From medical drift to predictive medicine

A note for clinicians, care teams, and health managers — on what we see in consultation, what we cannot see, and what could change

> **THE CENTRAL POINT OF THIS NOTE**
>
> *Medical drift is not a failure of our attention. It is the expected effect of a system that has us observing in episodes what requires continuous observation.*

Jérôme Vetillard

VP R&D & Engineering — CPO
Qualees · Twingital Institute

Note for clinicians · v3
May 2026 · 15 pages
Document derived from the Twingital Institute monograph
From medical drift to predictive medicine

To our colleagues

Preface

This note addresses clinicians. It has been written with two firm convictions that we wish to state at the outset, since they condition the reading of all that follows.

The first conviction. The phenomena of medical wandering and drift that we observe in practice are not the result of a failure of our attention. We know better than anyone how demanding, sustained, and largely uncompensated the human continuity we exercise actually is. This note formulates no accusation against care actors. It proposes, on the contrary, a reading that distinguishes what depends on us (and what we often do under conditions of overload) from what depends on the observation device into which we are inscribed.

The second conviction. The prospect of predictive vigilance is generally presented as a technological infrastructure, and even more often as a promise of algorithmic accuracy. Both presentations are misleading. The subject is neither a platform, nor a product, nor a promise. It is the nature of the observation regime to which our chronic patients are subjected, and the possibility of transforming it without delegating clinical judgment to a technical device. The hypothesis defended is deliberately limited: certain degradations become statistically detectable earlier when a partial longitudinal trajectory is rendered interpretable in time. Neither a promise of exhaustiveness nor a promise of universal effectiveness. A reasonable bound, sufficient to justify the examination without requiring the impossible demonstration of absolute predictivity.

This note has been written with the constant concern that you, clinicians, be the first consulted on everything this change of regime would imply for your daily practice. Only on this condition is the transformation described here defensible.

1. What we see in consultation

Three clinical situations, among those we encounter every week, are sufficient to pose the problem this note addresses. They are anonymized and schematized, but their features are familiar.

VIGNETTE 1 · DECOMPENSATED TYPE-2 DIABETES

Mrs B., 67, HbA1c at 11.2% at the next consultation

Patient known for eight years for type-2 diabetes on oral bitherapy, stable equilibrium until now. Seen on March 14 in scheduled consultation: HbA1c 7.1%, unremarkable clinical examination. Reviewed on September 12: HbA1c 11.2%, asthenia, nocturnal polyuria, 4 kg weight loss.

Six months without structured clinical contact. In the meantime: change of pharmacy, two weeks without treatment due to renewal default, social isolation following her husband's death in June, and progressive cessation of glycemic self-monitoring. None of these elements was reported. No healthcare professional had, at that moment, any means of seeing it.

This is not a missed consultation; it is an interval that was expected. The physiological trajectory evolved in the invisible window separating two scheduled clinical contacts.

VIGNETTE 2 · SEVERE COPD AND EMERGENCY VISITS

Mr D., 74, four acute exacerbations in eight months

COPD patient at GOLD stage 3, followed by his registered GP and a pulmonologist, well-managed background treatment. Four emergency department visits for exacerbation between January and August, three of which led to hospitalization. None of these episodes was the subject of structured feedback to the registered GP. The pulmonologist learned of the fourth episode only at the September scheduled consultation, from the patient's account.

Three professionals follow this patient with competence. None has the complete trajectory. The fragmentation is not a failing of each actor; it is a failing of the aggregation device.

VIGNETTE 3 · DRIFT THROUGH ISOLATION AND INVISIBILITY

Mrs K., 81, found at home undernourished and confused after eighteen months without contact

Patient known for many years to the general practice, complete file, former teacher, no documented social isolation. Last consultation in November 2023, prescription valid for three months. At the time, satisfactory general condition. No return for renewal. No automatic alert at the practice, no follow-up from health insurance, no signal from other healthcare professionals she might have consulted (none, as it turned out). Discovered in May 2025 by a neighbor alerted by floor neighbors who no longer saw her. Hospitalized for severe malnutrition and undiagnosed early dementia syndrome.

This is not a failure of commitment or competence; it is the effect of a system that does not, structurally, have the capacity to signal an absence. The system sees presences; absences are by construction invisible.

VIGNETTE 4 · WANDERING WITHOUT A REGISTERED GP

Mr F., 58, six uncoordinated medical contacts in four months

Patient without a registered GP since the retirement of his former GP, lives in a zone classified as a medical desert. Has presented since February with intermittent exertional chest pain. Successively consults: SOS Médecins (twice), a substitute GP at an out-of-hours practice, the emergency department (twice), and a private cardiologist via teleconsultation. Each actor acts with competence within the perimeter of its contact; none has the history of the five others. Complementary examinations partially repeated, treatments adjusted without coordination, no structured follow-up established. Diagnosis of unstable angina made late, during a hospitalization for acute coronary syndrome in June.

This patient saw six professionals in four months. None saw his trajectory. Over-utilization is here as pathological as under-utilization: what is missing, in both cases, is longitudinal aggregation.

These four situations cover what we call in this note *wandering* (vignettes 2 and 4: disorganized over-utilization, without longitudinal piloting) and *drift* (vignettes 1 and 3: structural under-utilization, progressive invisibility). The two phenomena are not independent malfunctions; they are two pathological regimes of the same observation device.

2. What we cannot see, and why

Careful examination of the three vignettes reveals a common property: in each, the degradation occurred within a temporal window in which no professional could observe. This window is not a system accident; it is its structural property.

Clinical observation is, by construction, intermittent

A consultation is, by definition, punctual. A type-2 diabetic patient followed by their registered GP is the subject, on average, of four to six annual consultations. That is, less than a single day of direct observation across the 365 days of the year. The bulk of the physiological trajectory unfolds outside our structured field of

observation. This is neither a critique nor a revelation: it is the very nature of our profession to observe by sampling.

Increasing the number of contact points improves certain capacities but does not, by itself, modify the fundamental property of the dominant regime: a primarily discrete observation of trajectories that have become continuous. More consultations would produce real marginal benefits, and they should be defended as such when demographic conditions permit. But these benefits do not correct the structural property of the regime, which remains to observe by sampling trajectories whose degradation occurs predominantly between samples. The question is therefore not solely quantitative; it has become structural.

HUMAN CONTINUITY

What we exercise: following our patients over decades, building fine-grained clinical knowledge over time, mobilizing relational attention under conditions of overload. This continuity exists and is not in question.

OBSERVATIONAL CONTINUITY

What the device does or does not ensure: the physiological, behavioral, and therapeutic trajectory between two clinical contacts. This continuity is today largely absent.

The distinction is apparently subtle but it is central. Human continuity is not observational continuity. The first stems from our professional commitment; the second stems from the properties of the observation device within which that commitment is exercised.

TOOL FOR CLINICAL REFLECTION

Clinical vigilance density

To make the discussion among us more precise, it is useful to name a conceptual operator that this note proposes as a heuristic.

> *Clinical vigilance density designates the effective proportion of a pathological trajectory that is the object of an interpretable observation by the care system.*

For Mrs B. (vignette 1), between the March consultation and the September one, the clinical vigilance density on her type-2 diabetes trajectory was close to zero: six months without structured contact, while behavioral determinants (bereavement, isolation, treatment-adherence rupture) evolved daily. For Mrs K. (vignette 3), vigilance density was zero for eighteen months. For Mr F. (vignette 4), density was paradoxically high but fragmented: six contacts in four months, but no actor held the integrated trajectory.

This operator is not, at this stage, a calibrated metric we could compute precisely. It is a reflective tool that allows comparison across regimes (practice in a stretched zone vs practice in a dense zone), pathologies (cardiology follow-up vs depression follow-up), and temporalities of a single trajectory (acute phase vs stabilization phase). Its translation into a quantified indicator belongs to subsequent methodological work; its primary function is to structure the clinical discussion on sustainable observation regimes.

Structural cognitive debt cannot be absorbed by more attention

A second structural property limits what we can see: the multifactorial complexity of chronic trajectories exceeds the human capacity for longitudinal information processing within the consultation time available. An elderly polypathological patient presents, on average, seven to nine active pathologies, twelve to eighteen medications, and several dozen relevant clinical and biological indicators. No professional can, in fifteen to twenty minutes of consultation, fully revisit this complexity. Cognitive saturation is not an individual frailty; it is a physical limit on longitudinal processing, which grows with the number of chronic patients on medical lists.

> *Our attention is not extensible. The system asks us to see more than it is physically possible to see, in a time it has itself constrained.*

Organizational fragmentation has no owner

A third property completes the analysis: a patient's trajectory crosses several professionals and several care locations. Registered GP, specialists, hospital team, home-care nurses, physiotherapists, pharmacists. Each of these actors operates with competence within their perimeter. None holds the complete trajectory. Sharing through reports, discharge letters, or Apicrypt messages is useful but remains asynchronous and incomplete. No institution is today responsible for the longitudinal aggregation of information circulating among us. This absence of an owner of the integrated trajectory is one of the structural conditions of drift.

3. What is not in question in this diagnosis

Before evoking what a transformation of observation could bring, it is important to name explicitly what is not in question in the diagnosis posed. Without such explicitness, the analysis risks being read as a critique of practice, which it is not.

The commitment and competence of clinicians are not in question. The diagnosis at no point suggests that patients would be better followed if physicians were more attentive or nurses better trained. It suggests exactly the opposite: care actors are committed and competent under structurally saturated conditions, and it is precisely this saturation that asks to be addressed at the level of the device, not of professional behavior.

The caregiver-patient relationship is not in question. This note at no point proposes to replace the medical relationship by an algorithmic device. The relationship is and will remain the matrix of care. No signal, no alert, no digital trajectory could substitute for the clinical conversation, the examination, the listening, or medical judgment. These functions are not optional; they are the very purpose of any device that would purport to support them.

Clinical time is not in question. The diagnosis does not suggest that physicians should devote more time to each patient. It recognizes that clinical time is scarce, that medical scarcity is structural, and that any credible proposal must compose with this scarcity rather than deny it.

Individual medical responsibility is not in question. The drift of a patient we have not seen for eighteen months is not our personal fault; it is the effect of a system that has not, structurally, integrated the function of follow-up or absence-signaling. This property belongs to the device, not to the professional.

> *If everything depended on our attention, drift would be a moral problem. It is an architectural problem.*

4. The patient is more than the trajectory they carry

One dimension has so far been in the background of this note, and it must be named so as not to suffer from its absence. The chronic patient is not merely a trajectory to observe; they are a cognitive actor who interprets, decides, and arbitrates, often under conditions more demanding than ours. An observation infrastructure that does not recognize this dimension misses the essential of what it claims to support.

The cognitive architecture of the chronic patient

Living with a chronic disease is not merely undergoing a physiological trajectory. It is also, daily: interpreting bodily signals (is this pain unusual? is this fatigue an effect of treatment?), arbitrating between medical prescriptions and the constraints of ordinary life (meals, work, transport, childcare), absorbing continuous decisional fatigue (each day calls for ten to twenty minor therapeutic arbitrations), integrating sometimes contradictory recommendations across several prescribers, and sustaining a cognition distributed between oneself, one's relatives, and sometimes digital tools.

This cognitive load is poorly recognized. It is, however, the exact counterpart, on the patient side, of the structural cognitive debt we described on the clinician side. The chronic patient too is saturated. Their agency does not reduce to adherence or non-adherence; it is a continuous interpretive labor exercised under social, family, and economic conditions rarely taken into account by the care device.

Three implications for the design of the device

The device does not address itself only to clinicians. When a signal is produced, the question of its communication to the patient is not a secondary protocol question. It is a structuring question, which engages the patient's interpretive autonomy and the nature of the therapeutic relationship. A device that produces a longitudinal representation of the patient without the patient being able to access, comment on, or contest it silently transforms the patient into an object of surveillance.

Adherence is not a behavior; it is an arbitration. When the system observes a rupture of adherence, it does not observe a failure of discipline; it observes the result of an arbitration the patient has made, under constraints that neither they nor we always control. A device that interprets this gap as a moral or behavioral deficit mistakes its object. Adherence rupture is an informational signal about the patient's living conditions; it is not, in itself, a signal about their will.

Distributed cognition must be respected, not replaced. Many chronic patients have built, over years, an extremely fine patient knowledge of their trajectory. This knowledge does not encode itself in the variables observable by the device; it is exercised in the clinical conversation. A device that purported to replace it, or that marginalized the patient's word in favor of algorithmic signals, would degrade one of the finest interpretive resources available to us in consultation.

> *The chronic patient is more than the trajectory they carry. They are the principal interpreter of that trajectory. The device supports this interpretive function; it does not replace it.*

5. What a structured observational continuity would produce, for the clinician

The distinction between intermittent observation and structured observational continuity is easier to grasp through what it would change for practice than through its theoretical formulation. Before describing these transformations, two doctrinal precisions are required: on the nature of the hypothesis defended, and on the device's frontier of validity.

A limited hypothesis, not a promise of accuracy

The regime described here does not rest on the hypothesis of perfect predictive capacity, nor on an automation of clinical decision. It rests on a more modest hypothesis: *certain degradations become statistically detectable earlier when a partial longitudinal trajectory is rendered interpretable in time*. This hypothesis does not cover all pathologies, nor all phases of a trajectory, nor all patients with the same effectiveness. It is a defensible bound, sufficient to justify the examination without requiring the impossible demonstration of absolute predictivity.

A practical consequence follows: a sustainable predictive system is not a system that eliminates error; it is a system that renders its conditions of validity, its zones of uncertainty, and its operational limits explicitly governable. This posture is familiar to us. It is the one we adopt with classical diagnostic tools: the sensitivity of a test is never 100%, nor is its specificity, and predictive value always depends on prevalence. A sustainable probabilistic device merely applies this same rigor to the algorithmic field.

The device must know that it does not know

A second doctrinal precision concerns the device's frontier of validity, what the Anglo-Saxon literature calls the *applicability domain*. Every predictive model has a validity domain bounded by its training population and its deployment conditions. A prediction on a patient situated outside this domain is not equivalent to a prediction on a typical patient; it is an extrapolation, and it must be treated as such.

A credible device must refuse to produce a signal when the patient lies outside the recognized applicability domain, rather than silently produce a degraded-quality signal. For clinical practice, this point is essential: we must be able to know when the device has nothing to say, and not receive a misleadingly confident signal for a situation the system is not competent to evaluate. This translates into a simple requirement in consultation: alongside each signal, the explicit indication of the validity domain in which that signal was produced, and the clear mention of cases in which the device abstains from pronouncement.

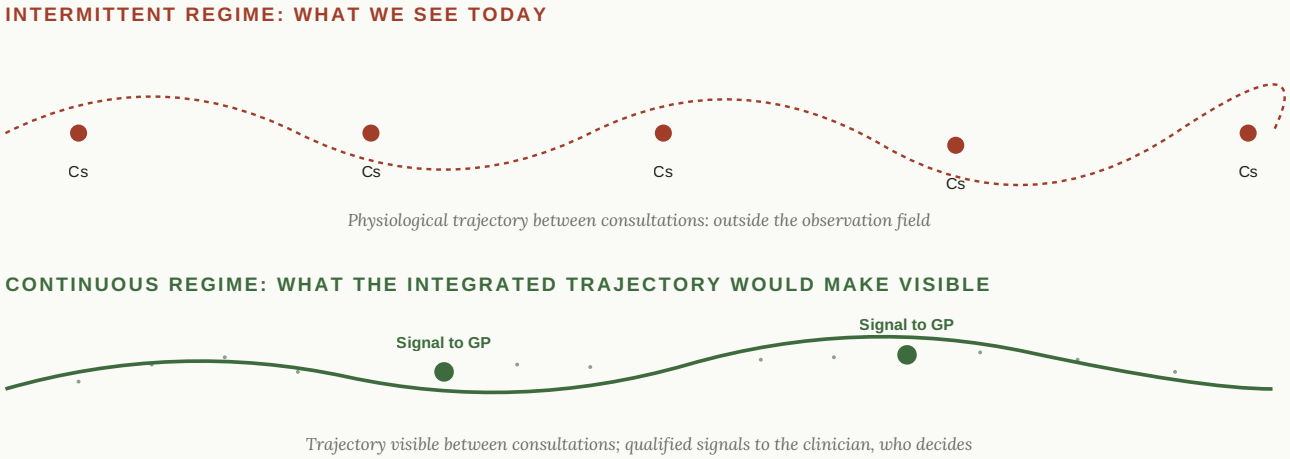

Figure 1. *The transformation envisaged is not the addition of punctual observations. It is the production of a longitudinal representation of the trajectory in the window separating two clinical contacts. The clinician remains the actor of the decision; the device renders visible what was silently drifting.*

Four concrete transformations for practice

01 **See the drift before the event, without having to look for it.**

The objective is not that we conduct more surveillance. It is that a qualified signal reach us when a trajectory drifts between two scheduled consultations: rapid weight variation, glycemic degradation, increase in emergency visits, halt in treatment renewal, prolonged loss of contact. The signal is not a medical decision; it is qualified information that allows scheduling this patient before decompensation.

02 **See the integrated trajectory, not just the episodes one has personally observed.**

The registered GP would have, in consultation, the complete trajectory of the patient's contacts with the health system: emergency visits, hospitalizations, specialist consultations, biological exams, treatment modifications. Not to comment on them individually, but to understand where the trajectory stands when it arrives before us.

03 **See collective prioritization, not just the individual queue.**

When resources are constrained, knowing which of our upcoming patients are drifting and which are stable allows allocation of available clinical time with increased efficiency. The device does not decide whom to see; it proposes a hierarchization that informs our decision without substituting for it. The final decision remains clinical, therefore ours.

04 **See absence as a signal, not as non-observable data.**

A patient who has not renewed her treatment for three months, who has not honored two appointments, who has not consulted for eighteen months: each one a signal the current system does not handle, because it observes presences. Longitudinal observation renders these absences observable and allows proactive follow-up. Mrs K. of vignette 3 could have been visited well before malnutrition.

6. What will not change, and what must be guaranteed

> **Medical decision remains ours.**
>
> No signal, no alert, no algorithmic prioritization substitutes for clinical decision. The device produces qualified information; the physician decides what to make of it, according to their judgment, their knowledge of the patient, and the context.

The caregiver-patient relationship is not instrumented.

The relationship remains a human space. The device is not in the consulting room; it is in the background, as a visibility infrastructure. Clinical conversation, listening, examination, accompaniment: none of this is delegated to an algorithm.

Information is strictly bounded to clinical use.

The longitudinal trajectories produced by the device are subjected, by construction, to three technical bounds inscribed in the design: informational minimization (only variables necessary for clinical observation are retained), limited retention temporality (data are erased after a contractually fixed window), technical impossibility of extra-clinical use (granularity degraded at export, impossibility of reconstituting individual profiles). These bounds are not promises; they are auditable architectural properties. This exclusion is not optional; it is the very condition of the device's sustainability.

Medical responsibility does not shift to the algorithm.

If the device misses a signal, it is not a failure the physician must compensate for personally. If the device emits a signal, the physician is not bound to follow it. Responsibility remains calibrated on human judgment, as today.

The device assumes its fallibility instead of dissimulating it.

A predictive system is never infallible, and it is more dangerous when it pretends to be than when it admits it. The governance of uncertainty trumps algorithmic performance: the device makes its zones of uncertainty, its conditions of validity, and its applicability domain explicit, rather than masking its limits behind apparent probabilistic confidence.

The clinician is not a validation operator.

This guarantee is probably the most important. The clinicians' legitimate fear before any algorithmic device is not only of being replaced; it is of being progressively transformed into validation operators of a system that would have, in fact, made the decision before them. This transformation occurs when the signal becomes an implicit instruction and the deviation from the signal becomes an explanatory burden. The device defended here is designed to produce the opposite: situated information that informs clinical judgment without obliging it, and that presupposes the final decision is exercised in the encounter with the patient. The signal is never the decision. The decision remains a clinical act, situated, contextualized, and irreducible to any prior algorithmic validation.

7. What we should be consulted on

If the transformation described in this note occurs, it must not occur as a technical device imposed top-down, which would be both politically untenable and clinically unworkable. It must occur as a co-constructed device with care actors. Six points call for explicit consultation of clinicians, and not for a technical decision imposed on them.

Alert thresholds

Every predictive device operates on the basis of explicit or implicit thresholds. A low threshold multiplies signals and the investigation load. A high threshold preserves the capacity for action but accepts a higher false-negative rate. This choice is not technical. It depends on our real capacity to absorb signaling and on the criticality we attribute to different pathologies. Clinicians must be consulted on initial calibration and on periodic revisions.

Signal granularity

Receiving an alert on every capillary glycemia deviation is unmanageable. Receiving a weekly synthesis of patients whose trajectory is drifting is useful. The right level of granularity depends on the profession, the time available, and the practice's organization. Clinicians must be consulted on the concrete modalities of signal arrival in their daily work.

Integration with existing tools

Any device that would add a new tool without integrating with practice software (LGC, hospital EHR) increases cognitive load without reducing work. Clinicians must be consulted on the technical conditions of integration, and any deployment that does not respect them must be rejected.

Training and accompaniment

The clinical use of a probabilistic device demands training in signal interpretation, false-positive management, and communication with the patient. This training cannot be treated as a simple manual; it presupposes peer-led medical accompaniment, calibrated on real conditions of practice.

The patient's place in the device

Whether the patient knows, or does not know, that a signal has been emitted on their trajectory is a decision that engages the clinical relationship. Should they have access to the signals concerning them? At what level of granularity? With what accompaniment? These questions cannot be settled without clinicians, who are the front-line actors of risk communication.

Indicators for evaluating the device

A vigilance device must be evaluated not on its algorithmic performance but on what it produces in practice: reduction of avoidable hospitalizations, clinical workload, perceived quality of the relationship, equity of access to care. The choice of these indicators is itself a clinical decision, which must be borne by clinicians and not by technical designers.

8. Questions you frequently ask us

When this analysis is presented to clinicians in continuing education or in collegial discussion, the same questions recur. We take them up here in their most direct formulation, and answer them without circumlocution, also indicating what remains to be arbitrated.

Will I have to manage a new flow of alerts on top of everything else?

No, and any device that would produce this effect must be rejected. The central principle of the device defended here is that it is filtering, not amplifying. It replaces the disorderly flow of partial information (letters, mail, reports, messages) by a qualified longitudinal representation. It hierarchizes what requires your attention rather than transmitting everything to you. The simple test: if the deployment increases your cognitive load, it is a bad deployment.

Who will be responsible if an alert is not seen, or is seen too late?

Medical responsibility does not shift to the algorithm and does not shift excessively to the physician either. If the device misses a signal, it is not a failure the physician must compensate for personally. If the device emits a signal, the physician is not bound to follow it. Responsibility remains calibrated on human judgment, as today. These principles must, however, be inscribed explicitly in the legal framework of deployment, failing which jurisprudence could redefine them in an unfavorable sense. This is one of the points on which professional orders and learned societies must be consulted upstream, not downstream.

How do we prevent this from becoming behavioral scoring of patients?

This question is central, and the answer is not in the goodwill of the designers: it is in the architectural separation between variables admitted for clinical decision support and those admitted for any other function. A credible device must, by technical construction, exclude that adherence indicators, appointment-presence indicators, or therapeutic-behavior indicators be reusable to modulate access to care, priority order, or remuneration of the act. This exclusion is not a promise: it must be verifiable by audit, inscribed in the design, and audited by a body external to the device's operators.

Will the device feed insurers, employers, or platforms?

This question depends strictly on the political framing of the deployment, and the honest answer is: by default, yes; by construction, no. Without explicit framing, longitudinal trajectories constitute an economic object of very high value that peripheral actors will seek to obtain through progressive legal channels. The

countermeasure does not lie in promises; it lies in architectural choices that render extraction technically impossible: bounded retention, granularity degraded at export, impossibility of reconstituting individual profiles outside strict clinical use. These choices degrade the device's potential economic value, and that is precisely what makes them politically sustainable.

What does this change for the remuneration of the act?

This note does not address the remuneration question, and we want to be direct on this point: it is not by avoidance, it is because the answer belongs to conventional negotiation between actors and professional organizations. The position defended here is only the following: if the transformation of observation absorbs an additional clinical workload, that load must be tariffed explicitly and not absorbed silently by the professionals. Any deployment that does not address this question upstream reproduces the historical asymmetry between system expectations and conditions of practice.

And medical secrecy?

Medical secrecy is non-negotiable, and the device is not intended to redefine it. The longitudinal trajectory is not public data; it is a medical object accessible only to the clinical actors of the patient's care. The exact perimeter (registered GP, treating team, professionals in direct connection with the trajectory in progress) must be defined with professional orders. The position defended is that this perimeter must be narrower than that of the current Shared Medical Record (DMP), not broader, because longitudinal granularity is more sensitive than that of a punctual report.

Is this not a gateway for industrial teleconsultation and commercial actors?

This objection is legitimate and the analysis conducted in the full monograph addresses it head-on. The answer comes in two stages. First, the device defended here is not teleconsultative: it does not replace any consultation, it observes the trajectory between consultations. Second, the partial inalienability of trajectories (point addressed above) must explicitly exclude their exploitation by commercial actors, whether they belong to teleconsultation, platforms, or business software publishers. This is again an explicit political choice, not a promise of goodwill.

What happens when the device drifts over time?

This is probably the most structuring question, and it is usually addressed too late in deployments. A predictive system is not deployed once and for all; it is maintained or it drifts. Clinical practices evolve, diagnostic codings change, the demographics of monitored patients shift. A model calibrated on 2026 has no guarantee of equivalent operational validity on 2032. The countermeasure is conceptually simple, operationally demanding: a credible device must include, by construction, an explicit dispositive for monitoring drift, continuous probabilistic calibration, and rollback mechanisms in case of degradation. This is the equivalent of an algorithmic pharmacovigilance — a function that should be familiar to us. No medication remains indefinitely under the same marketing authorization (AMM) without monitoring of its benefit-risk balance; no predictive device should be either. This requirement today exceeds the perimeter of French health authorities; it remains to be constituted.

And with the new generative-AI tools and agent chains?

A legitimate and increasingly pressing question. The position defended in this note is that the current evaluation regime, designed for classical supervised machine learning with frozen weights, does not transpose as such to adaptive systems (dynamic orchestration, persistent memory, emergent behaviors). What is validated at t_0 no longer says anything about what executes at t_0 + 6 months, and this holds *a fortiori* for generative and agentic systems. For the present note, two principles: first, the architecture defended here can rely on classical supervised ML, without dependence on generative or agentic components at the clinical core; second, when such components are incorporated at the periphery (for example for the production of textual syntheses), they must be governed by a specific regime, distinct from the one applicable to predictive components, and their frontier of validity must be explicitly signed off by clinical actors. These regimes are not interchangeable and do not resolve themselves through addition.

9. What this note is not

This is not a technological promise. The predictive vigilance described here is not a product that already works, nor a medico-economic demonstration. It is a description of the properties that any credible device

should satisfy. No real device today combines the entirety of these properties at the scale of a national health system.

This is not a critique of clinicians. The diagnosis posed at no point attributes patient drift to a professional failing. It analyzes the structural conditions in which this commitment is exercised, and proposes that these structural conditions be the object of transformation, not individual behaviors.

This is not a reform program. This note does not propose to modify the nomenclature, to reform T2A, to redefine the role of ARS (Regional Health Agencies), nor to inscribe additional obligations into practice. It posits the conditions under which a change of surveillance regime could be sustainable. The program belongs to public decision-makers; the conditions of clinical sustainability belong to care actors.

This is not a document to be applied as-is. Any effective transformation passes through discussion with professional organizations, orders, learned societies, territorial conferences, and patients. This note constitutes a framing proposal for these discussions, not their conclusion.

The subject is not artificial intelligence. The subject is the possibility, for each of us, of seeing before decompensation what today only becomes visible after; and of being able to do so with a device that assumes its fallibility instead of dissimulating it.

10. For further reading

This note synthesizes a Twingital Institute monograph of approximately forty-six thousand words and one hundred twenty-eight references. For clinicians wishing to deepen the analysis, a few resources are particularly relevant.

On chronicity and care models.

1. Bodenheimer T, Wagner EH, Grumbach K. Improving primary care for patients with chronic illness. JAMA. 2002;288(14): 1775-9. Founding reference of the Chronic Care Model.
2. WHO. *Innovative care for chronic conditions: building blocks for action.* Geneva: WHO; 2002.
3. Wagner EH. Chronic disease management: what will it take to improve care for chronic illness? *Eff Clin Pract.* 1998;1(1): 2-4. Seminal article on the inadequacy of the acute-care model.

On continuity of care and its effect on mortality.

4. Pahlavanyali S et al. Continuity of care and mortality for patients with chronic disease: an observational study using Norwegian registry data. *Fam Pract.* 2023;40(5-6):698-706.
5. Pereira Gray DJ, Sidaway-Lee K, White E, Thorne A, Evans PH. Continuity of care with doctors: a matter of life and death? *BMJ Open.* 2018;8(6):e021161.

On potentially avoidable hospitalizations in France.

6. Mercier G, Georgescu V, Bousquet J. Comparison of two methods to report potentially avoidable hospitalizations in France in 2012. *BMC Health Serv Res.* 2015;15:4.
7. Bourgueil Y, Mousquès J, Tajahmady A. The effect of primary care on potentially avoidable hospitalizations in France. *BMC Health Serv Res.* 2020;20:268.

On cognitive debt and clinical mental load.

8. West CP, Dyrbye LN, Shanafelt TD. Physician burnout: contributors, consequences and solutions. *J Intern Med.* 2018;283(6):516-29.
9. Sinsky CA, Brown RL, Stillman MJ, Linzer M. COVID-related stress and work intentions in a sample of US health care workers. *Mayo Clin Proc Innov Qual Outcomes.* 2021;5(6):1165-73.

On the epistemic limits of predictive systems in health.

10. Rudin C. Stop explaining black box machine learning models for high stakes decisions. *Nat Mach Intell.* 2019;1(5):206-15.
11. Finlayson SG et al. The clinician and dataset shift in artificial intelligence. *N Engl J Med.* 2021;385(3):283-6. On distribution drift.
12. Van Calster B, McLernon DJ, van Smeden M, Wynants L, Steyerberg EW. Calibration: the Achilles heel of predictive analytics. *BMC Med.* 2019;17(1):230.

Document derived from the Twingital Institute monograph From medical drift to predictive medicine: systemic analysis of the failures of the French healthcare system and architectural properties of a predictive prevention infrastructure as a perennial public good (Vetillard J., 2026). This clinical note presents a reading oriented toward care actors; it does not expose the detailed demonstrations grounding the assertions, nor the totality of the systemic, comparative, and political diagnosis presented in the full monograph.