

# On the Assumption of Distinguishability in Physical Theory

## A Constraint-First Perspective and the VERSF Programme

Keith Taylor VERSF Theoretical Physics Programme

---

### Abstract

Modern physical theories have achieved extraordinary success by constructing mathematical frameworks that describe and predict observable phenomena with high precision. These frameworks — quantum mechanics, quantum field theory, general relativity — share a common structural feature: they are formulated on spaces of distinguishable states, configurations, or events. The existence of such distinguishability is not typically derived within these theories; it is assumed as part of their foundational structure.

This paper is in two parts. Part I isolates and examines the assumption of distinguishability. It argues that the existence of distinguishable physical states is not a trivial or purely formal prerequisite but a restrictive structural condition: not all mathematical substrates can support persistent distinguishable records, and those that cannot are excluded as candidates for physical reality. We state a formal failure criterion — a substrate fails to support observation if there exists no mapping from its configurations to finite, retrievable observational records respecting differentiation, identifiability, and structural stability simultaneously — and give concrete examples of substrates that fail it: purely reversible dynamics, unbounded-precision continua, single-state manifolds. Standard physical theory, we observe, operates on the admissible side of this partition without testing admissibility; it assumes it.

Part II presents the Void Energy-Regulated Space Framework (VERSF) as one substantive pursuit of the constraint-first perspective. VERSF is characterised by three physical axioms — finite distinguishability, irreversible commitment, finite local capacity — each of which is defended as a physics requirement whose removal breaks a specific operational feature. The axioms are shown to be mutually irreducible and to substantially constrain the space of admissible substrates; VERSF is argued to occupy a specific, so-far-unoccupied position in that space. The programme's distinctive dynamical move is to treat commitment events as a source for a substrate field (the  $\kappa$ -field). The paper argues this move is forced rather than chosen, resting on the principle that physical existence and causal activity are not separable — where causal activity is stated precisely as appearance in the equations of motion, admitting no intermediate state. VERSF yields a no-go theorem on substrate geometry (whose bite depends, at each step, on the constraints of Part I), a discrete Noether theorem, a Single-Source Theorem for observables, and numerical predictions including the fine-structure constant, the tensor-to-scalar ratio, and a parameter-free experimental signature. Because a programme of this scope attracts a wide range of objections, Part II adopts a question-and-response format, treating eighteen

anticipated objections — including whether the admissibility constraints of Part I substantially collapse the space of viable substrate theories (we argue they do), whether facts require observers (we argue they require commitment, of which observation is one special case), whether facts require mathematically exact distinctions or only operational distinguishability (we argue the latter), and whether the relativity of observation undermines the notion of facts (we argue it does not, once fact occurrence is distinguished from fact reception). The epistemic register throughout distinguishes *proven*, *conditional*, and *conjectural* results, and each numerical prediction is accompanied by an explicit statement of what its failure would localise to.

The paper is intended both as a self-contained argument for the constraint-first perspective and as a consolidated entry point to the VERSF programme.

**Keywords:** foundations of physics; distinguishability; constraint-first approaches; substrate programmes; emergent time; falsifiability; VERSF.

---

# Part I. The Constraint-First Perspective

## 1. Introduction

The development of modern physics has been driven by a consistent methodological pattern: when confronted with unexplained phenomena, one seeks a deeper level of description. This approach has yielded atomic structure, subatomic particles, quantum fields, and the geometric formulation of gravity. Its success has strongly reinforced the expectation that fundamental explanations reside at increasingly microscopic scales.

At the same time, the foundational structures of contemporary theories share a striking commonality. Quantum mechanics is formulated on a Hilbert space of states. Quantum field theory extends this structure to fields defined over spacetime. General relativity models physical reality as a differentiable manifold endowed with a metric. In each case, the theory presupposes the existence of a space of distinguishable configurations.

This part of the paper focuses on a simple but underexamined question:

What must be true for distinguishable physical states to exist at all?

This question is not typically treated as a source of constraint in physical theory. Distinguishability is taken as given, and the task is to describe its dynamics. The aim of Part I is to examine whether this assumption can be elevated to a structural principle — not as a philosophical gesture, but as a restrictive condition that partitions the space of candidate substrates into those that can support physics and those that cannot. If the partition is real, the perspective is productive; if the partition is empty, the perspective collapses into a tautology. The argument of Part I is that the partition is real.

## 2. The Standard Starting Point in Physical Theory

Across its major formulations, physics begins with structures that encode distinguishability.

In *quantum mechanics*, the state of a system is represented by a vector in a Hilbert space, with observables corresponding to operators whose eigenstates define distinguishable outcomes. In *quantum field theory*, fields define local excitations that can, in principle, be differentiated and tracked. In *general relativity*, spacetime events are distinguishable points within a manifold, structured by a metric that defines causal and geometric relations.

These frameworks differ profoundly in interpretation and structure, yet they share a foundational feature: they presuppose the existence of distinguishable states, configurations, or events. Theories are then constructed to describe how these distinctions evolve, interact, and are related.

## 3. Distinguishability as a Restrictive Condition

The existence of distinguishable states is typically treated as a minimal or trivial requirement. We argue it is restrictive. The assumption decomposes into three implicit components, each of which identifies a class of substrates incapable of supporting observation:

*Differentiation.* There must exist physical configurations that are not identical. A substrate without differentiation reduces to a single state, incapable of encoding outcomes.

*Identifiability.* Configurations must be distinguishable in principle. A substrate that admits configurations indiscernible in principle collapses the distinction between distinct predictions; the mathematical framework on it would describe possibilities that do not correspond to separable physical situations.

*Structural stability.* Distinctions must persist long enough to be meaningfully referred back to. A substrate that produces records decaying faster than they can be consulted dissolves the operational content of measurement: the concept of a "measurement outcome" is not available on it.

Each failure mode removes a prerequisite that predictive physics silently relies upon. We can state the failure condition compactly:

A substrate *fails to support observation* if there exists no mapping — realised by finite physical interactions producing stable records — from its configurations to finite, retrievable observational records that respects differentiation, identifiability, and structural stability simultaneously.

The mapping is required to be physically realisable, not merely mathematically definable: a purely formal correspondence between substrate configurations and record-space elements does not count unless it is the image of some finite interaction capable of producing the record. The three components are therefore joint conditions on physically realised mappings from substrate

to record, not independent desiderata. A substrate may support some of them and not others; what is required for physics is that all three hold together, and the failure modes below are instances in which one or more break.

The failure modes can be instantiated concretely. A substrate whose dynamics are purely reversible as a complete description — for example a closed unitarily-evolving system with no mechanism anywhere by which commitment can occur — fails structural stability at the level required: any recorded configuration is in principle recoverable by the reverse dynamics, and the supposed record has no operational permanence. The point here is not that reversible dynamics are excluded from the constraint-first perspective (they are not: unitarily-evolving quantum mechanics is one of the theories the perspective is designed to accommodate), but that a substrate containing no commitment mechanism at any level cannot support records — a distinction further developed in §9.5 below. A substrate with unbounded local state density — a strict continuum admitting arbitrary precision in any finite region — fails operational identifiability: configurations can be formally distinct without being distinguishable by any finite measurement, so the mathematical distinguishability of the substrate does not translate into distinguishability of outcomes. A substrate without differentiation is the trivial limit of a single-state manifold. Each is a mathematical structure that has been, or could be, written down; each fails the condition stated above. The failure modes are not abstract.

The point can therefore be put as a concrete constraint rather than a philosophical observation:

Not all mathematical substrates are capable of supporting persistent distinguishable records, and those that cannot are excluded as candidates for physical reality.

The assumption of distinguishability therefore partitions the space of candidate substrates. Standard physical theory operates on one side of this partition without naming the partition; it selects substrates that support observation and studies their dynamics, but it does not test admissibility — it assumes it. A constraint-first perspective asks what the partition looks like and whether its shape imposes structural content that dynamics alone do not see. If the partition is trivial — if every conceivable substrate is on the side that supports physics — there is nothing here. If it is not, there is work to do upstream of existing frameworks.

## 4. Distinction vs. Description

The point can be sharpened by separating two levels.

The *descriptive* level comprises the mathematical structures used to represent and evolve distinguishable states: Hilbert spaces, manifolds, fields. The *existential* level comprises the conditions required for distinguishable states to exist at all.

Standard physical theories operate primarily at the descriptive level. The existential level is implicitly assumed. If the partition identified in §3 is real, then the existential level has structural content: the substrate must lie on the supporting side of the partition, and doing so imposes conditions that propagate into the mathematics. Physical theories describe distinctions; the constraint-first perspective asks what conditions make distinctions describable at all.

## 5. Relativity and the Locality of Distinction

The role of distinguishability becomes more subtle in the context of special relativity. Relativity eliminates global simultaneity, absolute ordering of events, and observer-independent slicing of spacetime. It does not, however, eliminate local structure. Physical theories in relativistic settings still rely on the ability to distinguish events within a local region, causal relationships, and invariant intervals.

Thus, while global descriptions of distinction may be frame-dependent, the existence of local distinguishability remains indispensable. This suggests a refined formulation:

Distinction need not be globally defined, but it must exist locally for physics to be meaningful.

This refinement matters. The constraint-first perspective is not committed to absolute or global facts; it is committed to the existence of local, invariant record structure. This commitment is compatible with relativistic physics as it stands, and it engages the strongest version of the frame-dependence objection that Part II addresses (§12 below).

## 6. A Constraint-First Perspective

The standard approach in physics can be summarised as:

Given distinguishable states, construct dynamics and derive predictions.

A constraint-first perspective reverses the order:

Given that distinguishable states must exist, what constraints does this impose on the underlying structure?

This shift does not reject existing theories. It asks whether their shared assumptions can be treated as sources of constraint rather than as background conditions. The partition identified in §3 is what makes the reversal productive: if not every conceivable substrate supports physics, then asking what supports physics is a non-trivial question with potentially restrictive answers.

## 7. Implications for Foundational Theory

Treating distinguishability as a structural condition suggests several avenues for further work. One is the identification of minimal conditions under which distinguishable states can exist. Another is the determination of whether these conditions restrict the allowable mathematical structures of physical theory. A third is the investigation of whether irreversibility, information, or entropy are necessary components of distinction. A fourth is the exploration of whether known constants or scales can be linked to such constraints.

Historically, progress in physics has often involved identifying and elevating implicit assumptions to explicit principles: the invariance of physical law under coordinate

transformations; the role of symmetry in conservation laws; the operational meaning of time and measurement. The assumption of distinguishability may represent a similar opportunity. The perspective developed here does not claim that current physical theories are incomplete in their domains of validity. It suggests that they begin from assumptions that themselves carry structural content, and that excavating that content is worth the effort.

## 8. From Perspective to Programme

The constraint-first perspective does not, by itself, determine a unique foundational theory. It is a methodological orientation: take the existence of distinction as the primary datum and ask what it requires. A number of substantive programmes can be, and have been, pursued in this spirit — Wheeler's "it from bit" [1], Zurek's decoherence and einselection programme [2], Deutsch and Marletto's constructor theory [3], the axiomatic reconstructions of quantum mechanics due to Hardy [4] and Masanes and Müller [5], Rovelli's relational quantum mechanics [6], the Connes–Rovelli thermal time hypothesis [7], Spekkens-class epistemic reconstructions [8], causal set theory [9], loop quantum gravity [10], and causal dynamical triangulations [11] all, in different ways, take operational or informational preconditions as constraining physical structure.

The remainder of Part I stages the transition from perspective to programme in four steps. First we state what the perspective requires of any candidate substrate. Then we present VERSF's three physical axioms and defend each as a physics requirement whose removal breaks a specific operational feature. Then we defend the most distinctive VERSF move — commitment-as-source — as forced rather than chosen. Finally we locate VERSF within the space of possible constraint-first programmes, making explicit what is specific to it and what is not.

### 8.1 What the perspective requires of any substrate

If the constraint-first perspective is taken seriously, a candidate substrate must support at least the following:

- *Differentiation*, since without it there is nothing to record.
- *Identifiability*, since indistinguishable records cannot ground predictions.
- *Structural stability*, since records that dissipate before reference cannot be referred back to.
- *Locality of record*, since relativity forbids global distinction but allows local distinction (§5).
- *Irreversibility in the record-producing process*, since the operational content of time — the thing that is counted in any clock — reduces to counting irreversible transitions that leave stable traces.

These are not equivalent to any specific set of technical axioms, but they are the requirements any specific axiomatisation must honour. A substrate failing any of them fails the perspective regardless of how it is formalised.

### 8.2 VERSF's three physical axioms

VERSF sharpens these requirements into three physical axioms which, together with the technical structure axioms A1–A6 used in the no-go theorem [Taylor, VERSF programme: No-Go Theorem], form the basis of the programme. Each physical axiom is stated here alongside a defence of its status as a physics requirement rather than a mathematical stipulation.

**A. *Finite distinguishability.*** Any finite region of the substrate supports a finite number of distinguishable configurations.

*If removed:* a region supports unboundedly many distinguishable configurations, so no finite observation can exhaust the record content of even a small volume. Holographic-style bounds on information density are violated; operational measurement loses content, because there is no fact of the matter about which of the infinitely many possible outcomes was recorded. The axiom is required for measurement to be operationally meaningful.

**B. *Irreversible commitment.*** The elementary transitions by which records come into being are ordered and cannot be undone.

*If removed:* records can be undone by the dynamics. The arrow of time loses its substrate-level representation and must be imported from outside. Operational time — the count of irreversible transitions in a clock — cannot emerge, because the transitions supposed to be counted are, in the underlying substrate, reversible. The axiom is required for any substrate-level derivation of emergent time.

**C. *Finite local capacity.*** Each local region supports a finite number of simultaneous commitments.

*If removed:* records can accumulate without bound at any point. The substrate supports any pattern of outcomes equally, which is to say it cannot constrain outcomes at all. Locality of physical influence, and the distinction between a universe with many records and one with few, both fail. The axiom is required for records to have operational cost and for locality to survive.

Each axiom corresponds to an operational feature of physics, and removing any one of them breaks a specific operational feature in a specific way. The axioms are not mathematically innocent, but they are not arbitrary either: they are the minimum commitments needed to turn the constraint-first perspective into a working programme.

The three axioms are mutually irreducible on their own content: each addresses a distinct class of failure mode identified in §3, and no subset of two suffices to cover all three. Finite distinguishability (A) addresses operational identifiability; irreversible commitment (B) addresses structural stability; finite local capacity (C) addresses the locality and cost structure without which differentiation becomes trivial. Full minimality in the stronger sense — that no other three-axiom set could cover the same ground — is not claimed, but minimality within the stated axioms is straightforward to verify by inspection: removing any one leaves at least one failure mode from §3 uncovered.

With these axioms in place, the space of admissible substrates appears to be substantially constrained. The irreversibility axiom excludes purely reversible substrate descriptions (in the sense clarified in §9.5 — reversible sub-dynamics are permitted within a substrate that contains commitment elsewhere); the finite distinguishability axiom excludes unbounded-precision continua; the finite local capacity axiom excludes substrates without operational cost structure. Even before the technical structure axioms A1–A6 are applied, the candidate space is substantially constrained. When the full axiom set is applied, the no-go theorem forces a specific geometry within the set: a triangular 2-complex with  $K=7$  minimal facts. VERSF's substrate is therefore not a selection from a broad pool of equally admissible candidates; it is the specific occupant, identified by the no-go theorem, of a constrained region of the candidate space. Whether this region has other occupants — whether alternative substrate structures could satisfy the same axioms — is a separate question, and one we intend to address in a future uniqueness analysis.

The relationship between the constraints and the theorem is worth making explicit: the constraints of §3 and the axioms of §8.2 are what make the no-go theorem bite. Without them, the theorem's restriction of the substrate space would be weaker or vacuous. The  $PGL(3,2)$  irreducibility argument on  $V_6$  operates inside a space that has already been narrowed by requiring irreversibility, finite local state, and finite capacity; if any of those requirements were dropped, candidate substrates outside the narrowed space would be admitted and the argument would not reach the  $K=7$  conclusion. The theorem is therefore not a self-contained mathematical curiosity about simplicial complexes — it is a derivation that depends, at each step, on the physical content of the constraints stated in Part I.

The no-go theorem operates on these three physical axioms together with the technical structure axioms, and appears to force — within this axiom set — that the substrate support the  $K=7$  simplicial structure. The uniqueness claim is of a specific form: VERSF is the unique substrate structure consistent with *this* axiom set, as established by the theorem. Whether the axioms themselves are the right ones remains open to challenge, and the theorem is offered as a tool for that challenge, not as a claim that closes it.

### **8.3 Why commitment must be a source, not mere bookkeeping**

The distinctively VERSF move is to give commitment events a dynamical role: they source the  $\kappa$ -field. This is the most original and the most attackable claim in the programme, and it deserves explicit defence before Part II rather than only as a response to an objection within it.

The defence rests on a general principle: physical existence and causal activity are not separable. A thing that exists physically affects the future evolution of the substrate. A thing that does not affect the future evolution of the substrate does not exist physically, whatever descriptive apparatus may be wrapped around it. This is not a stipulation; it is what "physical existence" means in the context of any theory whose content is given by equations of motion. A record that does not enter those equations — that does not appear in the Lagrangian, that does not influence what happens next — is at most an external label assigned to a substrate that would evolve identically whether the label were present or absent. It is not a physical record.

The constraint-first perspective requires records to exist physically (Part I, §§3–6). By the principle just stated, this requires records to be causally active. We can state what causal activity means without ambiguity:

A feature of the substrate is *causally active* if and only if it enters the equations of motion — that is, if the substrate's future evolution would differ depending on whether the feature is present.

This is not a philosophical definition; it is the operational content of causal activity in any theory whose dynamics are given by equations of motion, and it admits no intermediate state. Either a record appears somewhere in the action (as a source, a field, a modification of a coupling) or it does not. If it does, it is physical. If it does not, it is at best an external label assigned to a substrate that evolves identically with or without it.

For records to be causally active in this sense, the substrate must support a mechanism by which commitment events enter the equations of motion. Sourcing a field is the minimal such mechanism. Commitment-as-source is the move that implements it. The argument for the move is therefore not "this is one option that happens to work" but "this is what is needed for records to have the physical status the perspective demands of them."

Three specific failures follow if commitment is not a source.

First, irreversibility becomes bookkeeping. One can label transitions "irreversible" without making them elements of the dynamics, but such labelling imports the arrow of time from outside rather than deriving it. The axiom B of §8.2 then becomes a descriptive stipulation on an external ordering, not a structural feature of the substrate.

Second, records have no dynamical cost. Nothing in the substrate distinguishes a universe in which many commitments have occurred from one in which few have, because the substrate is unchanged by commitment. This conflicts directly with axiom C's requirement that records carry operational cost; it also conflicts with the pre-theoretic observation that the universe is evidently not the same with and without the measurements that have occurred in it.

Third, and decisively, there is no field whose equations of motion can be written down. Without such a field, there is no action principle connecting the substrate to observable physics; without an action principle, no emergent Lorentz invariance, no discrete Noether conservation laws, no recovery of effective field theory in the decoupling limit. Each of these results (see §§13.1, 13.2, 13.3) depends structurally on the source role.

Commitment-as-source is therefore not a free stipulation layered on top of the axioms. It is the move required to give records the physical status the perspective demands, and the results of Part II depend on it. A programme that adopts the constraint-first perspective but refuses the source role is entitled to do so, but it forgoes the machinery that makes the perspective a dynamics rather than a philosophy of measurement.

#### **8.4 VERSF as one construction, not the construction**

With these pieces in place, VERSF can be stated compactly. It is a specific axiomatisation of the constraint-first perspective, characterised by a specific triad of physical axioms (§8.2), a specific dynamical move (commitment-as-source, §8.3), a specific derived substrate geometry ( $K=7$  triangular 2-complex), a specific field content (the  $\kappa$ -field with derivable mass  $m = \sqrt{(4/3) \cdot \xi^{-1}}$ ), and a specific structural scale (the closure scale  $\xi$ , established as consequence of the axioms rather than as free parameter [Taylor, VERSF programme: Closure Scale]).

Each of these specifics could in principle be replaced by an alternative while preserving the constraint-first perspective. VERSF is one construction, not *the* construction. That said, to our knowledge it is currently the only construction in the published literature that simultaneously (i) satisfies the three physical axioms of §8.2, (ii) implements the commitment-as-source move defended in §8.3, and (iii) yields testable numerical predictions derived structurally rather than fitted. Alternative constructions meeting all three conditions may exist; we are not aware of them. The appropriate framing is therefore that VERSF occupies a particular and, so far, unoccupied position in the space of constraint-first programmes — a position whose specificity is protected by the no-go theorem and articulated by the numerical predictions. What distinguishes it within the family of constraint-first programmes is detailed in the Q&A that follows.

The programme is therefore open to challenge at three distinct levels, each of which can be engaged independently: the constraint-first perspective (challenge in Part I terms); the specific axiomatisation (challenge any of the three physical axioms); or the specific construction (challenge the technical derivations, the predictions, or the recoveries of standard physics). The epistemic register used throughout — *proven* results following from stated theorems within the axioms; *conditional* results following if certain axioms or structural identifications are accepted; *conjectural* results with substantial progress but no complete derivation — is maintained so that each level of challenge can be met where it lives.

One clarification before proceeding. Part II's defence of VERSF is not a defence of the constraint-first perspective as such. The perspective can be pursued without committing to VERSF's specific axioms, and a reader convinced by Part I but unconvinced by Part II loses nothing of Part I. Conversely, objections that successfully target VERSF do not automatically target the constraint-first perspective. The two theses are logically independent.

---

## Part II. Questions and Responses on VERSF

### 9. Foundational Objections

#### 9.1 Question 1: Physics does not need "facts" — it needs predictive models

**The objection.** Physics is an empirical, predictive discipline. Introducing a new primitive called a "fact" risks importing philosophical baggage without earning physical content. Standard measurement theory and statistical mechanics already handle whatever needs handling.

**Response.** The point is not to add a primitive but to make explicit what predictive physics already presupposes. Part I argues that this presupposition is restrictive rather than descriptive (§3): substrates failing to support differentiation, identifiability, or structural stability are excluded from supporting physics, and the exclusion is structural content. The "fact" terminology is a name for what is already in use; the move that makes it productive is treating the exclusion as a constraint to be exploited rather than a background condition to be assumed.

This is standard practice in foundational work. Decoherence theory names the structures that produce classical-looking outcomes from a quantum substrate [2]. Constructor theory names the tasks a substrate must support [3]. QBist reconstructions name the agent-relative status of probabilities. Every foundational programme introduces vocabulary for structures that predictive physics uses implicitly; the question is whether the new vocabulary earns its keep by generating consequences. VERSF does, and where those consequences are testable — the parameter-free protocols and numerical predictions discussed in §§11 and 14.2 — the work is physics by any reasonable standard.

## 9.2 Question 2: This is information theory restated

**The objection.** Wheeler's "it from bit" [1], Zurek's quantum Darwinism and einselection [2], Deutsch and Marletto's constructor theory [3], Spekkens-class epistemic reconstructions [8], and the axiomatic reconstructions of quantum mechanics due to Hardy [4] and Masanes and Müller [5] all proceed from informational or operational primitives toward physical structure. VERSF appears to occupy the same territory without acknowledging the lineage.

**Response.** The lineage is real and was acknowledged in §8. VERSF sits within the broader programme that treats informational or operational structure as constraining physical structure, and on the foundational questions of distinguishability and operational primacy of records, it shares commitments with the authors named.

The distinguishing claim is specific and technical, and was defended structurally in §8.3: in VERSF, commitment is not epistemic but *dynamical*. It sources the  $\kappa$ -field. This is the move the competing programmes do not make. Zurek's decoherence programme tracks which pointer states survive environmental monitoring; the environment plays a role, but irreversibility is bookkeeping rather than a source term in the dynamics. Hardy's reconstruction derives Hilbert-space structure from operational axioms without giving any operational process a dynamical role. Constructor theory classifies which transformations are possible without specifying a dynamics on the substrate itself. None of these programmes makes commitment an element of the action.

The argument of §8.3 is that the source role is not optional: without it, irreversibility becomes labelling, records have no dynamical cost, and no action principle can be written. The neighbouring programmes may choose to stop short of this move, but they pay specific costs for doing so — costs that Part II's technical results (§§13.1, 13.2, 13.3) would not be able to absorb. The honest framing is therefore that VERSF is a member of a family of substrate-constraining programmes, distinguished by its commitment-as-source mechanism and by the numerical and structural predictions that mechanism generates.

### 9.3 Question 3: Observations do not need "facts" — they are measurement outcomes

**The objection.** Standard measurement theory, decoherence, and the einselection programme already account for the emergence of stable outcomes. The "fact" terminology adds nothing operational.

**Response.** The objection concedes the substantive point. The moment one says "measurement outcome," one has presupposed the three components identified in §3: differentiation, identifiability, and structural stability. These are precisely what is meant by "fact" in the VERSF sense. The terminology is not new content; it is a name for the assumptions already in use.

VERSF goes beyond standard measurement theory in treating these requirements as restrictive constraints on the substrate (§3) rather than as background conditions. The decoherence programme tells us how classical-looking outcomes emerge given a quantum substrate with appropriate environmental coupling. VERSF asks what the substrate must be for that emergence to be possible at all. These are different questions at different levels, and answering one does not foreclose the other.

### 9.4 Question 4: Granting your constraints, why does that uniquely lead to VERSF?

**The objection.** Many possible structures could in principle support distinguishable, persistent, irreversible records. Asserting that one's preferred framework is the unique solution is a familiar overreach in foundational work.

**Response.** The objection is correct that uniqueness claims must be earned rather than asserted, and the defence must be placed at the right level.

At the level of the constraint-first perspective, uniqueness is not claimed. Part I explicitly acknowledges that the perspective is compatible with many programmes [1–11], and §8.4 states directly that VERSF is one construction, not *the* construction. At the level of the specific axiomatisation, each of the three physical axioms is defended in §8.2 as corresponding to a physics requirement whose removal breaks a specific operational feature; mutual irreducibility within the axiom set is argued at the end of §8.2; the axioms are shown to substantially constrain the space of admissible substrates before any technical machinery is applied. At the level of the specific construction — given the axioms — the *No-Go Theorem for Non-Simplicial Relational Substrates* derives that the substrate must support a triangular 2-complex geometry with  $K=7$  minimal facts [Taylor, VERSF programme: No-Go Theorem]. The  $PGL(3,2)$  irreducibility argument on  $V_6$  is what forces this conclusion within the axiom set; "minimality" and "stability" alone would not.

The scope of the claim, by tier:

- *Proven* (within the axiom set): given A1–A6 and the three physical axioms, the substrate is forced into the  $K=7$  simplicial structure. Non-simplicial relational substrates are excluded *within this axiom set*.
- *Conditional*: the axioms themselves carry physical commitments that can be independently challenged. If the axioms are accepted, the structural conclusion follows.
- *Conjectural*: that this structural skeleton, together with the BCB Lagrangian framework, is sufficient to recover the full content of known physics. A full uniqueness analysis — whether the constrained region of substrate space identified in §8.2 admits occupants other than VERSF — is reserved for a separate paper.

The correct response to a uniqueness challenge is therefore to point at the no-go theorem and invite scrutiny of its axioms, using the axiom-by-axiom requirements laid out in §8.2. If an axiom can be weakened or replaced while preserving the operational requirements of physics identified in Part I, the theorem will indicate how the conclusions change. This is an open invitation.

### 9.5 Question 5: Do the admissibility constraints merely exclude pathological cases, or do they substantially collapse the space of viable substrate theories?

**The objection.** The constraints identified in Part I — finite distinguishability, irreversible commitment, finite local capacity, and causal activity — may exclude certain pathological or unrealistic mathematical constructions. It does not follow that they significantly restrict the space of viable physical substrates. A large class of alternative structures might still remain admissible, in which case the constraint-first perspective would not meaningfully narrow the space of candidate theories, and the rhetoric of a "constrained region" in §8.2 would be doing more work than the argument can support.

**Response.** The constraints do not act in isolation; they act cumulatively. Their combined effect is not to exclude a few edge cases but to progressively eliminate entire classes of otherwise conceivable substrates. This process may be described as an *admissibility collapse*: the systematic reduction of the candidate substrate space under jointly imposed physical requirements.

The collapse proceeds by elimination of the distinct failure modes identified in §3. Finite distinguishability excludes substrates with unbounded local state density, such as strict continua admitting arbitrary precision; in such systems, formal distinctions need not correspond to operationally distinguishable outcomes, violating identifiability. Irreversible commitment excludes purely reversible dynamical systems *as complete substrate descriptions*; a substrate that lacks any mechanism for irreversible record formation at any level cannot support operational time or stable measurement outcomes. The exclusion is directed at the substrate as a whole, not at reversible dynamics as such: unitarily evolving quantum mechanics, Hamiltonian classical mechanics, and other reversible formalisms may appear as effective or sectoral descriptions within a larger framework that includes a commitment mechanism, and VERSF explicitly allows this — decoherence and einselection are standard examples of how classical-looking records emerge from unitarily-evolving sectors once the commitment structure is in place. What the axiom excludes is a substrate in which *no* commitment mechanism exists anywhere, not a

substrate that contains reversible sub-dynamics. Finite local capacity excludes substrates in which arbitrarily many records can accumulate in a finite region; such systems lack a meaningful cost structure for record formation and fail to preserve locality of physical influence. Causal activity — the commitment-as-source move defended in §8.3 — excludes substrates in which records do not enter the equations of motion; by the criterion stated there, a feature that does not affect the substrate's future evolution does not exist physically, and records that are dynamically inert cannot satisfy the requirements of the constraint-first perspective.

Each constraint removes not just a specific model, but a class of models. Taken together, they restrict the admissible substrate space to structures capable of supporting finite, local, irreversible, causally active record formation. The resulting region appears to be highly constrained rather than broad — a point that a full uniqueness analysis will need to quantify, and which is reserved for a future paper.

Within this constrained region, the technical machinery of VERSF applies. In particular, the no-go theorem for non-simplicial relational substrates operates on a space already narrowed by the admissibility constraints. As noted in §8.2, the theorem's force depends on this prior restriction: without finite distinguishability, irreversibility, and finite capacity, candidate substrates outside the admissible region would remain, and the theorem would not reach the  $K=7$  conclusion. The admissibility constraints are therefore not independent of the structural results; they are what make those results non-vacuous.

The appropriate conclusion is not that admissibility collapse uniquely selects a single theory, but that it substantially limits the space of viable alternatives. VERSF occupies a specific position within this constrained region, characterised by its axiom set, its commitment-as-source dynamics, and its derived structural and numerical results. Whether the admissible region admits more than one such construction remains open, and is the subject of the uniqueness analysis reserved for a future paper (§9.4). What the collapse establishes is that the space of viable substrate theories is not generic, but tightly structured by the operational requirements of physics itself.

## 9.6 Question 6: Do facts require observers, or do they exist independently?

**The objection.** In many interpretations of physics, particularly in discussions of measurement, facts are often associated with observation. This raises the concern that "facts" may be observer-dependent — requiring an observer to bring them into existence. If so, the use of facts as a foundational element risks reintroducing observer-centric assumptions that modern physics seeks to avoid, and inheriting the interpretational difficulties that come with them.

**Response.** Facts do not require observers. They require commitment.

A fact, in the present framework, is not defined by being observed but by being physically instantiated as a stable, distinguishable record through an irreversible process. This process — termed commitment — is the elementary transition by which a configuration of the substrate becomes fixed in a form that can be referred back to and can influence subsequent evolution. Commitment is a physical process, not an epistemic relation.

Observation is one way in which commitment can occur, but it is not the only one. Physical systems routinely produce committed records without observers: a photon scattering off an atom, a decay event recorded in surrounding matter, a macroscopic interaction leaving a persistent trace. In each case, a commitment event occurs whether or not an observer is present. The universe has been committing records since long before there were observers to notice them.

The role of commitment is therefore fundamental, while the role of the observer is secondary. An observer does not create facts; an observer becomes correlated with facts through additional commitment events. The framework is thus not observer-dependent but *interaction-dependent*: facts arise wherever irreversible, record-forming interactions occur. This distinction is essential — if facts required observers, the framework would inherit the ambiguities and interpretational difficulties associated with observer-centric formulations of quantum theory. By grounding facts in commitment instead, the framework locates them entirely within the physical dynamics of the substrate.

The requirement of commitment also prevents facts from being reduced to mere descriptive labels. A configuration counts as a fact only if it has undergone a physical process that makes it stable, distinguishable, and causally active — that is, only if it enters the equations of motion and affects the future evolution of the system, in the sense made precise in §8.3. Without commitment, there are configurations; with commitment, there are facts.

Facts in VERSF are therefore:

- *observer-independent* — they do not require observation;
- *locally instantiated* — they occur at specific spacetime locations (§5);
- *dynamically grounded* — they arise through commitment and influence subsequent evolution (§8.3).

This resolves the ambiguity. Facts are not brought into existence by observers, but neither are they passive descriptions of states. They are committed physical records, formed through irreversible processes and embedded in the dynamics of the substrate. The question "do facts require observers?" admits a clean answer: no — they require interactions of a specific kind, and observation is one special case of such interaction.

### **9.7 Question 7: Do facts require mathematically exact distinctions, or only operational distinguishability?**

**The objection.** Physical theories are often formulated on mathematically precise state spaces, including continua with arbitrarily fine distinctions. This suggests that facts may depend on exact mathematical differences between configurations. If so, the requirement of finite distinguishability (axiom A of §8.2) would appear too restrictive, excluding otherwise standard formulations. A system whose state differs by an infinitesimal amount from another is mathematically distinct; why should such distinctions not count as physically real facts?

**Response.** Facts do not require mathematically exact distinctions; they require operational distinguishability.

A fact, in the present framework, is a committed physical record that can be distinguished, in principle, by some finite physical process. The relevant criterion is therefore not whether two configurations are mathematically distinct, but whether they can be distinguished by a finite, physically realisable interaction that produces a stable record. This immediately limits the role of arbitrarily fine distinctions. A substrate that admits differences smaller than any physically realisable measurement cannot convert those differences into distinct records. Such differences exist at the level of mathematical description; they do not correspond to distinct facts.

A simple example illustrates the point. An object moving at an infinitesimally small velocity is mathematically distinct from a stationary object. If no finite measurement can resolve that velocity — because the displacement over any accessible timescale falls below the threshold required for record formation — then no commitment event distinguishes the two cases. In operational terms, the system produces the same record: it is treated as stationary. The distinction exists in the formalism but not in the space of facts.

This is not a limitation of measurement in practice; it is a structural feature of what it means for a fact to exist. Facts arise only where differences can be stabilised into records through irreversible processes. Differences that cannot be so stabilised do not enter the set of physically realised facts, regardless of their mathematical status.

The requirement of finite distinguishability follows from this principle. It does not deny the usefulness of continuous mathematical descriptions; rather, it asserts that only a finite subset of distinctions within those descriptions can be physically instantiated as facts. The constraint-first perspective therefore distinguishes between:

- the *descriptive space* of mathematically possible configurations, and
- the *admissible space* of physically realisable facts.

Physical theory may employ the former, but it is constrained by the latter. Facts live in the admissible space, not the full mathematical space. Infinitesimal distinctions can exist in the formalism without implying an infinite density of facts. What counts physically is not the existence of a difference in the equations, but the existence of a commitment process capable of recording that difference.

This connects the finite-distinguishability axiom to the admissibility collapse of §9.5 in a direct way. If facts require commitment and operational distinguishability together, then not every mathematically distinct state corresponds to a distinct physical reality — only those whose differences can be stabilised into records. Substrates admitting unbounded precision without a corresponding commitment mechanism fail operationally even if they are mathematically well-defined, and are excluded from the admissible region for the same reason that non-committing substrates are. The continuum is a useful descriptive tool; it is not, by itself, a generator of facts.

## 10. Status and Demarcation

### 10.1 Question 8: This is metaphysics, not physics

**The objection.** A framework that begins from "what must be true for observations to exist" is doing transcendental philosophy in the Kantian tradition. That can be valuable, but it is not physics.

**Response.** The distinction between physics and metaphysics is empirical: a claim is physical if it makes contact with measurement. VERSF crosses that line in several specific places. The coupled-temporal protocol predicts  $\sigma_{\tau}/\sigma_{\text{opt}} = \sqrt{(2\ln 2)}$  as a parameter-free relationship; a measurement deviating outside error bounds would falsify the structural derivation. The tensor-to-scalar prediction  $r \approx 0.008$  lies within current BICEP/Keck bounds [12], and the BICEP Array is positioned as the discriminating instrument. The  $\kappa$ -field mass derivation  $m = \sqrt{(4/3)} \cdot \xi^{-1}$  ties a particle-physics quantity to the closure scale. The ketosteroid isomerase prediction of an approximately  $9.3\times$  commitment lag ratio is testable in existing enzyme-tunnelling experimental traditions.

The ranked treatment of these falsifiers — which is primary, which follows, which cuts deepest — appears in §14.2. The point for the present objection is narrower: the framework starts from structural questions about the substrate, as Part I argues it should, but it does not stop there. It produces numerical predictions in cosmology, particle physics, and biochemistry that can be checked. That is the demarcation line.

## 10.2 Question 9: You are redefining time and entropy non-standardly

**The objection.** Time as "accumulated commitment events" is not the standard physical conception. Why should this redefinition be accepted?

**Response.** The framing as "redefinition" is misleading. The operational content of time in physics already reduces to counting irreversible transitions. This is not metaphor. An atomic clock produces a frequency standard by counting transitions between hyperfine levels of caesium-133; the count is physically meaningful only because each transition leaves a stable, distinguishable, recorded trace. A reversible dynamics, considered in isolation, contains no preferred direction and no operational time; it cannot realise a clock. The arrow of time, in every working clock, comes from irreversible commitment somewhere in the chain — which is precisely the operational feature that axiom B of §8.2 is designed to protect.

This view has established company. Rovelli's relational quantum mechanics [6] and the Connes–Rovelli thermal time hypothesis [7] share the orientation, as does the broader tradition treating time as emergent from irreversible processes. VERSF makes the connection precise by identifying the irreversible processes — commitment events — and giving them a dynamical role through the  $\kappa$ -field. The correct framing is therefore not that VERSF redefines time, but that it makes explicit and dynamical a connection between time and irreversibility that is implicit in operational physics and explicit in the Rovelli and Connes–Rovelli traditions.

## 11. Calculational Content

**Question 10: Even if the framework is consistent, does it calculate anything?**

**The objection.** Frameworks live or die by their numerical and structural predictions. Foundational reframing without calculational output is of limited interest.

**Response.** This is the right question, and it has concrete answers. We list each current prediction together with the derivation it depends on and the locus to which its failure would localise. The firewall structure matters: a prediction whose failure localises narrowly does not propagate to the rest of the programme, while a prediction whose failure localises to the commitment dynamics cuts to the core.

- **Fine-structure constant.**  $\alpha^{-1} \approx 137.034$ , derived from interface geometry [Taylor, VERSF programme: fine-structure constant derivation]. *Depends on:* the specific interface-geometry construction linking the substrate structure to the electromagnetic coupling. *Failure localises to:* the interface-geometry derivation, not the substrate axioms or the Lagrangian.
- **Tensor-to-scalar ratio.**  $r \approx 0.008$ , within current BICEP/Keck bounds. *Depends on:* the cosmological-sector derivation, which couples the  $\kappa$ -field dynamics to standard inflationary machinery. *Failure localises to:* the inflationary coupling or the cosmological-sector derivation, not the foundational results.
- **$\kappa$ -field mass.**  $m = \sqrt{(4/3)} \cdot \xi^{-1}$ , deriving the mass coefficient from PGL(3,2) irreducibility on  $V_6$  rather than fitting it. *Depends on:* the irreducibility argument and the closure-scale derivation. *Failure localises to:* the mass-coefficient calculation if  $\xi$  is independently validated; to the closure-scale derivation itself if  $\xi$  is not. Independent constraints on  $\xi$  discriminate these two cases.
- **Commitment barrier coefficient.**  $C^* = 3/8$ , derived as a structural quantity rather than a free parameter. *Depends on:* the commitment-barrier derivation. *Failure localises to:* the barrier calculation.
- **Cosmological constant.** Approximately 20% agreement via the Two-Planck Principle derivation. *Depends on:* the Two-Planck construction and the specific matching between substrate-scale and large-scale contributions. *Failure localises to:* the matching procedure; this is the weakest of the numerical claims in the programme and should be treated as such. Where standard QFT estimates miss by about 120 orders of magnitude, a framework achieving order-of-magnitude agreement is suggestive rather than confirmatory, and the result should be offered in that register.
- **Coupled-temporal experimental signature.**  $\sigma_\tau / \sigma_{\text{opt}} = \sqrt{2 \ln 2}$ , parameter-free. *Depends on:* the commitment dynamics directly, not on any sectoral application. *Failure localises to:* the commitment dynamics itself. This is the deepest-cutting prediction in the programme — a failure here propagates upward, not downward.
- **Enzyme catalysis.** An approximately  $9.3 \times$  commitment lag ratio for ketosteroid isomerase tunnelling. *Depends on:* the application of VERSF commitment dynamics to enzymatic tunnelling, which requires additional biochemical modelling assumptions beyond the core framework. *Failure localises to:* the biochemical application, not the foundational machinery.

The firewall structure is therefore asymmetric. Several predictions — fine-structure, tensor-to-scalar, barrier coefficient, cosmological constant, enzyme lag — fail locally if they fail. One — the coupled-temporal ratio — fails globally if it fails. A referee who wishes to find the deepest

test of the framework should look there, and a referee who wishes to find the most vulnerable prediction should look at the cosmological constant. Both are legitimate directions of scrutiny.

## 12. Relativity and Frame-Dependence

### 12.1 Question 11: Modern physics has dispensed with absolute facts

**The objection.** Observers in different inertial frames disagree on the simultaneity and ordering of spacelike-separated events. In quantum scenarios — Wigner's friend, Frauchiger–Renner [13] — different observers can apparently assign different outcomes to the same measurement. Facts are frame-dependent, observer-dependent, and not globally consistent. A framework that treats "facts" as a unified primitive imposing global constraints appears to assume a structure that physics has already abandoned.

**Response.** Part I already anticipated this concern: distinction need not be globally defined, but it must exist locally for physics to be meaningful (§5). The constraint-first perspective does not commit to global or absolute facts; it commits to the local, invariant existence of records. VERSF inherits this commitment without alteration.

Relativity tells us there is no global ordering of events and no observer-independent slicing of spacetime. It does not tell us that records fail to exist locally. Whenever a fact exists for any observer, it must be instantiated as a local physical record — a configuration of matter, a stored bit, a pointer position. The substrate must be capable of supporting such records locally, regardless of how observers in different frames reconcile them.

VERSF is natively relational on this point. The closure scale  $\xi$  and the structural constants of the framework are coordinate-invariant. The records themselves are local. The Single-Source Theorem expresses observables as functionals of the committed record density  $\rho(x,t)$ , where the density is the locally invariant object and observers' descriptions are the frame-dependent slices through it.

On Wigner's friend and Frauchiger–Renner specifically: VERSF treats commitment as a physical event, not an epistemic relation — a commitment event either occurs at a given location and time or does not. Observers' descriptions of when an outcome was "established" can differ without requiring a globally synchronised fact-set, and without weakening the structural constraints on the substrate that supports the local records. The framework's relational character is a feature here, not a cost.

### 12.2 Question 12: Does the relativity of observation undermine the notion of facts?

**The objection.** In relativistic settings, observers do not have access to events at the time they occur. Light from a star millions of light-years away, for instance, is only observed long after the emission event. This suggests that what the framework calls a "fact" may be observer-dependent or delayed, rather than a well-defined physical occurrence. If facts are not globally accessible or

synchronised, it is unclear how they can serve as a foundational element in a physical framework.

**Response.** The objection rests on a conflation between the *occurrence* of a fact and the *reception* of information about that fact. In VERSF these are distinct physical processes, and the distinction dissolves the apparent tension.

A fact occurrence is a local commitment event: a physical transition that produces a stable, distinguishable record at a specific location in spacetime. The emission of a photon from a distant star constitutes a fact at the location and time of emission, and contributes to the committed record density  $\rho(x,t)$  there. A fact reception is a *separate* commitment event: the local detection of that photon at another location, producing a new record which encodes information about the earlier event but is not identical to it. The two events are related by causal propagation, but they are not the same fact. Information does not transport facts; it generates new ones. The observation of a distant event is therefore not delayed access to a single global fact, but the formation of a new local fact correlated with a past one through the intervening light cone.

This distinction is fully consistent with relativistic physics and requires no additional machinery beyond what §5 already establishes. Relativity eliminates global simultaneity but preserves local causal structure. The constraint-first perspective requires only that facts exist as local, causally connected commitment events; it does not require global synchronisation or observer-independent simultaneity. Different observers may form different local records at different times, but each record is a valid physical fact at its own point of occurrence, and each enters  $\rho(x,t)$  at that location.

The apparent tension dissolves once the distinction is made explicit. Facts are not global entities distributed across spacetime; they are local events connected by causal structure. The framework does not treat observation as revealing pre-existing global facts, but as producing new local facts through interaction with propagating signals. This preserves both the operational role of facts in physics and the relativistic structure of spacetime, and it extends naturally to the stronger frame-dependence concerns addressed in §12.1: Wigner's friend and the starlight case are species of the same genus, and the same distinction resolves both.

## 13. Dynamical and Structural Machinery

### 13.1 Question 13: What is your action principle?

**The objection.** A physical theory worth taking seriously needs an action principle from which dynamics follow. Without a Lagrangian, claims about the framework's predictive content are unanchored — one cannot tell whether the predictions are derived or fitted, whether they are consistent with one another, or whether the theory has the right symmetry structure.

**Response.** The BCB Lagrangian unification paper provides the action structure [Taylor, VERSF programme: BCB Lagrangian]. The Lagrangian is constructed on the substrate such that commitment events source the  $\kappa$ -field (the move defended in §8.3); the action yields the  $\kappa$ -field

equations of motion, the discrete Noether conservation laws discussed in §13.3, and the appropriate continuum limit at scales well separated from the closure scale  $\xi$ . It is not a sketch.

By tier:

- *Proven* within the framework: the Lagrangian is well-defined on the substrate, generates the  $\kappa$ -field dynamics, and is consistent with the discrete Noether structure.
- *Conditional*: that this Lagrangian, together with the matter-sector coupling developed in the Standard-Model-from-hexagonal-geometry papers and the BCB/lepton mass model, is sufficient to recover the dynamical content of the Standard Model in the appropriate limit.
- *Conjectural*: full unification of the gauge-sector dynamics with the substrate Lagrangian beyond what is currently derived.

### 13.2 Question 14: How do you recover General Relativity and the Standard Model?

**The objection.** Any substrate-level framework must reduce to known physics in the appropriate limits: General Relativity in the classical, large-scale regime; the Standard Model in the high-energy regime. A framework that cannot demonstrate these reductions is not yet a candidate replacement for the existing structure, however elegant its foundational claims.

**Response.** VERSF is an emergence programme, not a reduction programme. Standard physics is recovered as an effective description in the appropriate decoupling regimes (see §14.1), not by rewriting it. The architectural question is not *what does VERSF reduce to* but *what emerges from the substrate in the appropriate limits, and is it what we observe?*

The honest position separates three recoveries at different stages of completion. *Lorentz invariance* is the most secure: the proto-time and emergent Lorentz invariance paper derives Lorentz symmetry as an emergent property of substrate dynamics rather than imposing it as a primitive [Taylor, VERSF programme: emergent Lorentz invariance]. *Standard Model content* is a partial recovery: the Standard Model from hexagonal geometry papers and the BCB/lepton mass model derive structural relationships and mass content from the substrate geometry, with the fine-structure constant derivation and the  $\kappa$ -field mass derivation as concrete pieces; the full gauge structure and complete mass spectrum are not yet derived. *General Relativity* is the most open: the closure scale  $\xi$  and the  $\kappa$ -field provide the structural ingredients for emergence of metric structure, but full derivation of the Einstein field equations as the appropriate-limit dynamics remains a conjectural target. The cosmological constant work via the Two-Planck Principle, with its approximately 20% agreement, is the area where contact with GR-scale physics is most concrete at present — and as noted in §11, it is the most vulnerable of the numerical claims.

Claiming more than this would invite dismissal of the parts that are solid alongside the parts that are not. The emergence framing matters: it sets the correct expectation that the programme recovers standard physics as an effective description, and it protects the derived results from being held to a reductionist standard they do not claim to meet.

### 13.3 Question 15: What are your conservation laws?

**The objection.** Conservation laws are the spine of physics. A framework that cannot specify, derive, or motivate its conservation structure is not yet a physical theory. Imposed conservation is unsatisfying; derived conservation is the standard set by Noether.

**Response.** The discrete Noether theorem and emergent time paper provides the formal basis [Taylor, VERSF programme: discrete Noether theorem]. Substrate symmetries — including the discrete time-translation symmetry that emerges from the ordering structure of commitment events — yield conservation laws via a discrete analog of Noether's theorem. The Bit Conservation and Balance (BCB) framework specifies these requirements at the substrate level, and the Ticks-Per-Bit (TPB) construction relates them to operational time.

The conservation structure derived within the framework includes bit conservation at the substrate level via BCB, which is the analog of unitarity in standard quantum mechanics and relates directly to the no-cloning and no-deletion theorems; energy-momentum analogs via discrete time- and space-translation symmetries on the substrate, recovered in the continuum limit as the standard conservation laws; and fact-momentum conservation, tied to the  $\kappa$ -field source structure and providing the link between commitment events and the dynamical role of records. The key claim — and it is a derivation rather than an imposition — is that conservation in VERSF emerges from substrate symmetries. The discrete Noether paper is among the submission-ready pieces of the programme and is targeted for *Foundations of Physics*.

Note that the derivation depends on the source role defended in §8.3. Without commitment-as-source, the symmetries are of an external ordering rather than of the substrate, and the Noether machinery does not engage. This is another concrete instance of the argument that commitment-as-source is forced rather than chosen.

## 14. Empirical Regime and Falsifiability

### 14.1 Question 16: Why has this not been seen?

**The objection.** Standard physics works extremely well across many orders of magnitude in energy, length, and time scale. Precision tests of QED agree with theory to better than ten significant figures. General Relativity passes Solar System tests, gravitational wave observations, and pulsar timing without strain. Any new substrate-level framework must explain why these measurements do not already see the deviations it predicts. A decoupling story that requires post-hoc parameter tuning to avoid existing constraints is not credible.

**Response.** The closure scale  $\xi$  provides the structural answer, and the answer is non-circular precisely because  $\xi$  is not a free parameter. VERSF effects are significant at scales comparable to  $\xi$ ; at scales well separated from  $\xi$  in either direction, the framework reduces to standard physics with corrections suppressed by the appropriate ratio.

At scales much larger than  $\xi$  (the cosmological regime), classical and effective-field-theory descriptions emerge with VERSF corrections suppressed. Where the framework predicts

deviations — the cosmological constant magnitude, the tensor-to-scalar ratio — they appear precisely in the regime where  $\xi$ -scale physics couples to large-scale observables.

At scales much smaller than  $\xi$  (the sub-Planckian regime), standard QFT calculations proceed with VERSF corrections suppressed.

At scales comparable to  $\xi$ , VERSF predictions diverge from standard expectations. This is where the  $\kappa$ -field mass  $m = \sqrt{(4/3)} \cdot \xi^{-1}$  lives, and where the parameter-free experimental signatures — the coupled-temporal protocol, the enzyme commitment lag — are designed to probe.

The non-circularity is essential. If  $\xi$  were a free parameter that could be retuned to evade any inconvenient measurement, the decoupling story would be empty. The closure scale paper establishes  $\xi$  as a structural consequence of the three axioms defended in §8.2 — not adjustable, not fitted [Taylor, VERSF programme: closure scale]. The decoupling regime is therefore a prediction of the framework, not a defensive manoeuvre. If existing high-precision measurements begin to show deviations of the predicted size in the predicted regime, that would confirm the framework. If they show deviations of the wrong size or in the wrong regime, the framework faces serious pressure.

## 14.2 Question 17: What single experimental result would falsify VERSF?

**The objection.** A theory's seriousness is partly measured by the sharpness with which its proponents can specify what would defeat it. Vague falsifiability is no falsifiability. A framework offering many qualitatively distinct predictions without ranking them, or without committing in advance to treating null results as falsifications, is unfalsifiable in practice.

**Response.** The primary falsifier is a cosmological measurement:

A sustained measurement of the cosmological tensor-to-scalar ratio above  $r \approx 0.05$  by the BICEP Array would falsify the current cosmological derivation, which predicts  $r \approx 0.008$ .

The BICEP Array is the discriminating instrument and the measurement is approaching the relevant sensitivity.

Secondary falsifiers, in order of discriminating power — using the firewall analysis of §11:

1. **Coupled-temporal protocol.**  $\sigma_\tau/\sigma_{\text{opt}} = \sqrt{(2\ln 2)}$ , parameter-free. A measurement showing the ratio deviating outside experimental error would falsify a structural derivation that cuts deeper than the primary, because  $\sqrt{(2\ln 2)}$  follows directly from the commitment dynamics rather than from a sectoral application. A failure here propagates upward to the commitment dynamics itself.
2.  **$\kappa$ -field mass relationship.**  $m = \sqrt{(4/3)} \cdot \xi^{-1}$ . Given an independently constrained closure scale  $\xi$ , the numerical prediction is violable; the firewall analysis of §11 shows how the localisation works.

3. **Ketosteroid isomerase commitment lag.** Approximately  $9.3\times$ , testable in the existing enzyme-tunnelling experimental tradition. Failure here localises to the biochemical application, which is narrower than the other falsifiers.

The framework therefore offers multiple parameter-free or structurally-constrained predictions, each independently falsifiable and each with a known localisation pattern on failure. The commitment to treat null results as falsifications rather than as opportunities for parameter retuning is enforced structurally by the closure scale paper, which establishes  $\xi$  as derived rather than free: there is no parameter available to adjust.

## 15. Relation to Neighbouring Programmes

### Question 18: How does this relate to loop quantum gravity, causal set theory, and causal dynamical triangulations?

**The objection.** Loop quantum gravity [10], causal set theory [9], and causal dynamical triangulations [11] are the nearest structural neighbours to VERSF — all discrete, all relational, all attempting emergence of spacetime from substrate-level structure. Causal set theory in particular shares the commitment to discrete, irreversibly-ordered events as the primitive. A framework that does not clearly distinguish itself from this established family risks being seen as a repackaging.

**Response.** The lineage deserves explicit acknowledgment, and the distinguishing features should be stated cleanly.

Causal set theory treats discrete events ordered by causal relations as the primitive and derives spacetime geometry from the statistics of the causal structure. Loop quantum gravity quantises spatial geometry via spin networks and evolves via spin foams. Causal dynamical triangulations regularises the path integral over geometries by restricting to causally well-behaved triangulations. All three share with VERSF the commitment to a discrete, relational, substrate-level description, and on this foundational commitment the programmes are cousins.

What distinguishes VERSF is the dynamical role of commitment, defended structurally in §8.3. In causal set theory, events are ordered but do not source a field; the programme is geometric and kinematic. In VERSF, commitment events source the  $\kappa$ -field — they have a dynamical role that feeds back into the substrate geometry via the BCB Lagrangian. That is the structural difference, and it is also what generates the specific numerical predictions ( $\alpha^{-1}$ ,  $r$ ,  $C^*$ , the  $\kappa$ -field mass) that the geometric-only programmes do not produce in the same form. VERSF and causal set theory can be seen as sharing foundational axioms on distinguishability and discrete ordering, and diverging at the point where VERSF introduces commitment-as-source.

---

## 16. Summary

The paper has argued two logically independent theses.

The first is general: the existence of distinguishable physical states is a *restrictive* foundational assumption. Not all mathematical substrates can support persistent distinguishable records, and those that cannot are excluded as candidates for physical reality. The assumption therefore partitions the space of candidate substrates, and the shape of the partition imposes structural content that dynamics alone does not see. This is the constraint-first perspective of Part I. It does not depend on any specific programme; it depends only on the observation that predictive physics presupposes what it does not derive, and that the presupposition has failure modes that exclude some substrates from supporting observation at all.

The second is specific. The Void Energy-Regulated Space Framework is one substantive pursuit of the constraint-first perspective, characterised by three physical axioms — finite distinguishability, irreversible commitment, finite local capacity — each of which is defended as a physics requirement rather than a mathematical stipulation, and by the dynamical move of treating commitment events as a source for the  $\kappa$ -field. Together with the causal-activity requirement on records, these constraints produce an *admissibility collapse*: the cumulative elimination of entire classes of candidate substrates, not merely pathological edge cases, leaving a substantially constrained region of substrate space within which VERSF occupies a specific, so-far-unoccupied position. The source-role move is defended as forced rather than chosen, on the grounds that without it irreversibility becomes mere labelling, records have no dynamical cost, and no action principle can be written. The framework yields foundational technical results — a no-go theorem fixing the substrate geometry within the axiom set (whose force depends, at each step, on the admissibility collapse), a discrete Noether theorem on conservation, a Single-Source Theorem on observables — and numerical predictions in cosmology, particle physics, and biochemistry that are independently falsifiable and that carry explicit firewall structure: some predictions fail locally, while the coupled-temporal signature fails globally if it fails. The closure scale  $\xi$ , established as structural rather than free, is what ties the foundational results to the empirical predictions and what makes the falsifiability non-circular. The primary falsifier is a BICEP Array measurement of the tensor-to-scalar ratio; the deepest-cutting falsifier is the coupled-temporal signature.

The programme is not complete. Full recovery of General Relativity remains conjectural; the Standard Model recovery is partial; several component papers remain under revision; the cosmological constant result is the weakest of the numerical claims. These gaps are stated because they are real, and because the epistemic register of the programme — *proven, conditional, conjectural* — is one of its architectural commitments. Engagement that pressure-tests the axioms, the derivations, and the numerical predictions is welcomed.

The compact version is this. Physics explains the behaviour of distinctions but rarely asks what makes distinction possible at all. Asking the question, and identifying what fails when distinguishability is not supported, is not a retreat into metaphysics; it is a source of structural constraint that, properly exploited, yields testable predictions. VERSF is one attempt to exploit it — a specific axiomatisation, a specific dynamical move, a specific derived structure, each offered alongside an account of what would break if the axiom or the move or the derivation

were removed. Part I makes the general case for the approach; Part II defends the specific attempt. The two stand or fall independently.

---

## References

- [1] Wheeler, J. A. (1990). Information, physics, quantum: The search for links. In *Complexity, Entropy, and the Physics of Information*, W. H. Zurek, ed., Addison-Wesley.
- [2] Zurek, W. H. (2003). Decoherence, einselection, and the quantum origins of the classical. *Reviews of Modern Physics* 75, 715.
- [3] Deutsch, D. and Marletto, C. (2015). Constructor theory of information. *Proceedings of the Royal Society A* 471.
- [4] Hardy, L. (2001). Quantum theory from five reasonable axioms. arXiv:quant-ph/0101012.
- [5] Masanes, L. and Müller, M. P. (2011). A derivation of quantum theory from physical requirements. *New Journal of Physics* 13, 063001.
- [6] Rovelli, C. (1996). Relational quantum mechanics. *International Journal of Theoretical Physics* 35, 1637.
- [7] Connes, A. and Rovelli, C. (1994). Von Neumann algebra automorphisms and time-thermodynamics relation in generally covariant quantum theories. *Classical and Quantum Gravity* 11, 2899.
- [8] Spekkens, R. W. (2007). Evidence for the epistemic view of quantum states: A toy theory. *Physical Review A* 75, 032110.
- [9] Bombelli, L., Lee, J., Meyer, D., and Sorkin, R. D. (1987). Space-time as a causal set. *Physical Review Letters* 59, 521.
- [10] Rovelli, C. and Smolin, L. Loop quantum gravity; Ashtekar, A. New variables for classical and quantum gravity. Standard references.
- [11] Ambjørn, J., Jurkiewicz, J., and Loll, R. Causal dynamical triangulations programme. Standard references.
- [12] BICEP/Keck Collaboration. Constraints on primordial gravitational waves from B-mode polarisation measurements. Most recent constraint paper.
- [13] Frauchiger, D. and Renner, R. (2018). Quantum theory cannot consistently describe the use of itself. *Nature Communications* 9, 3711.

## VERSF programme references

The technical results cited in this paper appear in the VERSF programme corpus hosted at versf-eos.com, including but not limited to: the *No-Go Theorem for Non-Simplicial Relational Substrates*; the *Closure Scale  $\xi$*  paper; the *BCB Lagrangian Unification* paper; the *Discrete Noether Theorem and Emergent Time* paper; the *Proto-Time and Emergent Lorentz Invariance* paper; the *Single-Source Theorem*; the *Fine-Structure Constant Derivation*; the  *$\kappa$ -Field Mass Derivation*; the *Commitment Barrier Coefficient* paper; the *Two-Planck Principle* paper; the *Coupled-Temporal Protocol* paper; and the *Ketosteroid Isomerase Application* paper. Specific bibliographic details are provided on the VERSF programme website.