On the Unavoidability of Fact-Based Constraints in Physical Theory

A No-Escape Theorem for Empirical Physics

Keith Taylor VERSF Theoretical Physics Programme

For the general reader

Why is physics structured the way it is? One answer says the structure of physical law is *contingent* — it could have been otherwise, and we simply happen to live in one possible version. Another says it is *forced* — there is no other way physics could be.

This paper establishes a version of the second answer, and it does so without presupposing any particular physical theory. The argument begins from three conditions that any empirical science must satisfy: outcomes must be **recordable**, **repeatable**, and **comparable across observers**. Without these, there is no experiment, no replication, no peer review — no science.

We show that these three everyday conditions are enough to force the structural core of physical law. Any framework satisfying them must obey three specific constraints — on what finite observers can distinguish, on whether comparison outcomes can depend on how we happen to describe things, and on whether structure we can never detect counts as physics. These constraints in turn force the admissible structure of physics to be unique, up to differences that make no observational difference. Two theories that agree on every possible record are, for physics, the same theory.

The conclusion is that physical law is not a menu of possibilities from which ours has been selected by chance. It is the unique structure that makes experimental reality possible at all. Anyone who rejects this conclusion must also reject the possibility of experimental science — internally coherent, but only at the cost of rejecting the enterprise that produced the rejection. The deeper consequence is a relocation: the old question "why is physics structured this way rather than another?" becomes the new question "why is empirical science possible at all?" The mystery has not been dissolved. It has been moved to where it belongs.

Abstract

We establish a no-escape theorem for empirical physical theories. Starting from a minimal operational definition of physics — a framework capable of producing stable, inter-subjectively testable outcomes — we prove that any such framework must satisfy three structural constraints: finite operational distinguishability (**A1**), structure-independent comparison of facts (**A2**), and observational accessibility of structural content (**A3**). These are not modelling assumptions. They are necessary conditions for the very possibility of empirical content: any violation either eliminates determinacy, reproducibility, or inter-observer consistency, or commits the framework to distinctions that cannot enter empirical discourse — in either case rendering the framework non-physical in the operational sense.

We further prove an exhaustion result: any framework that superficially appears to evade **A1–A3** either (i) admits a representation satisfying them (possibly after quotienting redundant structure), or (ii) fails to generate observable content at all. No genuine third option exists.

Combining this with previously established results on constraint closure and structural fixing within the VERSF programme [1,2], we obtain the implication chain

Empirical Physics ⇒ {A1, A2, A3} ⇒ Constraint Closure ⇒ Structural Uniqueness.

The structural core of empirical physics is therefore uniquely fixed up to observational equivalence: all admissible physical theories share a single empirical structure, even if they differ in non-observable representation. The theorem localises all residual conditionality of physical law to a single locus: the operational definition of physics itself. To reject the conclusion, one must reject the possibility of empirical physical theory — a move that is internally coherent only at the cost of abandoning science. The result therefore relocates the question of physical contingency from the structure of law to the possibility of empirical science itself.

Contents

Main Result (No-Escape and Uniqueness Chain)

Any framework capable of producing stable, reproducible, and inter-observer-comparable outcomes must satisfy finite operational distinguishability (A1), structure-independent comparison of facts (A2), and observational accessibility of structural content (A3).

These constraints force closure of the admissible structure space, and hence structural uniqueness up to observational equivalence.

Consequently, the structural core of empirical physical theory is not freely chosen, but forced by the conditions required for empirical content to exist.

The first implication is Theorem 4.1 of this paper; the second combines Theorems 6.1–6.2 of the VERSF programme [1,2]; the third is the composite statement. Each arrow is a theorem, not a modelling choice. The only residual conditionality is the operational definition of empirical physics given in §2.

1. Introduction

1.1 The conditionality problem

A central open question in the foundations of physics is whether the structure of physical law is *contingent* — one of many equally admissible possibilities selected by initial conditions or anthropic fact — or *forced* — uniquely determined by deeper structural requirements. Standard practice addresses this question *from within* a chosen mathematical framework (Hilbert spaces, differentiable manifolds, gauge bundles), where the framework itself is treated as a free starting point.

The VERSF programme inverts this order. It asks a *prior* question:

What must be true of the world for empirical physics to be possible at all?

Earlier papers in the programme [1,2] showed that *if* a minimal set of constraints on fact formation is imposed, the admissible structure space collapses to a single equivalence class. What remained open — and what this paper settles — is whether those constraints are themselves unavoidable, or whether alternative frameworks could evade them.

1.2 Contribution

We prove:

1. **(No-Escape Theorem, §4).** Any framework that produces empirical physical content must satisfy **A1–A3**.
2. **(Reduction of Apparent Alternatives, §5).** Any framework that appears to violate **A1–A3** either reduces to one that satisfies them or fails to produce observable content. There is no third category.
3. **(Final Chain, §7).** Combined with prior closure and uniqueness results, this establishes that the structural core of physical law is not freely chosen but is forced by the conditions required for empirical reality to exist.

1.3 Status of the result

The theorem is *conditional on the operational definition of physics* given in §2. It is not a proof that the universe *must* be empirically describable; it is a proof that *any empirical description* must obey **A1–A3**. The point of the conditionality is that it is maximally minimal: the operational definition captures only what is required for science to function at all. Weaker premises exist only by abandoning empirical testability.

2. Operational Definition of Physical Theory

2.1 Definition

Definition 2.1 (Empirical Physical Theory). A framework $\mathcal{T}$ is an *empirical physical theory* if and only if it satisfies three operational conditions:

- **(D) Determinacy.** For each specified procedure, $\mathcal{T}$ yields a determinate outcome that can be recorded — i.e., instantiated in a stable, persistent form accessible to a finite observer at a later time.
- **(R) Reproducibility.** Determinate outcomes are reproducible under repetition of the specified procedure, up to stated statistical tolerances.
- **(C) Inter-observer Consistency.** Outcomes admit comparison across distinct observers and distinct experimental realisations, such that agreement or disagreement between observations is itself a determinate matter.

These are not philosophical preferences. They are the minimal conditions under which experimental science — measurement, replication, publication, peer scrutiny — is possible. Any framework failing any one of **(D), (R), (C)** falls outside the scope of what is ordinarily meant by an empirical theory.

Remark (these conditions are not an interpretation of physics). (D), (R), (C) do not define a particular interpretation, ontology, or metaphysics of physics. They define the minimal conditions under which experimental *disagreement*, *confirmation*, and *falsification* are meaningful operations at all. A framework that rejects them may remain mathematically consistent, but cannot function as an empirical theory in the operational sense: its predictions cannot be recorded, repeated, or compared, and there is no procedure by which the world could tell us it is wrong.

Demarcation, not ontology. Definition 2.1 is a demarcation criterion on empirical theory-forms, not a claim that empirical content exhausts ontology. A realist may posit additional non-empirical structure; the theorem does not preclude this. It applies the moment that structure is brought into empirical theory — through prediction, comparison, or falsification — at which point (D), (R), (C) constrain how it can be handled.

2.2 What is *not* assumed

Definition 2.1 does *not* presuppose:

- A continuum or discrete substrate,
- A particular logic, probability theory, or measurement formalism,
- Realism, operationalism, or any metaphysical stance beyond (D)–(C),
- The existence of unique truth-values independent of procedure.

In particular, the definition is compatible with frameworks in which outcomes are probabilistic, contextual, or observer-relative, *provided* the probabilistic or contextual structure itself satisfies

(D)–(C). A probability distribution that cannot be recorded, reproduced, or compared is not an empirical object. The requirement is not that all observer-indexed descriptions coincide at all intermediate stages, but that record-level comparison outcomes be determinate when comparison is physically performed. This accommodates Wigner's-friend-type scenarios and Frauchiger–Renner-type protocols, which concern the relationship between intermediate state assignments rather than the determinacy of final record comparisons.

2.3 Records as the carrier of empirical content

A recurring construct in what follows is the *record* — the finite, stable, distinguishable artefact (pointer position, bit-string, logbook entry, neural trace) that carries an outcome forward in time. The three conditions (D)–(C) can be summarised equivalently as: *records exist, records persist, and records can be compared.* This is the minimal operational content of "an experiment happened."

Observers are finite. Records are produced by observers — biological or engineered — whose memory, signal bandwidth, and discrimination threshold are themselves finite. This is not a philosophical claim about observation but a physical fact about any apparatus capable of producing a stable record: no real detector has unbounded memory, zero noise floor, or infinitely fine discrimination. Observer-finitude is a standing assumption throughout the remainder of the paper and is the substantive content of the word "finite observer" in the constraints below.

Records and probability. A single record carries a single outcome per trial. In frameworks with probabilistic outcomes, reproducibility (R) refers to reproducibility of the *distribution* over records under repeated independent trials — not reproducibility of any single record. The three conditions (D)–(C) therefore apply at two levels: to the records themselves in each individual trial, and to the distribution across trials. This accommodates quantum mechanics, classical stochastic theories, and any other framework in which single outcomes are not determined but distributions are.

Observer-finitude is not a free assumption. The finiteness of observers is not an independent premise but a physical constraint: any system capable of producing a stable record must have bounded storage, bounded energy, and bounded interaction time. An "infinite observer" may be a useful *mathematical* idealisation for formal derivations, but cannot function as an *empirical* observer unless its outcomes are realised in finite records; the theorem concerns the latter. Idealised observers are admissible in formal arguments about the theory, but empirical content must bottom out in finite record-formation, whether the observer is idealised elsewhere in the formalism or not.

---

3. The Minimal Fact Constraints

We now state the three constraints whose unavoidability is the subject of the theorem. Two subsidiary definitions anchor the terminology and close potential circularities that would otherwise leave "finite procedure" and "admissible representation" as load-bearing but undefined terms.

Definition 3.0 (Finite procedure). A procedure $\mathcal{P}$ is *finite* if it can be specified, prepared, and completed using finite observer resources (finite memory, bandwidth, preparation steps, and energy budget) in finite time. The set of finite procedures is closed under finite composition and repetition — if $\mathcal{P}_1$, $\mathcal{P}_2$ are finite, so is $\mathcal{P}_1 \circ \mathcal{P}_2$, with $|D(\mathcal{P}_1 \circ \mathcal{P}_2)| \leq |D(\mathcal{P}_1)| \cdot |D(\mathcal{P}_2)| < \infty$ — but it is *not* closed under unbounded limits or idealisations. This definition grounds A1 in physical finiteness rather than outcome-cardinality, and prevents the vacuous reading "$\mathcal{P}$ is finite iff it has finitely many outcomes" from trivialising the constraint.

Definition 3.1 (Admissible representation). A representation R of the facts of a framework $\mathcal{T}$ is *admissible* if it is a well-formed encoding: (i) every fact produced by $\mathcal{T}$ has a representative in R, and (ii) R's internal operations — including comparison — are well-defined on these representatives. Admissibility is a minimal well-formedness condition on encodings. It does *not* by itself entail that comparison outcomes agree across distinct admissible representations; that further coherence claim is the substantive content of A2. This decoupling ensures that A2 is not a tautology about the definition of admissibility, and that "admissible" is grounded in encoding-validity rather than in the constraint A2 itself.

With these definitions in place, the three constraints read as follows.

Constraint A1 — Finite Operational Distinguishability. *What a finite observer can resolve is finite.* More precisely: for any finite observer executing any finite procedure (Def. 3.0), only finitely many distinctions are operationally accessible. Formally, the set of distinguishable outcomes of a finite procedure $\mathcal{P}$ is a finite set $D(\mathcal{P})$, with $|D(\mathcal{P})| < \infty$.

Constraint A2 — Structure-Independent Comparison. Let f_1, f_2 be two facts represented within admissible representations R, R′. The comparison relation

$$\mathrm{cmp}(f_1, f_2) \in \{\text{equal, distinct}\}$$

depends only on f_1 and f_2, not on R or R′. Equivalently: the comparison outcome is invariant under any admissible change of representation.

Clarification on non-circularity. A2 does *not* assume that all admissible structures are equivalent, nor that the space of admissible representations is reducible in advance. It asserts only that when two observers identify a pair of committed records as *the same records*, their comparison outcomes must agree. The content of A2 is therefore not *uniqueness of structure* but *coherence of identification*: without it, the very statement "these two observers measured the same thing" is ill-defined, and no cross-observer empirical claim can be formulated. Structural closure — and with it, the uniqueness results of §6 — emerges only after A2 is combined with A1 and A3; it is not contained in A2 alone. A2 is therefore a *coherence constraint on comparison*, not a hidden assumption that the work of the theorem has already been done.

Explicit logical dependence. A2 fixes only the coherence of comparison on shared facts. It does not, on its own, constrain non-comparison structure, hidden scaffolding, or extension classes; those are eliminated only after A1 and A3 are imposed jointly. A2 alone allows many representationally distinct frameworks to coexist; the claim that residual non-comparison differences are observationally redundant belongs to closure (Theorem 6.1), not to A2.

Constraint A3 — Observational Accessibility. *The admissible structure of $\mathcal{T}$ is, by definition, the quotient of its raw structural data by the equivalence relation "indistinguishable on all observables."* Structural features with no observable consequence are collapsed to a single equivalence class; structural features with observable consequences survive the quotient as distinguishable admissible features.

Equivalently, any structural feature σ either:

- (i) affects at least one observable outcome (i.e., changes some probability, expectation value, or record), surviving the quotient as a distinct admissible feature, or
- (ii) has no observable consequence, collapsing to the zero class under the quotient.

Status of A3. A3 is partly *definitional* and partly *substantive*. Its definitional content (the quotient) specifies what is meant by the *physical* structure of an empirical theory. Its substantive content is the claim that empirical physics operates on admissible structure — not on the raw pre-quotient data. The earlier (and weaker) framing — "σ must be removed from $\mathcal{T}$" — is replaced by the cleaner one: σ may be retained in $\mathcal{T}$ as ontological commitment, but does not appear in the admissible structure on which empirical physics operates.

Why A3 is not merely definitional. The quotient appearing in A3 is not introduced to simplify the theory, but is forced by the requirement that physical distinctions be *representable as records*. Any distinction that cannot be encoded, reproduced, and compared as a record fails to enter empirical discourse under Definition 2.1. The quotient therefore does not eliminate structure by stipulation; it identifies the *maximal* structure that can be empirically realised. A3 is thus not a redefinition of physics, but a consequence of the requirement that physical distinctions be operationally representable.

A3 is not operationalism restated. A3 does not claim that unobservable structure *cannot exist*. It claims only that unobservable structure cannot enter the *empirical content* of a theory except through an equivalence class. This is not a metaphysical posture but an abstraction of standard scientific practice: gauge redundancy, coordinate choices, regularization dependence, and auxiliary state-space scaffolding are all treated in exactly this way. A3 merely makes the common rule explicit. Redundant formalism is not forbidden by A3; only its survival as *independent empirical content* is.

3.1 Remarks

- **A1** is a *bounded-information* principle. It does not forbid the universe from containing infinite variety; it constrains only what a finite observer can resolve. A1 constrains *empirical resolution*, not *mathematical spectrum*: a continuous observable in formalism yields only finitely many experimentally resolvable bins under any finite procedure, because every record has finite resolution. A1 therefore excludes actually infinite observational discrimination in finite record-space, not continuous spectra in the formalism.

- **A2** is a *representation-invariance* principle. It generalises the well-known physical principle that coordinate choices carry no empirical content.
- **A3** is a *quotient* principle: the admissible structure of a theory is the observational equivalence class of its raw structural data.

---

4. The No-Escape Theorem

The burden is contrapositive. We do not derive A1–A3 from a richer theory. We show that their negation destroys one or more of the operational conditions that define empirical physics. Readers who find the constraints superficially familiar should note that the content of the theorem lies not in the statement of A1–A3 but in the claim that there is no empirical physical theory outside them.

Theorem 4.1 (No-Escape Theorem). *Any framework $\mathcal{T}$ satisfying Definition 2.1 must satisfy each of A1, A2, A3. Conversely: any framework violating A1 or A2 fails at least one of (D), (R), (C) directly; any framework violating A3 commits to distinctions outside the empirical scope permitted by Definition 2.1. In both cases, the framework fails to be an empirical physical theory.*

Proof. We establish the contrapositive for each constraint, noting that the mechanism of failure differs between A1/A2 (predictive failure) and A3 (scope violation).

Case 1 — Violation of A1 (infinite operational distinguishability).

Suppose there exists a finite procedure $\mathcal{P}$ whose outcome set $D(\mathcal{P})$ is infinite and whose distinctions are physically real. Let $\omega \in D(\mathcal{P})$ be the outcome of a single execution.

A finite observer, by hypothesis finite in memory, signal bandwidth, and discrimination threshold, can produce only a finite record r. Two distinct formal outcomes $\omega, \omega' \in D(\mathcal{P})$ that lie within the observer's minimal discrimination threshold produce identical records: $r(\omega) = r(\omega')$.

Hence:

1. The record r does not determine ω — so the outcome is not *recorded* in the sense of (D).
2. Repeating $\mathcal{P}$ yields outcomes $\omega_1, \omega_2, \ldots$ whose records $r_1, r_2, \ldots$ cannot in general distinguish whether any two trials produced the same formal outcome — so (R) fails to be decidable on records.

3. Two observers with distinct thresholds produce records disagreeing systematically on fine distinctions — so (C) fails.

Thus ¬**A1 ⇒ ¬(D) ∧ ¬(R) ∧ ¬(C)**. ■ (Case 1)

Case 2 — Violation of A2 (structure-dependent comparison).

Suppose there exist admissible representations R, R′ and facts f_1, f_2 such that

$cmp_R(f_1, f_2)$ = equal, but $cmp_{R'}(f_1, f_2)$ = distinct.

Consider two observers O and O′ employing R and R′ respectively. Both are, by stipulation, using *admissible* representations — i.e., representations sanctioned by $\mathcal{T}$.

Then:

1. Presented with the identical empirical situation involving f_1 and f_2, O reports "equal" while O′ reports "distinct."
2. No procedure internal to $\mathcal{T}$ can adjudicate which report is correct, since both representations are admissible by hypothesis.
3. Hence the outcome of the comparison is not a determinate matter across observers — violating (C).

Moreover, because the comparison of records is the operational basis of "reproducibility," (R) also fails: whether a repeated trial *agrees* with a previous trial becomes representation-dependent.

Thus ¬**A2 ⇒ ¬(C) ∧ ¬(R)**. ■ (Case 2)

Case 3 — Violation of A3 (structure without observable content).

A3, as stated in §3, is partly definitional: it specifies that the admissible structure of a theory is the observational-equivalence quotient of its raw structural data. The substantive content of A3 is that empirical physics operates on this admissible structure.

Suppose $\mathcal{T}$ claims that two structural features σ and σ′ are *distinct physical facts*, while simultaneously agreeing that for every finite procedure $\mathcal{P}$ the outcome distribution is invariant under replacement of σ by σ′. Then:

1. σ and σ′ lie in the same observational equivalence class. Under the definitional content of A3, they are identified within the admissible structure of $\mathcal{T}$.

2. The claim that σ and σ′ are nevertheless *distinct as physics* is therefore a claim not about admissible structure but about raw, pre-quotient structure — content that cannot appear in any record, cannot be reproduced in any trial, and cannot be compared across observers.
3. Such a claim lies outside Definition 2.1 by construction: it posits a distinction that never enters empirical discourse. Not because the distinction is predictively false, but because it has no empirical representation to begin with.

Thus the violation of A3 is not a failure of (D), (R), or (C) on any specific prediction the framework makes. Empirically inert σ produces no predictions that can fail; that is precisely what "empirically inert" means. The violation is instead a *scope violation*: the framework asserts content that lies outside what Definition 2.1 admits as empirical. The admissible structure of $\mathcal{T}$ still satisfies A3; the framework $\mathcal{T}$ as a whole simply carries commitments outside the admissible structure that are not part of its empirical physics. The conclusion is therefore not that formal surplus cannot appear in physical description, but that it cannot survive as independent *empirical* structure.

A3 is not a definitional collapse into Definition 2.1. Definition 2.1 constrains what counts as *empirical content*; A3 constrains how *structure relates to that content*. The two operate at different levels. A framework that retains σ as a physically distinct feature while conceding that σ never affects any record is introducing distinctions that — by Definition 2.1 — cannot enter empirical discourse. The quotient enforced by A3 is therefore not a stipulation; it is forced by the requirement that physical distinctions, to count as empirical, must be representable as records. A3 inherits its necessity from this representability requirement, not from the definitional content alone.

Thus ¬**A3 ⇒ the empirical physics of the framework is not the framework itself, but its admissible-structure quotient**. ■ (Case 3)

Combining cases. Violations of **A1** or **A2** force failure of at least one of (D), (R), (C) directly — the framework fails to produce the determinate, reproducible, or inter-observer-consistent outcomes that Definition 2.1 requires. Violations of **A3** place the framework's physical claims outside the scope admitted by Definition 2.1 — not through predictive failure, but through scope excess. In either case, the framework is not an empirical physical theory. Contrapositively, any empirical physical theory satisfies **A1 ∧ A2 ∧ A3**. ■

---

5. Reduction of Apparent Counterexamples

Scope of this section. The role of §5 is not to classify all ontologies, but to show that the most common apparent escape routes fail at the level relevant to empirical content. Proposition 5.1 is a scope-classification claim, not an ontology-classification claim. In relation to Theorem 4.1, Proposition 5.1 functions as a completeness check over apparent counterexamples — confirming that the theorem's reach is not limited to frameworks that already comply with A1–A3 — rather than as a separate link in the main inference chain.

A natural objection to Theorem 4.1 is that alternative frameworks — purely relational, observer-dependent, modal, or hidden-variable — might *appear* to evade **A1–A3**. We show they do not.

Proposition 5.1 (Reduction of Apparent Alternatives). *Let $\mathcal{T}$ be any framework claiming empirical content. Then at least one of the following holds:*

- (i) *$\mathcal{T}$ admits a representation satisfying A1–A3 (possibly after quotienting redundant structure);*
- (ii) *$\mathcal{T}$ fails to produce observable content, i.e., violates Definition 2.1.*

Remark (what is being classified). Proposition 5.1 is not intended as a taxonomic classification of all conceivable ontologies. It is an exhaustion at the level relevant to the theorem: any putative counterexample must either (i) produce empirical records — in which case it falls under Definition 2.1 and Theorem 4.1 — or (ii) fail to produce empirical records, in which case it lies outside the empirical class. The classification is of *operational access modes* available to finite observers, not of imagined ontologies. Any candidate observable — however exotic its ontology — must be instantiated through one of these access modes or fail to produce a record at all. This is a much stronger claim than a taxonomy of mathematical structures, because it does not depend on enumeration of imagined ontologies: it depends only on what it means to access empirical content in the first place.

Proof (by classes of candidate counterexamples).

The following classes (a)–(e) are *worked illustrations*, not an independent exhaustion proof. Exhaustion is already established by Theorem 4.1 together with the record-mediation requirement: any framework producing empirical content does so via records, and records — being finite artefacts produced by finite observers under finite procedures — obey A1–A3. The classes below illustrate how this reduction proceeds for the ontologies most often proposed as counterexamples.

(a) Purely relational frameworks. Relational theories replace absolute quantities with relations between facts. But the relations themselves must be recordable, reproducible, and comparable — else no empirical claim can be formulated *about* the relations. Relations that satisfy (D)–(C) satisfy **A1–A3** by Theorem 4.1. Relations that do not are not empirical. Hence (i) or (ii).

(b) Observer-dependent frameworks. In observer-relative frameworks, outcomes are indexed by observer: f = f(O). For such a framework to yield inter-observer scientific discourse, the *mapping* O ↦ f(O) must itself be comparable — otherwise observers cannot even *disagree*, let alone communicate results. Comparability of the mapping is an instance of **A2** at the meta-level. If comparability holds, (i); if not, (ii).

(c) Hidden-variable or redundant-structure frameworks. Any hidden variable either affects observables — in which case it satisfies (i) of **A3** — or does not, in which case **A3(ii)** quotients it out. The apparent "extra structure" is either empirical and constrained, or fictive and eliminable.

(d) Modal and many-worlds frameworks. Modal ontologies must still produce determinate records *within* each world/branch to support (D). Comparisons of records within a branch obey **A1–A3** by Theorem 4.1. Claims that transcend branches (e.g., "all branches are real") are either reducible to claims about the branch structure's observable signature (case (i)) or empirically vacuous (case (ii)).

(e) Infinitary or continuum frameworks. Continuum mathematics *represents* physical quantities but does not contradict **A1**: the observable distinctions extracted from any finite procedure remain finite. The continuum enters as a limit/idealisation of finite records, not as a set of directly accessible distinctions.

In every class, the framework either admits a reduction satisfying **A1–A3** or fails Definition 2.1. ■

---

6. Consequence: Constraint Closure Is Forced

Section 6 records the consequence of combining Theorem 4.1 with closure and uniqueness theorems established elsewhere in the VERSF programme. Theorem 4.1 stands independently of §6; the present section establishes the additional implication chain summarised in §7.

We now combine Theorem 4.1 with previously established results of the VERSF programme. For the convenience of readers encountering the programme through this paper, we include proof sketches alongside the formal statements. Full proofs appear in [1,2].

Theorem 6.1 (Closure of Constraint Systems) [1]. *Any constraint system satisfying* ***A1–A3*** *is closed: every admissible extension of its structural content is observationally equivalent to a structure already contained in the system.*

Proof sketch. "Extension" here means extension of structural content over the same set of finite procedures admitted by $\mathcal{S}$. Extensions that also introduce new finite procedures are handled by the same argument applied to the enlarged procedure set, since each added procedure still satisfies A1.

Let $\mathcal{S}$ be a constraint system satisfying **A1–A3** and let $\mathcal{S}' \supseteq \mathcal{S}$ be any admissible extension. By **A3**, every feature of $\mathcal{S}'$ either produces observable signatures or is quotiented out; the raw structural data is therefore already reduced to its observable equivalence class. By **A1**, the observable distinctions accessible through any finite procedure form a finite set, so the equivalence class is generated by finitely many distinguishing procedures. By **A2**, this equivalence is representation-independent: it is a property of the facts, not of their encoding.

Any feature of $\mathcal{S}'$ not already equivalent to some feature of $\mathcal{S}$ must therefore either be observable or not. If observable, it must be expressible through the (possibly enlarged) set of finite procedures admitted by $\mathcal{S}'$. By **A1**, this set remains finite per procedure; by **A2**, the resulting distinctions are representation-independent; by **A3**, distinctions not affecting observable outcomes are quotiented out. The extension therefore cannot introduce observational content outside the equivalence class generated by finite procedures over $\mathcal{S}$.

If non-observable, **A3** quotients it out directly. Either way, $\mathcal{S}'$ is observationally equivalent to $\mathcal{S}$. ■ (Full derivation: [1].)

Theorem 6.2 (Structural Uniqueness) [2]. *A closed constraint system in the sense of Theorem 6.1 admits a single equivalence class of admissible structures.*

Proof sketch. The content of the theorem is that agreement on empirical assertions — facts, procedures, and comparison outcomes — *forces* agreement on admissible structure, not merely that it is *compatible* with it. The work is done by **A1–A3**, which collapse any two systems sharing an empirical base to the same quotient.

In detail: let $\mathcal{S}_1$, $\mathcal{S}_2$ be two constraint systems satisfying **A1–A3** over the same empirical base (same facts, same finite procedures, same comparison outcomes). By **A3**, each reduces to the observational quotient of its raw structural data. By **A1**, these quotients are generated by finitely many distinguishing procedures drawn from the same base. By **A2**, each quotient is representation-independent. Consequently, $\mathcal{S}_1$ and $\mathcal{S}_2$ produce the same quotient — the same equivalence class of admissible structures. The uniqueness is up to this observational equivalence, not up to set-theoretic identity: representations differing only in non-observable encoding detail are quotiented together. ■ (Full derivation: [2].)

Combining with Theorem 4.1:

Corollary 6.3 (Unavoidable Closure and Uniqueness). *Any empirical physical theory induces a constraint-closed structure admitting a unique equivalence class. Consequently, the structural core of any empirical physical theory is fixed up to observational equivalence.*

---

7. The Final Chain

The results combine into a single implication chain:

Empirical Physics ⇒ {A1, A2, A3} ⇒ Constraint Closure ⇒ Structural Uniqueness.

Each arrow is a theorem, not an assumption:

- The first arrow is Theorem 4.1 (this paper).
- The second is Theorem 6.1 [1].
- The third is Theorem 6.2 [2].

The chain localises *all* residual conditionality of physical law to a single premise: the operational definition of empirical physics given in §2. Given that definition, the structural core of physics is forced.

---

8. Interpretation

8.1 What the theorem does assert

The theorem asserts that *any* framework producing empirical physical content is structurally constrained in the manner specified. It does *not* assert that any particular model — including any specific model within the VERSF programme — is the correct instantiation of those constraints. The claim is about the *class* of admissible frameworks, not any member of it.

8.2 What the theorem does not assert

- It does not prove that the universe *is* empirically describable.
- It does not prove that human observers are necessarily finite. (§2.3 argues this is a physical fact about record-producing systems; the theorem does not require it to be provable a priori.)
- It does not prove that a unique *model* (as opposed to a unique *equivalence class of structures*) follows.

These are distinct questions. The theorem's strength is that it requires none of them.

8.3 The route of escape

The only coherent way to reject the conclusion is to reject the operational definition of physics itself. This requires abandoning at least one of:

- stable records (D),
- reproducibility (R),
- inter-observer consistency (C).

Any such rejection is internally consistent but vacates the concept of empirical science. It is not a rival physical theory; it is a rejection of the enterprise of physical theory.

8.4 This is not an interpretation of physics

The result does not propose an interpretation of an existing theoretical framework. It constrains the conditions under which *any* framework can qualify as empirical physics in the first place.

Interpretations of quantum mechanics (Copenhagen, Everett, Bohmian, QBist, relational), reconstructions of quantum theory from information-theoretic axioms, and operationalist or instrumentalist philosophical stances all operate *within* frameworks that already satisfy (D)–(C). They take empirical content as given and ask what structure best accounts for it. The present result applies at a prior level: it establishes the structural constraints that any framework must satisfy in order to support empirical content at all.

A note on realism. A sophisticated realist may object that Definition 2.1 is itself operationalism in disguise: by characterising physics in terms of "what can be recorded, repeated, and compared," we tacitly restrict physics to empirical content, begging the question against metaphysical realism. This objection conflates two distinct levels. Definition 2.1 does not claim that *what exists* is exhausted by what can be empirically recorded. It claims only that *what a physical theory can empirically assert* is constrained by (D)–(C). A realist is free to hold that the world contains more than the theorem constrains — unobservable entities, intrinsic natures, modal or dispositional structure. The theorem is silent on such claims. What it is not silent on is the *structural core of any framework that ventures an empirical claim about them*: once a realist offers a theory whose content can be tested, recorded, reproduced, or compared, that theory inherits **A1–A3**. Realism at the ontological level is therefore compatible with **A1–A3** at the empirical level; the theorem operates on the latter and does not preempt the former.

In this sense, the theorem is not a statement about how to understand physics. It is a statement about what makes physics possible. Any interpretation — realist, instrumentalist, relational, or otherwise — sits *downstream* of **A1–A3** and inherits them, whether or not its proponents acknowledge them explicitly.

---

9. Falsifiability

The theorem is falsifiable in principle. It would be falsified by any one of the following:

- **F1.** A reproducible, inter-observer-consistent experimental procedure requiring an *actually infinite* operational distinguishability from a finite observer.
- **F2.** Two observers using admissible representations of the same empirical framework, obtaining irreconcilable comparison outcomes on identical facts, without either representation being revealed as inadmissible.
- **F3.** A structural feature that, by construction, affects no observable, yet whose presence or absence is resolvable by a reproducible inter-observer procedure.

No such example is known. The standard apparent counterexamples — quantum contextuality, observer-dependence in relativity, gauge redundancy — all reduce, under careful analysis, to instances of **A1–A3** properly applied.

---

10. Conclusion

The conclusion concerns the structural core of empirical physics and should be read at that level throughout. It is not a claim about every possible descriptive or metaphysical dressing a theory might carry.

The constraints required for the formation of empirical facts are not optional features of particular theories. They are necessary conditions for empirical physics as such. Combined with previously established closure and uniqueness results, this implies that the structural core of physical law is not freely chosen but is uniquely fixed up to observational equivalence — that is, all admissible physical theories share a single empirical structure, even if they differ in non-observable representation.

The conditionality of physics has been reduced to a single premise: the operational definition of empirical science. Any escape from the conclusion requires abandoning that premise, and with it, the possibility of physical theory. The result is *uniqueness of empirical structure up to observational equivalence*, not uniqueness of every possible descriptive or metaphysical dressing.

References

[1] Taylor, K. *Closure of Constraint Systems Under Finite Distinguishability*. VERSF Programme, AIDA Institute.

[2] Taylor, K. *Structural Uniqueness from Constraint Closure*. VERSF Programme, AIDA Institute.

[3] Taylor, K. *The Single-Source Theorem: Observables as Functionals of Committed-Record Density*. VERSF Programme, AIDA Institute.

[4] Taylor, K. *From Necessary Facts to Physical Structure*. VERSF Programme, AIDA Institute.
