Leaving the walls without dissolving authority

From governing cases to governing distributions, the decoupling of influence and responsibility, the integrity of execution conditions.

The public debate frames hospital-without-walls care as a liberation. Care leaves the building, the patient stays home, the system gains flexibility and capacity. This reading is not wrong. It is incomplete. Removing care from the walls does not abolish the clinical authority that was exercised within them; it redistributes the technical, informational, and operational conditions under which that authority continues to be exercised. The thesis of this text is precise. The risk is not the disappearance of a decisional center but the emergence of distributed architectures in which influence over a trajectory, the visibility of what tilts it, and assignable responsibility cease, gradually, to be aligned. This misalignment is not a scandal. It is a structural property, and it can be governed, on condition that it be named before it sediments.

I. What the wall made visible

This article is not about artificial intelligence in healthcare. It is not about the end of the hospital. It is about the conditions under which clinical authority is exercised, and about what happens to those conditions when the place that held them together dissolves. The clarification is non-negotiable, because most misunderstandings on this subject arise from confusing the transformation of a building with the transformation of a responsibility regime.

The hospital wall was not, primarily, a logistical constraint. It was a device of legibility. Within its walls, three things were in the same place without anyone having to think them separately: who knew, who saw, who answered. An adverse event in a ward had a place, a head of department, a medical commission, an institution that answered for it before the patient and before the regulator. Responsibility was not lighter than it is today. It was localizable.

Let us resist nostalgia, which would be a failure of analysis. The wall also enclosed. It produced nosocomial risk, immobilization, cost, loss of autonomy. Removing care from its walls answers real constraints, and no serious party proposes returning. The point is not that the wall was good. The point is that it made visible something we have not rebuilt elsewhere. The wall did not contain only beds. It rendered authority localizable.

This localization was procedural as much as administrative. The morbidity and mortality conference, the institutional medical commission, the on-call rotation, the hospital discharge report: as many devices as presupposed a place where a patient's trajectory was observed by an identified collective, and where an error could be replayed, discussed, imputed. Hospital medicine took more than a century to build these forms of shared responsibility. All of them rested on a tacit condition: that the trajectory and its observation take place in the same location. Distributed care removes this condition. It has not rebuilt the forms that depended on it.

Hence the distinction that opens the article. The wall did not hold care as if it had contained it. It made it legible, by spontaneously aligning three things on the same perimeter: influence, visibility, and responsibility. This alignment was free, because it was spatial. It ceases to be so as soon as care distributes itself.

Three propositions structure the demonstration that follows, and everything else is subordinated to them.

1. The first: orchestration governs probability distributions rather than cases.
2. The second: influence decouples from responsibility; the latter remains anchored to whoever decides, while influence shifts to whoever sets parameters.
3. The third: the integrity of execution conditions becomes a condition of clinical authority.

The other notions mobilized, namely architectures of influence, reconstructibility, attack surfaces, and institutional runtime, derive their value only from their relation to this triptych. It, and not the spatial motif of the wall, is the actual object of the text.

II. Distributed care already exists

A precaution before any diagnosis. Hospital-without-walls care is not a coming threat that could still be prevented. It is a reality deployed at scale, answering constraints that no serious party contests. Treating it as a future drift would be false, and defensive.

In France, home hospitalization (HAD) treated 184,400 patients in 2024, accounting for 7.7 million patient-days, delivered by more than 280 institutions, at a cost to the public health insurance of approximately 1.94 billion euros, that is to say about 2 % of the social security envelope, with an average daily price above 263 euros within territorial hospital groupings (ATIH; DREES). The dynamic is not marginal: patient-days are up 6.0 % and patients up 9.8 % over the single year 2024. Medical remote monitoring entered ordinary law through two decrees of March 3 and March 31, 2026 (HAS; Légifrance). And the coordination of complex pathways has, since the law of July 14, 2019, a dedicated public operator, the coordination support unit (DAC), whose orientation function is explicit. Distributed care is not waiting for the debate's permission. It is running.

The European picture is the same, at different governance speeds. England counted approximately 11,635 virtual ward beds in March 2025, against a target of 24,000 and an initial allocation of 450 million pounds over two years (NHS England; POST Parliament). The Nordic countries are scaling a Hospital@Home program in 2026. Germany has no national equivalent in acute care and remains centered on psychiatric home treatment. In the United States, the model is explicit and equipped: the federal Acute Hospital Care at Home waiver was extended by five years through the Consolidated Appropriations Act of 2026, deadline September 30, 2030, covering 366 programs across 139 health systems and 37 states (AMA; CMS).

Geography	**Governance rhythm**	**Deployment (dated)**	**Measurement and economics**	**Reliability**
France	Coordination instituted, alignment of responsibility not yet posed	HAD 2024: 184,400 patients, 7.7 M days, more than 280 institutions; remote monitoring in ordinary law (March 2026)	HAD ~1.94 B€, ~263 €/day, ~2 % of SS envelope	N1 institutional
United Kingdom	Mass deployment, late measurement	~11,635 virtual ward beds (March 2025), target 24,000, 450 M£ over two years	Regional savings documented; national minimum dataset expected only in 2026	N1 / N2
Nordics, Germany	Emerging in the North, near-absent in acute care in Germany	"Hospital@Home" at full scale in 2026; Germany oriented toward psychiatric home treatment	Not documented	N1 / N2
United States	Explicit and equipped orchestration	CMS waiver extended five years (deadline 2030), 366 programs, 139 systems, 37 states	Operator-claimed reductions, to be treated as allegations	N1 framework, N3 results

This panorama is not a ranking. It says one thing only, which is the diagnosis of this section. Operational scale advances faster than governance regimes. The British evidence is clear and limpid: 11,635 home care beds are being deployed, and the minimum dataset that will allow them to be evaluated is planned for 2026. That is to say, afterward. We deploy first. We measure later. We govern later still, if we govern at all.

It will be objected that this is simple maturation, and that governance will follow use as always. The objection is serious, but misses the point. It is not maturation that poses the problem, but the asymmetry of its speeds. When deployment matures without governance maturing at the same pace, the gap is not a transient lag. It becomes the structure within which the system operates. The medico-economic evidence itself forbids complacency: the available systematic review situates the effects of home hospital care on an interval ranging from savings above 8,000 euros per patient to additional costs above 2,000 euros (systematic review, PMC). This heterogeneity does not say that the model is bad. It says that it is not economical by nature, and that it is so only under organizational conditions. And organization is precisely what is distributing itself.

III. The perimeter of care becomes informational and operational

Care leaves the walls. Its perimeter does not disappear for all that. It changes nature. The wall delimited a physical perimeter, made of concrete and corridors, whose primary property was being visible to the naked eye. In its place an informational and operational perimeter installs itself, made of what is captured, seen, triggered, escalated, and prioritized. The word to avoid is instrumentation, which would reduce the subject to sensors. The question is not technical. It bears on visibility and on action: what the system sees, what it triggers, what it surfaces, what it leaves out of frame.

The inclusion boundary then shifts. The wall excluded by space: whoever was not admitted was not in the hospital, and everyone knew it. The informational perimeter excludes by absence of signal: whoever is not instrumented is not observed, and no one decided it. A perimeter bias, already named in our earlier work, takes here an unprecedented spatial form. The system sees only what it has rendered observable. The patient without a connected device, or whose pathology does not figure among the inscribed remote-monitoring activities, is not refused. They are invisible to the layer that orients. Care abandonment, which we have analyzed elsewhere as a silent underuse of the trajectory, finds here a new mechanism of emergence: one does not exit care by a decision, one exits by a default of capture.

The perimeter is also a regulatory boundary, not a pure affair of sensors. Not all connected objects are medical devices: a consumer wristwatch recording a heart rate is not, in itself, a reimbursed remote-monitoring device. The informational perimeter therefore composes itself from heterogeneous strata, some enforceable and reimbursed, others floating and unframed. The same patient may traverse these strata without anyone orchestrating the coherence of what happens to their observation.

An example fixes the stakes. A heart-failure patient under remote monitoring wears a connected scale and a blood-pressure monitor; their weight gain crosses a threshold, the

alert rises, a call is triggered, their incipient decompensation is caught in time. The following month, a battery has not been replaced, or the decompensating pathology does not fall within the inscribed activity. The same deterioration unfolds, but out of frame. Nothing was decided. The perimeter retracted through a material defect, and the trajectory exited observation without anyone pronouncing the exit. Within the walls, a patient ceases to be observed when they leave the hospital, and that departure is an act. Within the informational perimeter, they cease to be observed when the signal ceases, and that silence is no one's act.

In a distributed system, visibility ceases to be a simple clinical property. It becomes an informational dependency. A trajectory may cease to be observable not only when a patient exits the perimeter of care, but when a flow is interrupted, delayed, corrupted, desynchronized, or rendered unavailable. The capacity to observe then depends on the integrity of a technical execution chain, whose properties themselves become conditions for the exercise of clinical authority.

This dependency opens a second form of exclusion, more insidious than the absence of signal. The perimeter can exclude by silent alteration of the signal itself. In a perimeter governed by flows, seeing falsely sometimes becomes more dangerous than not seeing at all. Incomplete telemetry, a desynchronized threshold, an altered datum, a loss of availability do not produce only an IT incident. They modify the very conditions under which a trajectory becomes visible, and therefore orientable.

The distinction that closes this section thereby gains a third entry. The wall produced a decided exclusion, and therefore in principle contestable: one could ask why one had not been admitted. The informational perimeter produces an exclusion without decision, therefore difficult to contest: one does not litigate an absence of signal. And it now produces an exclusion without veracity, more difficult still to grasp: one does not litigate an altered signal one does not know is altered. The system does not see badly by accident. It sees according to a perimeter that no one has instituted as an act, and whose integrity no one guarantees as an act, although it produces all the effects of one.

IV. Orchestration, from governing cases to governing distributions

It remains to name what acts on this perimeter. Distributed care does not hold by magic. It is coordinated. Control tower, command center, orchestration platform, support unit: the names vary, the function describes itself without needing personification. And describing it correctly requires dismissing at once the polemical, false, and convenient formula according to which the platform would govern in the physician's stead. That is not what happens. Maintaining it would weaken the analysis.

Two too-seductive qualifications must be ruled out.

1. The first would speak of a "clinical function" of orchestration.
2. The second, more cautious, of a "function of structuring clinical trajectories".

Both carry the same defect: they pull the subject toward the clinical, and therefore toward the frameworks that already govern the medical act and the software. If orchestration is clinical, then the medical devices regulation, software qualification, the AI Act, and product liability seize it, and the question closes before being posed.

Yet the real subject is deeper, and it lies elsewhere: it is the architectures of influence. I therefore retain a deliberately less elegant and harder-to-attack qualification. *Orchestration is an operational orientation function with clinical effect*. Operational, because it performs no medical act. With clinical effect, because its settings determine care trajectories. This point, a clinical effect without a clinical act, is what escapes existing frameworks, which think responsibility as the attribute of a decision.

The thesis is cold, and difficult to refute. Orchestration architectures do not exercise direct medical decision. They fix the conditions under which certain trajectories become more probable, more visible, more rapid, or more accessible than others. They do not sign the prescription. They compose the terrain on which the prescription will, or will not, be written. Orchestration does not decide. It tilts.

The mechanism is without mystery, and is worth describing flatly. The device captures a flow of remote monitoring and medical objects. It qualifies this flow through thresholds, which transform a measurement into an alert. It prioritizes alerts according to a window of actionability. It routes the demand toward home, specialist, hospital, or wait. Each of these four operations falls under informational logistics. None has the form of a clinical decision. That is what makes them difficult to govern: non-decisional in their form, decisional in their effects. The French capacitive control function provides a measured illustration. In Valenciennes, a system predicts hospital activity at 48 hours with a reliability of about 95 % (esanum). It refuses no patient. It tilts flows, and a tilted flow, repeated at the scale of a territory, ends up determining who will be seen, when, and in what order.

This finding must be given its full weight, for it is here that the central concept of the text is lodged. A classical governance governs cases. It admits or refuses this patient, treats or abstains, orients here or there. Each act bears on an individual, and each act leaves a decision to display. Orchestration does not govern cases. It governs distributions of probability of access. It does not decide that a given patient will stay home; it adjusts a parameter that renders a fraction of the population more probably maintained at home, and another more probably summoned. The governed object is no longer the trajectory; it is its probability.

The concept deserves a settled definition, failing which it would remain an intuition.

A probabilistic governance acts not principally on identifiable individual decisions, but on the parameters that modify the distribution of probabilities of access, orientation, or care within a population.

This definition distinguishes itself from three neighboring notions with which it is confused. Capacitive triage sorts existing cases under bed constraint: it acts on individuals, through an explicit decision rule, however expeditious. Populational optimization pursues an aggregate objective, for example reducing avoidable hospitalizations: it is a finality, not a mode of action. Statistical governance steers through indicators and targets: it orients through measurement, but via identifiable decisions and policies. Probabilistic governance recognizes itself by a property these three notions do not have. It acts on the parameters that deform the distribution, and its effect cannot be attached to any individual decision. It is neither an objective, nor a sorting, nor a dashboard. It is a mode of government whose unit is not the case but the distribution, and whose signature is the absence of a contestable decision.

The clinical materialization is direct. Take two comparable patients, same pathology, same severity. The first sees their indicator slightly exceed a low threshold: they are summoned. The second, under a threshold raised the previous week to preserve beds, stays home. No physician arbitrated between them. The difference in trajectory is the effect of a parameter. At the scale of a cohort, the same setting shifts a few percentage points of admissions, that is to say, across a territory, hundreds of trajectories tilted without a single individual decision intervening. The threshold did not choose a patient. It deformed a distribution.

This displacement, from a governance of cases to a governance of distributions, is precisely what responsibility frameworks do not know how to grasp. They are built for acts. An act has an author, a date, a file, a recourse. A distribution has none of that. It is no one's act; it is the effect of a setting that presents itself as a technical operation. Probabilistic governance is, by construction, a governance without an individual decision to contest. That is its operational strength, and that is also what renders it, in the present state, ungovernable by the tools at our disposal.

There remains the question that technical discourse most willingly elides. Who sets this threshold, and in the name of what? The answer is never purely clinical. Parameter setting is decided in bed-management cells, by department managers, executives, platform editors, sometimes under the orientation of a regional agency or a payer. And the objectives it encodes are concrete, nameable, legitimate when taken one by one: reduce avoidable hospitalizations, protect emergency rooms from saturation, optimize counted nursing resources, hold a budgetary trajectory. None of these objectives is illegitimate. None is, however, clinical in the sense of the benefit to the patient in front of one. An orientation threshold does not encode a clinical truth. It encodes a trade-off, and this

trade-off is rarely purely medical. Orchestration therefore encodes institutional priorities. The misalignment does not arise on its own, by a natural entropy of systems; it is produced, setting by setting, by very concrete organizational arbitrations. Computation does not create these arbitrations, which preexisted in any constrained care organization. It makes them scalable, distributed, and removed from local clinical deliberation.

This analysis does not start from zero. The grammar of an operable chain, from observation to audit, has been laid down regarding predictive vigilance; the notion of optimization perimeter, which separates what a system claims to optimize from what it actually optimizes, has been laid down regarding coverage and care-regulation systems. I assume them acquired and refer to them. The point proper to this article lies elsewhere: this orientation function, because it produces an effect without posing an act, and because it acts on distributions rather than on cases, escapes by construction the frameworks that attach responsibility to a decision.

It will be objected that orchestration is merely a neutral support, an efficiency tool at the service of the clinician. The objection is serious; let us formulate it better than the opponent would. A tool that executed rules fixed and traced by the clinician would, indeed, be neutral. But a device that fixes thresholds, orders priorities, and filters visibility does not merely execute. It composes the conditions of the decision upstream of the decision. Neutrality ceases where the composition of the terrain begins. This thesis is a structural hypothesis, and it has an explicit falsifier: if, for an unfavorable trajectory, one can reconstruct which parameters tilted it and attach an enforceable responsibility to whoever set them, then orchestration remains governed, and my thesis falls.

A final property follows directly from what precedes. Insofar as these parameters govern distributions at scale, their alteration ceases to be a local technical incident. An orientation threshold is not a simple setting: it is a point of systemic modulation, and as such, an object of power. The layers that tilt trajectories become institutional attack surfaces. A silent modification of a threshold, a routing corruption, a desynchronization between rule versions, a partial unavailability of flows immediately displace probabilities of access to care, without any individual decision appearing irregular. Returning to Valenciennes: if the system predicts correctly but the parameter steering orientation is altered, mis-propagated, or executed differently from what had been validated, the effect is no longer computational. It is institutional and clinical, because it deforms a distribution of access. The cybernetic problem is therefore no longer only that of data confidentiality or of system availability. It becomes that of the integrity of conditions of visibility and orientation. Cybersecurity no longer protects an infrastructure alone; it protects the conditions of execution of an architecture of influence.

V. Misalignment and responsibility

The previous section established where influence is lodged, what it governs, and how it can be altered. It remains to follow what happens to responsibility when influence shifts thus. This is the heart of the problem. Within the walls, three things held together: influence over the trajectory, visibility of what tilted it, responsibility for what resulted. Distributed care does not abolish them. It desynchronizes them, and it is this desynchronization, not any dispossession, that constitutes the institutional risk.

The drift has a direction. Influence migrates toward layers that do not decide: rule engines, routing logics, queue prioritization, visibility filters. These layers sign nothing, yet they determine what surfaces, what becomes prioritized, what stays invisible. Monday morning, someone raises a threshold. Friday, a physician validates a maintenance at home without knowing that the bar has moved. Responsibility remains anchored to the signatory; influence has passed to the parameter setter. Whoever sets the threshold is not whoever answers for the trajectory. There lies the asymmetry, and it is human before being technical: it opposes those who bear the burden of the error to those who compose its conditions.

To asymmetry, dilution adds itself. In a distributed chain, an unfavorable trajectory has no single author. It is the composed effect of locally compliant links, none of which, taken alone, has committed a fault. The system may remain everywhere compliant with its specifications and produce, at scale, an attrition of access that no decision ordered. This point I have established elsewhere; I mobilize it for what it does here, rendering responsibility not absent but unfindable, because diffracted across a chain whose every segment refers to the next.

This dilution has a grammar that medicine knows poorly. It knew how to treat diagnostic error. It does not yet know how to treat the orientation error produced by a system. Four cases, nevertheless: the relevant orientation, the orientation in error, the justified absence of orientation, the unwarranted absence of orientation. This last case, the unwarranted silence, is the most fearsome. It leaves no trace in activity indicators: nothing was done, therefore nothing appears. A trajectory that should have tipped toward the hospital and did not is visible nowhere, except in the decompensation that follows, which one will then impute to anything but the threshold that let it pass.

Hence the resistance of clinicians, which efficiency discourse files too quickly among archaisms. The reading is lazy. When a clinician balks at validating orientations produced by an architecture whose thresholds and filters they master neither, they are not defending a territory. They perceive, often better than the designers, that they are being asked to bear a responsibility whose conditions escape them. Their resistance is not a brake on progress. It is a structural symptom, and an early indicator of misalignment.

Treating it as an obstacle to be removed, rather than as a signal to be read, amounts to unplugging the alarm because it bothers.

The strongest objection must be faced head-on, for it is just in its letter. Responsibility frameworks already exist, it will be said: prescriber, HAD coordinating physician, coordination support unit, data controller in the sense of the GDPR, operator of the remote-monitoring device. Each carries a defined responsibility. That is exact, for formal decision-makers. But these frameworks assign responsibility to whoever decides, when influence has migrated to whoever sets parameters. The gap is not the absence of responsibility. It is its decoupling from the actual locus of influence. It will be objected further that all of this is the lot of any complex organization, and that one is overselling a banal management problem. The answer is clear. What is proper to distributed orchestration is that the influential layer does not decide in its form, and therefore remains invisible to the frameworks that attach responsibility to a decision, and that it scales faster than it is measured. A management problem is corrected by the organization chart. A structural misalignment is corrected by doctrine.

It would, however, be dishonest to present every dissociation as a pathology to eradicate. A part of it is functionally necessary. It is what enables scale, continuity of follow-up, smoothing of pathways, coordination of a territory. Requiring that a clinician re-decide explicitly each threshold, each routing, each prioritization, for every patient, would amount to dissolving orchestration, and with it the gains that justify distributed care. Perfect alignment is not only impossible in a distributed system; it would be paralyzing. The target is therefore not total alignment, which is a regulatory myth, but a bounded misalignment whose effects remain visible, reconstructible, contestable. The problem has never been that a gap exists. The problem is that this gap has become uninspectable.

It remains to say what "reconstructible" requires, and what it does not. In distributed care, the integrity of execution conditions becomes a property of clinical governability. An auditable threshold has meaning only if one can demonstrate that the threshold executed was indeed the one that had been validated. A trace has value only if the conditions that produced it are reliable, synchronized, and enforceable. Without this execution trust, reconstructibility remains purely declarative. But one must guard here against excessive demand, which would turn against the thesis. Reconstructing does not mean reconstructing everything. Real distributed systems will remain partially opaque, probabilistic, emergent, incomplete, and to require total observability would be as unrealistic as it would be paralyzing. The correct measure is not the exhaustive reconstruction of reality; it is a sufficient reconstructibility to attribute and contest a significant influence. The criterion is not to see everything. It is to be able, when a trajectory has gone wrong, to trace back to the parameters that tilted it and to designate who answers for them.

The problem then changes nature. It is no longer only a matter of knowing who decides. It is a matter of knowing whether the environment that tilts trajectories can be inspected, attested, reconstructed, and contested to the extent necessary for a significant influence to remain attributable. And one must hold this thesis within its limits, on pain of caricaturing it. Misalignment is a tendency, not a law. Some architectures maintain alignment, by design or by regulation. The assertion that influence has, on average, migrated toward layers that do not decide is not measured at population scale, any more than the efficacy of the British deployment is. It remains a structural hypothesis, with a known falsifier. It is not an established fact, and presenting it as such would be to commit the error that this text reproaches in efficiency discourse.

VI. Re-aligning what must be, governing the rest

The consequence for the decision-maker, whether institutional director, regulator, or platform manager, is not to rebuild walls. The lost territoriality cannot be restored, and there is no nostalgia to entertain. Nor is it to aim at perfect alignment, which would be paralyzing and which the previous section ruled out. It is more exact, and more demanding: to re-align what must be re-aligned, and to render governable the misalignment that subsists.

Five enforceable principles fix this work. First: every critical orchestration function must be explicitly qualified, in the sense of an enforceable executory qualification, and not left in the state of an implicit technical parameter. Second: orientation thresholds must be auditable, for a threshold is an institutional act disguised as a setting. Third: causal visibility must be reconstructible, which presupposes at minimum the version of the model or rule, its execution configuration, the clinical context, the trace of its modifications, and the enforceable integrity of the conditions under which these parameters were executed. This reconstructibility is not total observability, impossible and paralyzing; it is the sufficient one that allows attribution and contestation of a significant influence. A reconstructibility that bore only on logs, without guarantee on the execution conditions that produced them, would not permit establishing a governable responsibility. Fourth: contestability must exceed simple logging, for a log is not a recourse. Fifth, which commands the other four: responsibilities must follow the actual capacities of operational influence, and not organization charts or signatures alone.

These principles translate into verifiable requirements. Qualifying an orchestration function means naming a responsible party for its parameter-setting and the class of decisions it tilts. Rendering a threshold auditable means preserving who set it, at what value, since when, and in the name of what trade-off. Rendering visibility reconstructible means being able to replay an orientation a posteriori to the extent necessary for its attribution, and not pretending to record everything. Effective contestability presupposes

that a patient or a clinician can not only consult a trace, but open a recourse that compels an answer.

These requirements call for a legitimate objection: how is all of this measured? The framework is today structural; it must become measurable, on pain of remaining a well-built intuition. Four observables can be proposed, as program more than as result. Detecting a probabilistic governance: identify the parameters whose variation displaces, at constant population, the distribution of orientations, and measure this sensitivity. Measuring the migration of influence: estimate the share of the variance of access trajectories explained by orchestration settings, compared to that explained by formal clinical decisions. Objectifying decoupling: for each critical parameter, verify whether the one who sets it figures in the responsibility loop, and report the number of influential parameters relative to the number of those for which a responsible party is identified. Detecting excessive misalignment: track the rate of unfavorable events that cannot be traced back to a parameter and a responsible party, and the rate of unwarranted silences not reviewed. None of these measures exists today in stabilized form. Their absence is not a detail of implementation; it is, in itself, an indicator of the present low governability.

A calendar discipline completes these principles. The British lesson is limpid: one does not measure after having deployed; one measures before extending. The minimum dataset awaited for 2026, while the beds have been deployed for years, is the example of what must not be reproduced. Measuring before scale is the condition for re-alignment to remain possible while it is still reversible.

PREDICARE, within the French perimeter, constitutes a relevant ground on which to test these principles, and not a proof that they work. The precision is doctrinal, not rhetorical. PREDICARE is a territorial experimentation of predictive vigilance on chronic trajectories at risk of decompensation [exact perimeter, target pathologies, partners, and calendar under fact-lock for the PREDICARE series, not lifted at this date]. What makes it an interesting ground is not that it predicts better, but that it explicitly poses, in its design, the separation of critical functions and patient contestability, that is to say the very conditions of re-alignment. The ground does not demonstrate the doctrine. It tests its feasibility.

A limit must be set without crossing it. The fine distribution of legal responsibility between suppliers and the coordinating operator, in home hospitalization as in remote monitoring, falls under specialized legal counsel that this text does not pretend to formulate. I signal the question; I do not settle the law. My object is to show that influence and responsibility have decoupled, and that this decoupling is governable.

This point displaces security beyond hospital cybersecurity in the classical sense. In distributed care, security no longer protects only data, workstations, or infrastructures. It protects the conditions of visibility, orientation, reconstructibility, and contestability of trajectories. The wall produced an implicit trust by physical co-location. Distributed care must produce an explicit trust, through execution proofs, enforceable traceability, and

reconstructible integrity. Without this, orchestration architectures remain capable of influencing trajectories without anyone being able to demonstrate, after the fact, under what conditions that influence was exercised.

What has just been described exceeds healthcare, and this must be said in one sentence rather than developed at length here. Any authority that distributes itself now depends on the execution properties of the chains that carry it, and there lies the seed of a runtime theory of the institution of which healthcare is only an advanced case. That is the object of another text; I limit myself to naming it.

At the end of the path, the transformation states itself without dramatization. The wall rendered authority localizable without one having to think about it. Distributed care does not render it unfindable at a stroke; it renders it gradually more difficult to reconstruct, as influence lodges in layers no one inspects. There is no single tipping point, no spectacular point of no return, but a slow drift: thresholds modified without trace preserved, routings whose responsible party is not identified, absences of alert that nothing revises. Between perfect alignment, which is a regulatory myth, and complete dissolution, which is a strawman, extends the actual space of governance: partial alignments, inspectable compromises, hybrid yet reconstructible governances.

The three propositions posed at the outset meet here. Governing distributions without a contestable decision, letting influence decouple from responsibility, and neglecting the integrity of execution conditions are not three distinct problems. They are three faces of the same renunciation. Leaving the walls does not compel us to dissolve authority. It compels us to re-index it on actual influence, and to render governable the gap that subsists. Failing which, authority does not disappear in one piece: it becomes, trajectory after trajectory, progressively unfindable.

It must then be stated more precisely still. Distributed care does not redistribute trajectories alone. It redistributes the technical conditions that render those trajectories visible, orientable, and contestable. From then on, cybersecurity ceases to be a peripheral computing layer. It becomes a constitutive property of the very possibility of a distributed clinical authority.