PREPRINT, WORKING DRAFT (EN) · MANIFESTO ARTICLE OF THE PROGRAM \"WASTE AS DYNAMICAL SYSTEMS\ · 2026-06-15

The Four Implicit Axioms of Static Life Cycle Assessment and Their Structural Failure in Non-Linear Environmental Systems

Theo Piens, Ph.D.

Independent environmental researcher, France · ORCID 0000-0002-4422-2786

Table des matières

Abstract 3
1. Introduction: a change of object 4
1.1 Every waste is a moving cycle, not a fixed unit 4
2. The implicit axiomatics of static evaluation 8
3. State of the art and positioning 10
3.1 The pedigree and the log-normal closure 10
3.2 Multimedia fate and transport 11
3.3 Dynamic LCA and coupling to system dynamics 11
3.4 The regulatory toxicology of the threshold 12
Positioning 12
4. A taxonomy of environmental non-linearity 12
5. The irreducible uncertainty floor (keystone) 14
5.1 When more data does not remove the surprise 15
5.2 Three corollaries (T2–T4) and the horizon by medium 15
6. Conclusion: a change of object, and a program 17
Appendix. Sketch proof of T2 and T4 (Fokker-Planck on the double-well, with the saddle-node as washout boundary) 18
References 19
Supplement 20

The Four Implicit Axioms of Static Life Cycle Assessment and Their Structural Failure in Non-Linear Environmental Systems

Theo Piens, Ph.D. · Independent environmental researcher, France
ORCID 0000-0002-4422-2786

> *Preprint, working draft (English version). Computations and figures are reproducible and openly deposited (Zenodo DOI at the first stable version). Scientific companion to the ChaoticLCA preprint (Piens 2026, Zenodo 10.5281/zenodo.20086772). No competing interests declared.*
>
> *This article belongs to the research program "Waste as dynamical systems" (Piens, in preparation). To keep the text light, section numbers (§x.y) follow the book-length manuscript from which the articles are drawn; references to sections not reproduced here point to the book or to the companion articles, namely the digester, the toxicology of emerging contaminants, environmental logistics, the ecology of critical transitions, and public decision-making.*

Abstract

A life-cycle assessment is, in the end, a quantitative argument that one waste-treatment process burdens the environment less than another, and people turn it into a decision. Yet the object the calculation rests on is not quite the one it assumes. Standard practice treats a waste as a *static, classified unit* (a fixed inventory, a quality score, a regulatory threshold, a log-normal margin closed around the whole), whereas once released a waste is no longer a unit but a **dynamical sub-system** nested in soils, groundwater, the atmosphere and the living world, whose output is a *regime* rather than a noisy number. We argue this substitution of object is a *structural* failure, and we name it: static evaluation rests on four implicit axioms (point estimation, linear characterization, path independence, categorical invariance), each of which fails, by a named mechanism, for a specific class of waste or pollutant. We give a taxonomy of eight failure modes and prove four theorems: a finite prediction horizon that more data cannot push back (T1), bimodal statistics in a bistable window that void the log-normal closure (T2), the impossibility of a single characterization factor under non-monotonic dose-response (T3), and the predictive blindness of a fixed threshold near a bifurcation (T4). The nine falsifiable vignettes that quantify these failures are developed in the companion articles of this program. We propose not a universal predictive model but a **change of object**: assess a waste's dynamical regime, not only its quantity, then report honestly the uncertainty that follows.

---

1. Introduction: a change of object

In the standard account of an environmental assessment, a waste-treatment process is a *unit*. One attaches to it a vector of inputs and outputs, a quality score, sometimes a regulatory threshold not to be crossed. The arithmetic that follows is of the same kind: inventories are aggregated, multiplied by characterization factors, and the result is compared with a limit. Everything is linear, stationary, and whatever ignorance remains is filed tidily into a log-normal margin, a "plus or minus" hung on a central value. This convention is remarkably productive, and it must be said at once: for a closed, controlled industrial process fed by measured flows, it works, and it works well. The problem is not that it computes badly. It is that it often computes on the **wrong object**.

For a waste is a unit only as long as one holds it. Once released, stored, spread or leached, its constituents enter soils, groundwater, catchments, the atmosphere and the living world, and there they do not wait politely to be characterized. They interact, through processes that are non-linear, coupled, driven by feedback. The output of such processes is not a mean quantity with Gaussian noise around it. It is a **regime**: a qualitatively distinct operating state in which the system *holds*, and which it can leave all at once (bifurcate) under a change of conditions that is itself nothing dramatic.

1.1 Every waste is a moving cycle, not a fixed unit

No major family of waste truly reduces to a static unit. The practitioner who lives with these flows knows it from experience; it is the model that has trouble learning it. Let us run through a few families, not to inventory them, but to show that in each the "fixed" character is an office convenience, not a property of the world.

Hazardous molecules have no intrinsic hazard. The same molecule does not act in the same way on a soil bacterium, a pollinating insect, a plant or a human being; its toxicity is a property of the encounter between molecule and organism, not of the molecule alone. Ecotoxicology has long measured this under the name of inter-species sensitivity variability, and it is large enough that entire distributions are devoted to it (the *species sensitivity distributions*). To speak of a "hazard unit" attached to a registry number is already to erase the fact that the effect depends on the receiving system as much as on composition. We return to this in the companion vignettes, where this dependence takes the sharp form of a non-monotonic dose-response (§7.2).

Inert wastes are inert by definition, not always in practice. A demolition aggregate, a crushed brick, a block of concrete are classified as inert because, in mass and at rest, they react little. But a heap is neither in mass nor at rest: it takes in water, it weathers, it becomes heterogeneous. A leachate loaded with minerals that penetrates the mass can mobilize within it water-soluble toxic species (mobile precisely because they are not fat-soluble) capable of migrating through porous concrete toward the groundwater downstream. The fines, above all, change in nature as soon as they are wetted and then drained: they release lead, arsenic and chromium, and can contaminate downstream groundwater (Townsend et al., 2004). The "inert" material then becomes the seat of processes that no category sheet anticipates. This is the subject of vignette §7.4.

Biogas substrates are indeed decomposable, but poorly contained over time. The process itself, in the tank, is rather well mastered: that is its purpose. What is less so are the external interactions over the life of an installation (methane leaks at the tank, at flaring, during batch operations) and above all the fact that no seal is perfect over twenty years. The climate credit of a digester assumes capture close to 100 %; as soon as one asks about the real leak rate, that credit becomes a function of time, maintenance and technology. And there is worse than a slow leak: the regime shift of the process itself, on which we dwell at length (§7.5).

Remediation residues inherit the history that produced them. This is the most path-dependent case. The hazard of an ultimate residue (should it go to a hazardous-waste landfill, to heavy incineration?) is not read in its instantaneous composition alone, but in the upstream chain that made it: which metals have already passed through, by which treatments, at what real incineration temperature, with what dust-removal efficiency. Two residues of similar composition may call for two different routes depending on their past. This is exactly the negation of the path-independence axiom we will name later (S3).

The common thread. None of these flows is really fixed, except in idea, and only *if* everything is hermetically controlled: sealed containers, closed conduits, an invariant context. As soon as the system breathes, as soon as it exchanges with a soil, a water, a climate, a living thing, the fixed unit dissolves. The image of the field says it plainly: over a hectare of soil, no instrument guarantees the chemical composition or the waste-related risks to 100 %; one can anticipate average tendencies, but "anticipate" already presupposes that one has entered a domain where there is no fixed unit. This is not a hypothesis of the article: it is a field certainty, and it is the starting point of everything that follows.

A case that makes the problem tangible: the luminaire and the assembly of average data

Take a banal, well-regulated object: a lighting luminaire, with its environmental declaration (a PEP, in the sense of the PEP Ecopassport program). Its stated carbon footprint is, most often, *correct on average*, and that is exactly what makes it misleading. For this central figure rests on an inventory assembled from dozens of generic datasets: steel and aluminium of non-European origin, electronic components, plastics, all informed by databases that never describe the real factory but a *category average*, often split by broad plant-size bracket, and part of which may be more than five years old.

Environmental assessment has an instrument to qualify this situation: the data quality rating, or DQR, of the Environmental Footprint method. It averages four criteria (technological representativeness, geographical representativeness, temporal representativeness and precision), each scored from 1 (very good) to 5 (very poor), the lowest being the best (Zampori & Pant, 2019). Generic, geographically extrapolated and slightly dated data mechanically push temporal and technological representativeness upward, and the overall score typically settles around **2 to 3 out of 5**: "acceptable" in the administrative sense, yet already in the zone where confidence becomes fragile.

The point is not that a score of 3 is "bad." It is that it is *blind* to two things, and the companion preprint to this article proves it formally (Piens, 2026). First, the log-normal closure that translates the score into an uncertainty margin is **self-truncated**: even scoring everything at its worst, the aggregated geometric standard deviation hits a ceiling (on the order of 2.3), whereas the real variability of low-quality flows (expert estimates, extrapolations from one continent to another, technologies known by analogy) routinely exceeds it. Second, and most important here, this margin is computed **flow by flow**: it does not see that dozens of medium-quality inventory lines share the *same* determinants, a single electricity mix, a single transport mode, a single upstream metallurgical chain. When individual uncertainties are large *and* correlated through common drivers, the uncertainty propagated through the coupled system does not remain a gently widened Gaussian: it can develop bifurcations, regime transitions, heavy tails that the log-normal envelope is, by construction, unable to produce. A "DQR 3" footprint is therefore not a correct number with a slightly wide error bar; it is, potentially, a correct central number set on a data substrate whose real uncertainty regime is invisible to the tool that scores it. The risk of a regime shift remains considerable, and it remains unspoken.

A calculation makes the first blindness tangible. The quality score is translated into a margin by the aggregation formula of the log-normal closure, $\sigma\_ILCD = \exp(\sqrt{\ln^2\sigma\_b + \Sigma_i \ln^2\sigma_i})$, where σ_b is the basic uncertainty of the flow and each σ_i is the geometric uncertainty factor attached to one of the pedigree criteria (Ciroth et al., 2016). These factors are bounded above: even pushing the five pedigree criteria to their worst level (5/5/5/5/5) and adding the basic uncertainty, the aggregate σ_ILCD does not exceed about 2.3, a ceiling we derive from the ecoinvent v3 table in the companion preprint (Piens, 2026, §1.2). Yet genuinely poor flows (a fugitive emission estimated by expert judgment, a technology known by simple analogy) show empirical coefficients of variation that cross this ceiling. The relevant comparison is like with like: grouping a same emission factor across independent processes of one product family, the within-family dispersion is precisely the variability a category datum hides when it is substituted for the real process (the luminaire case), and its median, of order 3.7, exceeds the 2.3 ceiling for close to two families out of three (Piens, 2026). The widest margin the tool can produce is then *narrower than reality*: the closure truncates itself.

The second blindness is just as easily quantified. Let two inventory lines x_1 and x_2 be fed by the same upstream determinant, the same electricity mix, say. Treated in isolation, flow by flow, their variances add: $Var(x_1+x_2) = Var(x_1) + Var(x_2)$. But if they share that driver with a correlation ρ, the true variance is $Var(x_1) + Var(x_2) + 2\rho\cdot\sqrt{(Var\ x_1 \cdot Var\ x_2)}$, larger by half as soon as $\rho = 0.5$ for two flows of comparable variance. The pedigree, scoring each line separately, never sees this cross term. And the stakes go beyond an under-estimated error bar: in an inventory solved by inverting a technology matrix, $s = (I - A)^{-1}\cdot f$ (Leontief, 1936; Heijungs & Suh, 2002), a shared determinant that drifts pushes *several* coefficients in the same direction at once, which can make the propagated result cross a regime boundary that the sum of per-flow margins rules out by construction (Piens, 2026, §6.5). That is the whole difference between an uncertainty one widens and an uncertainty that changes in nature.

To this *data* issue is added, in the luminaire case, a *moving-target* issue in the use phase, which must be named separately so as not to confuse it with the former. Lighting consumes electricity, and the carbon factor of electricity is not a constant: it is itself a regime quantity. In France, the average carbon intensity of electricity production settled at 21.7 gCO_2eq/kWh in 2024 for direct emissions (30.2 gCO_2eq/kWh life-cycle), its lowest level ever; nearly half the time it stayed below 10 g, and even at consumption peaks it did not exceed 70 g (RTE, 2024). This very low baseline rests on a largely nuclear and hydroelectric output; but it oscillates with fleet availability and weather (a cold snap calls on gas and imports, the intensity climbs), and from one year to the next it can degrade by a factor of several, as in 2022. To freeze this factor in an LCA is to freeze a quantity that, in reality, moves hour to hour and year to year. The number stays "correct" in the sense that it reflects an average; it becomes misleading the moment a decision (choosing a route, comparing two products a few percent apart) rests on the assumed stability of that average.

This case concentrates the thesis of the article in miniature: not a false datum, but an *ill-posed object*. The central value is defensible; what is missing is the reading of the regime (the real quality of the data substrate and the real volatility of the factors), without which one takes an irreversible decision on an appearance of stability.

---

2. The implicit axiomatics of static evaluation

Every method of calculation rests, in the last instance, on assumptions it does not prove: its axioms. Environmental assessment (LCA, characterization factors, regulatory thresholds) is no exception, except that its axioms are written nowhere. They are only the more powerful for it: one does not debate what one does not state. We isolate four. For each, we first show *why it is reasonable* (none is foolish, each solves a real problem in its domain of origin), then we exhibit the precise class of waste or pollutant where it becomes false, and the price its falseness costs. The stake is not terminological. A tacit and false axiom does not produce an error one corrects by measuring better; it produces a *category* error, invisible from inside the framework that posits it, and that is precisely why it must be brought into the open.

Axiom S1: Point estimation. *An inventory or impact quantity possesses a central value surrounded by a bounded variance; a pair (μ, σ), most often log-normal, suffices to summarize its uncertainty.* This is the assumption that structures the pedigree matrix of Weidema and Wesnæs (1996), operationalized in ecoinvent and the ILCD: five qualitative scores become a single geometric standard deviation, then propagated by Monte Carlo. In its domain of origin (well-instrumented industrial flows, of pedigree 1 or 2), the axiom is defensible: the residual variability is plausibly log-normal there (Limpert, Stahel & Abbt, 2001), and the composition of factors close to unity produces a well-behaved aggregate. It gives way at the other end of the scale, in two ways. The first we have already quantified: the closure is self-truncated, its ceiling (`σ_ILCD ≲ 2.3`) is narrower than the real dispersion of poor flows. The second is deeper. In a bistable window, the stationary density of the output is not unimodal but *bimodal*, as we prove in §5.2 (theorem T2) and in the Appendix. Now a two-peaked density has a mean lodged in the trough that separates them, that is, in a state the system never occupies: the pair (μ, σ) there is not imprecise, it is devoid of meaning. To summarize a digester on the edge of tipping by "500 ± 80 kg CO_2eq" is to describe the engine by a regime intermediate between idle and over-revving, the one regime in which, precisely, it is not running.

Axiom S2: Linear characterization. *Impact is proportional to quantity,* `impact = CF × inventory`*, with a constant characterization factor; and the contributions of distinct flows add.* The biological or chemical response is therefore assumed to grow steadily with dose, without a peak, without reversal, without a change of mechanism. The assumption is everywhere: a single factor per substance in USEtox (Rosenbaum et al., 2008), a single acceptable daily dose in regulatory toxicology, an additive aggregation of impacts in LCA. And it is reasonable in its domain: for many effects, at the high doses of the assays, the response *is* monotonic, and independent stressors *add* roughly. The counter-example is sharp and has a name: the non-monotonic dose-response (NMDR), frequent among endocrine disruptors (Vandenberg et al., 2012). If the effect function `g(d)` has an interior peak (`g` rises then falls), it is not injective over the interval containing that peak, and *no* constant `CF` can reproduce it: for any CF, there exist two doses $d_1 < d_2$ on either side of the peak such that $CF \cdot d_1$ and $CF \cdot d_2$ order the impacts the reverse of $g(d_1)$ and $g(d_2)$. We make this theorem T3 (§5.2) and measure it in §7.2. It is not an imprecision one tightens by widening the margin: it is an error of *shape*, which survives any improvement of the data.

Axiom S3: Path independence. *The consequence depends only on the present state, never on the history that led to it.* Three corollaries accompany it: no hysteresis, no amplifying feedback, and a *de facto* reversibility, reducing the load brings the system back to its prior state. This is the assumption that allows the modules A to D of an LCA to be aggregated independently of their order, and the fate of a pollutant to be modeled by a stationary transport. It is exact for linear systems, which are by construction memoryless. It becomes false as soon as a positive feedback creates two basins of attraction: then two systems in the *same* present state, but arrived by different paths, do not have the same future, because each is captive of a distinct basin. The textbook case is the lake that tips from clear to turbid and does not come back when phosphorus is reduced (§7.9): the consequence of a reduction depends on the history, on whether or not the system has already crossed the tipping threshold. To assume S3 in this regime leads to the costliest error of all in practice: believing that "it will be enough to reduce" in order to restore, when the system is locked.

Axiom S4: Categorical invariance. *Hazard is an intrinsic property of the substance, independent of its physical form, its speciation and the receiving medium.* This is the assumption of classification by CAS number (REACH), of the European waste list, of composition-based eco-labels. Its rationale is administrative, and it is legitimate: to manage tens of thousands of substances, one needs one sheet per substance, not one sheet per context. The axiom becomes false whenever it is the *form*, and not the molecule, that carries the hazard. Asbestos is the canonical example: the same fiber, bound in a matrix, is harmless to the touch; released, it is carcinogenic. The fluoride of an ash is immobile as fluorite at neutral pH, but the hydrofluoric acid it forms at low pH attacks silica and releases a soluble hexafluorosilicate (§7.4). Arsenic is mobile in its oxidized form As(V) and nearly immobile in its reduced form As(III), for the same total quantity (§7.8). In all these cases, the hazardous property belongs not to the molecule alone but to the couple *molecule* × *environment*: to classify it once and for all is to enter in the register half of the truth.

We can now gather these four fractures into a diagnostic table. It reads not as a nomenclature but as a *key*: given a waste, to identify the most exposed axiom is to anticipate the failure mode to fear, the vignette that documents it, the theorem that formalizes it, and (as we will see in Part V) the appropriate decision framework.

Assumed axiom	What it posits	Where it gives way (mechanism)	Class / vignette	Theorem	Cost of the error
S1 Point estimation	a bounded-variance mean summarizes all	bimodal density in a bistable window; self-truncated ceiling	AM2 digester, §7.5	T1, T2	false margin, misleading median
S2 Linear characterization	`impact = CF × quantity`, additive	non-monotonic dose-response	endocrine disruptors, §7.2	T3	under-estimation by an arbitrary factor
S3 Path independence	the future depends only on the present	hysteresis, two basins of attraction	lake, digester, §7.9, §7.5	T2	"reducing" does not restore
S4 Categorical invariance	hazard is intrinsic to the substance	form/speciation changes the hazard	asbestos, fluoride, arsenic, §7.4, §7.8	T4	"inert" that turns mobile

Each vignette of Part III is therefore not an isolated curiosity: it is the *proof*, on a quantified case, that a precise axiom is false for a precise class. And since these are axioms (founding assumptions, not parameters), no accumulation of data repairs them: it is the model, not its precision, that is at fault. Part V (§9) will show that these four assumptions are not mere habits of calculation but are *engraved* in regulatory texts, which explains why their fractures recur so regularly.

---

3. State of the art and positioning

Our proposal sits at the meeting of four bodies of work, none of which, taken alone, does the job we aim at. Running through them is not a bibliographic formality: it is by situating precisely what each can do, and what each lets slip, that the narrow interstice where this article's contribution lodges appears.

3.1 The pedigree and the log-normal closure

This is the closest body of work, and the one we extend. Since the introduction of the pedigree matrix by Weidema and Wesnæs (1996), operationalized in ecoinvent (Frischknecht et al., 2005) and codified by the ILCD handbook of the Joint Research Centre (JRC, 2010), the uncertainty of an inventory flow is expressed by five qualitative scores (representativeness, completeness, temporal, geographical and technological correlations) converted by a table into uncertainty factors, then aggregated into a geometric standard deviation σ_ILCD that becomes the dispersion of a log-normal law propagated by Monte Carlo. Justice must be done to this closure: it is *operational* (a non-specialist scores a flow in a few minutes), *international* (the table is shared and audited), *parsimonious* (one number per flow runs through the rest of the analysis), and empirically defensible at high data quality, where the residual variability of well-instrumented inventories is plausibly log-normal (Limpert, Stahel & Abbt, 2001). It does what it was designed for.

Its limits appear where it was never meant to operate: the low-quality tail. The companion preprint (Piens, 2026) demonstrates two, already met in §1.1: the closure is *self-truncated* (its ceiling $\sigma\_ILCD \lesssim 2.3$ is narrower than the real dispersion of expert-estimated flows) and it is *per-flow*, hence blind to the coupling of the inventory resolved by Leontief inversion. Above all, it remains a *static* description: a fixed dispersion around a fixed value, with no vocabulary to say that a system might possess two regimes, or approach a tipping point. Our article adds this vocabulary for the axis of environmental fate, where the preprint added it for the pedigree axis.

3.2 Multimedia fate and transport

To describe where a pollutant goes, the discipline has proven tools: multimedia fugacity models (Mackay, 2001), which distribute a substance among air, water, soil and sediment according to partition equilibria, and the advection-dispersion-reaction equations that follow its transport. These are the right tools for a stationary, linear regime: they close mass balances, they give equilibrium concentrations, they are indispensable. But their founding assumptions (stationarity, linearity, constant average parameters) are precisely those that fall as soon as a flow instability, a redox threshold, a feedback appears. An advection-dispersion model with average flux does not produce the moisture finger that brings the pollutant thirteen times sooner (§7.3); a constant partition factor does not produce the mobility switch of arsenic between As(V) and As(III) (§7.8). These models answer admirably the question "how much, at equilibrium?" and stay silent on "when, and in which regime?".

3.3 Dynamic LCA and coupling to system dynamics

We are not the first to want LCA to "think in systems." The coupling of system dynamics (stock-and-flow models) to LCA is a constituted field, the subject of systematic reviews: McAvoy, Grant, Smith and Bontinck (2021) survey its approaches in the *Journal of Cleaner Production*, and a more recent review (Science of the Total Environment, 2023) analyzes some fifty studies published between 2010 and 2023. This is genuine systems thinking, and it captures feedbacks and temporal trajectories that frozen LCA ignores. But, and this is where our niche lodges (rule G1), its object is the *temporal* trajectory of an industrial or policy system: fleet renewal, technology diffusion, the evolution of a mix. McAvoy et al. show, moreover, that the integration almost always runs one way: LCA results feed a system-dynamics model. None of these approaches asks the three questions that occupy us: to *classify* uncertainty by type of non-linearity, to identify an *irreducible* floor distinct from the epistemic, and to diagnose that the score-and-threshold apparatus is *structurally* blind to regime. The existing coupling simulates a trajectory; we propose a regime diagnosis.

3.4 The regulatory toxicology of the threshold

The fourth body of work is not a modeling tool but an institution: the threshold. No-observed-effect level (NOEL), acceptable daily dose, maximum admissible concentration, a single number, below which "safe," above which "hazardous." This paradigm works for classical toxicants with a monotonic response, and it has the immense virtue of being decidable. But it carries within it two of the axioms of §2: point estimation (S1, a threshold is a disguised mean) and linear characterization (S2, the monotonicity of the dose-response). It therefore fails exactly where these axioms fail: for a non-monotonic dose-response (Vandenberg et al., 2012), a threshold calibrated on high doses declares safe a low-dose regime that is not (§7.2); and for a substance whose toxicological understanding is evolving, the threshold itself moves by several orders of magnitude, like PFOA whose guideline value fell by a factor of 17,500 in six years (§7.1). The threshold is not neutral: it is an object that carries axioms, and it inherits their fractures.

Positioning

To our knowledge, no prior work applies a regime diagnosis drawn from dynamical systems (bifurcation, Lyapunov exponent, multistability) to the assessment of wastes and pollutants, in complement to the pedigree closure. The novelty is one of *framing and tooling*, not of new physical measurements: we do not claim to measure a concentration better, but to *qualify* better the system that produces it. Where the companion preprint installed this diagnosis on the data-quality axis, the present article installs it on the axis of environmental fate, the real journey of the waste through soils, groundwater, the atmosphere and the living world.

---

4. A taxonomy of environmental non-linearity

The central gesture of this article is a shift of question. Practice asks: "what is the uncertainty of this number?", and answers with a margin. We ask: "*what kind* of system produced this number?", because depending on the answer, the static model does not fail a little more or a little less, it fails in *qualitatively* different ways, and each type of failure calls for a different treatment. The table below is the inventory of these types. It reads as a bridge: on the left, a *phenomenology* (what one observes in the data); in the center, the *normal form* that is its minimal mathematical signature; on the right, the *axiom of §2* the mode defeats and the vignette that quantifies it. To give a symptom its normal form is already to know which mathematics governs it and which error the static account will commit.

Failure mode (right word)	**Normal form / signature**	**Defeated axiom**	**Class & vignette**
NMDR, non-monotonic dose-response	$g(d) = d \cdot e^{(-d/d_0)}$ (interior peak, g non-injective)	S2	endocrine disruptors, §7.2
Heavy-tailed persistence (fat-tail)	$P(X>x) \sim x^{(-\alpha)}$, unbounded variance	S1	PFAS, §7.1
Chaos (strict sense)	$\lambda_{max} > 0$; $t_{pred} \approx (1/\lambda_{max}) \cdot \ln(L/\delta)$	S1 (irreducible)	atmosphere, fingering, §7.3
Bistability (double well)	$dx = (a \cdot x - x^3 + b)dt + \sigma \cdot dW$ (cusp normal form, two minima)	S1, S3	anaerobic digestion, §7.5
Contextual regime transition	$f(x; c)$ changes structure when c crosses a threshold	S4	asbestos, fluorinated fines, §7.4
Critical slowing down	$\tau \sim \|r\|^{(-1/2)} \to \infty$ as the fold is approached	(warning signature)	early signals, §7.6
Bistable eutrophication	hysteresis: same load → two stable states	S3	lake, §7.9
Redox switch	$K_d = K_d(Eh, pH)$, jump between two regimes	S4 (and S1, S3)	metals As/Cr/Pb, §7.8

These eight rows can be grouped by the axiom they break, and the table then becomes the key of §2 again. Heavy tails and chaos break **S1**: the first because a law with infinite variance has no "±" that bounds it, the second because sensitivity to initial conditions makes the mean itself unpredictable beyond a finite horizon. Non-monotonicity breaks **S2**, by non-injectivity of the dose-response. Bistability and eutrophication break **S3**: the future state depends on the basin of attraction, hence on history. Contextual transitions and redox switches break **S4**: it is form or speciation, not the molecule, that carries the hazard. Critical slowing down, finally, is not a failure mode but its *useful counterpart*: the signature a system emits as it approaches a tipping point, which Part III will exploit as an early-warning system (§7.6).

Each row is thus a different mathematics, a different right word, a different reason why the static account is not merely imprecise but *structurally incapable*, and, each time, a vignette of Part III to prove it on a quantified case.

---

5. The irreducible uncertainty floor (keystone)

Should environmental uncertainty be seen as a mere lack of data, which the measurement effort will eventually fill? In part only, and it is that part which is deceptive. Two components of different nature must be separated, because they call for different answers.

Epistemic uncertainty is the one more data dissipates: a concentration not yet measured, a poorly known parameter, a route informed by analogy. It is real, often dominant, and the only one practice can name. One reduces it by measuring better.

Dynamic uncertainty is of another kind. It persists even with an exact model and perfectly measured initial conditions, because sensitivity to initial conditions amplifies any infinitesimal imprecision at an exponential rate. No measurement budget removes it: it has a floor. This floor is the central result of this part, and it can be proved.

Theorem T1 (finite prediction horizon). *Let a deterministic system have maximal Lyapunov exponent* $\lambda_{max} > 0$*. Even known exactly and started from an initial condition measured to precision* δ*, it allows an observable bounded by* L *to be predicted only up to a finite time* $t_{pred} \approx (1/\lambda_{max}) \cdot \ln(L/\delta)$.

Proof (sketch). Two trajectories δ apart at the initial instant separate as $\delta \cdot e^{\lambda_{max} \cdot t}$ while the dynamics stays in its linearized regime. Prediction loses its meaning when this gap reaches the amplitude L of the observable, that is $\delta \cdot e^{\lambda_{max} \cdot t} \approx L$, hence $t_{pred} \approx (1/\lambda_{max}) \cdot \ln(L/\delta)$. ∎

The whole lesson is in the *form* of the result. The horizon does not grow in proportion to the measurement effort; it grows as its **logarithm**. To divide δ by a thousand (to demand a thousand times more precision on the initial state) adds only $(1/\lambda_{max}) \cdot \ln(1000) \approx 6.9/\lambda_{max}$ to the horizon: a handful of characteristic times, never infinity. This is what rigorously grounds the claim of §1.1 about the one-hectare field. Not only can it not be characterized to 100 % (epistemic limit), but even if it could, that would not be enough (dynamic limit). The two coexist, and the second does not negotiate with a budget.

The canonical example is the atmosphere. Lorenz (1963) showed on his convection model that an exponent $\lambda_{max} \approx 0.05\ d^{-1}$ imposes a weather-predictability horizon on the order of two weeks, a bound that fifty years of progress in data assimilation and computing power have not pushed back, precisely because it is not a limit of means. Any pollutant transport coupled to meteorology, or to a heterogeneous subsurface flow, inherits a finite horizon of the same nature; we quantify these horizons by medium in §5.2, then examine their vectors one by one in §6.

5.1 When more data does not remove the surprise

There are two distinct ways for the measurement effort not to suffice, and practice often confuses them under the same "±".

The first is epistemic, but *unreported*. A considerable share of LCA data is of medium quality: sets scored around 3 out of 5 on the pedigree, accepted as is in regulatory studies for want of better. Behind such a score lies an inter-site and inter-operator variability that, on real processes (a local energy mix, a water consumption, a transport stage), routinely exceeds 20 % without showing in the official uncertainty, because the log-normal closure has folded it into its capped standard deviation (§1.1). That uncertainty *could* be reduced: one would only have to go and fetch the primary datum. As long as one does not, it remains present and silent, a surprise in reserve, not an impossible surprise. It belongs to the epistemic, but to an epistemic the tool masks instead of flagging.

The second is of an entirely different nature: the quantity one freezes is a *moving target*, a regime, and no amount of past data pins its future value. The carbon intensity of French electricity illustrates it exactly. It settled at 21.7 gCO_2eq/kWh on average for direct emissions in 2024, its lowest level ever, below 10 g nearly half the time, and never above 70 g even at consumption peaks (RTE, 2024), but this very low baseline oscillates with the availability of the nuclear fleet and the weather, and it had clearly degraded in 2022. An LCA that enters "the" carbon factor of electricity freezes a quantity that, in reality, varies by a factor of several from one hour and one year to the next. This is not noise around a constant: it is the regime dynamics of the electrical system, and to freeze it is to mistake a photograph for the trajectory. Here, measuring the past with perfect precision says nothing certain about next winter, we have left the domain where "more data" is the answer.

5.2 Three corollaries (T2–T4) and the horizon by medium

T1 bounds the horizon; three corollaries specify *how* the static description fails when the system is non-linear. We state and prove them here; Part III makes quantified use of them, vignette by vignette.

Theorem T2 (no single value in a bistable window). *In a bistable window, under input noise, the stationary density of the observable is not Gaussian: it is bimodal.* The normal form of bistability (the cusp, a double well) under additive white noise is $dx = (a\cdot x - x^3 + b)dt + \sigma\cdot dW$, with potential $U(x) = -a\cdot x^2/2 + x^4/4 - b\cdot x$; the Fokker-Planck equation gives the stationary density $\rho(x) \propto \exp[(2/\sigma^2)\cdot(a\cdot x^2/2 - x^4/4 + b\cdot x)]$ (development in the Appendix). For $a > 0$ the potential has **two minima**, so ρ is **bimodal**

and normalizable: the quartic term makes the potential rise on both sides and bounds the integral. The consequence is sharp: the log-normal closure of the pedigree assumes a single mode and a bounded variance; in a bistable window it is not merely imprecise, it is *false in form*. This is the tool that dismantles the so-called "process-unit value" of the anaerobic digester (§7.5), where the mean falls into the trough between the healthy state and the acidified state, a state the reactor never occupies. The saddle-node fold $\dot{x} = r - (x-x^*)^2$ plays a distinct role: it marks the **washout**, the loss of the healthy branch at the edge of the window, where the density is metastable (the cubic potential is unbounded) and the escape time vanishes, tying the washout to the T1 floor.

Theorem T3 (non-injectivity). *If the dose-response $g(d)$ is not monotonic, no constant characterization factor reproduces it.* The proof holds in one line: a constant factor imposes $\text{impact} = CF \cdot d$, a strictly monotonic function of d; it therefore cannot coincide with a g that rises then falls over an interval containing its peak, where g takes the same value twice for two distinct doses $d_1 \neq d_2$. The operational corollary is severe: a factor calibrated on the range of regulatory assays (high doses) can under-estimate the real impact in environmental exposure (low doses) by an arbitrarily large factor. We give the measure of it in §7.2.

Theorem T4 (threshold blindness). *Near a bifurcation, a fixed threshold has no predictive value over the regime the system will occupy.* In the bistable window above, the two stable states $x_- < x_+$ (the two minima of the double well) can both lie on the same side of a fixed threshold $x = c$, so classifying "safe" or "hazardous" by whether x crosses c does not separate the two regimes; and the tipping from one to the other obeys a control parameter crossing a fold that destroys one of the wells, not the crossing of c by x. The threshold lives on the state axis x; the bifurcation, on the parameter axis. To classify by a fixed number then amounts to monitoring the wrong variable. T4 underlies two vignettes: the moving "safe" threshold of PFAS, which shifts by 17,500-fold without the molecule changing (§7.1), and the reclassification of inert wastes, where the class tips with speciation and not with a crossed content (§7.4).

It remains to say *where* T1's floor really bites, for the horizon t_{pred} depends on the medium through its exponent λ_{max}. The orders of magnitude below bound deterministic prediction (perfect model included) for the main transport vectors, which §6 will then examine one by one.

Medium	`λ_max` (order)	`t_pred` (order)	Origin of the estimate
Free atmosphere	~0.05 d^{-1}	~2 weeks	Lorenz (1963), convection model
Heterogeneous vadose soil	~0.003–0.01 d^{-1}	~1–4 years	front instability, fingering (§6.2)
Soil redox systems	~0.001–0.002 d^{-1}	~3–5 years	oxidation-reduction kinetics (§6.5)
Lakes / surface waters	~0.01–0.1 d^{-1}	~1–5 months	seasonal mixing, stratification

The reading matters more than the decimals: for the subsurface, the dynamic floor sits at the year scale, that is, *well inside* the horizons over which storage and remediation decisions are made, which count in decades. The uncertainty floor is therefore not an asymptotic curiosity relegated to infinity; it falls in the middle of the management calendar, where it is most in the way.

---

6. Conclusion: a change of object, and a program

Environmental assessment of waste was built on a fertile simplification: treat the waste as a classified unit, and uncertainty as a margin around a number. This convention has rendered immense service, and it continues to do so wherever the process stays confined, controlled, fed by measured flows. We have argued that it fails elsewhere, and that it fails not at the third decimal place, but at the level of the *kind* of object.

We established this in two stages. First by naming the four implicit axioms on which the whole edifice rests, point estimation, linear characterization, path independence and categorical invariance, and showing that each, reasonable in its domain of origin, becomes false for a precise class of waste. Then by proving four theorems that turn this observation into results: a finite prediction horizon that data do not push back (T1), a bimodal statistic in a bistable window that voids the log-normal margin (T2), the impossibility of a single characterization factor under non-monotonicity (T3), the blindness of a fixed threshold near a bifurcation (T4).

This theoretical foundation calls for its field proofs. The nine falsifiable vignettes that quantify these fractures, from the PFAS threshold that moves by 17,500-fold to the digester that is a sink or a source depending on its history alone, from preferential flow that over-estimates a safety delay by a factor of thirteen to the bifurcated collapse of a bee colony, are the subject of the companion articles of this program, as is the inscription of these axioms in law (REACH, CLP, EN 15804) and the decision frameworks that follow. What we propose here is not a universal predictive model; we have none, and we doubt one exists. It is a *change of object*: assess a waste's dynamical regime together with its quantity, distinguish within its uncertainty what is reducible ignorance from what is the irreducible floor, and match the decision to the diagnosed regime rather than to a mean that, near a tipping point, describes no real state.

For three centuries, science promised that to measure better was to predict better; and the promise was kept, as long as one studied systems willing to stay linear. Wastes are not willing. To recognize the boundary beyond which prediction is no longer bought with precision is not an admission of weakness: it is the beginning of an honesty.

Appendix. Sketch proof of T2 and T4 (Fokker-Planck on the double-well, with the saddle-node as washout boundary)

T2 concerns a bistable window, whose normal form is the cusp (a double well), not the saddle-node. Consider

$$dx = (a\cdot x - x^3 + b)\, dt + \sigma\, dW,$$

with potential $U(x) = -a\cdot x^2/2 + x^4/4 - b\cdot x$. The density $\rho(x, t)$ obeys the Fokker-Planck equation

$$\partial\rho/\partial t = -\partial_x[(a\cdot x - x^3 + b)\, \rho] + (\sigma^2/2)\, \partial_{xx}\, \rho.$$

At steady state ($\partial\rho/\partial t = 0$) with zero flux, one integration gives

$$\rho(x) \propto \exp[\, (2/\sigma^2)\, (a\cdot x^2/2 - x^4/4 + b\cdot x)\,].$$

For $a > 0$, the potential has **two minima** (at $\pm\sqrt{a}$ when $b = 0$): ρ is **bimodal and normalizable**, because the quartic term makes U rise on both sides and the integral converges. The mean falls in the trough between the two modes, a state the system never occupies. The log-normal closure of the pedigree (a single pair μ, σ) is therefore *structurally* unable to represent the output statistics in a bistable window: this is the statement of T2.

The saddle-node is the boundary, not the bimodality. At the edge of the window one well merges with the barrier; this is the fold $\dot{x} = r - (x - x^*)^2$, whose cubic potential $U(x) = -r\cdot x + (x - x^*)^3/3$ is **unbounded**, so the associated density is **metastable**, not stationary. As the fold is approached ($r \to 0^+$), the relaxation time diverges as $\tau \sim |r|^{(-1/2)}$ (critical slowing down) and the escape time out of the remaining well vanishes: this is the **washout**, the loss of the healthy branch, and it ties that loss to the finite-horizon floor of T1.

For T4, the argument is geometric. In the double well the two stable states are $x_- < x_+$ (the two minima). A fixed threshold $x = c$ separates them only if $c \in (x_-, x_+)$; when both minima fall on the same side of c, the threshold classifies the two regimes identically ("safe" or "hazardous" together), while the tipping is governed by a control parameter (here a, or the load) crossing a fold that destroys one of the wells, independently of c. The threshold lives on the state axis, the bifurcation on the parameter axis: this is the statement of T4.

(Background references: Strogatz 2018 for the cusp and saddle-node normal forms; the Fokker-Planck stationary-density literature for the double-well potential.)

References

- Anderson, P. W. (1972). More is different. *Science*, 177(4047), 393–396.
- Berkowitz, B., Cortis, A., Dentz, M., & Scher, H. (2006). Modeling non-Fickian transport in geological formations as a continuous time random walk. *Reviews of Geophysics*, 44(2), RG2003.
- Ciroth, A., Muller, S., Weidema, B., & Lesage, P. (2016). Empirically based uncertainty factors for the pedigree matrix in ecoinvent. *International Journal of Life Cycle Assessment*, 21(9), 1338–1348.
- Frischknecht, R., Jungbluth, N., Althaus, H.-J., et al. (2005). The ecoinvent database: overview and methodological framework. *International Journal of Life Cycle Assessment*, 10(1), 3–9.
- Heijungs, R., & Suh, S. (2002). *The Computational Structure of Life Cycle Assessment*. Springer.
- JRC, European Commission (2010). *ILCD Handbook: General guide for Life Cycle Assessment, Detailed guidance*. Joint Research Centre, EUR 24708 EN.
- Leontief, W. (1936). Quantitative input and output relations in the economic system of the United States. *Review of Economics and Statistics*, 18(3), 105–125.
- Limpert, E., Stahel, W. A., & Abbt, M. (2001). Log-normal distributions across the sciences: keys and clues. *BioScience*, 51(5), 341–352.
- Lorenz, E. N. (1963). Deterministic nonperiodic flow. *Journal of the Atmospheric Sciences*, 20(2), 130–141.
- Mackay, D. (2001). *Multimedia Environmental Models: The Fugacity Approach* (2nd ed.). CRC/Lewis Publishers.
- McAvoy, S., Grant, T., Smith, C., & Bontinck, P.-A. (2021). Combining life cycle assessment and system dynamics to improve impact assessment: a systematic review. *Journal of Cleaner Production*, 315, 128060.
- Rosenbaum, R. K., et al. (2008). USEtox, the UNEP-SETAC toxicity model: recommended characterisation factors for human toxicity and freshwater

ecotoxicity in life cycle impact assessment. *International Journal of Life Cycle Assessment*, 13(7), 532–546.

- Strogatz, S. H. (2018). *Nonlinear Dynamics and Chaos* (2nd ed.). CRC Press.
- Vandenberg, L. N., Colborn, T., Hayes, T. B., et al. (2012). Hormones and endocrine-disrupting chemicals: low-dose effects and nonmonotonic dose responses. *Endocrine Reviews*, 33(3), 378–455.
- Weidema, B. P., & Wesnæs, M. S. (1996). Data quality management for life cycle inventories, an example of using data quality indicators. *Journal of Cleaner Production*, 4(3–4), 167–174.
- Zampori, L., & Pant, R. (2019). *Suggestions for updating the Product Environmental Footprint (PEF) method*. JRC Technical Report EUR 29682 EN, European Commission.
- RTE (2024). *Electricity report / eco2mix: CO_2 emissions of electricity generation in France*. rte-france.com
- Piens, T. (2026). Chaotic dynamics in life-cycle inventory pedigree (preprint). Zenodo 10.5281/zenodo.20086772.

Supplement

Code and figures are reproducible (Python; numpy/scipy/matplotlib) and openly deposited (Zenodo DOI at the first stable version). The factors quoted are not cherry-picked points but representative values over a plausible parameter space, with robustness sweeps.