The Substrate Response Principle

Loading, Capacity Response, and the Recovery of Committed Information

Keith Taylor — VERSF Theoretical Physics Programme

General Reader Summary

Across the VERSF programme, several different phenomena — gravity, the way clocks slow near mass, the "collapse" that happens when a quantum measurement is made, and the formation of black-hole horizons — have each been explained as the substrate (the informational fabric underlying physics) reacting to how much committed structure it is carrying. The more committed structure piles up in a region, the harder it becomes to register anything new there. That reaction has been *assumed* in many places but never studied directly. This paper makes the reaction itself the object of study and calls the principle behind it the **Substrate Response Principle**.

The paper is deliberately modest about what it proves. It does not write down the exact law of how the substrate responds, it does not calculate any numbers (such as the strength of gravity), and it does not claim to have finished any of the programme's open problems. What it does is two things.

First, it pins down the general *shape* any such response law must have. Using only a few assumptions already part of the programme — that a finite region can hold only a finite amount of distinguishable information, that committed facts cannot later be erased, and that the "response" really does track what the substrate is holding rather than being an independent free-floating quantity — it proves a small theorem: the response must always rise (never fall) as loading increases, must be capped at a finite ceiling, and must level off toward that ceiling. In plain terms, *the substrate must eventually run out of room, and it cannot un-commit what it has already committed.* A useful by-product is a "no-go" result: any rival theory that lets the response grow without limit is ruled out.

Second, it sharpens the programme's deepest remaining question about how committed information is *read back* from the substrate. There are two genuinely different possibilities, and the paper shows they come down to a single yes-or-no question — does committed information carry an intrinsic "twist" (a curvature) or not? If it does not, reality keeps one consistent global record that every observer shares. If it does, there is no single global record, only local comparisons that need not agree. The paper proves this is the whole of the choice, shows that the rest of the programme does not yet settle it, and identifies what kind of measurement could. It does not claim to know the answer.

A few cross-cutting honesty notes run through the paper. One assumption (that the response faithfully tracks the substrate state) is doing real work, and the paper defends it rather than hiding it: rejecting it amounts to saying the "response" is not really a response at all. And one important question is flagged but left open — whether the "twist" studied here is the same object as the sevenfold closure structure studied elsewhere in the programme, or a different one. The paper's purpose, in short, is to isolate the right question and prove cleanly the few things that genuinely follow, rather than to claim more than the programme has earned.

Table of Contents

1. Introduction
2. The Forward Map
3. Constraints on the Response Law
4. The Response Saturation Theorem
    - 4.0.1 Why faithful registration is not optional
    - Theorem 4.1 — Response Saturation
    - Corollary 4.2 — No oscillatory response
    - Corollary 4.3 — Saturation limit
    - Theorem 4.4 — No-Go for unbounded amplification
    - Lemma 4.5 — Record recoverability forces saturation
5. The Minimal Response Law
    - 5.1 The Universality Conjecture
6. Gravity as Capacity Response
    - Corollary 6.1 — Horizons before geometry
7. Decoherence as Capacity Response
8. Horizons as Capacity Response
9. The Reverse Map — A Fork, Not a Filter
    - 9.1 The deciding object — reverse-map holonomy
    - 9.2 What the corpus does and does not settle
    - 9.3 The deciding observable
    - 9.4 An identification left open — Ω and the closure sector
10. The Unification Conjecture
11. Open Questions

Abstract

Much of the VERSF programme rests on a common but largely implicit assumption: that the substrate responds in a systematic way to the accumulation of committed distinctions. Gravity, decoherence, horizon formation, Ticks-Per-Bit scaling, and record stabilization all depend on some form of capacity response, yet the response itself has not previously been isolated and studied as a primary object.

This paper proposes the **Substrate Response Principle (SRP)**: the hypothesis that finite distinguishability, admissibility, closure, and record formation jointly constrain the allowable response laws of the substrate. Rather than treating gravity, measurement, and horizon phenomena as separate mechanisms, we argue that they can be read as different manifestations of a single response function linking committed informational loading to local distinguishability capacity.

The paper has two goals. The first is to determine the most general admissible form of substrate response consistent with the existing VERSF programme. The second is to characterise how committed information is operationally recovered from the resulting substrate state. Together these define the forward and reverse maps connecting loading, response, and observation.

The paper is explicitly a problem-isolation paper, not a derivation. Its central results are a constraint catalogue on the forward map (qualitative, not quantitative) and a sharp statement of the reverse-map question as a two-branch fork whose resolution is left open. No microscopic response law is derived, no saturation scale is fixed, and no numerical constant is claimed. Each such item is recorded as an open problem.

Epistemic markers used throughout: ***(established)*** *for results inherited from prior VERSF papers;* ***(constraint)*** *for restrictions the existing programme forces on any admissible response law;* ***(conjecture)*** *for the unifying hypotheses this paper proposes but does not derive;* ***(open)*** *for questions explicitly deferred.*

---

1. Introduction

The VERSF programme has developed a number of apparently distinct phenomena:

- gravity as a response to committed structure;
- proper-time suppression as suppression of local bit formation;
- horizons as capacity saturation;
- decoherence as local stabilization of distinguishability;
- Ticks-Per-Bit as a measure of commitment cost.

These developments share a common theme. Each requires a substrate whose ability to support distinctions changes as distinctions accumulate. In effect, each of these treatments assumes the existence of a response law without making that law its object.

The purpose of this paper is to elevate that response law to a primary object of study. The question is simple to state:

Given a local loading of committed distinctions, how must the substrate respond?

and conversely:

Given a substrate state, how is committed information recovered from it?

These two questions define what we shall call the **forward map** and the **reverse map** of the framework. The forward map has a relatively settled qualitative shape, constrained by results already established elsewhere in the programme. The reverse map does not: as §9 makes precise, it presents a genuine fork, and which branch the substrate realises is — on the present assessment of the programme — its central unresolved structural question.

---

2. The Forward Map

The forward map sends loading to response:

$L(x) \to R(x)$

where

- $L(x)$ — local committed-distinction density at x;
- $R(x)$ — local substrate response at x.

The interpretation of $R(x)$ is deliberately left open at this stage. Across the programme it has appeared as reduced distinguishability capacity, increased TPB cost, proper-time suppression, geometric response, and decoherence tendency.

The organising claim of this paper **(conjecture)** is that these are not distinct objects but distinct manifestations of one underlying response. §§6–8 develop that reading for gravity, decoherence, and horizons respectively; §10 states the unification claim in full and marks it as conjecture.

---

3. Constraints on the Response Law

Not every response law is admissible. The existing programme already imposes strong restrictions, and they are sharp enough to fix the qualitative form of $R(L)$ even though they do not fix it quantitatively. Each constraint below is labelled by the prior result that forces it.

Constraint 1 — Locality (constraint)

The response at a location cannot depend on arbitrarily distant loading, on pain of violating the locality structure inherited from the record-dynamics layer. Therefore $R(x)$ may depend on nearby loading but not instantaneously on arbitrarily distant regions:

$R(x) = R[\, L(y) : y \in \mathcal{N}(x) \,]$

for some bounded neighbourhood $\mathcal{N}(x)$. The precise content of "neighbourhood" is itself a substrate-level question and is not settled here.

Constraint 2 — Finite Distinguishability (constraint)

If loading increases without bound while response remains linear, finite distinguishability eventually fails. Response must therefore saturate: there exists a finite capacity

$C\_max < \infty$

beyond which additional loading no longer produces proportional distinguishability support. This excludes unbounded linear response laws.

Constraint 3 — Admissibility Preservation (constraint)

The response law must preserve the possibility of global admissibility. A law that amplified distinctions without bound would eventually permit incompatible records to coexist. Response must therefore act as a regulating mechanism rather than a runaway amplifier: the substrate must become increasingly resistant to further distinguishability as loading grows.

$dR/dL > 0$ (away from saturation); $d^2R/dL^2 \leq 0$ (near saturation only)

A clarification on curvature, since it is easy to over-claim. Monotone non-decrease is what admissibility preservation forces, and near capacity the response must turn over ($d^2R/dL^2 \leq 0$ in the saturating regime). But **global** concavity from the origin is *not* forced: a sigmoidal response — convex at low loading, concave near saturation — is monotone, bounded, and saturating, satisfying everything §4 proves, while violating $d^2R/dL^2 \leq 0$ at low L. We therefore record near-saturation concavity as a defining feature of the regulating regime and explicitly decline to assert concavity globally; §4's perimeter likewise does not derive it. An earlier formulation wrote the concavity condition unqualified; that overstated the constraint and is corrected here.

Constraint 4 — Record Stability (constraint)

Committed facts must remain recoverable. Loading may suppress new commitments; it may not erase existing ones. This asymmetry is fundamental: the substrate may resist additional loading, but it may not invalidate already-committed facts. Formally, if a record r is committed at loading L_0, it remains recoverable for all $L \geq L_0$ — meaning r stays readable along reconstruction paths, retaining a non-zero distinguishability margin from its alternatives. We intend recoverability here in this *persistence* sense, **not** in the sense of single-valued global recovery: r remaining readable is compatible with reconstruction being path-dependent (different paths may return r-relative-to-path), which is the open question of §9. The distinction matters because the whole paper turns on it — (RS) asserts that committed records persist readably, and is deliberately silent on whether their readings compose into one global value. The forward map is therefore constrained not only in its shape but in what it is forbidden to do to the structures the reverse map depends on: it may not drive any committed record's margin to zero.

4. The Response Saturation Theorem

The constraints of §3 are stated as restrictions the programme makes reasonable. Three of them can be sharpened into a structural result: monotonicity, finite capacity, and asymptotic saturation are not merely reasonable to impose — they are forced, from finite distinguishability and record stability together with faithful registration. We isolate that result here, prove the three properties that genuinely follow, locate the fourth named property (record preservation) as the premise it is rather than a conclusion, and leave the remaining qualitative features (§5) as labelled assumptions defining the admissible class rather than as theorem conclusions.

One honesty point about where the "force" comes from, stated before the proof rather than after. The theorem's conclusions follow tightly, and the tightness originates in the *strength of the premises* — in particular in (FR), and specifically in its strict clause (strictly greater realised support gives strictly greater response, until the ceiling). It would be easy to dismiss that strictness as a convenient modelling choice. We argue in §4.0.1 that it is not optional: it is a minimal *observability* condition on what it means for R to be a substrate response at all, and the cost of rejecting it is admitting that R carries information absent from the substrate state — i.e. that R is a hidden field, not a response. A reader who grants (FR) will find (i)–(iii) nearly immediate; a reader who rejects it is committed to R not being a response variable. We flag this so the result is not mistaken for deriving saturation from thin premises — it derives saturation from finite distinguishability plus an explicitly-defended registration condition.

4.0.1 Why faithful registration is not optional

Faithful registration is not an arbitrary extra assumption. It is the condition that prevents the response variable from becoming physically detached from the substrate state it is meant to describe. (We use the notation fixed in the Setup just below: L for loading, D(L) for realised distinguishability support, R(L) for the substrate response.) The premise splits into two claims, and it is worth separating them because they defend against different objections.

Identity of physical state. If R(L) is to be called a substrate *response*, then two loadings that realise the same distinguishability support must give the same response:

$D(L_1) = D(L_2) \Longrightarrow R(L_1) = R(L_2)$.

Otherwise R would contain information not present in the realised substrate configuration — it would be an additional hidden degree of freedom rather than a response to the substrate state. Rejecting this half is not a disagreement about dynamics; it is a concession that R is not what the paper says it is.

Non-degeneracy of registration. If one loading realises strictly more distinguishability support than another while preserving all earlier records, the registered response cannot be smaller — and, below the ceiling, must be strictly greater:

$D(L_2) > D(L_1)$ and $R(L_1) < C\_max \Longrightarrow R(L_2) > R(L_1)$.

A smaller response under strictly greater committed structure would mean the substrate had acquired additional committed structure while its response to committed structure decreased — breaking the intended reading of R as a capacity-response variable. The ceiling qualifier $R(L_1) < C\_max$ is essential and is what reconciles (FR) with saturation (Theorem 4.1(iii)): the strict clause is a claim about the regime *below* capacity. Once the ceiling is reached, further attempted loading may fail to increase realised support, or may be absorbed by coarse-graining, but it cannot produce unbounded response — so (FR) does not falsely demand that R keep rising after saturation.

The upshot. (FR) is best understood not as a free modelling postulate but as a minimal observability condition: the response must register exactly the distinguishability support the substrate actually sustains — no hidden response variable, and no silent loss of committed structure. The premise is therefore hard to reject on its merits: either one accepts (FR), or one accepts that R is not really "substrate response" but an extra unexplained variable. This is the sense in which the saturation theorem rests on a defended premise rather than a stipulated one — though we are clear that "defended" means "shown to follow from what R is meant to be," not "derived from more primitive substrate axioms," which remains the deeper open direction.

Why rejecting (FR) costs more than it saves. A critic may still reply: very well, let R be a hidden field. The following argument shows that this reply is not free — it does not remove an assumption, it adds one to the ontology.

Suppose (FR) fails. Then there exist substrate states with identical realised distinguishability support but different response:

$D_1 = D_2, R_1 \neq R_2$.

The response therefore carries information not contained in the realised support. Denote that extra information H, so that

$R = R(D, H)$, with H not recoverable from D.

But R was introduced precisely to describe how the substrate responds to committed structure. If R depends on H as well as D, then the physical state relevant to response is the pair (D, H), which carries strictly more information than D alone. Any theory rejecting (FR) must therefore introduce H as an additional hidden state variable beyond distinguishability support.

So (FR) is the unique minimal description of substrate response: it is the choice on which the response is a function of the realised support and nothing more. **Rejecting it does not eliminate assumptions; it increases ontology** — it forces a new hidden degree of freedom into the theory. This is the form of defence physicists weigh most readily: a strong assumption that *reduces* what must be posited is preferable to a weaker one that *adds* an unexplained variable. (FR) is the ontology-minimising option, and that is a stronger standing than mere plausibility.

One boundary must be stated honestly, because it is the only reply left to a determined critic. The minimality argument is minimality *relative to taking the substrate state to be fully specified by its realised distinguishability support* — the distinguishability-first ontology inherited from the broader programme. A critic could insist that (D, H) was the true state all along and that the D-only description was the lossy one; against that move, (FR) is not minimal but simply false. We do not refute this: it is a challenge to the distinguishability-first ontology itself, not to (FR) within it, and whether that ontology is correct is precisely one of the programme's deeper open commitments, not something this paper settles. Within the distinguishability-first setting, however — the setting the whole programme operates in — (FR) is the unique ontology-minimal response description, and rejecting it strictly enlarges the state space. That is as strong as (FR) can be made without deriving it from more primitive substrate axioms, which remains the open direction noted above.

Setup

Let the substrate at a location support a finite distinguishability budget. Write D(L) for the realised distinguishability support at loading L — the number of mutually distinguishable committed structures the substrate sustains at that loading — and let the response be R(L), the substrate's registered reaction to loading L. We take as given:

- **(FD) Finite distinguishability.** $D(L) \leq D\_max < \infty$ for all L. The substrate sustains at most finitely many mutually distinguishable committed structures at any loading. *(established; [VERSF–CHS], record-dynamics layer.)*
- **(RS) Record stability.** If a structure is committed at loading L_0, it remains among the distinguishable, recoverable structures at every $L \geq L_0$. Loading never removes an already-committed distinction. *(established; Constraint 4 above, fact-production layer.)*
- **(FR) Faithful registration.** R is a non-degenerate function of realised distinguishability support: equal realised support gives equal response, and strictly greater realised support gives strictly greater response *until the local response ceiling is reached* (i.e. wherever R is still below its supremum). R registers what the substrate actually sustains; it is not an independent hidden field. *(observability condition; see §4.0.1. The ceiling qualifier prevents the false demand that R keep rising after saturation; the supremum is referred to here as a primitive, and §4.1(ii) proves it is finite and equal to C_max.)*

Theorem 4.1 — Response Saturation (proven, from FD + RS + FR)

Under (FD), (RS), and (FR), any admissible response law R(L) possesses:

(i) ***Monotonicity*** *— R is non-decreasing in L ($dR/dL \geq 0$);*

(ii) ***Finite capacity*** *— R is bounded above by a finite $C\_max < \infty$;*

(iii) ***Asymptotic saturation*** *— $R(L) \to C\_max$ as L grows without bound.*

Equivalently: an admissible response cannot decrease as loading accumulates, cannot grow without bound, and must approach a finite ceiling.

A fourth property — **record preservation**, that loading never erases an already-committed distinction — completes the qualitative picture, but it is **not** a conclusion of this theorem. It is premise (RS) itself, and is listed among the four named properties only so the reader can see all four in one place. We flag this explicitly rather than re-derive (RS) from (RS): record preservation is the most secure of the four, precisely because it is foundational (it is the definition of commitment, not a consequence of the response law), and presenting it as a theorem output would be circular. The theorem proves the three properties that genuinely follow; the fourth is the assumption they are proven under.

Proof.

(i) Monotonicity. Suppose, for contradiction, that R is not monotonic non-decreasing: there exist loadings $L_1 < L_2$ with $R(L_1) > R(L_2)$. By (FR), R is a function of realised distinguishability support D and is non-decreasing in it (the strict clause below the ceiling sharpens this, but only the non-decreasing direction is needed here). So $R(L_1) > R(L_2)$ requires $D(L_1) > D(L_2)$ — strictly more distinguishable committed structure supported at the lower loading L_1 than at the higher loading L_2. But every structure committed at or below L_1 is, by (RS), still recoverable at L_2; the support at L_2 therefore includes all of the support at L_1. Hence $D(L_2) \geq D(L_1)$, contradicting $D(L_1) > D(L_2)$. No such pair (L_1, L_2) exists, so R is monotonic non-decreasing: $dR/dL \geq 0$.

(ii) Finite capacity. By (FR), R is a function of the realised support D(L), and by (FD) that support satisfies $D(L) \leq D\_max < \infty$ for all L. Since R is monotonic in D (by FR) and D is bounded above by D_max, R is bounded above by its value at the maximal sustainable support:

$$R(L) \leq R(D\_max) < \infty \text{ for all } L.$$

An unbounded response would require unbounded realised support, which (FD) forbids. Define

$$C\_max := \sup\_L R(L) \leq R(D\_max) < \infty.$$

Then R is bounded above by the finite capacity C_max. (Given the locality constraint of §3, C_max is a per-location ceiling — the capacity of the local distinguishability budget at a point — not a single global constant; we write it without a location index only to lighten notation, and nothing in the theorem requires it to be uniform across the substrate.)

(iii) Asymptotic saturation. By (i), R is monotonic non-decreasing along L; by (ii), it is bounded above by C_max. A monotonic non-decreasing function bounded above converges to its supremum as its argument grows. Therefore

$$\lim(L\to\infty) R(L) = C\_max.$$

Divergence is excluded by (ii), and oscillation is excluded by (i); the only remaining behaviour is convergence to the ceiling. The response saturates.

Parts (i)–(iii) establish the theorem. Note that the argument uses only (FD), (RS), and (FR): finite distinguishability supplies the bound, record stability supplies the non-decrease, and faithful registration ties the response to the support that both constrain. No appeal to locality or to admissibility-as-regulation is required — those play roles elsewhere in the programme but do no work here, and we do not list them among the premises of this theorem.

What the theorem does and does not establish

The theorem proves three of the qualitative properties used downstream — **monotonicity**, **finite capacity**, and **asymptotic saturation** — and correctly locates a fourth, **record preservation**, as the premise (RS) it is rather than a derived conclusion. The three proven properties are now structural results, not stipulations: they follow from finite distinguishability, record stability, and faithful registration by the arguments above, and any purported admissible response violating any of them is inadmissible.

The theorem does **not** establish two further properties that appear in the broader qualitative picture of §5:

- **Coarse-graining stability** — that the response law retains its form under rescaling of the loading description — does not follow from (FD), (RS), (FR). It is an additional regularity condition on how R behaves across scales, and we record it as a *defining assumption* of the admissible class (§5, item 4), not a theorem conclusion. Whether it can be derived from a substrate renormalisation argument is left open (§11).
- **Bounded support** in the strong sense — a sharp cutoff in the loading variable beyond which no response is registered — is not proven; the theorem establishes boundedness of R (finite capacity), not compactness of the loading domain. The distinction matters: a finite ceiling on R is weaker than a hard cutoff in L, and only the former is forced here.

This is the honest perimeter of the result. The theorem elevates monotonicity, finite capacity, and saturation from observation to structure; it names record preservation as the foundational premise it rests on; and it does not silently annex coarse-graining stability or strong bounded support, which remain assumptions.

Corollary 4.2 — No oscillatory response (immediate from 4.1(i))

No admissible response law can be oscillatory in loading. Concretely, there is no admissible R with loadings $L_1 < L_2 < L_3$ such that $R(L_2) > R(L_1)$ and $R(L_3) < R(L_2)$: a rise followed by a fall.

This is not a new result but the reader-facing form of monotonicity. Theorem 4.1(i) establishes $dR/dL \geq 0$; an oscillation requires a stretch where $dR/dL < 0$, which 4.1(i) forbids directly. We state it separately because "the response cannot oscillate with loading" is the physically vivid statement, and because it forecloses a tempting class of models in which response peaks and then relaxes as structure crowds the substrate — such models are inadmissible, not merely disfavoured. The strengthening here is one of exposition, not logical content: the work is done entirely by 4.1(i).

Corollary 4.3 — Saturation limit (immediate from 4.1(ii)–(iii))

As loading accumulates, the marginal response falls to zero: for any $\varepsilon > 0$ there is a loading L_ε beyond which any further loading ΔL produces a change in response smaller than ε. Formally, since R is monotonic (4.1(i)) and bounded above by C_{max} (4.1(ii)), it converges to C_{max} (4.1(iii)), so $dR/dL \to 0$ as L grows.

Interpretation, conditional. This says the substrate possesses a limiting regime in which additional loading produces vanishing additional response — a saturation limit. **If** gravity is a manifestation of substrate response (the conjecture of §6), this limiting regime is the natural home of horizon behaviour: the point past which loading no longer buys distinguishability support is precisely where §8 locates horizon formation. We flag the conditional explicitly: what is *proven* is that an admissible response has a saturation limit; that this limit *is* a horizon is the §8 reading, not a theorem. We avoid the term "fixed point," which would suggest an iterated map converging to a stationary state; the object here is the asymptotic ceiling of a monotone bounded response, which is a limit, not a dynamical fixed point.

Theorem 4.4 — No-Go for unbounded amplification (proven, contrapositive of 4.1(ii))

No admissible substrate response law can simultaneously satisfy finite distinguishability (FD), record stability (RS), and faithful registration (FR), together with unbounded amplification ($R(L) \to \infty$ as $L \to \infty$).

Proof. By Theorem 4.1(ii), (FD), (RS), (FR) jointly force R bounded above by a finite C_{max}. Unbounded amplification asserts R unbounded. The two are contradictory; no R satisfies all three premises plus unbounded amplification.

This is a no-go result in the strict sense: it excludes an entire family of candidate substrate theories — those positing a response that amplifies distinguishability without limit as structure accumulates — on the same premises that ground the rest of the programme. It is the contrapositive of finite capacity, stated as exclusion because that is the form in which it bears on rival accounts: any framework with finite distinguishability and stable records *cannot* have runaway response, regardless of its other commitments.

Lemma 4.5 — Record recoverability forces saturation (proven; largely independent route to 4.1)

Theorem 4.1 derives saturation from (FD), (RS), (FR). There is a second, *largely* independent route to the same conclusion that runs through *recoverability*. We say largely, not wholly, independent and are exact about the overlap: both routes need R tied to realised distinguishability (the coupling that is (FR)'s content), so 4.5 does not dispense with (FR). What 4.5 adds — the genuinely new ingredient absent from 4.1 — is the **margin-counting bound**: that a finite distinguishability budget partitioned into records each holding a minimum non-zero margin can contain only finitely many records. That bound, not independence from (FR), is the lemma's contribution.

We introduce one notion the theorem did not need: the **distinguishability margin**. A committed record is recoverable only if it retains a non-zero margin of distinguishability from every incompatible alternative — a record separated from its alternatives by zero margin cannot be read out, compared, or preserved, since nothing distinguishes it from what it is not. This sharpens (RS): record stability requires not merely that a record persist as a token, but that it persist *with a readable margin*.

Lemma 4.5. *Under (FD), (FR), and recoverability-with-margin, the response must saturate. Equivalently: if response does not saturate, then finite distinguishability and recoverability cannot both hold.*

Proof. Let a finite region have distinguishability capacity $C_D < \infty$ (this is (FD); C_D is a budget on the support D, distinct from the response ceiling C_{max} of §4.1). Let N(L) be the number of committed distinctions registered at loading L. Suppose, for contradiction, response does not saturate: then registered distinguishability grows without bound with L — here we use (FR), which ties the response to the realised distinguishability it registers, so non-saturating response means unbounded registered distinguishability — and hence $N(L) \to \infty$.

A finite capacity C_D cannot hold unboundedly many records *each separated from the others by a fixed non-zero margin* — the margins partition a finite distinguishability budget, so their number is bounded by C_D divided by the minimum margin (this is the margin-counting bound, the new ingredient). As $N(L) \to \infty$ within fixed C_D, therefore, one of two things must occur:

- **(A)** older records are overwritten or erased — which violates record stability (RS) directly; or
- **(B)** the distinguishability margin between records shrinks toward zero — which makes records unrecoverable, since a record with vanishing separation from its alternatives cannot be reliably read out.

Both horns are fatal: (A) destroys the records, (B) destroys their recoverability. So non-saturation is incompatible with the conjunction of finite distinguishability and recoverability. Contrapositively, finite distinguishability together with recoverability forces saturation.

Why this matters. Theorem 4.1 and Lemma 4.5 reach saturation by routes that overlap in their use of (FR) but differ in their decisive ingredient — 4.1 closes through faithful registration of bounded *support*, 4.5 closes through the margin-counting bound on recoverable *records*. Saturation is therefore not an artifact of any single *downstream* assumption: it follows whether one reasons from bounded support (4.1) or from record margins (4.5), given the shared premises (FD) and (FR). This is weaker than full independence — both routes rest on (FR) — but it is still a genuine robustness result: the conclusion does not hinge on the particular bridge (support-boundedness vs margin-counting) one chooses, only on the common premises plus either bridge. The residual concern a referee may still press is that (FR) underlies both; we own that explicitly in §4's perimeter and do not claim 4.5 escapes it.

This lemma also has a direct consequence for the relationship between the two maps, developed in §9: saturation is not an optional feature of the forward map but a precondition for the reverse map's records to be readable. The reverse map's recoverability requires the forward map's saturation — a one-way dependence, not a symmetric pairing.

---

5. The Minimal Response Law

The constraints of §3 jointly suggest a simple qualitative structure. At low loading,

$R(L) \approx k L$ (linear regime)

and at high loading,

$R(L) \to C_{max}$ (saturation regime)

The detailed function is not known and is not derived here. What admissibility fixes is the qualitative class: any admissible R(L) is

1. monotonic non-decreasing — *proven, Theorem 4.1(i)*;
2. positive — *immediate, since response registers non-negative support*;
3. saturating, with finite C_{max} — *proven, Theorem 4.1(ii)*;
4. stable under coarse-graining — *defining assumption of the class, not proven (Theorem 4.1, perimeter)*;
5. record-preserving in the sense of Constraint 4 — *premise (RS)*.

Items 1 and 3 are now structural results; items 4 and 5 are stipulations that define which laws count as admissible. This defines a **universality class** of admissible substrate responses rather than a unique law. Whether the class collapses to a single member under further substrate constraints is recorded as an open problem (§11).

A caution on status. The pairing "linear at low loading, saturating at high loading" is the simplest representative of the class, not a derived result. It is consistent with §3 but not forced by it; other members of the class (for example, responses that are sub-linear from the outset) are not excluded by the constraints as stated. The claim of this section is only that the class is non-empty and qualitatively constrained, not that its minimal representative is the physical one.

5.1 The Universality Conjecture (conjecture)

The constrained class invites a stronger hypothesis, which we state as a conjecture and are careful not to present as a theorem.

Universality Conjecture. All admissible response laws — those satisfying locality, finite distinguishability, record stability, and faithful registration — flow under coarse-graining toward a common infrared response structure. Microscopic differences between admissible laws are washed out at large scale, so that gravity, decoherence, and horizon behaviour depend only on the shared infrared form and not on the microscopic details of R(L).

If established, this would be the substrate-response analogue of universality in critical phenomena: just as microscopically distinct systems share critical exponents because they flow to the same renormalisation-group fixed point, microscopically distinct admissible response laws would share their observable large-scale consequences because they flow to the same infrared structure. The downstream phenomena would then be **insensitive to substrate microphysics** — a property that would explain why a single effective description (gravity, the measurement transition, horizon thermodynamics) succeeds without knowledge of the underlying law.

We must be explicit about the gap between this and §4. Theorem 4.1 and the class definition above establish that admissible laws are monotone, bounded, and saturating — a *constrained class*. The Universality Conjecture asserts that this class *flows to a common fixed point* under coarse-graining, which is a far stronger and entirely separate claim: it requires a defined coarse-graining operation on response laws, a flow generated by it, and a proof that admissible laws converge under that flow. None of this machinery is constructed here. The analogy to thermodynamics and the renormalisation group is offered as the right *language* for the conjecture and as motivation for pursuing it — not as evidence that it holds. Establishing it (or finding the admissible law that fails to flow) is recorded as an open problem (§11, Q6).

---

6. Gravity as Capacity Response

Gravity can be read within this picture as follows.

Mass corresponds to persistent committed structure **(established)**. Persistent structure loads local capacity. Capacity response modifies local distinguishability support, and the resulting gradients produce proper-time suppression. Worldlines extremize proper time, and the observed result is what we call gravitational attraction.

On this reading gravity is not a separate mechanism but one manifestation of the substrate response law. This is the qualitative, infrared content already developed in the gravity papers of the programme; it is restated here to exhibit gravity as a *case* of R(L), not to add to it. In particular, this section derives no field equation, fixes no coupling, and does not bear on the open status of Newton's constant (§11, Q3).

Corollary 6.1 — Horizons before geometry (conditional on the §6 reading)

If gravity is a manifestation of substrate response, then horizon-like structures are fixed by the response law alone — they are the saturation limit of distinguishability support (Corollary 4.3), reached where loading exhausts capacity — and do **not** require an independent geometric postulate to exist.

The conceptual force of this is that horizon-like behaviour is *expected before any geometry enters the description*. In the standard account a horizon is a geometric object: a surface in a metric manifold where a null congruence fails to escape. Here, under the §6 reading, the horizon regime is already present at the level of the response law, as the regime $R(L) \to C_{max}$ — prior to, and independent of, the emergence of the metric in which a geometric horizon would be defined. Geometry, when it emerges, inherits a horizon where the substrate was already saturated; it does not create one.

This does not derive general relativity, and it does not derive the horizon area law or its thermodynamics — those remain open (§8, §11). What it establishes, conditional on the substrate-response reading of gravity, is the weaker but non-trivial claim that the *existence* of horizon-like structure is a consequence of saturation rather than an independent feature of the emergent geometry. The conditional is essential and we keep it visible: the corollary is a consequence of the §6 conjecture, not an independent result about gravity.

---

7. Decoherence as Capacity Response

The same response appears in measurement.

An unmeasured quantum system corresponds to unresolved distinguishability — distinctions held provisional, below the commitment threshold **(established)**. Environmental interaction increases local loading. Increasing loading pushes the local substrate toward saturation, and as saturation is approached, stable record formation becomes increasingly favoured. The result is decoherence.

Measurement is therefore, on this reading, a local manifestation of the same response that appears at large scale as gravity. The two are presented here as regimes of one law, not as a derived identity; the identity itself is the conjecture of §10.

---

8. Horizons as Capacity Response

The strongest response occurs when loading approaches capacity:

$R(L) \rightarrow C_{max}$

At this point the substrate can no longer support additional distinguishability. This is horizon formation. On this reading horizons are not singular boundaries but regions where distinguishability support reaches saturation — the same law that governs weak gravity and measurement, extended to its limiting regime.

This is a structural interpretation, not a derivation of horizon geometry. The relation between the saturation surface $R(L) = C_{max}$ and the standard horizon area law is not established here and is part of the open problem set.

---

9. The Reverse Map — A Fork, Not a Filter

The second half of the problem is the reverse map:

$R(x) \rightarrow I\_committed$

This asks how committed information is read from the substrate. A physical observer never accesses the substrate directly; an observer accesses records. The reverse map is therefore a reconstruction process — the extraction of stable, admissible structures from the underlying substrate state.

Before describing its structure, we record a result that ties the reverse map to the forward map. It is, in proof-content, a restatement of Lemma 4.5 — corollary-strength relative to 4.5, not an independent theorem — and we state it separately not because it proves something new but because the dependence it makes explicit is the conceptual hinge of the paper: it is what licenses treating forward and reverse phenomena as downstream of one constraint. Throughout this theorem, "recoverable" means **recoverable-with-margin in the sense of Lemma 4.5** — records persist with non-zero distinguishability margin from their alternatives — and "the reverse map is non-degenerate" means it *has such recoverable records as input*. Neither phrase here means *path-independent* or *single-valued*; that stronger property is Branch A only (§9, below), and the present theorem is deliberately neutral between the branches. The distinction is load-bearing: the theorem is about whether the reverse map has readable input at all, not about whether its output is globally single-valued.

Theorem 9.0 — Saturation is Necessary for Readout (proven, from Lemma 4.5 under FD)

Given finite distinguishability (FD), the reverse map can be non-degenerate (records recoverable-with-margin) only if the forward map saturates:

Recoverability-with-margin ⟹ Saturation (under FD).

That is: saturation is a necessary condition for the reverse map to have recoverable input. A non-saturating forward map leaves the reverse map no recoverable input. The theorem asserts necessity, not sufficiency — saturation does not by itself produce recoverable records — and it says nothing about path-dependence; it holds identically on both branches of §9.

Proof. Suppose the reverse map is non-degenerate: records are recoverable-with-margin. By Lemma 4.5, recoverability-with-margin together with (FD) forces the forward map to saturate. Contrapositively, a non-saturating forward map drives records (under FD) to either erasure or zero-margin unreadability, leaving the reverse map no recoverable input. Hence recoverability requires saturation.

The converse does **not** hold and is not claimed: saturation is necessary but not sufficient for recoverability. A saturating forward map whose records were never laid down with margin — or whose margins were lost for reasons other than non-saturation — satisfies saturation while failing recoverability. Record persistence is carried by (RS), not produced by saturation;

saturation is the condition that *permits* recoverable records, not the cause that *guarantees* them. The result is therefore a one-way dependence, which is all Lemma 4.5 licenses and all the paper needs.

The force of Theorem 9.0 is that it makes the paper's title principle genuinely *one* principle. The reverse map's readable records cannot exist without the forward map's saturation ceiling; the two are not independent facts that happen to co-occur but are linked by necessity under finite distinguishability. Consequently gravity (forward, §6), measurement (forward, §7), horizons (forward saturation limit, §8), and record recovery (reverse, §9) are downstream of a single constraint rather than four independent posits. The shared premise (FD) is essential and we keep it visible: without a finite distinguishability budget the margin-crowding argument of Lemma 4.5 does not run, and the dependence dissolves. With it, the reverse map's existence is contingent on the forward map's saturation — a one-way lock, not a symmetric equivalence, but enough to tie the two maps to one underlying constraint.

The temptation is to declare the reverse map a *filter*: a process that simply selects the stable structures surviving the response field. We resist that declaration, because it silently answers what is in fact the open question of the programme. The reconstruction process admits two structurally different forms, and they are not equivalent.

Branch A — Global ledger (path-independent readout)

Readout consults a single global record that all admissible observers share. Committed information is recovered by reference to one consistent ledger, and any two reconstruction paths that reach the same event return the same committed content. Formally, the reverse map is **path-independent**:

$I_{committed}(x)$ is independent of the reconstruction path γ reaching x.

Under Branch A the reverse map is well-defined as a global object, the "filter" reading of the naïve picture is correct, and the leftover closure-memory channel of the substrate is **occupied** — there is a global structure for it to register.

Branch B — Local comparisons (path-dependent readout)

Readout relies only on local comparisons between neighbouring committed structures, which need not reconcile into a single global picture. Reconstruction is then **path-dependent**: distinct paths reaching the same event may return distinct committed content, with the discrepancy a holonomy of the reverse map around closed loops.

$I_{committed}(x; \gamma) - I_{committed}(x; \gamma') = H(\ell) \neq 0$ in general,

where $\ell = \gamma \circ \gamma'^{-1}$ is the loop formed by the two paths and H is the reverse-map holonomy defined in §9.1.

Under Branch B there is no global ledger to consult, the "filter" reading fails as a global statement, and the leftover closure-memory channel is **empty** — there is no global structure for it to register, only locally reconciled fragments.

Why this is the fork

Branches A and B are not two descriptions of one thing; they are different physics with different downstream consequences. Whether the closure-memory channel is occupied or empty — a question that bears on the readout side of measurement as much as on gravity — is exactly the question of which branch the substrate realises. The present paper does **not** select a branch. It claims only that the reverse map must be one or the other, that the choice is not free once the substrate's readout structure is fixed, and that characterising each branch and its distinguishing signature is the proper content of the reverse-map problem.

This is the regime-characterisation stance: state both branches, state what each implies, and identify what would distinguish them — rather than assume the branch one would prefer. A confirmed path-dependence in committed-record reconstruction would select Branch B and exclude Branch A; a demonstrated global path-independence would do the reverse. Which obtains is open.

9.1 The deciding object — reverse-map holonomy (proven characterisation)

The fork is not a vague pair of pictures. The two branches are distinguished by a single well-defined object, and we can say exactly what it is.

Local comparison of committed structures at neighbouring substrate points defines a transport rule — a connection Γ on the bundle of committed-record data over the substrate. This connection is the reverse-map analogue of the gauge connection A_μ established for the forward comparison of rays in the gauge sector of the programme: it specifies how a committed record at one point is carried to a comparison at a neighbouring point. Given Γ, reconstruction along a path γ from a reference point to x is the transport of record data along γ, and the recovered content $I_{committed}(x; \gamma)$ is the endpoint of that transport.

Define the **reverse-map holonomy** around a closed loop ℓ as the net transport accumulated on traversing ℓ:

$H(\ell) := \oint_\ell \Gamma.$

Proposition 9.1. *The reverse map is path-independent (Branch A) if and only if its holonomy vanishes on every contractible loop:*

Branch A $\Leftrightarrow H(\ell) = 0$ *for all contractible* $\ell \Leftrightarrow \Gamma$ *is flat (zero curvature).*

Conversely, the reverse map is path-dependent (Branch B) if and only if $H(\ell) \neq 0$ *for some loop.*

Proof. If two paths γ, γ' reach the same point x, their concatenation $\gamma \circ \gamma'^{-1}$ is a closed loop ℓ. The difference in recovered content is the transport around ℓ:

$I\_committed(x; \gamma) - I\_committed(x; \gamma') = H(\ell)$.

If $H(\ell) = 0$ for all contractible loops, every such difference vanishes and reconstruction is path-independent — the defining property of Branch A. If $H(\ell) \neq 0$ for some loop, the two paths bounding it return different content — path-dependence, the defining property of Branch B. Vanishing holonomy on all contractible loops is equivalent, by the standard connection–curvature correspondence, to flatness of Γ.

So the fork reduces to one question with a precise answer: **is the reverse-map connection Γ flat or not?** Branch A is the flat case; Branch B is the curved case. This is a genuine sharpening — the choice between a global ledger and irreconcilable local comparisons is exactly the choice between zero and non-zero curvature of a single connection.

The local form of this is worth naming, because it turns the fork into a recognisably physical question. In the abelian (or linearised) case, define the **closure-memory curvature** as the exterior derivative of the reverse-map connection:

$\Omega := d\Gamma$ (abelian / linearised),

with $H(\ell) = \oint\_\ell \Gamma = \int\_\Sigma \Omega$ by Stokes' theorem on the loop ℓ bounding a surface Σ, so that non-zero holonomy on contractible loops is equivalent to non-zero Ω.

A note on generality, since Γ is introduced by analogy to the gauge connection and the gauge sector is non-abelian ($SU(3)\times SU(2)\times U(1)$). If Γ is non-abelian, the correct curvature is $\Omega = d\Gamma + \Gamma\wedge\Gamma$ and the holonomy is the path-ordered exponential $\mathcal{P}\exp(\oint\_\ell \Gamma)$, not the plain integral; the formulae above are the abelian linearisation. We state the abelian form because it suffices to define the fork — the characterisation that *decides the branches* is flatness, and "Branch A $\Leftrightarrow$ trivial holonomy on all contractible loops $\Leftrightarrow$ Γ flat" survives intact in the non-abelian case (a flat connection has trivial path-ordered holonomy on contractible loops just as in the abelian case). What changes non-abelianly is the explicit expression for Ω and H, not the flat/curved dichotomy that the fork rests on. A reader needing the non-abelian curvature should read $\Omega = d\Gamma + \Gamma\wedge\Gamma$ throughout; the fork is unaffected. The branch statement is therefore:

Branch A $\Leftrightarrow$ Γ flat (trivial holonomy on contractible loops) $\Leftrightarrow$ $\Omega = 0$ in the abelian form,
Branch B $\Leftrightarrow$ Γ curved (non-trivial holonomy) $\Leftrightarrow$ $\Omega \neq 0$ in the abelian form.

Stated this way the question is: **does committed information possess intrinsic curvature?** That is a physical question about the substrate's record structure, not a metaphysical one about ledgers — and it is the form in which the fork is most likely to connect to a substrate-level computation, since Ω is a local object that a substrate model could in principle yield. We name Ω to make the question concrete; we do not claim to know its value. Whether Ω vanishes is precisely the open content of the fork (§9.2–9.3).

Curvature-Unification Conjecture (conjecture)

Ω is not merely *analogous* to curvature; it is the substrate-level instance of curvature itself, of which the gauge field strength and the curvature of emergent geometry are the forward-sector and geometric-sector instances. On this conjecture the programme's three connection-like objects — the gauge connection A_μ (forward transport of rays, §1), the emergent metric connection (geometry, the gravity papers), and the reverse-map connection Γ (readout, §9.1) — would each carry a curvature, and physics would read as a set of curvature statements at three levels: gauge curvature governs force, geometric curvature governs gravity, and closure-memory curvature Ω governs readout. The fork of §9 would then be the special case, at the readout level, of the general question "is this connection flat?"

We state this as a conjecture and nothing more. What the paper establishes is that Ω is *a* curvature (of the connection Γ); what the conjecture asserts is that it is *the same kind of object* as the curvatures in the gauge and gravitational sectors, unified at the substrate level. That identification would require showing the three connections descend from a common substrate structure — which the paper does not do and the programme has not done (it is part of the continuum-limit and closure-algebra problems recorded elsewhere). The conjecture is offered because it is provocative and, if true, structurally deep: it would make "flatness of a connection" the single question underlying force, gravity, and readout alike. It is not offered as a result.

9.2 What the corpus does and does not settle

It is tempting to think the gauge and magnetism results already decide this — that because the programme establishes comparison as *local* (a connection is required; there is no primitive global lookup), the global ledger of Branch A is thereby excluded. This inference does not go through, and it is worth being exact about why, because the error is natural.

The gauge sector establishes that **comparison is local**: reconstruction proceeds by transport along a connection, not by consultation of a primitive nonlocal table. That excludes "global lookup" as a *mechanism*. But Branch A is not global lookup as a mechanism. Branch A is the case in which the local connection Γ is *flat*, so that local transport happens to reconstruct to a path-independent global ledger as an *emergent* property. A flat connection is fully local in its dynamics and still path-independent — general relativity supplies the standing example of a local theory supporting flat (path-independent) sectors alongside curved ones. The gauge corpus therefore fixes that transport is local; it does **not** fix whether that local transport is flat. Locality of comparison is consistent with both branches.

Hence:

- The corpus **excludes** primitive nonlocal readout (a global table consulted instantaneously). Neither branch requires it, so nothing is lost.
- The corpus **does not exclude** Branch A, because Branch A needs only a flat local connection, not nonlocal lookup.
- The corpus **does not force** Branch B, because local comparison permits a flat connection (zero holonomy) as readily as a curved one.

Neither branch is ruled out by what is presently established. This is not a failure to look hard enough; it is the precise content of the open question. Whether Γ is flat is a fact about the substrate's readout structure that the existing results constrain (transport is local) but do not determine (flat or curved).

9.3 The deciding observable

Because the branches differ by H(ℓ), they are in principle distinguishable by observation rather than only by stipulation. A confirmed dependence of a reconstructed committed record on the reconstruction path — equivalently, a non-zero H(ℓ) around some physically realisable loop — would establish Branch B and exclude Branch A. A demonstration that reconstruction is path-independent across loops that would otherwise register curvature would establish flatness and hence Branch A. The holonomy is the signature; its vanishing or non-vanishing is the experiment.

The present paper does not compute H(ℓ) — that computation is a substrate-level problem the programme has not closed, and is recorded among the open questions (§11, Q4). What this section establishes is narrower and real: the reverse-map fork is controlled by a single object, the flatness of the reverse-map connection; that object is not fixed by the existing corpus, which constrains transport to be local without constraining it to be flat; and the object is observable in principle as reconstruction path-dependence. The fork has been narrowed from two pictures to one number, without that number being invented.

9.4 An identification left open — Ω and the closure sector

The phrase "closure-memory channel" used above should not be read as a settled pointer to any specific structure elsewhere in the programme, and one cross-corpus question must be flagged explicitly rather than left to implication. The reverse-map curvature Ω defined here is a continuous object — a curvature 2-form valued (via the gauge analogy) in a Lie algebra. The closure sector studied in the Gate-3 line of the programme is a *discrete* object — a cohomology class valued in $\mathbb{Z}_7$. Whether these are the same channel in two notations or two genuinely distinct objects is **open**, and the paper asserts neither.

If they are the same channel, the discrete-versus-continuous valuation must be reconciled: one would need to know whether the physically meaningful holonomy is $\mathbb{Z}_7$-valued (a discrete closure residue) or Lie-algebra-valued (a continuous curvature), and how a continuum limit relates them — the same continuum-limit question that recurs across the programme. If they are distinct, the "closure-memory" language here is merely suggestive and carries no commitment to the Gate-3 sector. We do not resolve which holds; we flag the identification as open (§11, Q7) so that the present fork and the Gate-3 closure-residue question are neither silently conflated nor silently severed. Determining their relationship is a cross-paper coherence check that is prior to circulating this paper alongside the Gate-3 line.

---

10. The Unification Conjecture

The central conjecture of this paper **(conjecture)**:

Gravity, decoherence, TPB scaling, horizon formation, and record recovery are manifestations of a single substrate response law. They are not separate physical mechanisms but different regimes of one capacity-response principle.

If correct, the divide between gravitational physics and measurement theory is largely one of scale: both become special cases of a deeper relationship between committed distinctions and the substrate that supports them, with low loading giving quantum behaviour, local saturation giving classical definiteness, and global saturation giving gravity and horizons.

This is stated as a conjecture, not a result. §§6–8 exhibit each phenomenon as *consistent with* a single R(L); they do not prove that the same R(L) governs all of them with the same parameters. The unification is a proposal for how the pieces fit, and its discharge depends on the open problems below — in particular on whether a unique admissible response law exists (Q5) and on which branch the reverse map takes (§9).

The conjecture has one proven leg, which distinguishes it from a bare hope: Theorem 9.0 establishes that reverse-map recovery is not independent of the forward-map phenomena — a readable reverse map requires forward-map saturation under finite distinguishability. So the unification is not conjectural everywhere — the forward/reverse seam is secured by a necessity result (recovery depends on saturation); what remains conjectural is that the *forward* phenomena share one law with one parameter set.

Three conjectures appear in this paper and should not be conflated. This one (§10) concerns *phenomena* — that gravity, decoherence, and horizons are regimes of one response law. The Universality Conjecture (§5.1) concerns *laws* — that all admissible laws flow to one infrared structure. The Curvature-Unification Conjecture (§9.1) concerns *connections* — that gauge, geometric, and readout curvatures are one object at three levels. They are complementary claims at three different levels of the framework, and each could hold or fail independently of the others.

---

11. Open Questions

The paper isolates the problem; it does not solve it. The following are recorded as open **(open)**:

1. **Microscopic form.** What is the exact microscopic form of the response law R(L), beyond the universality class of §5?
2. **Saturation scale.** Can C_{max} be derived from substrate structure rather than inferred?
3. **Newton's constant.** Can the gravitational coupling — and ultimately Newton's constant — emerge directly from the response function, rather than being matched? (This is the same closure-normalisation gap recorded as open elsewhere in the programme; the present paper does not close it.)

4. **Reverse-map resolution.** §9 reduces the Branch A / Branch B fork to a single object — the flatness of the reverse-map connection Γ, equivalently the vanishing of the closure-memory curvature $\Omega = d\Gamma$ or its holonomy $H(\ell)$ — and shows this is not fixed by the existing corpus. What remains open is to compute Ω from substrate structure, or to determine it observationally via reconstruction path-dependence. Does committed information carry intrinsic curvature?
5. **Uniqueness.** Does a unique admissible response law exist, or only the universality class of §5?
6. **Universality flow.** Does the Universality Conjecture (§5.1) hold — do all admissible response laws flow under coarse-graining to a common infrared structure? Establishing this requires constructing the coarse-graining operation and proving convergence; finding an admissible law that fails to flow would refute it.
7. **Closure-sector identification.** Is the reverse-map curvature Ω (§9.1, §9.4) the same object as the Gate-3 closure sector $\kappa \in H^1(\Gamma\_vac; \mathbb{Z}_7)$, or a distinct one? If the same, how do the continuous (Lie-algebra-valued) and discrete ($\mathbb{Z}_7$-valued) valuations reconcile under the continuum limit? This determines whether the §9 fork and the Gate-3 closure-residue question are one open problem in two notations or genuinely two.

These define the next stage of the programme. The contribution of the present paper is narrower and prior: to isolate the substrate response law as a first-class object rather than an implicit background assumption, to catalogue the constraints the existing programme already places on it, and to state the reverse-map question in its true form — as a fork between a global ledger and local comparison, whose resolution determines whether the framework's last open structural channel is occupied or empty.