

# Four “dimensions,” one arbitration layer

**Under operational constraint, the orchestration layer does not merely render care, place, time, and cost comparable: it renders certain trajectories calculable and others invisible.**

The hospital that steps outside its walls presents itself with a compact promise, imported from the vocabulary of *value-based care*: the right care, in the right place, at the right time, at the right cost. *Wherever needed, whenever needed, the right care at the right cost.* The formula is serious. It is no empty slogan: French *hospitalisation à domicile* was already treating more than 184,000 patients in 2024 (ATIH), medical telesurveillance entered mainstream coverage through the decrees of 3 and 31 March 2026 (HAS, Légifrance), and British *virtual wards* counted close to 11,635 beds by March 2025 (NHS England, POST Parliament). Distributed care exists, works, and keeps part of its promise.

But the promise rests on a presupposition it never names: that its four “rights” are separate objectives, which one could tune one by one, like four independent sliders. This text defends the opposite thesis. Right care, right place, right time, and right cost are not four sliders. They are four effects of one and the same parameterization layer, and the mechanism that produces them is not arbitration between competing objectives. It is deeper, and less visible: **under operational constraint, the orchestration layer constitutes the space within which these objectives become calculable. It renders certain trajectories observable, and others invisible. A distributed system ends up optimizing what it can render visible; the cost of what it does not see never appears in its accounts.**

Three precautions, laid down once.

1. This text is not a critique of value-based care, nor of payers: it describes an architecture, not an intention.
2. It does not address accounting cost, nor the appropriateness of the hospital-outside-its-walls model, which is not at issue.
3. It advances a falsifiable structural hypothesis, not a demonstration: the arbitration functions of the platforms are proprietary, and no one can produce, at this stage, direct evidence of what plays out within their runtime. The thesis is worth what its falsifier is worth, stated below.

## I. The promise of the four “rights” presupposes their separability

The genealogy of the formula is well known. It descends from the Institute for Healthcare Improvement’s *Triple Aim* (better population health, better care experience, controlled cost), extended into the *Quadruple Aim* by Bodenheimer and Sinsky in 2014 (adding clinician well-being), then into the *Quintuple Aim* by Nundy, Cooper and Mate in 2022 (adding equity), and reformulated by market discourse as an orchestration imperative: a performant system would be one that orchestrates interoperable data, digital clinical programs, access experience, and payment structures.

The promise “right care, right place, right time, right cost” is its consumer-facing version.

Its strength lies in its plausibility. Each of the four “rights” corresponds to a real lever, already deployed.

1. **The right place:** arbitration between conventional hospitalization, hospital-at-home (HAD), community care, and the medico-social sector.
2. **The right time:** early detection through telesurveillance, which moves intervention upstream of decompensation.
3. **The right care:** matching need to resource.
4. **The right cost:** substituting a day of conventional hospitalization with home-based care, whose order of magnitude is documented (in the UK, Medway reports £187 per day on virtual ward versus £657 in hospital, NHS data; in France, the HAD day within a Hospital Group exceeds €263, ATIH 2024).

Presented this way, the four rights look like a dashboard. Four dials, four dimensions of value, a pilot who adjusts at the margin. The metaphor reassures because it presupposes the dials independent: one would pull on cost without deforming care, advance the moment without displacing the place, and the only difficulty would be one of weighting.

It is false, and for a technical, not ideological, reason. The four rights do not live in the same *measurement space*. And that space is not given in advance: it is produced by the very layer that claims merely to measure it.

## II. Separability is fragile, because each “right” depends on a perimeter of observability

Let us pose the question the promise sidesteps: what defines “right,” and on what does that definition depend?

- **The right care** is defined relative to an observed need. But need is not observed in itself: it is inferred from what the system captures. A degradation emitting no instrumented signal is not a need for the system; it is a silence.
- **The right place** is defined relative to a map of resources, and that map is not neutral: it references what is contracted, instrumented, connected.
- **The right time** is defined relative to a detection, hence to a threshold and a sampling frequency.
- **The right cost**, finally, is defined relative to a perimeter of counted costs: a cost that is just over a narrow perimeter, the avoided day, may be unjust over the full pathway, the deferred readmission.

None of these four perimeters is a fact of the world. All are fixed by the parameters of orchestration. Hence a distinction the dashboard erases: that between the *objective* and the *measurement space* in which the objective is defined. Optimizing an objective presupposes a stable measurement space. Yet the measurement space of distributed care is not stable: it is itself a product of parameterization.

Metrology has known this for a long time. Three facts, by now laboratory commonplaces: the instrument modifies the phenomenon it measures; every measurement space already encodes arbitrations in the choice of what counts; comparability is never neutral. What changes here is not the principle but the scale. No longer an isolated instrument in a laboratory, but a continuous clinical orchestration, under constraint, over an entire population, with direct consequences for patient addressability.

Rendering the four rights comparable has a learned name, *commensuration*, and a more practical working name: the manufacture of the *shared decision space*. To arbitrate between treating here rather than there, now rather than tomorrow, at this cost rather than another, the system must first project these heterogeneous dimensions into that shared space, where they become confrontable. The operation is not a neutral prelude to the calculation. It is already the calculation, and it decides most of the answer.

### III. Under constraint, this shared space is a dynamic compression that manufactures visibility

If this shared space were only a static projection, it would remain a conceptual elegance, and one would be right to see in it a floating abstraction. The decisive point lies elsewhere. **In execution, under constraint, the measurement space is not built once: it is continuously reduced, deformed, recomputed. It is not a projection. It is a dynamic compression.**

Five runtime forces perform this compression. None appears on the dashboard.

1. **The feedback loop:** what the system rendered visible yesterday conditions what it can detect today. Triage models retrain on already-observed trajectories, and a class underrepresented in the history remains underdetected in the present. The observable reproduces itself.
2. **Recalibration temporality:** a model drifts. The question is not whether it drifts, but how often it is recalibrated. Between two recalibrations, the system measures the present with the measurement space of the past.
3. **Missing data:** what is not instrumented is not calculable, and what is not calculable does not enter into arbitration. A missing data point is not a neutral hole in a matrix. It is a trajectory that exits the field.
4. **Capacity dependence:** routing follows capacity as much as need. When the downstream bed is missing, the effective escalation threshold shifts, whether or not the written rule says so. Capacity-planning tools, such as bed-occupancy forecasting deployed in French hospitals, make this dependence explicit: availability becomes an input to the routing calculation.
5. **Saturation:** under load, priorities do not narrow; they reconfigure. The same modeling that, in nominal regime, jointly optimizes triage, staffing, and beds (the hospital-operations digital twins described in the 2026 literature, JMIR, claim emergency-department wait-time reductions of 20% to 40% and throughput gains of roughly 20%, figures with strong vendor coloration) produces, under saturation, a different hierarchy. Saturation does not degrade uniformly. It redraws the map of what counts.

A scene, to embody the mechanism. A heart-failure patient is monitored at home by telesurveillance. One evening, the downstream beds of the reference cardiology unit are saturated. No written rule states it, but the effective escalation threshold rises: for lack of a place to refer to, the system reclassifies as “stable” what it would, in nominal regime, have flagged as “to be monitored.” The patient’s weight and heart-rate variations remain beneath the threshold thus displaced. No alert fires. The patient is not poorly monitored: they are correctly monitored in a measurement space that has tightened, without anyone having decided so. Decompensation arrives a few days later. Late. One will conclude it was not predictable. It was. The trajectory had simply never become calculable.

The system did not observe a stable patient: it *instituted* the patient’s stability, as a side effect of a capacity constraint. This is the hardest point in the text, and it lends itself to two misreadings that must be set aside. It is not the sociotechnical commonplace that every system shapes what it

measures: here the object shaped is not the measurement, but the set of trajectories that become clinically actionable. Nor is it a constructivism denying the real: the patient's degradation is real; what the system constitutes is its *calculability*, that is, whether it becomes an object of decision. **The clinical status of a trajectory depends on its regime of calculability.** This is a property of execution, measurable and falsifiable, not a philosophical position.

Hence the exact name of the mechanism, more accurate than "arbitration": the layer *manufactures institutional visibility*, and since a system arbitrates only over what it computes, some patients become structurally more addressable than others, without any explicit decision having posed this. The shortest formulation is also the most disquieting: **tuning a threshold does not move a value on the scale, it moves the scale itself.**

The objection must be given its full force. An engineer will say that all of this is a multi-objective optimization problem already solved: a *Pareto front* handles competing objectives with explicit trade-offs, and sharing sensors is not sharing an objective. The objection is correct on its own terrain, and it is the terrain that is in question. A Pareto front presupposes objectives defined on a stable measurement space. Here, the space is endogenous to the optimized parameters: moving a threshold moves both the objective's value and the set of trajectories over which that value is defined. One does not arbitrate on a fixed front. One deforms the front by moving along it.

The terrain of this text is the hospital outside its walls, but the mechanism is in no way specific to hospital-at-home. The same layer governs emergency-room triage, waitlist prioritization, imaging routing, primary-care–hospital coordination, remote monitoring platforms, oncology coordination, critical-flow management. Wherever a distributed clinical architecture arbitrates under constraint, it manufactures the observable before measuring it. The structure overflows even the clinical: any distributed decision system under constraint may exhibit it. But it is in health that it becomes most visible, because saturation there is permanent, trajectories are vital, and arbitrations are clinically contestable.

#### IV. Therefore the right cost is, first, a cost of non-observation

Once it is granted that the layer manufactures the observable, "right cost" changes nature. It is not a financial quantity measured after the fact. It is pre-inscribed in the thresholds, routes, priorities, and filters that decide who will be seen, when, where, and with what level of escalation: a clinical arbitration prior to any measurement.

What remains is to count a cost that classical economic analysis does not know how to see, because it bears on what the system did not observe: **the cost of non-observation**. It is not a bag in which to dump every negative externality. It has a causal structure, in three stages that follow the path of a signal.

1. **First stage, the cost of non-detection:** the degradation that never crossed the threshold, or crossed it too late. A claimed sensitivity above 80%, consistent with a literature predominantly North American and second-hand on remote monitoring, leaves at most one degradation in

*Jérôme Vétillard, VP R&D + Engineering / Chief Product Officer TweenMe® by Qualees®*

five beneath the threshold. It measures what the system catches; it also names, in the negative, what escapes it. Visibility delay belongs to the same stage: a detection arriving after the point where it would have changed the outcome is a non-detection unaware of itself.

2. **Second stage, the cost of non-escalation:** the detectable trajectory that the system did not escalate, for lack of downstream capacity or because saturation had reconfigured priorities. This is the cost of the preceding scene.
3. **Third stage, the cost of non-reconstructibility:** after an unfavorable outcome, the impossibility of reconstituting which parameters tipped the trajectory. The error becomes an event without assignable cause.

None of these three stages appears in an ex-post cost-effectiveness analysis, by construction: the analysis measures the trajectories the system has rendered visible. This illuminates an awkward fact in the data. A systematic review of hospital-at-home care (2018 indexed prices) reports differences ranging from savings of more than €8,000 to additional costs of more than €2,000 per patient. An economist will see case-mix there, and will be partially right: severity explains part of the dispersion. But case-mix does not predict a change of sign. An independent cost slider would produce a tight distribution around an average saving; what one observes is a swing, from savings to overruns, consistent with a cost that emerges from local parameterization rather than one piloted as an autonomous quantity. Consistent is not proven: the thesis remains structural.

The same contradictor will object that all this is mere operational *tuning*, the art of adjusting parameters for efficiency. Tuning adjusts quantities within a stable real. Here, the object shaped is the *observable real* itself. This is not finer tuning: it is a change in the nature of the object tuned.

One asperity, to close, because the mechanism rarely presents itself as a degradation. It presents itself as a success. An orchestration system can genuinely improve its indicators, fewer readmissions, shorter lengths of stay, as some North American programs claim, while increasing the share of trajectories it does not render calculable. The two movements do not contradict each other: it performs on what it can see, and what it does not see degrades no indicator, since it does not figure there. **The more a system performs on its observable, the less the cost of what it does not observe is reconstructible. The efficiency curve and the silence curve can rise together.**

This observation does not call for denunciation; it requires an engineering discipline, in four design requirements.

1. **The arbitration function must be explicit and inspectable:** an unreadable objective function does not suppress the clinical decision, someone did set it, it renders it unassignable, and soon irreconstructible as a clinical act. The problem is not the absence of decision; it is the loss of its author.
2. **The perimeter of observability must be declared in the specification,** including the missing-data policy and the recalibration temporality: a system that conceals how it handles data absence conceals half of what it decides.

3. **Every cost or capacity lever must be traceable to its effect on the observable**, failing which tuning becomes an unqualified clinical act; the doctrinal corpus this text extends names this requirement *opposable executory qualification*, which places a function with clinical effect among clinical functions and not among technical utilities.
4. **Tuning must be reconstructible**: not recording everything, but being able to replay a trajectory together with its parameters and the state of saturation in which it played out. A log attests that things happened. It does not say how they were rendered calculable.  
**Reconstruction is not logging.**

These requirements are not compliance, whose fine-grained analysis falls outside this text's scope. They are the condition for the governance question to remain *posable*: no one is held responsible for an influence buried in an implicit observable. The governance of distributed care is not a regulatory supplement. It is a property of execution, or it is nothing.

## V. Limits

Three limits bound what this analysis authorizes one to conclude.

1. **It is structural, not empirical**: platform arbitration functions are proprietary, and no direct evidence of the mechanism is produced here. Hence the falsifier, which serves as guardrail. *The thesis falls if one exhibits, in a real architecture, a cost or capacity lever whose modification leaves the distributions of access, routing, and escalation invariant.* If the observable does not move when one reparameterizes, it is not manufactured, and the analysis is wrong.
2. **It holds only under constraint**. A well-instrumented system, at ample capacity, recalibrated often, manufactures little differential visibility. The thesis bites as distribution, constraint, and recalibration debt increase.
3. **It describes a tendency, not a fatality**. That the layer institutes the observable does not require doing so in silence. This is the entire purpose of the four requirements.

## Conclusion

The promise of the four “rights” is tenable, though not necessarily in the terms in which it formulates itself.

It presupposes four independent sliders. It operates through a single layer that, under constraint, decides what becomes calculable.

**The true subject of distributed care is neither cost, nor place, nor time, nor quality. It is the manufacture of the observable: which trajectories a system knows how to render visible, and at the cost of what invisibility for the others.**

So long as this manufacture remains implicit, the efficiency measured after the fact will appear neutral. It is not. It measures only what the layer was willing to render observable.

The rest figures in no dashboard. It becomes a blind spot, then an out-of-perimeter item, without anyone having had to decide it should be so.

**What a system cannot render calculable, it eventually ceases to treat, for lack of ever having decided to.**

Making this mechanism explicit is not a moral demand. It is the first technical specification of a distributed care for which one could still answer.

## Sources

- **Hospitalisation à domicile (HAD), 2024 activity:** 184,400 patients managed and 7.7 million days in France; growth primarily driven by palliative care and patients aged 80 and over. ATIH – Analyse de l’activité hospitalière 2024 (HAD); ATIH – Analyses activité hospitalière 2024; DREES, Établissements de santé – édition 2025. Avoid the isolated assertion of an “HAD day above €263 within a Hospital Group,” methodologically under-stabilized and ambiguous as to its tariff perimeter.
- **Medical telesurveillance under mainstream coverage:** integration into mainstream coverage since 2023; regulatory updates by decrees of 3 and 31 March 2026; HAS methodological and clinical doctrine on eligibility, evidence and organization conditions. Légifrance; HAS – Télésurveillance médicale.
- **Virtual wards (England):** approximately 11,635 operational beds across the NHS in March 2025; the economic comparison frequently cited for the Medway programme (£187 versus £657 per day) should be interpreted as a local use case and not as a national generalization. NHS England / NHS Digital – Virtual ward capacity statistics; UK Parliament POSTnote 744 – Virtual wards and hospital at home; NIHR Evidence – Hospital at home and virtual wards: what works?
- **Digital twins applied to hospital operations:** claimed reductions in emergency-department wait times of 20% to 40% and throughput gains close to 20%, drawn primarily from industrial use cases, local simulations, or publications with a strong demonstrative component; caution on external generalizability. Journal of Medical Internet Research (JMIR), 2026.
- **Remote clinical monitoring:** several studies report sensitivities above 80% for certain alerting or clinical-degradation detection systems, but the literature remains heterogeneous, often North American, and centered on specific devices that are difficult to transfer at scale; these figures should not be presented as a consolidated consensus. JMIR – Remote patient monitoring and early warning systems.
- **Cost-effectiveness of hospital-at-home models:** systematic reviews show a general tendency toward cost reduction and higher patient satisfaction, but with strong methodological heterogeneity, variable economic perimeters, and moderate evidence quality; some studies report savings that can exceed several thousand euros per patient depending on organizational context. PMC – Systematic review on hospital-at-home cost effectiveness.
- **Readmission and average length-of-stay reductions:** results drawn primarily from North American programs operated within the CMS Acute Hospital Care at Home waiver; the reported effects should be treated as program or operator outcomes, not as universally demonstrated effects. CMS – Acute Hospital Care at Home initiative; POSTnote 744 – Virtual wards and hospital at home.

- **Triple Aim, Quadruple Aim and Quintuple Aim frameworks:** Institute for Healthcare Improvement – Triple Aim; Bodenheimer T., Sinsky C. (2014), From Triple to Quadruple Aim; Nundy S., Cooper L.A., Mate K.S. (2022), The Quintuple Aim for Health Care Improvement.