The Cost of the Contraction

Political economy of digital twins in healthcare: what a paradigm costs, what it costs not to have others

Introduction

This note extends a prior article which established that the term *digital twin* covers (at minimum) four families of objects (organ, trajectory, territorial, pathway…) governed by four independent dimensions: referent, coupling, temporal scale, validation regime. It also established that the semantic contraction of the concept to the sole high-fidelity organ twin produces a temporal framing error whose consequences are structurally allocative.

The first article closed the terminological-architectural question. The present one opens the medico-economic question: what does a paradigm cost, what does it cost not to have others, and who pays the consequences of this asymmetry?

The thesis of this note is threefold.

First, current French digital twin programs in healthcare, MEDITWIN foremost, are investments in real options on a future medico-economic benefit whose magnitude remains to be demonstrated on hard outcomes. This qualification is not pejorative for a pre-competitive R&D program. It is, however, pejorative for a public communication that sometimes presents these investments as directly generating savings for the health system.

Second, several families of digital twins (trajectory, pathway, compact organ twin) operate on documentable addressable pools at modeling unit costs one to two orders of magnitude lower. None benefits, to date, from a national envelope comparable to that dedicated to the high-fidelity organ twin.

Third, the opportunity cost of this preferential allocation does not lie in what is funded. It lies in what is not funded, and in the temporal window during which these unfunded families remain industrially behind.

This note does not seek to attack MEDITWIN. It seeks to name an allocation logic that has not been publicly examined, and whose consequences can be quantified in order of magnitude from public sources.

1. Why medico-economic analysis is inseparable from architectural analysis

A taxonomy that does not extend into allocative analysis remains an intellectual exercise. An economic analysis that does not rest on a rigorous taxonomy confuses objects and produces false comparisons. The two registers are jointly necessary.

The first article established that the four families of digital twins in healthcare (organ, trajectory, territorial, pathway) are not substitutable, that they fall under distinct validation regimes, and that they call for funding architectures adapted to objects whose cost profiles and return horizons differ structurally. The present note operationalizes this finding by applying to it the standard medico-economic grid: direct costs, direct benefits, indirect benefits, opportunity costs.

The exercise carries a major methodological difficulty that this note assumes head-on: the budgetary envelope allocated to MEDITWIN is not public. Official communications from the consortium, from industrial partners, and from the seven University Hospital Institutes (Instituts Hospitalo-Universitaires, IHU) involved mention France 2030 support without specifying an amount. This opacity is not specific to MEDITWIN; it characterizes public communication on France 2030 health programs. It nonetheless has a direct consequence: a rigorous cost-benefit analysis of a public program of this magnitude would presuppose that its budgetary envelope be public. As of the date of this analysis, it is not. The article therefore proceeds in orders of magnitude, mobilizing typical ranges for comparable structuring programs.

A public digital health policy must ask itself simultaneously what it funds, what that costs, and what it does not fund as a consequence. The third question is the hardest. It is also the most important.

2. Methodology: assumed asymmetry of evidence

The medico-economic analysis of emerging objects such as digital twins is constrained by a structural asymmetry: costs are partially observable, benefits are mostly projected, and inter-family comparisons rest on heterogeneous maturity levels. For MEDITWIN specifically, public figures essentially cover direct costs, and only partially. Benefits remain largely projective and, to date, not demonstrated in French real-world conditions. This feature is not a weakness specific to MEDITWIN; it is structural for digital twin programs. But it forbids exact cost-benefit analysis and allows only order-of-magnitude reasoning.

Three operating rules are therefore set for this note.

1. First, orders of magnitude are mobilized solely to compare classes of objects, with an uncertainty margin of a factor of two to three, and not with decimal precision. Conclusions drawn must remain robust to this imprecision.
2. Second, benefits not demonstrated in real-world conditions are treated as options, not as acquired savings. Whenever a figure is not publicly available, the article states so explicitly and fabricates none by unsourced extrapolation.
3. Third, inter-family comparisons are interpreted as orientation indications and not as definitive arbitrations. Where the literature presents a range or contrasting results, the note retains the conservative bound for projected benefits and the median bound for costs. This asymmetry is deliberate: it shields the analysis against the charge of having maximized the anticipated benefits of counterfactual reasoning.

This methodological posture is defensible because it is explicit. It would become indefensible if it were implicit. Its caution reduces the normative reach of the argument. It strengthens its robustness.

3. The costs of a high-fidelity organ twin

The cost structure of a high-fidelity organ twin decomposes into three distinct strata, whose orders of magnitude are cumulative and whose industrial determinants differ.

3.1 The programmatic cost

MEDITWIN was announced in December 2023 for a five-year duration (2024-2029), supported by France 2030, led by Dassault Systèmes as industrial prime, and associating seven IHUs, the Nantes University Hospital Center (Centre Hospitalo-Universitaire, CHU), the French National Institute for Research in Digital Science and Technology (Institut National de Recherche en Sciences et Technologies du Numérique, Inria) through eleven project-teams, and four startups (inHEART, Codoc, Qairnel, Neurometers). The stated objective is the production of seven virtual health products in cardiology, neurology and oncology.

The official communications consulted (Inria, Dassault Systèmes, IHU ICAN, IHU LIRYC, IHU Strasbourg, Institut Imagine, Institut du Cerveau, France 2030 press kit December 2023) mention no explicit amount allocated to the program. The systematic phrasing is: *the investment of the partners in this project will be financially supported by the State within the France 2030 framework*.

For France 2030 programs of this size (public-private consortium over 5 years with a large industrial group as prime, partner IHUs, associated startups), the typically observed ranges fall within an order of magnitude of 50 to 150 million euros, all components combined (public subsidy and industrial self-funding cumulated). This order of magnitude is consistent with the envelopes of Priority Research Programs and Equipment (Programmes et Équipements Prioritaires de Recherche, PEPR) of the same register, generally endowed with 30 to 60 million euros per program over 5 years for the public share. It is also consistent with recent IHU envelopes: Gustave Roussy obtained 30 million euros for its IHU PRISM in the third France 2030 call, and several comparable sector-specific digital health programs fall within the same range.

This order of magnitude remains an indirect estimate. It must be underlined: a public investment whose envelope is not publicly detailed cannot be publicly evaluated.

MEDITWIN's budgetary opacity is not a fault of the consortium; it reflects a current practice of France 2030 communication. But it mechanically limits the capacity of public debate to arbitrate, in full knowledge, the allocation choices between families of digital twins.

3.2 The marginal cost per patient

This is where the medico-economic question becomes operative. For a high-fidelity organ twin to be sustainable at scale, its cost per clinical use must be known. The scientific literature provides indications without yielding a ready-made figure in euros per patient in a French clinical environment. Bhagirath and colleagues, in the reference review published in December 2024 in *EP Europace*, document that the typical 400 µm spatial resolution used in several recent clinical applications (ventricular tachycardia (VT) stratification, VT ablation, atrial fibrillation (AF) ablation) mobilizes several CPU-days per patient on advanced high-performance computing (HPC) resources, and that doubling the resolution multiplies the cost by a factor close to ten. The HPC cost alone, expressed in euros, thus lies in a range of a few hundred to a few thousand euros per patient depending on resolution.

To this HPC cost are added engineer time (segmentation, meshing, calibration, quality control), interpreting-clinician time, and the cost of the underlying imaging (typically cardiac MRI with late gadolinium enhancement, LGE). In full-cost terms, the order of magnitude commonly cited in industrial presentations and in cost-effectiveness literature currently in publication falls within a range of a few thousand euros per patient for current clinical cases, to be compared with the hospital cost of an AF ablation in France (a stay of the order of 8,000 to 15,000 euros according to the Homogeneous Stay Group classification) and with the cost of a reimbursed LGE cardiac MRI (of the order of 200 to 400 euros).

The organ twin thus adds, to the existing interventional procedure, a marginal surcharge that may represent, case by case, 15 to 30 % of the procedural cost, for an expected benefit in terms of ablation precision, reduced operative duration, and potentially reduced recurrence rate. This surcharge is not in itself problematic. It becomes problematic only if the expected benefit on hard outcomes is not demonstrated at the level that justifies the surcharge. This is precisely the point that the next section addresses.

3.3 The cost of industrialization and scale-up

This is where the architectural critique of the first article articulates with the economic analysis. An organ twin in a research environment is not industrially equivalent to an organ twin in regulated clinical deployment. The passage from one to the other presupposes an architectural redesign (hexagonal architecture with isolated input and output ports, dedicated reliability ports, continuous instrumentation of quality control) and a dynamic governance framework adapted to an evolving system (Predetermined Change Control Plan, PCCP, formalized by the FDA in December 2024).

Without this redesign, the maintenance and update cost of an industrial twin becomes prohibitive, because every modification calls for full re-validation. With this redesign, the marginal cost of evolution is contained, because every modification is traced, documented and authorized within a predefined framework.

Many digital twin prototypes in healthcare documented in the scientific literature are not designed with this hexagonal isolation. They are research artifacts intended to demonstrate feasibility, not industrial systems intended to evolve in a regulated framework. The passage from one state to the other is not a simple improvement; it is a change of nature that may require substantial redesign. This transition carries a cost which, as of today, is not integrated into the budgetary announcements of public programs. It is a hidden cost of the paradigm.

4. Documented benefits and projected benefits

The cost-benefit analysis of the organ twin in 2026 mobilizes two registers of evidence that it is useful to separate rigorously.

4.1 Documented benefits on hard outcomes

Before April 2026, no prospective study documented the effect of a cardiac twin guiding ablation on hard clinical outcomes. The cohort of Hwang and colleagues published in 2024 in *Circulation: Arrhythmia and Electrophysiology* validated the technical accuracy

of the VT twin (sensitivity 81.3 %, specificity 83.8 %, negative predictive value 98.8 %) on 18 patients, but did not measure the effect on survival or recurrence.

The publication of the TWIN-VT study by Chrispin, Prakosa, Trayanova and colleagues in the *New England Journal of Medicine* of April 2, 2026 altered this landscape. On 10 post-infarction patients enrolled in an FDA investigational device exemption trial (NCT03536052), VT was non-inducible in 100 % of patients after twin-guided ablation; at a mean follow-up of 13 months, 8 patients out of 10 were free of recurrence without antiarrhythmic therapy. The historical comparator arm of standard-of-care ablation in ischemic scar VT sits at approximately 60 % success at one year.

These figures do not constitute a randomized multicenter demonstration. The cohort is small, the study lacks a randomized control group, one co-author (Barcelon) belongs to Johnson & Johnson MedTech. These limits are real. They do not erase the structuring conclusion: the clinical feasibility of the organ twin paradigm in ablation cardiology is now supported on hard outcomes, at an evidence level that exceeds mere technical validation.

For the French perimeter, as of the date of this note, no randomized MEDITWIN trial on hard outcomes has been published. The prospective multicenter cohorts carried by the cardiac strands of the consortium and by the associated University Hospital Research programs (Recherche Hospitalo-Universitaire, RHU, notably RHU-TALENT at Bordeaux University Hospital) are in the constitution phase; their results will not be available before 2027 or 2028. The scientific debate on feasibility is beginning to be settled; the allocative debate on the comparative relevance of the investment remains, for its part, entirely open.

4.2 Projected benefits at scale

International academic estimates suggest that an accurate electrophysiological twin could improve the success rate of persistent AF ablation by 5 to 15 points (from the current order of 60-65 % to a projected 70-80 %), reduce reinterventions, reduce procedural duration. These benefits, *if* they are confirmed on multicenter randomized phase III trials, would generate substantial savings. The key phrase is *if* they are confirmed.

What distinguishes projected benefits from documented benefits is not their *a priori* probability, which may be high. It is their epistemic status. So long as benefits are not measured on a large cohort in routine clinical environment, they remain a promise. The medico-economic consequence is direct: the French MEDITWIN investment is, at this stage, an investment in real options on a future medico-economic benefit whose magnitude remains to be demonstrated.

This qualification is not pejorative. It is accurate for any pre-competitive R&D program. A real option is, by definition, an asset whose value depends on future arbitrations. But it

forbids a simple accounting reading: so long as the option is not exercised by a demonstration on hard outcomes, the investment remains a cost; it becomes a benefit only retrospectively, and only if the favorable arbitration is realized.

The medico-economic study conducted by the University of Rotterdam within the European inEurHEART project (cohort of 112 patients across 16 centers spread between Germany, Austria, France and Switzerland, comparing standard ablation versus inHEART) will constitute the first European published reference on the cost-effectiveness ratio of an *industrial* cardiac twin in real-world environment. Its results are expected in the short term and will be decisive. They do not, however, cover the MEDITWIN perimeter as a whole.

5. Heart failure as a case study

To assess the comparative allocative relevance of the organ twin, it must be confronted with families whose savings pool is already documented. Heart failure (HF) constitutes the most telling case study in France, because epidemiological and cost data are public, robust and recent.

5.1 The order of magnitude of the burden

According to data published by Santé Publique France in the 2025 special issue of the *Bulletin Épidémiologique Hebdomadaire*, HF concerned approximately 1,376,692 adult prevalent cases in France in 2022, or 2.6 % of the adult population, reaching 23.7 % above age 85. The same year, 181,178 adults were hospitalized for acute HF, making this pathology the leading cause of hospitalization of the elderly in France. In 2021, 24,645 deaths had HF as initial cause, and 73,121 deaths mentioned it as initial or associated cause.

The average hospital cost of an acute HF stay is estimated by the French National Health Insurance Fund (Caisse Nationale de l'Assurance Maladie, Cnam) at 8,338 euros (2019 expenditure mapping). On the basis of 181,178 annual hospitalizations, the acute hospital line item thus represents, in order of magnitude, 1.5 billion euros per year for acute decompensation alone. The total annual cost of chronic HF in France, aggregating hospitalizations, ambulatory care, medications and long-term illness coverage, is estimated according to the broadest sources at approximately 13.6 billion euros, or of the order of 12.6 % of expenditures linked to long-term illnesses, and 1.5 to 2 % of total health expenditures.

5.2 The allocative computation

Even a modest reduction in avoidable hospitalizations for HF decompensation would have major economic consequences. A 10 % reduction would represent approximately 18,000 hospitalizations avoided per year, or of the order of 150 million euros per year on the acute hospital line alone. A 15 % reduction would represent approximately 225 million euros annually. Over five years, the cumulative order of magnitude would reach a billion euros.

These amounts do not constitute acquired net savings. They define the gross addressable ceiling of the acute hospital line. Real savings would depend on the cost of deploying the monitoring and decision-support infrastructure, on patient adherence rate, on the cost of clinical coordination between ambulatory care, hospital and medico-social sectors, on the false-positive rate and the cost of unnecessary alerts, on the organizational capacity to actually avoid stays rather than displace costs toward ambulatory or medico-social care, and on the share effectively capturable by the National Health Insurance or by the hospital establishment under the current tariff system. None of these factors is negligible. Several of them may transform a significant gross addressable ceiling into modest, even null, net savings under unfavorable configurations.

This qualification is essential. The argument of this note does not bear on acquired savings; it bears on the order of magnitude of an addressable pool that the current allocation does not put in a position to be captured. The distinction is not accounting, it is doctrinal: an unaddressed pool remains theoretical; an instrumented pool begins to become economic.

To be compared with the public investment projected in HF trajectory twin infrastructures at the national scale: as of the date of this note, no dedicated France 2030 envelope exists. HF telemonitoring exists in France through the ETAPES framework (Expérimentation de Télémédecine pour l'Amélioration des Parcours En Santé or Telemedicine Experimentation for the Improvement of Healthcare Pathways) instituted in 2018, with limited coverage and progressive ramp-up; it is not backed by a structured program of trajectory digital twins in the sense of the first article's taxonomy.

The order of magnitude of accessible performance is, moreover, documented. Pavon and colleagues, in a non-randomized prospective study covering 150 HF patients in post-discharge over one year, trained a long short-term memory (LSTM) network (a type of recurrent neural network suited to time series) exploiting telemonitored vital signs and patient profile to dynamically predict 30-day readmission risk. The model detects emerging readmissions with sensitivity above 71 %, specificity above 75 %, and an area under the ROC curve (AUROC) of the order of 80 %, performance maintained above 78 % AUROC even when reducing telemonitoring frequency to every other day.

These figures do not prove that a trajectory model universally outperforms structured clinical follow-up. They prove that, at a defensible performance level, on a real clinical cohort, an object of the trajectory family exists, is validated, and is deployable on lightweight infrastructure.

A conceptual precision is required here: a trajectory twin is not a model. It is a loop *prediction → decision → action → re-evaluation*. This distinction is not terminological. It directly conditions the reading of medico-economic results: the effect of an isolated predictive model is by construction limited; the effect of a complete clinical loop mobilizes four distinct links each of which can be the limiting factor. Recent meta-analyses on HF telemonitoring remain contrasted on net effect in routine practice, precisely because that effect depends as much on care organization as on the predictive model itself. But this mixedness does not disqualify the model; it contextualizes its function as an instrument inserted into an organizational chain.

5.3 The allocative consequence

The gross addressable ceiling related to modeling unit cost appears, in order of magnitude, more favorable for the HF trajectory twin than it is, in the state of public data, for the high-fidelity organ twin. This does not disqualify the latter. It qualifies the former as an under-invested object, on condition that the conversion from addressable ceiling to net savings can be documented by a dedicated organizational validation program.

The structural reason is threefold.

1. First, the eligible population is massive: 1.3 million prevalent patients and 181,000 annual hospitalizations, against a few thousand patients per year for the peri-interventional indications concerned by high-fidelity organ twins.
2. Second, the modeling unit cost is one to two orders of magnitude lower: a compact LSTM model runs in near-real-time on lightweight infrastructure, against several CPU-days per patient on HPC for a multi-physics solver.
3. Third, avoided costs are direct and quantifiable: an avoided hospitalization for decompensation represents 8,338 euros immediately, against a longer chain of effects for an optimized interventional gesture.

A national digital health strategy that preferentially funds the family with the lowest expected cost-benefit ratio, because that family is the most visible and the most narratable, does not maximize the impact of public investment.

This formulation is not a moral judgment. It is a direct consequence of placing public orders of magnitude side by side.

6. Informed de-indication of the defibrillator in primary prevention

A second case study illustrates the *specialized indication* regime of the first article, and demonstrates that the organ twin itself can have a very high cost-benefit ratio on the condition that it be deployed in the right configuration.

6.1 The clinical context

Current European Society of Cardiology (ESC) guidelines for the implantation of an implantable cardioverter-defibrillator (ICD) (in French, *défibrillateur automatique implantable* or DAI) in post-infarction primary prevention rest largely on a left ventricular ejection fraction threshold of 35 % or below. This threshold, validated by the MADIT-II and SCD-HeFT trials, presents limited specificity: a significant number of implanted patients never trigger the defibrillator over the lifetime of the device. These patients bore the procedural risk, the psychological cost, the economic cost, without demonstrated individual benefit.

6.2 The economic order of magnitude

According to the September 2022 French National Authority for Health (Haute Autorité de Santé, HAS) evaluation on the reassessment of implantable cardiac defibrillators, approximately 10,000 ICD implantations are performed each year in France, of which a significant share in primary prevention. The cost of an ICD over 10 years (device, procedure, follow-up, replacements) is estimated according to several European medico-economic studies in a range of 25,000 to 40,000 euros per patient.

The total annual line item thus represents, in order of magnitude, on the order of 200 to 400 million euros on the basis of 10,000 implants at 30,000 euros amortized over 10 years.

6.3 The predictive lever as a program hypothesis

The Hwang 2024 cohort documents a negative predictive value of 98.8 % for the *in silico* prediction of VT critical sites. This metric has a precise scope: it qualifies the capacity of an electrophysiological twin to identify ablation targets in the context of ischemic scar VT ablation. It does not, as such, bear on the long-term sudden death risk that structures the ICD primary indication. To bridge the two clinical objects in one step would be a shortcut.

The operative hypothesis to retain is therefore more modest and more rigorous: an electrophysiological twin could, on the condition of dedicated prospective validation, become a component of an *integrated de-indication score*, to be articulated with the fibrotic substrate assessed in LGE-MRI, the rhythmic risk assessed by classical electrophysiological methods, the competing risk of non-rhythmic death, the temporal evolution of the substrate, the medical therapy optimization, and comorbidities. This

multifactorial integration presupposes a longitudinal cohort of several hundred to several thousand patients, multi-year follow-up, and a demonstration of clinical non-inferiority under de-indication intent.

If such an informed de-indication strategy succeeded in validating, on a larger cohort, the identification of 10 to 15 % of current candidates as at sufficiently low risk to justify non-implantation, the order of magnitude of the addressable pool would be 25 to 100 million euros per year, in gross ceiling, on the primary ICD line alone. This calculation remains a hypothesis of a prospective program, not a near-demonstrated lever. It defines a target for medico-economic investigation, not an available acquisition. But the order of magnitude justifies that a prospective program be considered.

6.4 Why this case is instructive

This projection illustrates that the organ twin can be a serious candidate for a high expected cost-benefit ratio under two explicit architectural conditions, provided that prospective validation is conducted with appropriate rigor.

1. The first condition is *compact form*. An electrophysiological twin intended for de-indication must be deployable beyond expert centers equipped with dedicated HPC and a modeling-engineer team. This presupposes a *reduced-order model*, a model that preserves the essence of the simulated physiological behavior while drastically reducing computation cost compared to the reference solver.

   Two families of approaches are documented today in the electrophysiology literature.

   - The *hybrid eikonal model* approximates the propagation of the cardiac activation front by an eikonal equation (which computes arrival times rather than detailed electrical potentials), while preserving an explicit representation of substrate heterogeneities; this approximation can move a computation from several CPU-days to a few minutes per patient.
   - The *differentiable surrogate* is a statistical learning model trained to reproduce the outputs of a reference multi-physics solver, whose differentiability moreover enables fast optimization operations; it typically operates in real time on standard clinical infrastructure.

One or the other approach rests on fidelity-cost trade-offs that must be clinically validated for each indication.

2. The second condition is the *specialized indication regime* rather than pure peri-interventional. Here the twin does not serve to optimize a gesture, but to prevent it from taking place. This inversion of finality changes the nature of the economic computation: the saving is no longer marginal (a procedure better performed), it is radical (a procedure not performed).

The organ twin, in compact form and indication regime, is therefore a serious candidate for sound public allocation, subject to dedicated prospective validation. In high-fidelity form and pure peri-interventional regime, it remains an investment in real options.

The first article's taxonomy makes this distinction explicit; the economic analysis makes it operative as an investigation target, not as an available acquisition.

7. The CAPEX allocation logic

The bias is not first scientific. It is institutional. Public instruments more readily fund a capital-intensive, visible, productizable object carried by an identifiable consortium than a populational infrastructure that is distributed, organizational, and multi-actor.

Allocation thus follows not only the clinical promise; it follows the administrative form of the fundable project. This section makes this mechanism explicit.

It remains to understand why the concentration of public investment on the high-fidelity organ twin structurally occurs, even though the other families present higher expected cost-benefit ratios in order of magnitude.

The answer is not in any strategic intent of the actors. It is in the mechanical logic of existing funding vehicles.

7.1 The fundability of the organ twin

The high-fidelity organ twin presents a cost profile compatible with the industrial and institutional funding architectures that structure French public innovation investment. It naturally fits the France 2030 large-program format: identified large-group industrial prime, unified HPC platform, public-private alliance over 5 years, deliverables in the form of commercializable products, deployment on sovereign cloud. Each of these elements is fundable in existing funding architectures, because it corresponds to the vehicles designed for them: Bpifrance i-Démo calls, IHU endowments, Inria agreements, CIFRE contracts.

7.2 The structural non-fundability of populational families

The trajectory twin and the territorial twin present a cost profile dispersed across public actors none of which alone carries the modeling. Data governance simultaneously mobilizes National Health Insurance, the Regional Health Agencies (Agences Régionales de Santé, ARS), ambulatory medicine, the medico-social sector, learned societies and territorial collectivities. Structuring investment struggles to find its natural sponsor: it is neither a medical device (hence not the domain of the Directorate-General for Health in the first instance), nor an industrial product (hence not the domain of Bpifrance large-program), nor a research infrastructure (hence not the domain of PEPR).

These families are fundable in theory, lightly funded in practice, because they fall in the interstices between funding vehicles. The pathway twin moreover mobilizes actors (National Health Insurance, Hospitalization Information Agency, National Cancer Institute, learned societies) whose culture is statistical-epidemiological rather than industrial-modeling. Prospective modeling is read there as a retrospective evaluation tool rather than as a closed-loop decision-support infrastructure. This is not a funding problem, it is a conceptual frame problem: data corpora exist (Hospital Discharge Database, Programme de Médicalisation des Systèmes d'Information, PMSI ; FRANCIM registries; REIN database; chained National Health Data System, Système National des Données de Santé, SNDS), but they are not today mobilized as digital twins in the strict sense.

7.3 An unexamined convergence

The consequence is clear. The concentration of public investment on the organ twin is not the result of a deliberate medico-economic choice that would have concluded to its allocative superiority. It is the result of a convergence of funding structures, industrial narratives, and existing platforms. This convergence is neither illegitimate nor strategic in the pejorative sense. It is simply unexamined.

An allocation logic can be rational for each of the actors composing it and collectively suboptimal for the system that results from it. This is precisely what makes public coordination necessary. Without coordination, allocation amounts to stacking individually defensible decisions that collectively produce a structural bias. With coordination, allocation becomes a choice whose opportunity costs are named and explicitly arbitrated.

As of today, this coordination does not exist. There is, in 2026, no published and arbitrated French national strategy for digital twins in healthcare that would distinguish the four families, compare their expected cost-benefit ratios, and allocate public investment accordingly. There are individual programs (MEDITWIN, RHU-TALENT, ETAPES, PEPR Santé Numérique) each of which is defensible, and whose aggregation is not itself publicly examined.

8. Conclusion: from semantic cost to real cost

The first article established that the semantic contraction of the digital twin concept to the sole high-fidelity organ twin produces a temporal framing error with structurally allocative consequences.

The present article has shown how these consequences can be quantified in order of magnitude from public sources.

Three operative findings emerge.

1. First, in the absence of a detailed public budget, the exact cost of MEDITWIN cannot be publicly evaluated. The reasoning of this note therefore does not rest on an amount attributed to the program. It rests on a comparison of economic profiles between families: high capital cost and deferred benefit for the high-fidelity organ twin; lower modeling unit cost and potentially more immediate populational impact for trajectory and pathway twins, under condition of organizational integration. Clinical validation on hard outcomes of the former is undergoing international initiation (TWIN-VT, NEJM 2026); French results from MEDITWIN and associated RHU programs are not expected before 2027-2028. The doctrinal question is not within that perimeter. It is within the absent perimeter.
2. Second, several families of digital twins operate on addressable pools whose order of magnitude, in gross ceiling, reaches several hundred million euros per year on the sole examples examined here: 150 to 225 million euros annually in gross addressable ceiling on the acute hospital HF line for a 10-15 % reduction in hospitalizations, 25 to 100 million euros annually in gross addressable ceiling on the primary ICD line under a hypothesis of validated prospective program. These amounts are not acquired savings. They define the addressable target, whose conversion to net savings depends on the organizational, tariff, and operational factors enumerated in §5.2. But they establish the order of magnitude of the pool, which justifies that an explicit public arbitration be conducted.
3. Third, the absence of explicit public arbitration between families does not mean that no arbitration takes place. It means that the arbitration is delegated to the existing funding vehicles, which mechanically privilege the high-fidelity organ twin because it aligns with their architectures. The argument of this note does not demonstrate a misallocation; it demonstrates the absence of the public framework allowing one to know whether the current allocation is sound.

This note does not plead against the high-fidelity organ twin. It pleads for the national digital health strategy to explicitly name the four families, their validation regimes, their expected cost-benefit ratios, and their reciprocal opportunity costs.

This explicitation is technically feasible: it presupposes a comparative medico-economic mapping exercise that the HAS, National Health Insurance, Inserm or France 2030 itself could pilot.

A national digital health strategy needs several paradigms to be effective. A strategy that preferentially funds a single one may do so for sound industrial reasons. But it must then name what it does not fund, and accept the associated opportunity cost.

Not to name this cost is to count by halves. To count by halves is to arbitrate without knowing it. To arbitrate without knowing it is to fund one objective while believing one funds another.

Principal references

Baillargeon, B., Rebelo, N., Fox, D. D., Taylor, R. L., & Kuhl, E. (2014). The Living Heart Project: A robust and integrative simulator for human heart function. *European Journal of Mechanics A/Solids*, 48, 38–47.

Bhagirath, P., Strocchi, M., Bishop, M. J., Boyle, P. M., & Plank, G. (2024). From bits to bedside: entering the age of digital twins in cardiac electrophysiology. *EP Europace*, 26(12), euae295.

Chrispin, J., Prakosa, A., Kholmovski, E., Kolandaivelu, A., Barcelon, A., Aronis, K. N., Berger, R. D., Calkins, H., & Trayanova, N. (2026). Digital Twin-Guided Ablation for Ventricular Tachycardia. *New England Journal of Medicine*, 394(13), 1345–1347. doi: 10.1056/NEJMc2517822. NCT03536052.

Cnam (2023). *Cartographie des dépenses de santé 2019*. Caisse Nationale de l'Assurance Maladie, annual report.

FDA (2024). *Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence-Enabled Device Software Functions*. U.S. Food and Drug Administration, Final Guidance, December 2024.

HAS (2022). *Réévaluation des défibrillateurs cardiaques automatiques implantables avec sonde(s) endocavitaire(s)*. Haute Autorité de Santé, evaluation report, September 2022.

Hwang, M., Han, S., Park, M.-S., Yu, H. T., Kim, T.-H., et al. (2024). Heart Digital Twins Predict Features of Invasive Reentrant Circuits and Ablation Lesions in Scar-Dependent Ventricular Tachycardia. *Circulation: Arrhythmia and Electrophysiology*.

Pavon, A. G., et al. (2021). A Vital Signs Telemonitoring Programme Improves the Dynamic Prediction of Readmission Risk in Patients with Heart Failure. *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 2021, 1968–1972.

Santé Publique France (2025). *Épidémiologie de l'insuffisance cardiaque en France*. *Bulletin Épidémiologique Hebdomadaire*, special issue.

See also in this corpus

Vetillard, J. (2026). *The Digital Twin Is Not an Organ Model. Genealogy of a contraction, epistemic asymmetry, and the distinction between peri-interventional regime, indication regime, and predictive regime*. Twingital Institute.

Vetillard, J. (2026). *AI in Healthcare Is Not Solely an Algorithm Problem. It Is First and Foremost an Architecture Problem*. Twingital Institute.

Vetillard, J. (2025). *Do Buzzwords Dream of Clearer Substance? The Journey of Digital Twins*. Twingital Institute / LinkedIn.