems. The European problem is therefore not to imagine a trajectory logic *ex nihilo*, but to translate it into an opposable prudential regime. The difficulty is regulatory and institutional, not conceptual.

But the parallel has an asymmetry that must be named. What distinguishes the medical context is not the technical management of the trajectory, which the software industry has mastered, but **the nature of the proof required**. This proof is *clinical* (it is established on subjects, not on SLAs), *populational* (it is validated at the scale of a cohort, not of an A/B test), and *legally opposable* (it engages product liabilities that survive the update). None of these three characters is required in an ordinary Kubernetes deployment. Translating the industrial paradigm therefore requires not the importation of tools, but their articulation with an evidentiary regime that the software industry has never had to produce. That is the work that remains ahead.

Pseudo-structure of a holding authorization dossier.

The dossier would integrate, from the initial submission, four sections.

A. Authorized state (current MDR equivalent)

- description of the device,
- target population, indication, reference performance,
- initial clinical evaluation.

B. Qualified evolution envelope (new)

- admissible bounds of variability (performance drift, input distribution, clinical subgroups),
- declared and authorized retraining cadence,
- sources and qualification criteria for new training data,
- non-regression tests on a reference cohort.

C. Surveillance and controlled withdrawal mechanisms (new)

- metrics monitored continuously, alert thresholds,
- controlled rollback procedure in case of envelope exit (suspension, return to the previously qualified version, under human validation),
- conditions of notification to the regulator (deadline, format),
- observability contracts with clinical user sites.

D. Requalification conditions (reformed)

- modification within the envelope → ordinary compliance cycle (no referral),
- modification outside the envelope → requalification of the envelope itself,
- the multi-year PSUR ceases to be the trigger; the trigger is a surveillance event.

Application to the monthly radiology case. Monthly retraining is no longer qualified event by event. It is qualified *ex ante* in section B at the moment of initial authorization. As long as it remains within the envelope (measured drift, thresholds respected, data from the same clinical distribution), it triggers nothing. If it exits, it is the envelope that is requalified (by notification, within a defined deadline), and not each update that goes through a notified body procedure.

The notified body evaluates the pipeline and its envelope. The provider operates within the envelope and notifies the exits. Liability ceases to be ambiguous: the notified body is responsible for what it has qualified; the provider is responsible for what it operates within the qualification.

Institutional consequence. The displacement is not only procedural. It qualifies the role of the notified body differently: it ceases to be a validator of state and becomes a qualifier of trajectory. The mutation has a direct consequence on the required competences. The dominant mastery is no longer principally metrological and clinical; it becomes algorithmic and statistical (capacity to evaluate a learning pipeline, drift bounds, non-regression protocols, observability contracts). The shortage of notified bodies capable of operating in this regime is today the principal operational lock, more constraining than the legal text itself. It is also for this reason that the reform does not reduce to a regulatory amendment: it presupposes the reconstruction of a prudential competence that does not yet exist, in the required quantity, within the European ecosystem.

This is not a utopia. It is a targeted rewriting of MDR Annex IX, doubled by a policy of support to notified bodies.

9. Empirical indication: PCCP

The Food and Drug Administration framework, finalized in January 2026:

- initial authorization + declared modification protocol,
- beginning of integration of {state + trajectory}.

But incomplete:

- does not cover structural representational biases,
- nor the relevance of training data on target populations.

The PCCP is not the European model to transpose. It is a **proof of existence**: a partial holding regime is technically constructible. That is enough to forbid treating the holding regime as doctrinal utopia.

10. Counter-theses

First objection. Removing the AI Act creates a void.

Response. Cumulation was not additive, coverage was already incoherent, one does not lose a coverage that had never existed. The real question is not the HRAIS exit but the qualification of what an amended MDR can absorb, of what remains structurally orphaned, and of the distinct prudential infrastructure to be built for that zone.

Second objection. Incoherence is protective: it creates redundancy, hence safety, as in any critical system.

Response. Redundancy is protective only if it bears on the same object. A double validation of a state increases reliability. A validation bearing on the state combined with a validation bearing on the trajectory do not reinforce each other: they overlap partially and leave the non-common zones without coverage. Authentic redundancy requires a categorial compatibility that is lacking here. What presents itself as «double coverage» is, on analysis, an incomplete coverage paid at the price of a complete one. This is precisely what non-commutativity (§4) demonstrates: the order of application changes the result, a property that no authentic redundancy possesses.

Qualification of the orphaned zone.

Two HRAIS requirements resist any plausible MDR absorption. The first: traceability of training corpora derived from third-party foundation models, whose exact composition escapes by construction the provider of the final device; an amended MDR can demand documentation, it cannot produce the transparency of an upstream it does not control. The second: evaluation of structural representational biases when the device is deployed on target populations whose distribution differs from the initial training corpus; the MDR knows how to certify clinical evaluation on an indication, it does not know how to certify the robustness of a model to a distribution shift. These two zones are not particular cases: they are the generic characteristics of a trained model. They are, at this stage, structurally orphaned.

Question of coverage.

Identifying the orphaned zone is not enough. Three coverage architectures are conceivable, with distinct institutional costs.

- **Dedicated regulator.** Create a specialized authority for transversal algorithmic components, by analogy with the EDPS for personal data. Institutionally heavy solution, long calendar; potentially the strongest doctrinal coherence.
- **Transverse normative layer.** Extend the role of existing bodies (EMA, ENISA, the future European AI Office) to handle upstream requirements (corpora, biases, distributional robustness) without new institutional creation. Faster solution, uncertain operational capacity, risk of doctrinal fragmentation.
- **Insurance mechanism.** Move coverage of the orphaned zone toward a mandatory product insurance device, indexed on the qualification of the trajectory and on the history of incidents. Operationally functional solution, but one that transfers the prudential arbitration to reinsurance companies rather than to the regulator, which amounts, in practice, to delegating the political decision to the market.

None of these options is neutral. None resolves the problem on its own. But their identification is necessary for the critique of cumulation to cease being a critique and become a proposal. The thesis defended here does not arbitrate between the three; it requires only that the arbitration be posed.

11. Empirical test

Within the 24 months following any adoption.

Case 1. Convergence

- single chain of compliance,
- falling marginal cost of algorithmic update,
- emerging holding regime.

Case 2. Divergence

- persistent double chain,
- litigation on the order of application,
- the thesis fails.

Case 3. Pseudo holding regime

- double chain maintained but articulation internalized,
- marginal cost stabilized,
- industrialization absorbs what the law does not.

Case 3 is dangerous for the thesis because it makes the holding regime technically superior but operationally circumventable. Industrialization there absorbs the problem, at the price of a transfer of cost from regulators to manufacturers that becomes bearable once amortized.

This is the economic definition of a success. It is also that of an admission.

Methodological precision. The test does not bear on the persistence of cumulation until 2028 (the transition period being, by construction, the period in which the cumulation is suspended) but on the post-adoption trajectory, when the new MDR / AI Act articulation enters into force on medical devices. The 2026 to 2028 period neither validates nor invalidates the thesis: it merely shifts the horizon of observation.

12. Conclusion

The European movement is not a decision.

Nor is it a symptom. A symptom presupposes a disease one is trying to cure.

It is a mode of functioning.

The regulatory system of medical AI devices is not a system in breakdown. It is a system that has discovered it can dispense with coherence, by making it paid elsewhere.

Stabilization by absorbed incoherence.

Regulation is not coherent. The institutions know it. Resolution is deferred. The cost is transferred. And everyone carries on.

This stabilization needs no malevolent actors to hold. It holds because each party has an interest in its holding:

- the regulator avoids the political cost of an overhaul,
- industry avoids the risk of a poorly calibrated overhaul,
- manufacturers absorb (and bill) the residual complexity,
- patients endure a system they do not see, until the event that makes it visible.

Real resolution presupposes that someone agrees to stop functioning this way. This pertains neither to technical diagnosis nor to doctrinal prescription. It pertains to a political decision that no one, at this stage, has the mandate to take.

The 2028 postponement is not a step toward resolution. It is the extension of the period during which one does not resolve. As long as there is someone to absorb the cost, there will be no institutional reason to stop.

Which leaves the only question that remains, after this one: at what moment does absorption become impossible?

Typology of rupture shocks

Absorption does not exhaust itself by moral wear of manufacturers, which has never determined a European policy. It yields by the crossing of a threshold, in one of four forms. These classes are not alternatives: they may cumulate, and they reinforce each other when they coincide.

Class	Mechanism	Time to effect	Observable precursor signal
Legal shock	Member State case law requalifying an externality as a product obligation.	12 to 36 months after a founding litigation.	Multiplication of disputes on the same type of algorithmic failure.
Economic shock	Insurance crisis: premiums incompatible with margins, rational exit of absorbing operators.	24 to 60 months depending on reinsurance cycles.	Progressive tightening of warranty exclusions on AI devices.

Class	Mechanism	Time to effect	Observable precursor signal
Clinical shock	Cohort harm originating from the orphaned zone, transiting from the statistical register to the media register.	Indeterminate, structurally incompressible.	Aggregation of weak pharmacovigilance signals on indications outside HRAIS.
Political shock	Legislative window opened by a change of majority, an institutional crisis, or alignment with another major regulatory dossier.	European electoral cycles.	Resumption of the medical AI dossier within a non-sectoral political agenda.

Concretely: when an operator in bankruptcy reveals, through its proceedings, what regulation was concealing. When a Member State case law forces a clarification that had not been made. When an insurance crisis renders premiums incompatible with margins. When a political change opens a legislative window that did not exist. When a cohort of patients sustains harm from the orphaned zone, and that harm ceases to be a statistical externality to become a political fact.

At that moment, or by accumulation of events rendering the cost politically unbearable, the collision will cease to be a doctrinal abstraction. It will become a decision.

Until then, it remains an equilibrium.

Footnote

1. «Ontological collision» designates the divergence of categorial constitution of the regulated object according to the regime that qualifies it: regulation does not describe the artifact, it constitutes it as an object for itself. The term is preserved for its impact. The technical demonstration of §5 specifies it as *categorial qualification of the regulated object*.