The Incompleteness of Incompleteness

Why Gödel's Theorems Do Not Constrain Physically Admissible Systems

Keith Taylor VERSF Theoretical Physics Program

Core Result

Gödel's incompleteness theorems are mathematically correct but do not constrain physically interpretable propositions.

Any physically admissible system — one constrained by finite information capacity, irreversible thermodynamic commitment, and bounded recursion depth — cannot express full Peano arithmetic. It therefore falls outside the scope of Gödel's theorems entirely.

Even when we use infinite mathematical frameworks to *model* physical systems, the sentences that Gödel's theorems render undecidable are exclusively those requiring unbounded quantification over domains no physical system can instantiate. These sentences fail the Admissibility Filter: they admit no finite operational test and cannot be stabilised as irreversible committed facts. They have no physical referent.

Gödel's reach into our formalism is real. His reach into nature is empty.

The limitations that physically admissible systems actually face are not logical barriers but admissibility boundaries — resource constraints imposed by thermodynamics. These boundaries are qualitatively different from Gödelian incompleteness: they are dynamic rather than permanent, physical rather than logical, and reflect the cost of being finite rather than the paradoxes of being infinite.

For the General Reader

What did Gödel prove?

In 1931, the mathematician Kurt Gödel proved something remarkable: any mathematical system powerful enough to do basic arithmetic will contain true statements that the system itself can never prove. This is called the incompleteness theorem, and it is often interpreted as placing a permanent, unbreakable limit on human knowledge — the idea that there will always be truths about the universe that no theory can capture.

What does this paper argue?

We argue that this extrapolation — from abstract mathematics to physical reality — doesn't follow. Gödel's proof has a critical requirement: it only works on systems that can handle *infinite* arithmetic. It needs a system that can count without limit, encode infinitely many formulas, and search through infinitely many possible proofs. Every step of the proof depends on this infinity.

Why doesn't that apply to the real world?

Because the physical universe is not infinite — at least not in the way Gödel's proof requires. Every physical system, including the entire observable universe, can store only a finite amount of information (roughly 10^{122} bits, set by the laws of thermodynamics). It can make only a finite number of distinctions. And every time it produces a fact — every time something goes from undecided to decided — it must pay an irreversible energy cost. There is no free lunch in physics, not even for logic.

Because no physical system meets Gödel's requirements, his theorem simply doesn't apply to nature.

But physicists use infinite mathematics all the time. Doesn't that bring Gödel back in?

It does — but only at the level of the mathematical model, not at the level of reality. We show that even within an infinite mathematical framework, the specific sentences that Gödel makes undecidable are always about the infinite scaffolding of the model. They are never about anything you could physically measure, test, or observe. Think of it this way: the model has more structure than nature does, and Gödel's limitations live entirely in the excess.

So what limits *does* physics face?

Real ones — but different ones. Some truths may be inaccessible because proving them would cost more energy than the universe can supply. That's a thermodynamic constraint, not a logical paradox. It's the difference between a door that is locked by the laws of logic (Gödel) and a door that is too expensive to open (physics). Our paper argues that every door physics encounters is the second kind.

In one sentence?

Gödel proved that infinite systems have inherent logical limits — but reality isn't infinite, so those limits don't apply to the physical world.

Table of Contents

Abstract

Gödel's incompleteness theorems demonstrate that any consistent formal system capable of expressing Peano arithmetic contains true statements that cannot be proven within the system. This result is mathematically unassailable within its domain. However, we argue that Gödel's theorems do not constrain physical reality — not because the theorems are wrong, but because physically admissible systems cannot physically instantiate or interpret the preconditions those theorems require. Drawing on the Physical Admissibility Framework (PAF), Finite Commitment Mathematics (FCM), and the principle of finite distinguishability from the VERSF programme, we show that any system constrained by finite information capacity, irreversible commitment, and bounded distinguishability falls below the threshold at which Gödel's theorems engage. We further show that the finitude of physical systems is not a philosophical preference but a thermodynamic necessity, grounded in the Bekenstein bound and the irreversibility of fact-production. The incompleteness theorems are theorems about a class of mathematical structures that nature does not — and cannot — instantiate.

---

1. Introduction

Gödel's 1931 incompleteness theorems are among the most celebrated results in mathematical logic [1]. The first theorem states that any consistent, effectively axiomatised formal system capable of expressing basic arithmetic contains statements that are true but unprovable within the system. The second theorem strengthens this: such a system cannot prove its own consistency.

These results are often interpreted as placing fundamental limits on knowledge, proof, and even physical theory. If the mathematics describing reality is subject to Gödel's constraints, then there are truths about the universe that no finite theory can capture. This extrapolation has influenced philosophy of mind [11], foundations of physics [6], and even popular conceptions of what science can achieve.

We argue that this extrapolation — from formal arithmetic to physical reality — is unjustified. The argument is not that Gödel's theorems are false. They are impeccable within their scope. The argument is that their scope does not include the systems nature actually uses to produce facts, and that any mathematics adequate to describe a physically admissible universe inherits structural constraints that prevent Gödel's preconditions from being satisfied.

The finitude of physical systems is not a new observation. Ultrafinitist programmes [12, 13] have long questioned the applicability of infinite arithmetic on philosophical grounds. What the VERSF programme contributes is a *physical* grounding: the claim that finitude is not a preference but a consequence of thermodynamic law. Nature does not merely happen to be finite — it is *required* to be finite by the structure of information, entropy, and distinguishability.

The key insight is simple: Gödel's incompleteness requires unbounded arithmetisation of syntax — quantification over infinite domains that no physical system can supply. The remainder of this paper makes that claim precise and traces its consequences.

1.1 Related Literature and Positioning

The question of Gödel's relevance to physics and computation has been addressed from several directions. We briefly situate our contribution relative to existing work.

Franzén [18] provides a careful analysis of common misapplications of Gödel's theorems, demonstrating that many popular extrapolations (to human cognition, physics, and artificial intelligence) rest on misunderstandings of what the theorems actually require. Our analysis is consistent with Franzén's: we are not claiming that Gödel's theorems are wrong, overrated, or poorly understood. We are making a specific structural argument about the relationship between their preconditions and physical admissibility.

Feferman [19] has argued extensively that the mathematics required for physics is predicatively reducible — that physically relevant mathematics does not require the full strength of set theory or Peano arithmetic. Our argument is complementary: where Feferman asks "what mathematics does physics *need*?", we ask "what mathematics can physics *ground*?" The answer, in both cases, is less than the full apparatus of classical arithmetic.

The Lucas-Penrose argument [11, 20] claims that Gödel's theorems show human minds transcend Turing machines. Our framework implies that this argument rests on a category error: if physical systems (including brains) are PAC-compliant, then they fall outside the scope of Gödel's theorems entirely, and the question of whether minds "see" the truth of Gödel sentences is not a question about transcending computation but about the relationship between physical semantics and formal overhead. Shapiro [21] provides a thorough philosophical analysis of these issues.

Aaronson [22] has explored the connections between computational complexity, physics, and the Bekenstein bound, arguing that physical constraints impose real limits on computation. Our treatment of admissibility boundaries as the physically relevant successor to Gödelian incompleteness aligns with this perspective: the limitations physics faces are resource constraints, not logical barriers.

---

2. What Gödel's Theorems Actually Require

To identify precisely where the theorems lose purchase on physical systems, we must state their preconditions with full rigour and examine the structural dependencies of the proof.

2.1 The Three Preconditions

Gödel's first incompleteness theorem applies to a formal system **F** if and only if:

1. **Consistency:** F does not derive a contradiction (there is no formula φ such that $F \vdash \varphi$ and $F \vdash \neg\varphi$).
2. **Effective axiomatisation:** The set of axioms of F is decidable — there exists an algorithm that determines, for any formula, whether it is an axiom.
3. **Sufficient expressive power:** F can represent all primitive recursive functions — equivalently, F can express the arithmetic of the natural numbers $\mathbb{N} = \{0, 1, 2, 3, \ldots\}$.

Condition (3) is the critical one. It requires that the system contain, or faithfully represent, the entire structure of Peano arithmetic — an infinite domain with unbounded successor operations, unbounded multiplication, and the full apparatus of recursive self-reference via Gödel numbering.

2.2 The Structural Dependencies of the Proof

Clarification. Our claim is not that "infinity" is the only ingredient in incompleteness, nor that every bounded arithmetic theory is complete. Incompleteness-like phenomena can arise in various weak systems depending on their interpretability strength, and the boundary between "Gödel bites" and "doesn't" is a subtle question in mathematical logic. The point we are making is narrower: Gödel's original undecidable sentence — and the standard proof method that constructs it — requires a theory to support unbounded arithmetisation of syntax and an unbounded proof quantifier. Those are precisely the features that fail the Admissibility Filter in PAC-compliant systems (Section 3.4). Our argument targets this specific construction and its physical interpretability, not the full landscape of incompleteness phenomena in abstract mathematics.

Every step of Gödel's construction depends on infinitary structure. It is worth tracing these dependencies explicitly, because each will fail independently under physical admissibility constraints.

Gödel numbering assigns a unique natural number to every symbol, formula, and finite proof sequence in the system. This encoding is injective: distinct syntactic objects receive distinct numbers. Let $\sigma_1, \sigma_2, \ldots$ be the symbols of the system, and let $\ulcorner \varphi \urcorner$ denote the Gödel number of formula φ. The encoding requires:

For any well-formed formula φ, there exists a unique $n \in \mathbb{N}$ such that $\ulcorner \varphi \urcorner = n$.

Since the set of well-formed formulas is countably infinite (any finite alphabet generates infinitely many strings under concatenation), this requires an unbounded supply of distinct integers. If the available integers are bounded above by some M, then formulas whose Gödel numbers exceed M simply cannot be encoded.

The diagonal lemma (also called the fixed-point lemma) is the engine of self-reference. It states that for any formula $\psi(x)$ with one free variable, there exists a sentence σ such that:

$$F \vdash \sigma \leftrightarrow \psi(\ulcorner \sigma \urcorner)$$

The construction of σ requires composing substitution functions over the Gödel numbering — operations that traverse the full domain of natural numbers. The lemma's proof relies on the totality of these functions: for *every* input, the function must produce a well-defined output. In a bounded domain, substitution functions become partial (undefined for inputs exceeding the bound), and the fixed-point construction breaks.

The undecidable sentence G encodes the claim: "There is no proof of G in F." Formally, G asserts:

$\neg\exists y\ (\text{Proof\_F}(y, \ulcorner G \urcorner))$

where Proof_F(y, x) is the arithmetised proof predicate — "y is the Gödel number of a proof of the formula with Gödel number x." The quantifier ∃y ranges over all natural numbers, i.e., over all possible proofs. In a bounded system, this quantifier ranges over a finite set, and the statement becomes decidable by exhaustive search.

The essential role of infinity. Remove the infinite domain, and:

- Gödel numbering becomes a finite map, unable to encode all formulas.
- The diagonal lemma loses the total functions it requires.
- The undecidable sentence becomes decidable by finite inspection.

This is not a minor technical point. It is the structural load-bearing wall of the entire result. Gödel's theorems are, at their core, theorems about what happens when a system is powerful enough to mirror its own infinite complexity. Systems that lack this power are simply outside the theorem's scope.

---

3. Physical Admissibility: The Constraints Nature Imposes

The VERSF programme, through its constituent frameworks — Bit Conservation and Balance (BCB), Ticks-Per-Bit (TPB), and the Physical Admissibility Framework (PAF) — identifies three non-negotiable constraints that any physically realisable system must satisfy [5, 8, 9]. These are not approximations, modelling choices, or engineering limitations. They are structural features of any system that produces definite, distinguishable outcomes.

3.1 Finite Distinguishability

Any physical system can encode only a finite number of distinguishable states. The Bekenstein bound [2] provides a rigorous upper limit: the maximum entropy of a region of space with radius R and energy E is:

$S_{max} \leq 2\pi kRE / (\hbar c)$

This bounds the information content of any finite region. A system of finite spatial extent and finite energy contains a finite number of bits — and therefore a finite number of distinguishable configurations. For the observable universe as a whole, this yields an estimated upper bound of approximately 10^{122} bits [14], corresponding to at most $2^{(10^{122})}$ distinct states.

This is a staggeringly large number, but it is finite. It means no physical system — not even the entire observable universe — can instantiate the full natural numbers. There is always a largest number the system can represent. Beyond that boundary, "the next number" is not merely unknown — it is physically meaningless. The system lacks the degrees of freedom to encode it.

A note on cosmological subtleties: the Bekenstein bound is rigorously established for systems in asymptotically flat spacetime with finite energy. In de Sitter space — our actual cosmological setting — observer-dependent horizons mean that different observers have access to different finite regions, and the global entropy associated with the cosmological constant introduces subtleties about whether "the universe's information capacity" is well-defined as a single number. These refinements do not threaten the present argument: the bound applies to any causally accessible region, and observer-dependent horizons refine but do not eliminate the finiteness constraint. Every observer, in every cosmological setting, has access to at most finitely many distinguishable states.

Formally: for any bounded physical system S, there exists a finite integer N such that the total number of distinguishable configurations of S is at most 2^N. The "arithmetic" such a system can perform is therefore bounded arithmetic over the domain $\{0, 1, \ldots, M\}$ for some M determined by N.

A further structural point deserves brief note. The apparatus of distinguishability presupposes a reference ground — a null state against which difference is defined. In physically admissible systems, this ground state is not a member of the state space but the pre-informational substrate from which the state space is carved: zero is the platform for distinction, not a participant in it. This contrasts with Peano arithmetic, where $0 \in \mathbb{N}$ is the initial element of the domain. The distinction does not, by itself, invalidate Gödel's theorems — but it deepens the structural gap between PA and physical systems, reinforcing that the arithmetic of physically admissible systems is not merely bounded PA but a differently structured calculus rooted in the physics of distinction.

3.2 Irreversible Commitment

In PAF, the production of a fact — a definite, distinguishable outcome — requires an irreversible thermodynamic commitment [5]. This claim requires careful statement, because it is sometimes conflated with a stronger claim than is warranted.

Landauer's principle [15] establishes that erasing one bit of information dissipates at least $kT \ln 2$ of energy as heat. Bennett [16] showed that computation itself can, in principle, be carried out reversibly — near-dissipationlessly — provided the computing system retains all intermediate information. These results are not in tension: reversible computation avoids dissipation by deferring erasure, not by eliminating the thermodynamic constraint.

The relevant claim for our purposes is narrower and more defensible than "all computation costs entropy." It is this: *the stabilisation of a definite record — a committed fact that can be read, copied, and used as input to further processes — requires entropy export*. A reversible computation may avoid dissipation during processing, but the moment it produces a stable, reusable output — a record that persists independently of the computation that generated it — some degrees of freedom must be reset, and Landauer's bound applies to that reset. The cost may be deferred, redistributed, or minimised, but it cannot be eliminated entirely in any finite system that produces and retains definite facts.

More precisely: a physically admissible system that commits n bits of stable fact must, somewhere in its causal history, export at least $n \cdot kT \ln 2$ of entropy to its environment. This is a lower bound; realistic systems dissipate far more.

This has direct consequences for proof verification. In classical logic, checking whether a sequence of symbols constitutes a valid proof is treated as a costless operation — a function from strings to {True, False} that exists timelessly. In a physically admissible system, each step of proof verification is a physical process that produces a committed record ("valid" or "invalid"), and each such commitment exports entropy.

The implication is this: verification of unboundedly many proofs requires unbounded entropy export. Let P(n) be the entropy cost of verifying a proof of length n and stabilising the result. In any physically admissible system:

$\Sigma_{i=1}$^∞ P(i) → ∞

But the system's total available entropy throughput is finite:

S_total < ∞

Therefore, only finitely many proof-verification results can ever be committed as stable facts. The quantifier ∀y in Gödel's proof predicate cannot be physically instantiated as a sequence of committed verification outcomes.

3.3 Bounded Recursion Depth

The Ticks-Per-Bit (TPB) framework [9] establishes that every distinguishability operation requires a minimum temporal cost — at least one "tick" per bit of information resolved. This is not a contingent engineering fact but a consequence of the structure of information: resolving a binary distinction requires at least one irreducible temporal step.

Self-referential constructions that require unbounded recursion depth are therefore temporally unbounded. In a universe of finite age — and, more fundamentally, of finite total entropy throughput — only finitely many levels of recursion can ever be physically instantiated.

Gödel's diagonal construction requires a system to:

1. Represent its own proof apparatus as arithmetic (first level of encoding).
2. Construct a formula that refers to properties of this encoding (second level).
3. Apply substitution to generate a self-referential fixed point (third level and beyond).

Each level requires the system to maintain coherent representations of the previous levels while performing operations at the current level. The total information required grows with each recursive step. In a bounded system, there exists a maximum recursion depth d_{max} beyond which the system cannot maintain coherence — not because of a logical prohibition, but because it runs out of distinguishable states to encode the meta-levels.

The depth d_{max} is determined by the system's total information capacity N:

$$d_{max} \approx N / \log_2(|\Sigma|)$$

where $|\Sigma|$ is the size of the system's symbol alphabet. For any finite N, this is finite, and the tower of self-reference that Gödel's proof requires is truncated.

3.4 Emergent Physical Semantics: The Admissibility Filter

The constraints of Sections 3.1–3.3 establish that physically admissible systems are finite in capacity, irreversible in commitment, and bounded in recursion. But these are *structural* constraints — they tell us what a physical system *cannot do*. We now need a *semantic* principle: a criterion that determines which formal sentences are *about* anything physical at all.

In physics, meaning is not primitive. A sentence acquires physical content only when it can be connected to a finite physical procedure that produces a definite, distinguishable outcome. This is not a philosophical preference — it is how physics has always worked. The empirical content of general relativity is not the Einstein field equations as formal syntax; it is the set of finite, operational predictions those equations generate (deflection angles, redshift values, orbital precessions). Sentences that cannot be connected to any such procedure — however well-formed they may be syntactically — have no physical content.

We formalise this through two admissibility principles that together constitute the **Admissibility Filter (AF)**.

Principle 1: Operational Closure. A sentence φ is physically interpretable in a system S only if there exists an admissible operational test M_φ such that:

(a) M_φ terminates within S's admissibility budgets (finite entropy, finite energy, finite time),

(b) M_φ partitions the physically distinguishable outcomes of S into two disjoint sets corresponding to "φ" and "$\neg\varphi$."

Physical semantics is not given in advance — it *emerges* only when the loop from formal symbols → physical procedure → finite outcome can be closed. A sentence for which no such loop exists is not false; it is physically meaningless — a statement about the formalism's internal structure, not about nature.

Principle 2: Fact-Commitment Semantics. A proposition φ has physical truth-value in a system S only insofar as its truth-value can be stabilised as a committed fact — an irreversible record produced by a thermodynamic process. Formally: φ has physical truth if there exists a physically realisable process that produces a stable record R such that:

(a) R is robust under admissible perturbations (it does not spontaneously revert),

(b) The theory maps R to φ (and not to ¬φ).

In physics, "truth" is not a timeless model-theoretic property. It emerges when a system pays the thermodynamic cost to make a distinction irreversible. A proposition whose truth-value cannot be stabilised as a committed record — because the required distinctions span an unbounded domain — has no operationally definable physical truth-value. It is a meta-proposition about an idealised infinity, not a fact about nature.

The Admissibility Filter (AF). A formal sentence φ has physical meaning in a PAC-compliant system S only if it satisfies operational closure (Principle 1) and admits fact-commitment (Principle 2).

The Admissibility Filter is not an ad hoc restriction. It is the formalisation of what makes physics *physics*: the requirement that claims about nature must, in principle, be testable by finite procedures and settleable by irreversible processes. It upgrades finite distinguishability from a capacity bound into a semantic emergence criterion — a filter that separates the physically meaningful content of a theory from its formal overhead. The justification for AF — including paradigm cases, invariance properties, and its relationship to verificationism — is developed in Section 3.5.

The consequences for Gödel's theorems are immediate. A Gödelian sentence G asserts the non-existence of a proof across an unbounded domain. Its truth conditions quantify over objects — arbitrarily large proof-numbers — that cannot be instantiated, encoded, or operationally closed under PAC constraints. No admissible operational test can settle G, because its truth conditions range over a domain exceeding the system's finite state space (failure of Principle 1). No finite thermodynamic process can commit the result as a stable record, because the required distinctions span an unbounded range of proof objects (failure of Principle 2). Gödelian sentences therefore fail the Admissibility Filter when interpreted as candidates for physical content. They exist within the formal scaffolding of a theory but have no emergent physical meaning.

3.5 Justifying the Admissibility Filter

A natural objection to the Admissibility Filter is that it is engineered to exclude Gödel sentences — that the conclusion ("Gödel sentences lack physical content") is smuggled into the criterion. This section demonstrates that AF is not tuned to any particular logical result. It is a general criterion for physical content, already implicit in the practice of physics, that correctly classifies paradigm cases of physical and non-physical sentences. The Gödel result is a corollary, not the motivation.

3.5.1 Paradigm Inclusions: What AF Admits

AF must correctly classify the sentences that physics unambiguously treats as having physical content. Consider:

- "The electron's anomalous magnetic moment is $g/2 = 1.00115965218128 \pm 0.00000000000018$." This corresponds to a finite-precision measurement outcome, produced by a terminating experimental procedure (Penning trap + quantum electrodynamics prediction), yielding a committed record. AF passes it.
- "The cosmic microwave background power spectrum has a first acoustic peak at multipole $\ell \approx 220$." This corresponds to a finite angular resolution measurement by instruments (WMAP, Planck), producing a committed data record. AF passes it.
- "The cross-section $\sigma(pp \to H + X)$ at $\sqrt{s} = 13$ TeV is 55.6 ± 2.5 pb." This is a tolerance-band claim about a finite-precision observable, verified by a terminating procedure (collider event counting + statistical analysis), yielding a committed record. AF passes it.

In each case, the physical claim is a finite-precision, tolerance-band statement — not an exact real-number identity. The claim "$\alpha = 1/137.035999084...$" *as an exact equality over* $\mathbb{R}$ would fail AF (no finite measurement can verify infinite decimal precision). But "α lies in the interval $[\alpha_0 - \delta, \alpha_0 + \delta]$" passes AF, because it partitions outcomes into a finite distinction. This is precisely how metrology works: physical measurements produce interval estimates, not exact real numbers. The real-number continuum is a representational convenience; AF attaches physical content to finite-precision outcome equivalence classes, not to the ideal continuum itself.

3.5.2 Paradigm Exclusions: What AF Rejects

AF must also correctly classify sentences that physics treats as lacking physical content — syntactically well-formed but physically redundant. Consider:

- **Gauge potentials vs field strengths.** In electrodynamics, the vector potential A_μ is gauge-dependent: different choices of A_μ correspond to the same physical situation. The field strength $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ is gauge-invariant. Statements about specific values of A_μ fail AF: no admissible measurement can partition outcomes according to A_μ, because gauge-equivalent potentials produce identical observations. Statements about $F_{\mu\nu}$ pass AF.
- **Coordinate singularities.** The Schwarzschild metric in standard coordinates appears singular at $r = 2GM/c^2$. But this is a coordinate artefact — the spacetime is perfectly

regular there, as shown by transforming to Kruskal-Szekeres coordinates. The sentence "there is a physical singularity at $r = 2GM/c^2$" fails AF: no local measurement can distinguish the coordinate singularity from a smooth region. AF correctly classifies this as representational surplus.
- **Global phase of a quantum state.** The state $|\psi\rangle$ and $e^{i\theta}|\psi\rangle$ are physically indistinguishable: no measurement can partition outcomes according to global phase. Sentences attributing a specific global phase fail AF. Observable quantities — expectation values, transition probabilities — pass AF.

In every case, AF reproduces the judgement that physics already makes by other means (gauge invariance, coordinate independence, projective Hilbert space structure). This is strong evidence that AF is not an ad hoc construction but a formalisation of existing physical practice.

3.5.3 Invariance Under Representational Redundancy

The paradigm cases above reveal a unifying principle: *physical content is whatever is invariant under representational redundancy that preserves all admissible observational procedures.* Gauge transformations, coordinate changes, and global phase rotations are all representational redundancies — they change the formal description without changing any observable outcome.

We can state this as a robustness condition on AF:

(AF-Inv) A sentence φ has physical content only if its truth-value is invariant under all representational transformations that preserve the system's admissible observational procedures.

AF-Inv is not an additional axiom; it is a consequence of Principles 1 and 2. If two representations R_1 and R_2 yield identical outcome partitions for all admissible tests, then any sentence whose truth-value differs between R_1 and R_2 cannot be operationally closed (Principle 1) and cannot correspond to a unique committed record (Principle 2). AF-Inv is therefore built into AF, not added to it.

A caveat is warranted. AF-Inv identifies physical content as what is invariant under representational redundancy — but which transformations count as "representational" rather than "physical"? In the paradigm cases above, this is settled by the mathematical structure of the theory: gauge groups, diffeomorphism invariance, and projective Hilbert space structure provide independent criteria for identifying redundancy. AF-Inv is therefore most rigorous when applied to theories whose redundancy structure is already mathematically characterised. For theories where the boundary between representational and physical is contested, AF-Inv provides a criterion (invariance under whatever preserves all admissible observations) but does not by itself resolve the contest. In practice, this is not a limitation for our argument: the cases where AF-Inv matters most (gauge theory, general relativity, quantum mechanics) are precisely the cases where the redundancy structure is well understood.

This invariance property further distinguishes AF from naive verificationism (which typically lacks a representational-invariance criterion) and connects it to the structural realist tradition in philosophy of physics.

3.5.4 Why AF Is Not Tuned to Gödel

The Admissibility Filter was introduced in the VERSF programme to address the general *surplus structure problem* in physics [5, 8]: the fact that our mathematical formalisms routinely contain more structure than the physical systems they describe. Gauge degrees of freedom, coordinate charts, global phases, and unphysical sectors of Hilbert space are all instances of surplus structure. AF provides a uniform criterion for identifying and quarantining this surplus.

Gödelian sentences are simply another instance of surplus structure. They arise because a theory T, in order to be mathematically powerful, embeds arithmetic that goes beyond what any physical system can ground. This excess arithmetic carries its own limitations (incompleteness), but those limitations pertain to the surplus, not to the physical content. AF was not designed to handle this case specifically; it handles it as a corollary of a general principle.

To put it sharply: a critic who rejects AF must explain how physics should distinguish gauge artefacts from observables, coordinate singularities from physical singularities, and global phases from measurable quantities — because AF is the formalisation of exactly those distinctions. Rejecting AF to save the physical relevance of Gödel sentences requires rejecting the conceptual apparatus that physics uses to identify its own content.

3.5.5 Relation to Verificationism

AF shares a family resemblance with verificationist criteria of meaning, and intellectual honesty requires acknowledging this lineage. But AF differs from classical verificationism in two key respects.

First, *AF is not a thesis about meaning in general.* It is an admissibility criterion for *physical content* — the subset of a theory's sentences that are intended to correspond to measurable or recordable facts about nature. AF makes no claim about the meaningfulness of mathematical, ethical, aesthetic, or metaphysical sentences. Gödel sentences are meaningful in the model-theoretic sense; AF's claim is that they lack *physical* content, not that they lack *all* content.

Second, *AF allows indirect operational closure.* A theoretical term need not be directly observable to have physical content. The electron, for example, is not directly observed in any single measurement, but it sits inside a stable inferential network (quantum electrodynamics) that yields precise, finite-precision predictions for admissible measurement outcomes. The electron's properties pass AF because the predictive network in which they participate cashes out in operationally closable claims. This is essentially a Ramsey-sentence or structural-realist move: theoretical terms acquire physical content through their role in generating observable predictions, not through direct verification.

What AF *does* exclude is sentences that cannot be connected to any admissible outcome partition, even indirectly — sentences whose truth conditions require distinctions that no finite physical system can instantiate. This is not too strict (it admits electrons, quarks, and dark matter via their predictive networks) and not too loose (it excludes gauge artefacts, coordinate singularities, and — as a corollary — Gödelian sentences about unbounded proof domains).

An important consequence: AF's classification is *theory-relative*. A sentence can pass AF under one physical theory and fail under another, because the set of admissible measurement procedures depends on the theory's structure. For example, "the electron has a definite position and momentum simultaneously" is syntactically well-formed and involves an entity (the electron) that passes AF in many other contexts. But quantum mechanics provides no admissible measurement procedure that simultaneously partitions outcomes according to both position and momentum at arbitrary precision — complementarity prevents it. The sentence fails AF relative to quantum mechanics, even though it would pass relative to classical mechanics. This theory-relativity is a feature, not a bug: it means L_phys is determined by the full structure of the physical theory, not by formal properties of sentences alone. The confinement of Gödel sentences to L(T) \ L_phys should therefore be read as relative to the physical theory that defines L_phys — which is exactly the right reading, since the question "does incompleteness constrain physics?" is always asked in the context of a specific physical framework.

The classic dilemma for verificationism — that the verification criterion cannot verify itself — does not arise for AF, because AF is not a self-applying criterion of meaning. It is a criterion for physical content within a theory, stated from outside the theory at the level of metatheoretic methodology. Its justification is not self-verification but classificatory adequacy: it correctly identifies what physics already counts as physical content, and it does so via principles (operational closure, fact-commitment, representational invariance) that are independently motivated by thermodynamics and information theory.

---

4. The Core Theorem: Gödelian Incompleteness Does Not Reach Physical Content

We can now state the central result precisely. We proceed in three steps: a formal definition, a bridging proposition that identifies the structural partition, and the main theorem.

4.1 Definition

Definition (Physical referent). A sentence φ has *physical referent* in a PAC-compliant system S if and only if φ passes the Admissibility Filter: there exists an admissible operational test M_φ satisfying Principles 1 and 2 of Section 3.4 — that is, M_φ terminates within S's admissibility budgets, partitions S's distinguishable state space into "φ" and "¬φ," and the outcome can be stabilised as an irreversible committed fact.

Sentences without physical referent — those whose truth conditions require distinctions that exceed the system's state space or commitment budget — are not "about" any physically realisable state of affairs. They are properties of the formalism, not of nature.

4.2 The Physical-Interpretation Restriction

The key structural move of this paper is not to attack any formal theory T, but to identify a partition within it. Let T be any formal theory intended to describe a PAC-compliant system S. The language of T, denoted L(T), will in general contain sentences that go beyond what S can physically ground — because T may embed Peano arithmetic as part of its mathematical scaffolding.

We define:

L_phys ⊂ L(T) : the set of sentences in T that pass the Admissibility Filter with respect to S.

L_phys is the *physically interpretable sublanguage* of T — the fragment whose truth conditions can be connected, via operational closure and fact-commitment, to distinguishable states of S.

Proposition (Physical-Interpretation Restriction). Let T be any formal theory describing a PAC-compliant system S, and let L_phys ⊂ L(T) be defined as above. Then L_phys cannot interpret unbounded arithmetic sufficient to construct Gödel sentences with physical referent. That is: any Gödel sentence constructible in T belongs to L(T) \ L_phys.

Proof sketch. Gödel's construction requires (a) unbounded successor, (b) total multiplication, (c) a proof predicate quantifying over all ℕ, and (d) a diagonal fixed-point over the full domain. Each of these requires truth conditions that range over an unbounded domain of natural numbers. But every sentence in L_phys, by definition, has truth conditions that partition only the finitely many distinguishable states of S (PAC-1), via procedures that terminate within S's entropy budget (PAC-2), producing committable records (Principle 2 of AF). No sentence in L_phys can have truth conditions ranging over an unbounded domain. Therefore the arithmetical resources required for Gödel's construction — unbounded successor, total multiplication, unbounded quantification, and diagonal self-reference — are available in L(T) but not in L_phys. Any Gödel sentence G constructed in T resides in L(T) \ L_phys. □

This proposition makes the paper's thesis model-theoretically precise. We are not claiming that T is defective, or that Gödel's theorems fail. We are claiming that Gödel's theorems operate exclusively in the complement of L_phys — the formal overhead of T that has no physical referent. The restriction is not merely cardinal (finite vs infinite), but structural: the interpretability strength of L_phys is bounded by the finite state space and commitment constraints of S, preventing the internal coding of syntax required for Gödel's fixed-point construction. The partition L(T) = L_phys ∪ (L(T) \ L_phys) cleanly separates physical content from formal scaffolding, and Gödelian incompleteness is confined entirely to the latter.

4.3 The Main Theorem

Theorem (Gödelian Non-Constraint for Physical Observables). Let S be any system satisfying the Physical Admissibility Constraints:

- (PAC-1) Finite distinguishability: S has at most 2^N distinguishable states for some finite N.
- (PAC-2) Irreversible commitment: Each distinguishability operation in S requires entropy expenditure bounded below by $\varepsilon > 0$.
- (PAC-3) Bounded recursion: S can instantiate at most d_{max} levels of self-referential encoding.

Let T be any formal theory describing S, and let $L_{phys} \subset L(T)$ be the physically interpretable sublanguage (sentences passing the Admissibility Filter). Then:

(a) L_phys cannot support the unbounded arithmetisation of syntax required for Gödel's construction.

(b) Gödel's incompleteness theorems do not yield undecidable propositions in L_phys.

(c) Any undecidable sentence produced by Gödel's theorems in T belongs to $L(T) \setminus L_{phys}$ and lacks physical referent.

Proof.

(i) Unbounded successor is not available in L_phys.

By PAC-1, S can represent at most M distinct values, where $M \leq 2^N$. Any sentence in L_phys has truth conditions that partition S's finite state space. The successor axiom of PA:

$$\forall n \, \exists m \, (m = n + 1)$$

requires truth conditions ranging over an unbounded domain. No sentence with such truth conditions can pass the Admissibility Filter, because no admissible operational test can partition S's states according to a distinction (the existence of a successor for every n) that exceeds S's representational capacity. Therefore no physically interpretable fragment of any theory about S can instantiate the unbounded successor structure required by PA. □ (for this sub-claim)

(ii) Total multiplication is not available in L_phys.

PA requires closure under multiplication: for all $m, n \in \mathbb{N}$, the product $m \times n$ exists. In the physically interpretable domain, if $m \times n > M$, the product has no corresponding distinguishable state in S. Multiplication, interpreted physically, is a partial function — undefined when its outputs exceed the representable domain. Since Gödel's encoding of proof predicates relies on the totality of multiplication (products of prime powers encode sequences), the encoding cannot be grounded in L_phys. □

(iii) Unbounded quantification is not available in L_phys.

Gödel's undecidable sentence G asserts:

$\neg\exists y\ (\text{Proof\_F}(y, \ulcorner G \urcorner))$

The quantifier ∃y ranges over all natural numbers. In L_phys, every quantifier is implicitly bounded by S's representational capacity: truth conditions can range only over distinctions that S can physically instantiate. The sentence G, interpreted in L_phys, would reduce to:

$\neg\exists y \in \{0, \dots, M\}\ (\text{Proof\_F}(y, \ulcorner G \urcorner))$

But this bounded sentence is not the Gödel sentence — it is a different, weaker claim. The actual Gödel sentence, with its unbounded quantifier, has truth conditions that require surveying an infinite domain of proof objects. These truth conditions cannot be operationally closed (Principle 1: no finite test can settle a claim whose domain exceeds S's state space) and cannot be committed as fact (Principle 2: the required distinctions span an unbounded, non-instantiable range). The Gödel sentence therefore belongs to L(T) \ L_phys. □

(iv) The diagonal lemma cannot be grounded in L_phys.

The diagonal lemma requires the substitution function sub($\ulcorner \varphi(x) \urcorner$, n) to be total: for every formula φ(x) and every number n, it must return the Gödel number of $\varphi(\bar{n})$. In the physically interpretable domain, when $\ulcorner \varphi(\bar{n}) \urcorner > M$, the function has no corresponding output in S's state space. The fixed-point construction — which requires this function to be total over the entire domain — cannot be grounded in L_phys, because the self-referential sentence it produces may have a Gödel number exceeding S's representational capacity. □

(v) The semantic gap is irreducible.

Parts (i)–(iv) establish that L_phys lacks the arithmetical resources for Gödel's construction. This is not a contingent limitation that might be overcome by cleverness in encoding. It is a structural consequence of the Admissibility Filter: any sentence whose truth conditions require unbounded quantification, total functions over an infinite domain, or self-referential encoding exceeding S's capacity is, by definition, outside L_phys. The gap between L_phys and L(T) is precisely the gap between physically groundable claims and formal overhead, and Gödel's construction lives entirely in the latter.

A note on time-dependence: the bound M on S's representational capacity is itself a physical quantity that can, in principle, change over time as the system's coherence horizon expands (see Section 5.2). This means L_phys is not a fixed set but a time-dependent one — L_phys(t) at time t. The proof holds at every time: L_phys(t) is finite at every t, and a growing finite set is still finite at every point. Gödel's construction requires *unbounded* resources, not merely *large* ones, so no finite expansion of L_phys ever brings Gödel sentences within reach. The gap is irreducible not because L_phys is static, but because it is always finite while Gödel's requirements are always infinite.

Moreover, by PAC-2, even the bounded approximation of G (the sentence with quantifier restricted to {0, …, M}) may exceed S's entropy budget for verification. Such a sentence is logically decidable relative to the bounded domain but may be physically inadmissible — an admissibility boundary limitation rather than a logical one. □

Since Gödel's construction requires arithmetical resources available only in L(T) \ L_phys, the first incompleteness theorem yields undecidable sentences only in the non-physical fragment of T. The second theorem, being a strengthening of the first with shared preconditions, likewise produces no physically referential undecidable propositions. ■

Remark. This result is not a loophole or a technicality. It is the recognition that Gödel's standard construction targets features — unbounded arithmetisation and unbounded proof quantification — that the physically interpretable sublanguage of any theory about a PAC-compliant system cannot support. The mathematical observation that certain bounded arithmetic theories do not support Gödel's standard construction is well established [3, 4], though the full landscape of incompleteness phenomena in weak systems is richer than a simple "finite vs infinite" dichotomy. What the present argument adds is the physical and semantic claim: the constraints that force finitude are not a modelling choice but thermodynamic necessities, and the L_phys / L(T) \ L_phys partition makes the confinement of incompleteness to formal overhead structurally explicit.

Scope clarification. Nothing in this argument denies that a formal theory describing S may syntactically contain Peano arithmetic and hence Gödel sentences. The claim is narrower and precisely stated: L_phys, the physically interpretable sublanguage of such a theory — the fragment connected by the Admissibility Filter to measurable distinctions — cannot support Gödel's unbounded arithmetisation of syntax. Incompleteness remains a property of the formal scaffolding; it does not propagate into the theory's physically referential core.

---

5. What Replaces Gödelian Incompleteness?

If Gödel's theorems do not apply to physically admissible systems, what does constrain them? The answer, developed across the VERSF programme, is that physical systems face a different and more fundamental kind of limitation: admissibility boundaries.

5.1 The Finite Commitment Boundary

In Finite Commitment Mathematics (FCM) [6], every mathematical operation has a cost. A proof is not a timeless abstract object — it is a sequence of distinguishability operations, each requiring entropy expenditure. The "theorems" of a physical system are those statements whose proofs fit within the system's entropy budget.

This means that some true statements are inaccessible — not because they are logically unprovable (in the Gödelian sense), but because their proofs exceed the system's physical resources. This is a resource limitation, not a logical one, and it is qualitatively different from incompleteness:

	Gödelian Incompleteness	Admissibility Boundary
Source	Self-referential logical structure	Finite physical resources
Nature	Permanent, structural	Resource-dependent
Domain	Infinite formal systems	Finite physical systems
Character	∃ true unprovable statement	Some proofs cost more than the system can pay
Resolution	Cannot be resolved within the system	May be resolved by expanding entropy throughput

5.2 The Coherence Horizon

A physically admissible system has a **coherence horizon**: the boundary of what it can currently distinguish, verify, and commit to fact. Beyond this horizon, statements are not "undecidable" in the Gödelian sense — they are beyond the system's current resolution capacity.

As the system processes more entropy (in cosmological terms: as the universe ages and its total entropy throughput increases), the coherence horizon can expand. Statements previously beyond reach may become accessible. This captures, in physically precise language, an intuition that has recurred throughout the history of mathematics: that truth unfolds over time, and that what is provable depends on the resources available for proof.

But the coherence horizon is not guaranteed to recede indefinitely. As the system's expressive capacity grows, new statements also become formulable. The frontier expands, but the territory beyond it does not shrink to nothing. There is always a boundary — and for some statements, the entropy cost of proof may exceed the total entropy throughput remaining in the system's future. Such statements are not logically undecidable. They are *cosmologically inaccessible* — a limitation that is, in any operationally meaningful sense, permanent, even though it is in principle dynamic.

This is an important nuance. The VERSF programme does not naively promise that all truths will eventually become accessible. It distinguishes between:

- **Gödelian undecidability:** structural, permanent, independent of resources.
- **Admissibility inaccessibility:** resource-dependent, in principle dynamic, but potentially permanent in practice for systems with bounded future entropy.

The difference is one of kind, not merely of degree. The first is a theorem about logic. The second is a consequence of physics.

5.3 Consistency as a Monitorable Condition

Gödel's second theorem states that a sufficiently powerful consistent system cannot prove its own consistency. But a finite system is logically capable — in principle — of verifying its own consistency, because the space of derivations is finite (though potentially physically infeasible to search exhaustively). In PAF terms, consistency is not an abstract logical property but a structural integrity condition: the system's committed facts do not contradict each other within the current distinguishability resolution.

However, "in principle" must be qualified by the system's own constraints. The space of derivations, even in a finite system, can be combinatorially vast. For a system with M available states and rules of inference that allow derivation chains of length L, the number of candidate derivations can scale as M^L — a number that may itself exceed the system's entropy budget for verification.

Thus, self-consistency verification becomes an *admissibility boundary problem* rather than a logical impossibility. A physically admissible system can monitor its coherence incrementally — checking local consistency as facts are committed — even if global exhaustive verification exceeds its resources. This is precisely how physical systems maintain structural integrity in practice: not by surveying all possible contradictions, but by enforcing consistency constraints at the point of commitment.

This transforms the philosophical import of Gödel's second theorem. The question is not "can the system prove its own consistency?" (a question about logical power) but "can the system maintain its own consistency?" (a question about physical architecture). The answer, for PAF-compliant systems, is yes — through the irreversibility of commitment, which ensures that once a fact is produced, it cannot be silently altered to create a contradiction.

---

6. The Relationship to Classical Mathematics

6.1 Why Classical Mathematics Works

The success of classical mathematics is not threatened by this analysis. Classical mathematics, including Peano arithmetic and its Gödelian limitations, remains a powerful and internally consistent framework. Infinite mathematics works spectacularly well in physics because the finite systems it approximates are very large — so large that the boundary effects of finitude are negligible in almost all practical calculations.

The situation is precisely analogous to fluid dynamics. The Navier-Stokes equations treat water as a continuous medium, and they are extraordinarily effective. But water is made of discrete molecules. The continuous model works because the discretisation scale (molecular spacing ≈ 10^{-10} m) is vastly smaller than the scales of interest (≥ 10^{-6} m for typical flows). The model is an idealisation — a brilliant and useful one — but it does not imply that water is truly continuous.

Similarly, Peano arithmetic idealises the structure of counting by extending it to infinity. This idealisation is effective because the bound on physically realisable numbers (≈ $2^{(10^{122})}$ for the observable universe) is so large that finite-size effects are invisible at the scales where we normally do mathematics. But the idealisation does not imply that nature is truly infinite, any more than the Navier-Stokes equations imply that water is truly continuous.

6.2 Convergence of Bounded and Unbounded Arithmetic

For any fixed computation that terminates within the bounds of a finite system, bounded arithmetic and full PA give identical results. The two diverge only at the boundary — when computations approach or exceed the system's capacity. In the language of limits:

For any PA-theorem φ involving only numbers less than M and proofs of length less than L, there exists a bounded system S(M, L) in which φ is also a theorem.

This is a well-known conservation result in the theory of bounded arithmetic (cf. Parikh [23]; see also [3, 4] for systematic treatments). It means that the vast body of mathematical results used in physics — results about differential equations, group theory, probability, and analysis — survive intact in the bounded setting, provided only that the numbers and proof lengths involved remain within the (astronomically large) physical bounds. Gödel's theorems, which specifically concern statements that exist only at the boundary of infinite self-reference, are precisely the results that do *not* survive.

6.3 The Map and the Territory

One might argue: even if the physical system is finite, our *reasoning* about the system uses full Peano arithmetic, and Gödel's theorems apply to our reasoning. This is true — but it is a statement about the mathematical framework we choose, not about nature.

We are free to reason about finite systems using infinite mathematics, and the limitations we discover in our reasoning tools (including Gödelian incompleteness) are limitations of the *map*, not of the *territory*. A cartographer's inability to draw a perfectly accurate map of England does not imply that England contains unmappable regions. The Gödelian limitations of our mathematical tools do not imply Gödelian limitations of the physical systems those tools describe.

This distinction extends to the question of nonstandard models. A theory about a physical system, formalised in full PA, may admit nonstandard models containing "numbers" larger than any physically realisable quantity. These nonstandard elements are artefacts of the theory's excess expressive power — ghosts of the infinite idealisation. They do not correspond to any physical state and impose no physical constraint.

---

7. Addressing Objections

7.1 "But we use infinite mathematics in physics."

We do — as an approximation. The real number continuum, infinite-dimensional Hilbert spaces, and unbounded arithmetic are modelling tools. They are extraordinarily effective because the finite systems they approximate are very large. But the map is not the territory. The effectiveness of infinite mathematics in physics does not imply that nature *is* infinite, any more than the effectiveness of continuous fluid dynamics implies that water is not made of atoms.

The relevant question is not "do we use infinite mathematics?" but "does nature require infinite mathematics?" The VERSF programme answers: no. Every physical prediction that has ever been confirmed experimentally involves only finite quantities, finite precision, and finite computation. The infinities of our mathematical models are scaffolding, not structure.

7.2 "Gödel's theorems apply to any system that can do arithmetic, even bounded arithmetic."

This conflates bounded and unbounded arithmetic. Gödel's standard construction specifically requires the ability to represent *all* natural numbers and to quantify over them. Bounded arithmetic — arithmetic with an upper limit — is a well-studied subject in mathematical logic [3, 4], and it is known that sufficiently weak systems of bounded arithmetic do not support Gödel's standard construction of undecidable sentences in the same way, and in any case such sentences fail the Admissibility Filter when interpreted physically.

The critical technical point is this: Gödel's first theorem requires the system to represent all primitive recursive functions as *total* functions. In bounded arithmetic, multiplication and exponentiation become partial — undefined when their outputs exceed the bound. The proof predicate Proof_F(y, x), which must be a total primitive recursive function for Gödel's proof to work, becomes partial or undefined for large arguments. The self-referential construction collapses.

The theory of feasible arithmetic and the study of sub-Peano systems confirm that unbounded arithmetisation of syntax is not a dispensable convenience but a necessary precondition for Gödel's standard incompleteness construction.

7.2.1 Bounded Independence Does Not Automatically Imply Physical Incompleteness

A sophisticated critic may press further: incompleteness-like phenomena do occur in weak and bounded arithmetics. For example, the theory $I\Delta_0$ (arithmetic with induction restricted to bounded formulas) cannot prove the totality of the exponential function. Paris and Harrington [17] showed that a strengthened finite Ramsey theorem is independent of PA despite being "finitary" in flavour. Are these not counterexamples?

They are not, but the reason is instructive. The existence of a syntactically bounded independent sentence does not automatically imply the existence of a physically referential independent sentence. The bridge from syntactic form to physical content requires an explicit mapping via the Admissibility Filter — and this mapping is not automatic.

Consider the $I\Delta_0$ independence result: "exponentiation is total" is a statement about the representability of a function *within a formal system*. It is a claim about proof-theoretic strength, not about any physical observable. To become physically referential, it would need an explicit interpretation: a physical system S, an admissible measurement procedure, and a demonstration that the measurement outcome partitions according to "exp is total" vs "exp is not total." No such interpretation is forthcoming, because the statement is about the internal architecture of a formal calculus, not about the behaviour of any physical system.

More generally, independence results in bounded arithmetic typically concern:

- Totality of rapidly growing functions (proof-theoretic growth rates),
- Reflection principles (a system's ability to assert its own soundness),
- Combinatorial statements whose "natural" formulation requires quantification at or beyond the system's proof-theoretic ordinal.

These are statements about the structural capabilities of formal systems — they live, in the L_phys / L(T) partition, squarely in L(T) \ L_phys unless someone constructs an explicit AF-satisfying interpretation. We do not claim this is impossible in all cases. We claim that the burden of proof is on the critic: *show* the AF-satisfying measurement procedure, or the independence result remains a property of the formalism.

The Paris-Harrington case deserves particular candour, because it is the hardest for our framework. Unlike $I\Delta_0$ totality-of-exponentiation, which is transparently about proof-theoretic strength, the Paris-Harrington theorem is a statement about *finite combinatorial structures* — it asserts the existence of certain finite colourings with specific properties. This makes the AF burden harder to dismiss: one might imagine a physical system whose behaviour corresponds to a finite Ramsey-type combinatorial structure. We maintain, however, that the independence of the Paris-Harrington statement from PA is driven by the *growth rate* of the required witnesses (which exceeds any primitive recursive function), not by the finiteness of the objects involved. The statement's independence reflects a mismatch between the combinatorial complexity of the witnesses and the proof-theoretic strength of PA — a formal-system limitation — not a physical fact that resists measurement. Constructing an AF-satisfying interpretation would require a physical system whose observable behaviour partitions according to the existence of a specific Ramsey-type colouring, with the critical feature being that no PA-provable bound suffices. We are not aware of any such physical interpretation, but we flag this as the point where the AF burden is least trivially discharged.

If such an interpretation were provided — if a bounded independent sentence could be shown to have genuine physical referent via AF — then the resulting limitation would be an *admissibility boundary* (the proof exceeds the system's entropy budget), not a Gödelian incompleteness (a truth with no proof in principle). The distinction between "too expensive to prove" and "logically unprovable" would remain.

7.3 "This just pushes the problem to the meta-level."

The meta-level objection deserves careful treatment. It runs: "Even if the physical system is finite, any *theory* we construct about it — if it's powerful enough to be useful — will express enough arithmetic for Gödel to apply. So incompleteness returns at the level of theory."

We accept the premise but not the conclusion. Our claim concerns physical constraint: whether incompleteness produces undecidable propositions with physical referent — that is, propositions that pass the Admissibility Filter of Section 3.4. The theory may indeed be subject to Gödelian limitations. But those limitations are properties of the theory, not of the physical system the theory describes. Physics does not ask: "What can our theory prove?" It asks: "What does nature do?" If the physical system is finite and its behaviour is determined by finite information, then the limitations of an infinite theory about that system are artefacts of the theory's excess power.

To put it concretely: suppose a finite physical system S has some definite behaviour B. A Gödelian theory T about S might contain a statement about B that is true but unprovable in T. This is a limitation of T, not a limitation of S. The physical system has already "decided" B — it has produced a definite fact via irreversible commitment. The theory's inability to prove that fact is a failure of the map, not an inscrutability of the territory.

We can make this precise, using the L_phys / L(T) partition from Section 4.2:

Lemma (Finite-Semantics Lemma). Let T be a formal theory describing a PAC-compliant system S, and let L_phys ⊂ L(T) be the physically interpretable sublanguage. Every sentence in L_phys has truth conditions ranging only over distinctions realisable within S's finite state space and commitment budget. Gödel sentences — whose truth conditions quantify over unbounded proof objects — therefore belong to L(T) \ L_phys and have no physical referent.

Proof sketch. By the Admissibility Filter, every sentence φ ∈ L_phys must admit operational closure (a finite test that terminates within S's budgets) and fact-commitment (an irreversible record stabilising its truth-value). Since S's state space is finite (PAC-1), the operational test partitions only finitely many distinguishable outcomes. Since each commitment step costs entropy (PAC-2), the total verification is bounded. The truth conditions of φ therefore involve only finitely many distinctions. But a Gödel sentence G requires ¬∃y(Proof_F(y, ⌜ G ⌝)) where y ranges over all ℕ — an unbounded domain that cannot be instantiated, encoded, or operationally closed under PAC constraints. The truth conditions of G quantify over proof objects that exceed S's representational capacity (failure of Principle 1: its truth conditions range over a domain exceeding S's finite state space). The resulting truth-value cannot be stabilised as an irreversible record, because the required distinctions span an unbounded range not realisable in S (failure of Principle 2). Therefore G ∉ L_phys; it belongs to L(T) \ L_phys and lacks physical referent. □

This lemma converts the intuition — "Gödel lives in the overhead" — into a structural claim grounded in the L_phys partition. With it in hand, we can say something sharper. The specific statements that Gödel's theorem renders undecidable are precisely those in L(T) \ L_phys — the sentences without physical referent. A Gödelian sentence G in theory T asserts the non-existence of a proof — quantifying universally over all natural numbers, including numbers vastly exceeding the information capacity of any physical system. This universal quantifier ranges over a domain that no bounded system can instantiate. The sentence G is therefore not *about* any physically distinguishable state of affairs. It is a statement about the infinite overhead of the formalism — the scaffolding, not the structure it supports. Any statement whose truth value depends on quantifying over the full, unbounded natural numbers is, by the finite distinguishability constraint, a statement without physical content: there is no configuration of any physical system that could correspond to its truth or falsity. Conversely, any statement that *does* describe a physically distinguishable outcome — a measurement result, a committed fact, a resolved bit — lives in L_phys, where bounded quantification suffices and Gödel's construction cannot gain purchase. The incompleteness theorems apply to T, but the sentences they render undecidable are confined to L(T) \ L_phys — the formal overhead, not the physical core.

7.3.1 Cross-Boundary Derivations and Deductive Adequacy

A further sharpening of the meta-level objection deserves attention. Even granting the L_phys / L(T) \ L_phys partition, a critic might ask: what if a physically referential proposition φ ∈ L_phys can only be *derived* via lemmas that pass through L(T) \ L_phys? In that case, Gödelian incompleteness in the scaffolding could block the derivation of a physically meaningful result. The incompleteness would not produce an undecidable physical sentence directly, but it could obstruct the *proof pathway* to a decidable one.

This is a genuine conceptual possibility, and we address it as follows. When a derivation of a physically referential claim requires passage through non-physical surplus structure, the obstruction is a property of the chosen formalisation, not of the physical proposition's referent. Physics already treats such situations as representational pathology: we seek gauge-invariant, coordinate-free, and manifestly covariant reformulations precisely to avoid dependence on non-physical intermediaries. The derivation of the perihelion precession of Mercury does not require a specific coordinate chart; if a coordinate choice creates a spurious obstruction, we change coordinates.

We therefore adopt the following as a methodological requirement for physically adequate theories:

Physical Derivability Principle (PDP). For any physically referential proposition φ ∈ L_phys, a physically adequate formalisation of T should admit a derivation of φ (or its negation) in which all intermediate lemmas can be chosen within L_phys or within a conservative extension that introduces no new physical commitments.

PDP is not a theorem; it is a conjecture about the structure of physically adequate theories — one motivated by strong precedent but not yet established as a general principle. Gauge-invariant reformulations, coordinate-free derivations, and manifestly covariant formulations all instantiate PDP in specific domains, achieved through substantial mathematical work (fiber bundles, BRST cohomology, effective field theory). But there is no guarantee that an analogous "physically grounded derivation" exists for every physically meaningful claim in every domain. If a particular theory T entangles physical claims with surplus structure so deeply that no L_phys-internal derivation exists, PDP says the appropriate response is to seek a better formalisation — but we acknowledge this may be a programme rather than a fait accompli.

The upshot: even if cross-boundary derivations exist in some formalisations, this is — on the evidence of physical practice to date — a representational artefact, not a physical limitation. Gödelian incompleteness can obstruct proof pathways that pass through L(T) \ L_phys, but PDP requires that physically meaningful claims should have physically grounded derivations. Where such derivations have been sought, they have been found. Whether this pattern extends universally remains an open question — one that we flag as a substantive conjecture rather than a settled principle.

7.4 "This is just ultrafinitism rebranded."

There is genuine affinity between the position defended here and the ultrafinitist programme [12, 13]. Both reject the actual infinite as a feature of mathematical reality. The difference is one of grounding.

Ultrafinitism is typically motivated by philosophical scepticism about infinite objects: the claim that we have no epistemic access to completed infinities, and therefore no warrant for asserting their existence. This position has been philosophically productive but has struggled to gain traction in mainstream mathematics, partly because the motivation is epistemological rather than structural.

The VERSF programme provides a different and, we argue, stronger foundation. The claim is not "we cannot know the infinite" but "nature cannot *instantiate* the infinite." This is a physical claim, grounded in thermodynamics, information theory, and the Bekenstein bound. It does not require philosophical scepticism about infinite objects — one can accept that Peano arithmetic is a perfectly legitimate mathematical structure while maintaining that it describes more structure than any physical system provides.

This shifts the argument from epistemology to ontology. The question is not whether we can think about infinity (we obviously can, and fruitfully do). The question is whether the physical systems we inhabit and study are members of the class of structures to which Gödel's theorems apply. The answer, given the constraints of physical admissibility, is no.

7.5 "What about the universe as a whole?"

Even the observable universe has a finite information capacity. Current estimates place the entropy of the observable universe at approximately 10^{122} bits [14], dominated by the entropy of supermassive black holes. This means the universe can represent at most $2^{(10^{122})}$ distinct states.

This is an unimaginably large number, but it is finite. The entire observable universe is a bounded system. Its total capacity for representing numbers, encoding formulas, and constructing proofs is finite. Gödel's theorems, which require quantification over an infinite domain of proof objects, do not apply to it.

One might speculate about the universe beyond the observable horizon, or about a hypothetical infinite multiverse. But such speculations do not rescue the applicability of Gödel's theorems to physics, for two reasons. First, any region causally accessible to a finite observer has finite information content. Second, even a spatially infinite universe composed of finite-density matter has, in any bounded region, finite entropy — and it is bounded regions, not infinite totalities, that produce physical facts.

8. Implications

8.1 For the Foundations of Mathematics

This argument does not invalidate mathematics. Classical mathematics, including Peano arithmetic and its Gödelian limitations, remains a powerful and internally consistent framework. But it is a framework that describes more structure than nature provides. The natural numbers, in their full Peano glory, are an idealisation — a useful one, but an idealisation nonetheless.

The mathematics that describes physical reality is not Peano arithmetic. It is something more constrained: a finite, commitment-bound, resolution-limited calculus. Gödel's theorems are theorems about the idealisation, not about the physics. Recognising this does not diminish mathematics — it clarifies the relationship between mathematical structure and physical reality.

8.2 For Physics

If the mathematical framework describing nature is not subject to Gödelian incompleteness, then the limitations physics faces are admissibility boundaries — resource constraints — not logical barriers. Some of these boundaries may be, in practice, permanent: the entropy cost of certain computations may exceed the total remaining thermodynamic budget of the universe. But the nature of the limitation is different. It is a constraint of cost, not of logic. It is the price of being physical.

This is a more nuanced picture than naive optimism. We do not claim that all physical truths will eventually be discovered. We claim that the *reason* some truths may remain inaccessible is thermodynamic, not logical — and that this distinction matters for how we understand the structure of scientific knowledge.

8.3 For Computation

The relationship between computational complexity and physical admissibility becomes clearer in this framework. Computational intractability (e.g., the conjectured $P \neq NP$) is not a Gödelian phenomenon — it is a resource phenomenon. Some computations require more entropy expenditure than is physically available, not because of self-referential paradox, but because of combinatorial scaling against finite thermodynamic budgets.

This connects naturally to the VERSF treatment of P vs NP as a physical law rather than a purely mathematical conjecture [7]: the boundary between tractable and intractable computation is, at bottom, a boundary between physically admissible and physically inadmissible entropy expenditure.

9. Conclusion

Gödel's incompleteness theorems are correct, elegant, and profound — within their domain. That domain is the class of consistent, effectively axiomatised formal systems powerful enough to express the full arithmetic of the natural numbers.

Physical reality is not in that class.

Any system constrained by finite distinguishability, irreversible commitment, and bounded recursion depth — which is to say, any physically admissible system — cannot express full Peano arithmetic. It falls below the threshold where Gödel's machinery engages. And even where the formal theories we use to *describe* such systems do engage Gödel's theorems, the undecidable sentences those theorems produce are confined to the theories' infinite overhead — they have no physical referent and constrain no physical observable.

This does not diminish Gödel's achievement. It contextualises it. Gödelian incompleteness is a property of a mathematical idealisation — the infinite. Nature, built from finite distinctions and irreversible commitments, operates under a different regime. Its limitations are real but they are admissibility boundaries, not logical barriers. They are dynamic, not permanent. They are the cost of being physical.

The incompleteness of incompleteness is this: Gödel proved that infinite systems have inherent limits. But he could not have proven — and it is not the case — that reality is infinite.

References

[1] Gödel, K. (1931). Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I. *Monatshefte für Mathematik und Physik*, 38, 173–198.

[2] Bekenstein, J. D. (1981). Universal upper bound on the entropy-to-energy ratio for bounded systems. *Physical Review D*, 23(2), 287.

[3] Buss, S. R. (1986). *Bounded Arithmetic*. Bibliopolis.

[4] Pudlák, P. (2013). *Logical Foundations of Mathematics and Computational Complexity*. Springer.

[5] Taylor, K. (2025). The Physical Admissibility Framework (PAF): Finite Distinguishability, Irreversible Commitment, and the Cost of Facts. VERSF Theoretical Physics Program.

[6] Taylor, K. (2026). Finite Commitment Mathematics (FCM). VERSF Theoretical Physics Program.

[7] Taylor, K. (2026). Why P vs NP May Be a Law of Physics, Not Just Mathematics. VERSF Theoretical Physics Program.

[8] Taylor, K. (2025). Finite Distinguishability and Irreversible Commitment: A No-Go Result for Infinite-Precision Physics. VERSF Theoretical Physics Program.

[9] Taylor, K. (2025). Introducing the Ticks-Per-Bit Framework. VERSF Theoretical Physics Program.

[10] Nagel, E. & Newman, J. R. (1958). *Gödel's Proof*. NYU Press.

[11] Penrose, R. (1989). *The Emperor's New Mind*. Oxford University Press.

[12] Nelson, E. (1986). *Predicative Arithmetic*. Princeton University Press.

[13] Sazonov, V. (1995). On feasible numbers. In *Logic and Computational Complexity*, Springer, 30–51.

[14] Egan, C. A. & Lineweaver, C. H. (2010). A larger estimate of the entropy of the universe. *The Astrophysical Journal*, 710(2), 1825.

[15] Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3), 183–191.

[16] Bennett, C. H. (1973). Logical reversibility of computation. *IBM Journal of Research and Development*, 17(6), 525–532.

[17] Paris, J. & Harrington, L. (1977). A mathematical incompleteness in Peano arithmetic. In *Handbook of Mathematical Logic*, North-Holland, 1133–1142.

[18] Franzén, T. (2005). *Gödel's Theorem: An Incomplete Guide to Its Use and Abuse*. A K Peters.

[19] Feferman, S. (1998). *In the Light of Logic*. Oxford University Press.

[20] Lucas, J. R. (1961). Minds, machines and Gödel. *Philosophy*, 36(137), 112–127.

[21] Shapiro, S. (1998). Incompleteness, mechanism, and optimism. *Bulletin of Symbolic Logic*, 4(3), 273–302.

[22] Aaronson, S. (2005). NP-complete problems and physical reality. *ACM SIGACT News*, 36(1), 30–52.

[23] Parikh, R. (1971). Existence and feasibility in arithmetic. *Journal of Symbolic Logic*, 36(3), 494–508.