Physical Admissibility: A Constraint-Based Foundation for Physics

Finite Distinguishability, Irreversible Commitment, and the Architecture of Physical Law

---

Table of Contents

Abstract

We develop a minimal admissibility framework for physically grounded mathematics. Beginning with a pre-physical process ontology (the tick layer), we introduce two universal operational constraints—finite distinguishability (FD) and irreversible commitment (IC)—that define the admissibility layer governing physical law. From these constraints alone, we derive exclusion theorems establishing necessary conditions on physically realizable theories: no unbounded information density, unavoidable thermodynamic costs for information erasure, and bounds on operationally accessible soft excitations. We demonstrate that the fundamental equations of physics—the Schrödinger equation, Heisenberg uncertainty principle, Maxwell's equations, Einstein field equations, and the second law of thermodynamics—are each consistent with this admissibility structure, with some structural features (unitarity, entropy monotonicity) enforced by admissibility alone and others (specific field equations) requiring additional symmetry inputs. Finally, we show that the tick layer provides a minimal process substrate for meta-computation, illustrating how "outside-theory" reasoning in the sense of Gödel can be operationally instantiated.

What distinguishes this work: Unlike prior operational reconstructions that aim to derive quantum theory from informational axioms, we isolate a pre-physical process substrate (tick layer) and a resource-relative admissibility layer (FD/IC over $\mathcal{O}$), yielding a constraint-complete "building code" that cleanly separates necessity constraints from dynamical selection.

Clarification on scope: By "theory of everything" in this work, we mean a complete specification of necessary admissibility constraints that any physical theory must satisfy—not a unique dynamical law that determines all physical phenomena. The result is a **completeness theorem for constraints, not a completeness theorem for dynamics.** The framework provides constraint-completeness at the level of necessary conditions while remaining agnostic about dynamical completeness.

Assumptions Ledger

Before proceeding, we explicitly list the minimal assumptions employed beyond the two core constraints (FD and IC):

1. **Physical Church-Turing Thesis** — Used in Theorem 4.1. Every physically realizable input-output procedure is Turing-simulable to the same operational accuracy. If violated by future physics, FD would require revision.
2. **Operational State Space Axioms** — Used in unitarity derivation. States form a convex set with continuous reversible transformations and pure states (standard operational quantum axioms; cf. Hardy 2001, Chiribella et al. 2011). These yield Hilbert space representation.
3. **Operational Distinguishability Measure** — Used throughout. The measure D is: (i) operationally definable via optimal discrimination error, (ii) contractive under physical

channels (data-processing inequality), and (iii) metrizes operational distinguishability. Examples: trace distance, Helstrom optimal error probability.
4. **Continuity and Composition** — Used for Stone's theorem. Time evolution forms a continuous one-parameter group: $T_s \circ T_t = T_{s+t}$.
5. **Basic Thermodynamic Consistency** — Used for Landauer's principle. Systems can exchange heat with thermal reservoirs at well-defined temperatures.

These assumptions are standard in operational approaches to physics. We flag them explicitly so that "you assumed things!" becomes "yes, and here they are."

For the General Reader: What This Paper Is About

Imagine you wanted to write down all the rules that any possible universe must follow—not the specific rules of *our* universe (like "gravity pulls things down" or "light travels at 299,792 km/s"), but the deeper rules that constrain what kinds of rules are even possible.

This paper argues that there are exactly two such meta-rules:

1. **You can't distinguish infinitely many things with finite effort.** No matter how clever you are, if you only have limited time, energy, and equipment, you can only tell apart a finite number of different states.
2. **Some changes can't be undone.** There exist physical processes that lose information permanently—you can't always rewind the tape.

From just these two constraints, we show that many features of physics we usually take as separate assumptions—like quantum mechanics' wave equation, the uncertainty principle, and the law that entropy always increases—actually *follow* as logical consequences.

The paper doesn't claim to explain *everything* about physics (why is the speed of light what it is? why do electrons have the mass they do?). But it does claim to identify the universal constraints that any physically possible law must satisfy.

Part I: Architectural Foundations

Chapter 1: The Tick Layer (Pre-Physical Process Ontology)

1.1 Motivation and Definition

The key idea in plain language: Before we can talk about space, time, energy, or any physical laws, we need something even more basic—the simple fact that *things happen*. The "tick layer" is our name for this most primitive level: a sequence of elementary changes, like the individual frames of a movie before we've decided what the movie is about.

To avoid category errors regarding fundamentality, it is essential to distinguish between a pre-physical substrate of change and the admissibility layer at which physical constraints properly reside. This distinction strengthens, rather than weakens, the constraint-based nature of the framework developed here.

The tick layer represents the most primitive level of description: discrete events of change. At this level there is no spacetime, no fields, no symmetry, and no notion of physical law. A **tick** is defined only as a minimal update event—an atomic transition—sufficient to induce a partial ordering among events.

Definition 1.1 (Tick): A tick is an elementary event of change, representing the minimal unit of transition from one configuration to another. Ticks carry no intrinsic duration, spatial extent, or energy—these concepts emerge only at higher levels of description.

For the general reader: Think of a tick as the most basic possible "something happened"—not *what* happened, not *how long* it took, not *where* it occurred, just the bare fact of change. It's like asking "what's the simplest thing a computer could do?" The answer is: flip a single bit from 0 to 1 (or vice versa). A tick is even more basic than that—it's just "a change occurred," period.

The tick layer provides a **process ontology**—the raw substrate of change—not the laws governing aggregates of events. It answers the question "what is the most primitive thing that happens?" without yet addressing "what rules govern happenings?"

1.2 Mathematical Structure

For the general reader: We need a way to say "event A might have influenced event B" without yet having a notion of space or time. Mathematicians have a tool for this called a "partially ordered set" or "poset"—it's a collection of things where some pairs are ordered (A comes before B) but not all pairs need to be comparable (maybe C and D happened "independently" with neither before the other).

Mathematically, the tick layer may be modeled as a locally finite event graph or causal poset.

Definition 1.2 (Causal Poset): A causal poset is a pair $(P, \preccurlyeq)$ where P is a set of primitive events (ticks) and $\preccurlyeq$ is a partial order satisfying:

1. **Reflexivity:** $e \preccurlyeq e$ for all $e \in P$
2. **Antisymmetry:** if $e \preccurlyeq f$ and $f \preccurlyeq e$, then $e = f$
3. **Transitivity:** if $e \preccurlyeq f$ and $f \preccurlyeq g$, then $e \preccurlyeq g$

In plain language:

- Reflexivity: Every event is "before or simultaneous with" itself (trivially true)

- Antisymmetry: If A is before-or-equal-to B, and B is before-or-equal-to A, then A and B are actually the same event (no circular time)
- Transitivity: If A is before B, and B is before C, then A is before C (causation chains)

The partial order $\preccurlyeq$ represents potential causal influence: $e \preccurlyeq f$ means that event e could, in principle, influence event f.

Definition 1.3 (Local Finiteness): A causal poset $(P, \preccurlyeq)$ is locally finite if for any $e, f \in P$ with $e \preccurlyeq f$, the **causal interval**

$$[e, f] = \{g \in P : e \preccurlyeq g \preccurlyeq f\}$$

is a finite set.

For the general reader: "Local finiteness" means that between any two causally-connected events, only *finitely many* things can happen. You can't squeeze an infinite number of events between breakfast and lunch. This might seem obvious, but it's actually a substantive assumption that rules out certain exotic mathematical structures (like having infinitely many moments packed into a finite time interval).

The local finiteness condition is physically crucial: it ensures that between any two causally related events, only finitely many intermediate events occur. This prevents the existence of infinitely dense tick structures within bounded causal intervals—a first hint of the finite distinguishability constraint that will be formalized at the admissibility layer.

Remark 1.1: Local finiteness does not impose an upper bound on the total number of ticks, nor does it require discreteness of any emergent spacetime. It merely ensures that causal chains are countable and that no finite causal interval contains infinite structure.

1.3 What the Tick Layer Does Not Contain

Crucially, **no assumption of reversibility or irreversibility is imposed at the tick level**. The partial order $\preccurlyeq$ encodes potential causal influence but does not specify:

- Whether transitions can be "undone"
- Whether information is preserved or lost
- Any dynamics or evolution law
- Any notion of energy, momentum, or conservation
- Any spatial or temporal metric

For the general reader: The tick layer is deliberately "content-free"—it's like having a blank canvas before you've decided what to paint. We know that changes happen and that some changes can influence later changes, but we haven't yet said anything about *what* is changing, *how fast*, or whether you can reverse a change. All of that comes later, when we add constraints.

These concepts emerge only when ticks are aggregated under constraints at the admissibility layer.

Remark 1.2: The tick layer is deliberately minimal. One might object that such a sparse structure cannot ground rich physical phenomena. This objection mistakes the role of the tick layer: it provides the **ontological substrate** for change, while the **nomological content** (laws, constraints, regularities) resides at the admissibility layer.

Important note on incompleteness: The tick layer does not "escape" Gödelian limitations. Any sufficiently powerful formal description of the tick layer would itself be subject to incompleteness theorems. The tick layer provides a process substrate, not a privileged formal system.

1.4 Alternative Formalizations

While we have presented the tick layer as a causal poset, alternative mathematical structures could serve the same foundational role:

- **Causal sets (causets):** Locally finite posets with additional conditions ensuring Lorentz invariance in the continuum limit
- **Event graphs:** Directed acyclic graphs where edges represent direct causal links
- **Process theories:** Category-theoretic frameworks where morphisms represent transitions

The specific formalization matters less than the conceptual content: a pre-physical process ontology with partial causal ordering.

Chapter 2: The Admissibility Layer

2.1 Physical Grounding

The key idea in plain language: Not every mathematically possible system can actually exist in the physical world. A "physically grounded" mathematical model is one where every operation it describes could actually be performed by a real physical process—using finite time, finite energy, and finite resources. This rules out mathematical fantasies like "compute the exact value of pi to infinite decimal places" or "check infinitely many possibilities instantly."

The admissibility framework does not reside at the tick layer. Instead, it emerges one level above, at the **admissibility layer**, where aggregates of ticks are subject to operational constraints.

Definition 2.1 (Physical Grounding): A mathematical model is **physically grounded** if every operation it requires can be implemented by a finite physical procedure using finite energy, time, memory, and precision.

This definition excludes:

- Operations requiring infinite precision (exact real number computation)
- Operations requiring infinite time (non-halting procedures treated as completed)
- Operations requiring infinite energy (unbounded force applications)
- Operations requiring infinite memory (storing infinite data)

For the general reader: Think about what your computer can actually do. It has finite memory (maybe a terabyte), finite processing speed, and runs on finite power. Any calculation that would require infinite memory or infinite time isn't something your computer—or any physical device—can actually perform. Physical grounding is the requirement that our theories only describe operations that could, in principle, be carried out by some physical system.

Physical grounding is not a limitation but a **consistency requirement**: any mathematics claiming to describe the physical world must be implementable within that world.

2.2 The Operational Process Class

Before stating the core constraints, we define the class of operations they quantify over.

Definition 2.2 (Operational Process Class): Let $\mathcal{O}(R, E, \tau, M, \delta)$ denote the class of physically implementable operations within a specified finite resource budget:

- **R:** Spatial support (bounded region)
- **E:** Energy budget (maximum total energy)
- **τ:** Time budget (maximum duration)
- **M:** Memory budget (available storage)
- **δ:** Error tolerance (maximum acceptable error probability)

All subsequent definitions of FD and IC are stated relative to this operational class. This makes precise that our constraints concern what can be done with finite resources, not what is mathematically conceivable.

2.3 Core Constraints: Finite Distinguishability

Definition 2.3 (Finite Distinguishability — FD): For any fixed finite resource budget (R, E, τ, M) and fixed error tolerance δ, the maximal size of a set of mutually distinguishable states—states that can be reliably discriminated with error probability $< \delta$ using operations in $\mathcal{O}(R, E, \tau, M, \delta)$—is finite.

For the general reader: Imagine you have a box, and you're trying to figure out what state it's in. FD says that with limited time, limited equipment, and limited energy, you can only reliably tell apart a *finite* number of different possibilities.

Here's an everyday example: Suppose you're trying to identify colors. With your naked eye, you might distinguish a few million colors. With a good spectrometer, maybe billions. But no matter how good your equipment, if it's finite, you can only distinguish finitely many colors. You can't build a device that reliably distinguishes infinitely many different shades.

Critical clarification (capacity, not metaphysics): FD is a **capacity constraint**, not a metaphysical claim about the state space. If resources scale (more time, more energy, better equipment), the bound may increase. There is no claim that the underlying state space is intrinsically finite—only that operational access to it is bounded under any fixed finite resources. FD does not exclude arbitrarily fine distinctions in the limit of unbounded resources; it excludes their availability at fixed operational capacity.

For the general reader: This is an important subtlety. In quantum mechanics, physicists often work with mathematical structures (called Hilbert spaces) that have infinitely many dimensions—infinitely many possible states in principle. FD doesn't say this is wrong. It says that even if there are infinitely many *mathematical* states, only finitely many can be *operationally distinguished* with finite resources.

Analogy: The real number line has infinitely many points between 0 and 1. But if your ruler only has millimeter markings, you can only distinguish about 1000 positions. The infinite mathematical structure exists; your finite ability to probe it is limited. FD is about operational access, not metaphysics.

FD does **not** claim that Hilbert spaces are finite-dimensional in general. It claims that the operationally accessible, reliably distinguishable portion of any state space is finite under finite resources.

2.4 Core Constraints: Irreversible Commitment

Definition 2.4 (Irreversible Commitment — IC): There exist physical processes $\Phi \in \mathcal{O}$ such that for some initially distinguishable states $\rho_1 \neq \rho_2$:

$$D(\Phi(\rho_1), \Phi(\rho_2)) < D(\rho_1, \rho_2)$$

and for all recovery operations $\Psi \in \mathcal{O}(R, E, \tau, M, \delta)$:

$$D(\Psi(\Phi(\rho_1)), \Psi(\Phi(\rho_2))) < D(\rho_1, \rho_2)$$

For the general reader: Some things can't be undone. If you scramble an egg, you can't unscramble it. If you burn a book, you can't unburn it. IC is the formal statement that such irreversible processes genuinely exist—that there are physical operations where information is lost and cannot be recovered, no matter how clever you are (given finite resources).

This might seem obvious from everyday experience, but it's actually philosophically controversial. Some physicists have argued that at the fundamental level, all physical laws are reversible—that the appearance of irreversibility is just an illusion arising from our limited perspective. IC takes no position on this debate; it just says that *operationally*, with finite resources, some processes cannot be reversed.

In plain language: There exist processes that make two previously distinguishable states harder to tell apart, and no amount of subsequent processing (with finite resources) can fully restore the original distinguishability.

Critical clarification (operational, not ontological): IC is defined relative to the operational process class $\mathcal{O}$, not to the set of all mathematically imaginable operations. If the universe is microscopically unitary, a "recovery" map Ψ might exist in principle but require resources far exceeding any finite budget. IC makes no claim about microscopic ontology—only about what can be achieved operationally.

Agnosticism on origin: IC is stated as an operational constraint, compatible with multiple interpretations of its origin:

1. **Fundamental irreversibility:** Some processes genuinely destroy information at the most basic level
2. **Emergent irreversibility:** All microscopic dynamics are reversible, but coarse-graining over environmental degrees of freedom produces effective irreversibility
3. **Resource-bounded irreversibility:** Reversibility exists in principle but requires resources exceeding any finite budget

For the general reader: Physicists disagree about *why* things are irreversible. Some think information is truly destroyed. Others think it's just hidden in countless environmental particles where we can't practically recover it. Still others think it's technically recoverable but would require godlike resources. The admissibility framework doesn't take sides—it just notes that *operationally*, irreversibility is real.

The admissibility framework is agnostic among these interpretations. What matters operationally is that certain maps cannot be inverted with available resources—not whether this reflects fundamental physics or practical limitation.

Empirical support for IC:

- Second law of thermodynamics (coarse-grained entropy increase)
- Measurement as operational irreversibility (amplification + record formation), consistent with multiple interpretations of quantum mechanics
- Decoherence (environmental entanglement producing effective irreversibility)
- Landauer's principle (erasure requires dissipation)
- Horizon thermodynamics strongly suggests FD and IC; the status of fundamental information loss at black hole horizons remains debated

2.5 The Arrow of Time

For the general reader: Why does time seem to flow in one direction? Why do we remember the past but not the future? Why does a dropped glass shatter but shattered glasses never spontaneously reassemble? The "arrow of time" is physics' name for this asymmetry.

In our framework, the arrow of time *emerges* from IC: time's direction is defined by the direction in which irreversible commitments accumulate.

The monotone ordering induced by IC is identified with the **thermodynamic arrow of time**. This provides an operational definition of temporal direction:

Definition 2.5 (Temporal Order): Event A precedes event B (written $A < B$) if there exists a chain of irreversible commitments connecting A to B but not B to A.

This definition:

- Does not assume a background time parameter
- Derives temporal asymmetry from operational irreversibility
- Is compatible with relativity (different observers may disagree on simultaneity but agree on causal/commitment ordering)
- Explains why we remember the past but not the future (memory formation involves irreversible commitment)

Clarification: The "macroscopic time metric" emerging from counting irreversible records is a **candidate operational time parameter**, not necessarily *the* fundamental time. This framing is compatible with relativistic settings where different observers may define different operational clocks.

For the general reader: When you form a memory—say, of eating breakfast—that's an irreversible physical process. Neurons fire, proteins fold, information gets encoded. You can remember breakfast because that memory-formation was a one-way process. You can't remember lunch (which hasn't happened yet) because no irreversible commitment has connected you to it. Time's arrow *is* the direction of irreversible commitment.

2.6 Relation Between Layers

The tick layer and admissibility layer are related as follows:

Aspect	Tick Layer	Admissibility Layer
Content	Primitive events	Constraints on event aggregates
Structure	Causal poset	FD + IC relative to $\mathcal{O}$
Reversibility	Unspecified	Distinguished (IC defines operational irreversibility)
Distinguishability	Unspecified	Bounded (FD)
Time	Partial order only	Full temporal direction via IC
Role	Process ontology	Necessary conditions for physical law

For the general reader: Think of it this way:

- **Tick layer:** The raw material—individual changes, with some causal connections

- **Admissibility layer:** The building codes—rules that any structure made from ticks must follow
- **Effective physics:** The actual buildings—specific physical laws that obey the codes

A common misconception is that a complete physical framework must specify the deepest ontological layer. This is incorrect. A complete admissibility framework specifies **necessary conditions** that any physically realizable law must satisfy. The tick layer supplies the minimal process ontology; the admissibility layer supplies the constraint structure.

Important limitation: The admissibility framework provides **constraint-completeness at the level of necessary conditions**, not dynamical completeness. FD and IC constrain what theories are physically admissible; they do not uniquely determine which admissible theory is realized. Additional inputs (symmetries, initial conditions, perhaps contingent facts) are needed to select among admissible possibilities.

For the general reader: This is crucial. We're not claiming to derive all of physics from FD and IC alone. We're claiming that FD and IC are *necessary* constraints—any valid physics must satisfy them. But they're not *sufficient* to determine exactly which physics we get. That's like saying any house must follow building codes (necessary), but building codes don't tell you exactly what your house will look like (not sufficient).

Chapter 3: Emergence of Time and Law

3.1 From Ticks to Macroscopic Time

For the general reader: How do we get from the primitive tick layer (where time isn't even defined yet) to the familiar flow of time we experience? The answer involves "coarse-graining"—ignoring fine details and looking at the big picture.

When ticks are aggregated under FD and IC, a partial order on equivalence classes of event-histories is induced. This ordering is identified with **macroscopic time**.

The mechanism is as follows:

1. **Coarse-graining:** Define equivalence classes of tick configurations that are operationally indistinguishable under finite resources (per FD)
2. **Commitment tracking:** Among equivalence classes, identify those related by irreversible commitments (per IC)
3. **Temporal ordering:** The IC-induced partial order on equivalence classes defines macroscopic before/after relations
4. **Metric emergence:** The "amount" of elapsed time between equivalence classes is determined by the number of independent irreversible commitments required to transition between them

For the general reader: Here's an analogy. Imagine watching a movie at different resolutions:

- At maximum resolution, you see every pixel of every frame
- At low resolution, many different high-res frames look identical—they blur together
- The low-resolution version has fewer distinguishable states

Time "emerges" when we go from the tick layer to a coarse-grained description:

- Many different tick configurations look the same (because FD limits what we can distinguish)
- We group them into equivalence classes
- The irreversible commitments between these classes define temporal order
- The "amount of time" between events relates to how many irreversible steps separate them

Proposition 3.1: The macroscopic time ordering derived from IC is consistent with the causal ordering inherited from the tick layer: if tick e causally precedes tick f ($e \preccurlyeq f$), then any equivalence class containing e does not temporally follow any equivalence class containing f.

Proof sketch: Causal precedence at the tick level constrains which irreversible commitments are possible. If $e \preccurlyeq f$, then commitment events influenced by f cannot precede commitment events influenced only by e. ■

3.2 Physical Laws as Fixed Points

For the general reader: Where do physical laws come from? In this framework, they're not handed down from on high—they *emerge* as stable patterns that survive coarse-graining.

Here's the intuition: Imagine looking at a beach from various distances. Up close, you see individual grains of sand. From far away, you see "the beach"—a stable, recognizable pattern that persists regardless of exactly which grains are where. Physical laws are like "the beach"—patterns that remain stable even when you ignore microscopic details.

Physical laws arise as **stable fixed points** of coarse-graining over tick histories subject to admissibility constraints.

Definition 3.1 (Coarse-Graining Family): Let $\{C_\lambda\}$ be a family of coarse-graining maps parameterized by resolution scale λ, mapping microscopic descriptions to macroscopic equivalence classes while respecting FD and IC.

Definition 3.2 (Induced Effective Dynamics): For each λ, let L_λ denote the effective dynamical law at scale λ.

Definition 3.3 (Fixed Point Law): L is a fixed point of coarse-graining if L_λ is invariant (up to reparametrization) under $\lambda \to \lambda'$. That is, the law's form does not depend on the choice of coarse-graining scale.

For the general reader: A "fixed point" is something that doesn't change when you apply some operation to it. For example, if you keep pressing the √ button on a calculator starting from any positive number, you eventually get to 1 (since √1 = 1). The number 1 is a fixed point of the square root operation.

Physical laws are fixed points of coarse-graining: no matter how you adjust what counts as "microscopic" versus "macroscopic," the laws stay the same. That's what makes them laws—they're the robust, universal patterns.

Clarification: In this paper we use "fixed point" in the RG-inspired structural sense; a fully formal RG construction is beyond scope. Such a construction would require specifying the space of effective theories, defining explicit coarse-graining operators, and proving convergence—a program for future work.

Remark 3.1: This notion of stability parallels the renormalization group concept in statistical mechanics and quantum field theory. Effective theories at different scales are connected by coarse-graining flows, and physically relevant theories correspond to fixed points or slow flows under such transformations. The laws of physics are precisely those regularities that survive coarse-graining—that are robust against changes in how we divide the world into "relevant" and "irrelevant" degrees of freedom.

3.3 What Admissibility Does and Does Not Determine

The admissibility constraints (FD and IC) impose **necessary conditions** on physical laws but do not uniquely determine them:

Enforced by admissibility alone (+ operational axioms + continuity):

- Structure preservation of reversible dynamics
- Entropy monotonicity under irreversible processes (second law structure)
- Finite operational information density

Constrained but not determined (requires additional symmetry inputs):

- Maxwell's equations: require U(1) gauge symmetry + Lorentz covariance + locality + variational principle
- Einstein's equations: require diffeomorphism invariance + locality + second-order dynamics + Lovelock uniqueness
- Specific Hamiltonians: require specification of degrees of freedom and interactions

Not determined by admissibility:

- Which gauge groups are realized
- Which matter content exists
- Values of coupling constants
- Initial/boundary conditions

For the general reader: Here's what this means in plain terms:

Admissibility tells you:

- Quantum systems must evolve in a specific mathematical way (unitarily) between measurements—no choice about this
- Entropy (disorder) can never decrease overall—this is required, not optional
- You can't pack infinite information into finite space—hard limit

Admissibility doesn't tell you:

- What particles exist (electrons, quarks, etc.)
- How strong gravity is
- Why the speed of light is 299,792 km/s
- How the universe started

The admissibility layer constrains the **form** of physical law more than its **content**.

3.4 Architectural Summary

The resulting structure is hierarchical:

```
LAYER 2: EFFECTIVE PHYSICS
─────────────────────────
Content: Dynamical laws (Schrödinger, Maxwell, Einstein...)
Status: Fixed points of coarse-graining under admissibility
Requires: Admissibility + symmetry inputs + initial data

[What we calculate with; the familiar equations of physics]

                        ▲
                        │ constrains

LAYER 1: ADMISSIBILITY LAYER
────────────────────────────
Content: FD (finite distinguishability) + IC (irreversible
         commitment) relative to 𝒪(R,E,τ,M,δ)
Status: Universal necessary conditions
Role: Constraint-complete specification of admissibility

[The "building codes" all physics must follow]

                        ▲
                        │ aggregates
```

```
LAYER 0: TICK LAYER
───────────────────
Content: Primitive events of change
Structure: Locally finite causal poset (P, ≼)
Role: Minimal process ontology

[The raw "stuff happens" before physics]
```

This stratification preserves mathematical rigor, prevents category errors (confusing ontology with constraint), and clarifies how constraint-based unification coexists with emergent physical law.

Part II: Physical Admissibility Theorems

For the general reader: Now we get to the payoff. From just FD and IC, we can prove several important theorems about what's physically possible and impossible. These aren't assumptions—they're logical consequences of the two core constraints.

From the two core constraints FD and IC, we derive exclusion theorems that establish necessary conditions on any physically grounded mathematics. These theorems are best understood as **reductio arguments**: if the constraint were violated, physical realizability would fail in specific ways.

Chapter 4: No Unbounded Operational Information Density

4.1 Statement

For the general reader: This theorem says: "You can't store infinite information in a finite space with finite resources." This might seem obvious, but the proof shows *why* it's true—and what would go catastrophically wrong if it weren't.

Theorem 4.1 (No Unbounded Operational Information Density — Physical CT Form): Assume the Physical Church-Turing Thesis: every physically realizable input-output procedure is Turing-simulable to the same operational accuracy. If a bounded system admits an unbounded family of mutually distinguishable states with uniform preparation and uniform readout within a fixed finite resource budget, then Physical Church-Turing is violated.

For the general reader: The "physical Church-Turing thesis" says that any computation a physical system can perform can also be performed by a standard computer (given enough time and memory). It's one of the most well-supported principles in computer science—no one has ever found a physical process that computes something a regular computer couldn't.

The theorem says: if FD were false, the physical Church-Turing thesis would also be false. Since we strongly believe Church-Turing holds, we should believe FD holds too.

Assumptions (Uniform Prepare/Read): There exist procedures Prep(n) and Read(·) in $\mathcal{O}(R, E, \tau, M, \delta)$ such that:

- Prep(n) prepares state σ_n for any $n \in \mathbb{N}$ within the same fixed resource budget
- Read discriminates σ_n from σ_m for any $n \neq m$ with error $< \delta$ within the same fixed budget

Assumption (Uniform Configurability): The device admits a reconfigurable control/program register (itself physically realizable within the bounded system) such that for any finite N and any bit pattern $b_0,\ldots,b_n$, one can configure the device so that on inputs $n \leq N$ it outputs b_n with error $< \delta$, while preserving the same uniform Prep/Read bounds. This captures the operational notion of a programmable memory/lookup device.

4.2 Proof

Proof (by contradiction):

Assume such $\{\sigma_n\}$ and uniform Prep/Read exist. Then the bounded device implements a physically realizable, uniform input-output map $f: \mathbb{N} \to \mathbb{N}$ given by: input n, run Prep(n), then Read to output n.

Now consider decision procedures parameterized by subsets $S \subseteq \mathbb{N}$. By Uniform Configurability, the device can realize arbitrarily long finite prefixes of membership functions χ_S via physical configuration: configure outputs to match χ_S on $\{0,\ldots,N\}$ for any N.

If unboundedly long prefixes are physically realizable in a single bounded device without scaling resources, the induced input-output behavior is not Turing-simulable uniformly to the same operational accuracy. Physical Church-Turing implies the set of physically preparable configurations is Turing-enumerable; unbounded uniform addressability would permit non-Turing-enumerable configuration-spaces (since arbitrary $S \subseteq \mathbb{N}$ can be encoded), yielding contradiction. ■

For the general reader: The key insight is that if you could store unlimited information in a finite device, you could "pre-load" the answers to problems that no computer can solve—like which programs will run forever. The device wouldn't compute the answers; it would just look them up. But physical Church-Turing says this can't happen, so unlimited storage must be impossible.

4.3 Interpretation

The contradiction does not require "computing" a non-recursive set during preparation; it arises because FD failure plus uniform addressability permits physically realizing bounded configurations whose induced input-output behavior corresponds to non-recursive sets. Physical Church-Turing rules these out; therefore FD must hold (as an operational capacity constraint).

Remark (Halting oracle specialization): If one additionally assumes a physically realizable mechanism for embedding a halting oracle into initial conditions, the above construction specializes to a direct halting decider. This stronger assumption is not required for the theorem.

4.4 On Uniform Readout

Remark (Uniformity as the operational hinge): Uniform readout is intentionally strong: it posits a hypothetical system in which increasing the number of distinguishable states does not force a corresponding scaling of discrimination resources. In realistic physical systems, discrimination typically scales with dimension, time, energy, or sample complexity (e.g., tomography and hypothesis-testing bounds). That scaling behavior is itself a concrete manifestation of FD. Here, Uniform Readout is introduced only to make explicit which "escape hatch" prevents the reductio: namely, that unbounded distinguishability can only be physically meaningful if it is uniformly accessible.

Remark (Assumption transparency): This theorem is conditional on the physical Church-Turing thesis, which is an empirical claim, not a mathematical theorem. If future physics were to establish violations of physical Church-Turing (e.g., through novel quantum gravitational effects or other mechanisms), then FD would require corresponding revision rather than the argument being invalid. The theorem shows: Physical Church-Turing ⇒ FD. The contrapositive—¬FD ⇒ ¬Physical Church-Turing—is equally valid.

4.5 Corollaries

Corollary 4.1 (Operational Entropy Bounds): Under FD, the operationally accessible entropy of any bounded system with bounded resources is finite:

$$S\_operational(R, E, \tau, \delta, M) \leq k\_B \ln N(R, E, \tau, \delta, M) < \infty$$

where N is the number of operationally distinguishable states.

For the general reader: Entropy measures how many different microscopic arrangements correspond to the same macroscopic state—roughly, how much hidden information a system contains. This corollary says that hidden information is finite, not infinite.

Corollary 4.2 (Consistency with Bekenstein-Type Bounds): FD is consistent with—and provides operational grounding for—Bekenstein-type entropy bounds:

$$S \leq 2\pi k\_B RE / (\hbar c)$$

For the general reader: In 1981, physicist Jacob Bekenstein discovered that black holes place an upper limit on how much information can be stored in any region of space. The limit depends on the region's size (R) and energy content (E). FD explains *why* such bounds exist: they're necessary to prevent infinite information storage.

Corollary 4.3 (Holographic Consistency): The holographic principle (entropy bounded by boundary area rather than volume) is consistent with FD. Holography provides a gravitational mechanism enforcing FD at the Planck scale.

For the general reader: The "holographic principle" is a deep idea from quantum gravity: the maximum information in a 3D region is proportional to its 2D surface area, not its 3D volume. It's as if reality is like a hologram—all the information is somehow encoded on the boundary. FD is compatible with and helps explain this strange fact.

4.6 Relation to Existing Physics

FD is not a new physical postulate but a formalization of constraints already implicit in established physics:

Domain	FD Manifestation
Quantum mechanics	Finite operationally distinguishable states for bounded systems
Statistical mechanics	Finite accessible phase space volume (h^n per state)
Black hole physics	Bekenstein-Hawking entropy bounds
Quantum gravity	Holographic principle
Information theory	Finite channel capacity under resource constraints

---

Chapter 5: Irreversible Cost of Information Erasure

5.1 Clarification: Branching vs. Erasure

For the general reader: Before diving into the theorem, we need to clarify an important distinction that even many physicists get confused about.

Branching means exploring multiple possibilities—like a chess computer considering many possible moves. This can be done *reversibly*, keeping track of all the branches.

Erasure means forgetting which branch you took—resetting to a blank state. This is *irreversible* and has an unavoidable energy cost.

The theorem below is about erasure, not branching.

A critical distinction: **branching** (exploring multiple computational paths) and **erasure** (discarding information about which path was taken) are different operations with different thermodynamic consequences.

- **Reversible branching:** A computation can explore k distinguishable branches reversibly, maintaining full information about the branching history. Quantum computers exploit this via superposition.

- **Irreversible erasure:** Returning to a standard initial state after branching—discarding the record of which branch was taken—is irreversible and carries unavoidable thermodynamic cost.

The admissibility theorem concerns **erasure**, not branching per se.

5.2 Statement and Proof

Theorem 5.1 (Irreversible Compression Cost): Any computation that produces $\log_2 k$ bits of mutually distinguishable branch record and then returns to a standard initial memory state must dissipate at least k_B T ln k of heat (or export equivalent entropy to the environment), regardless of whether the intermediate branching was performed reversibly.

For the general reader: In plain English: "Erasing information costs energy." Specifically, erasing information about k possibilities costs at least k_B T ln k in energy (where k_B is Boltzmann's constant and T is temperature). At room temperature, erasing one bit costs at least about 3×10^{-21} joules—tiny, but not zero.

This is called "Landauer's principle," and it has been experimentally verified.

Proof:

Step 1: Consider a computation that explores k distinguishable branches, producing a branch record—a memory state encoding which branch was taken. The k possible records are mutually distinguishable states of the memory system.

Step 2: Information content of the branch record: $\log_2 k$ bits (minimum bits needed to specify one of k possibilities).

Step 3: Returning to a standard initial state (reset/erasure) means transforming any of the k distinguishable record states to a single fixed state. This is a many-to-one map on the memory system.

For the general reader: If you have k different possible starting states and they all end up in the same final state, that's a "many-to-one" map. Information about which starting state you had is lost.

Step 4: By IC, many-to-one maps on distinguishable states cannot be implemented without irreversible commitment—information about which input state was present is lost.

Step 5: By **Landauer's principle**, erasure of one bit requires dissipation of at least:

$W_{1\beta it} \geq k\_B T \ln 2$

where T is the temperature of the thermal environment.

Step 6: Erasure of $\log_2 k$ bits requires:

$$W\_total \geq k\_B T \ln 2 \times \log_2 k = k\_B T \ln k$$

This is a lower bound; actual implementations typically dissipate more. ■

5.3 Landauer's Principle: Derivation and Status

For the general reader: Landauer's principle connects information theory to thermodynamics. It says that information isn't free—erasing it has a physical cost. This principle, proposed by Rolf Landauer in 1961, was controversial for decades but has now been experimentally confirmed.

Derivation: Consider a single-bit memory that can be in state 0 or 1 with equal probability. Erasing means resetting to state 0 regardless of initial state.

- Initial entropy of memory: $S_i = k\_B \ln 2$
- Final entropy of memory: $S\_f = 0$

For the general reader: Entropy measures uncertainty or "disorder." Initially, the bit could be 0 or 1 (uncertainty = ln 2). After erasure, it's definitely 0 (uncertainty = 0). The entropy decreased.

By the second law, total entropy (system + environment) cannot decrease:

$$\Delta S\_total = \Delta S\_memory + \Delta S\_environment \geq 0$$

$$-k\_B \ln 2 + \Delta S\_environment \geq 0$$

$$\Delta S\_environment \geq k\_B \ln 2$$

For the general reader: If the memory's entropy went down, the environment's entropy must go up by at least the same amount. The total entropy can't decrease.

For a thermal environment at temperature T:

$$Q = T \times \Delta S\_environment \geq k\_B T \ln 2 \approx 2.9 \times 10^{-21} \text{ J at room temperature}$$

Experimental status: Landauer's bound has been experimentally verified in multiple systems:

- Colloidal particles in optical traps (Bérut et al., Nature 2012)
- Single-electron boxes (Koski et al., PNAS 2014)
- Nanomagnetic bits (Hong et al., Science Advances 2016)

5.4 Implications for Computation

Proposition 5.1: Any physical computer that explores k candidate solutions and returns to a ready state must dissipate energy at least $k_B T \ln k$, regardless of intermediate computation strategy.

For the general reader: This has profound implications for computing. Every time your computer clears its memory, it must pay an energy cost. Modern computers dissipate far more than the Landauer minimum (for engineering reasons), but the minimum itself is a fundamental physical limit.

Remark 5.2: Quantum computers do not violate this bound. Quantum speedup comes from maintaining coherent superposition (reversible branching), exploiting interference, and measuring only at the end. The quantum advantage is structural, not thermodynamic evasion.

For the general reader: You might wonder: "Don't quantum computers beat classical computers by exploring many possibilities at once?" Yes, but they do it reversibly—they don't erase the branch information until the very end (measurement). The Landauer cost is paid at measurement, not during the quantum computation itself.

Remark 5.3: This theorem does not resolve P vs NP (a mathematical question). It establishes that any **physical implementation** of exponential search pays thermodynamic cost at least linear in the exponent—a physical constraint, not a logical one.

Chapter 6: Structure Preservation Between Commitments

6.1 Operational Metric Assumption

Assumption (Operational Metric): The distinguishability measure D is operational in the sense that:

- (i) It is non-increasing under admissible channels (data-processing/contractivity): $D(\Phi(\rho), \Phi(\sigma)) \leq D(\rho, \sigma)$ for all $\Phi \in \mathcal{O}$
- (ii) Equality $D(\Phi(\rho), \Phi(\sigma)) = D(\rho, \sigma)$ for all pairs characterizes reversibility with respect to $\mathcal{O}$

Examples include trace distance, Helstrom optimal error probability, and other contractive distinguishability measures.

6.2 Reversibility Lemma

Definition (Operational Reversibility): $\Phi \in \mathcal{O}$ is operationally reversible if there exists $\Psi \in \mathcal{O}$ such that $\Psi \circ \Phi = \text{id}$ on the relevant state set.

Lemma 6.1 (Reversible ⇒ Isometry of D): If $\Phi \in \mathcal{O}$ admits an inverse $\Psi \in \mathcal{O}$, then contractivity of D implies $D(\Phi(\rho), \Phi(\sigma)) = D(\rho, \sigma)$.

Proof: By contractivity, $D(\Phi(\rho), \Phi(\sigma)) \leq D(\rho, \sigma)$. Applying contractivity to Ψ: $D(\rho, \sigma) = D(\Psi(\Phi(\rho)), \Psi(\Phi(\sigma))) \leq D(\Phi(\rho), \Phi(\sigma))$. Combined: equality. ■

For the general reader: If a process can be undone, it can't lose information—the distinguishability of states must be preserved exactly.

6.3 Structure Theorem

For the general reader: This theorem answers the question: "What kind of evolution is allowed between irreversible events?" The answer: only transformations that preserve distinguishability—which in quantum mechanics means unitary evolution.

Theorem 6.1 (Quantum-Case Structure Theorem): Given the operational quantum axioms (convex state space, pure states, continuous reversible transformations) implying a Hilbert-space representation, any continuous family of operationally reversible transformations is implemented by a unitary one-parameter group.

Proof:

Step 1: By Lemma 6.1, operationally reversible maps preserve D (are D-isometries).

Step 2: On pure states, taking distinguishability as transition probability, Wigner's theorem implies that any bijection preserving transition probabilities is unitary or antiunitary. More generally, for affine bijections on the convex state space preserving an operational metric (e.g., trace distance), Kadison-type theorems yield implementation by unitary or antiunitary conjugation.

For the general reader: These are famous results from quantum physics saying that the *only* transformations preserving the distinguishability structure are unitary (or antiunitary). There's no other option—it's a mathematical uniqueness result.

Step 3: Continuous one-parameter evolution connected to the identity must be unitary throughout (antiunitary maps cannot form continuous groups connected to I).

Conclusion: Continuous reversible dynamics form one-parameter unitary groups. ■

6.4 From Unitarity to Schrödinger Evolution Form

For the general reader: Now comes the punchline. A famous mathematical theorem (Stone's theorem) says that *any* continuous family of unitary transformations can be written in a specific form. This form *is* the Schrödinger equation.

In other words: we didn't assume the Schrödinger equation—we derived its form from FD and IC!

Stone's Theorem: Every strongly continuous one-parameter unitary group U(t) on a Hilbert space $\mathscr{H}$ has the form:

$U(t) = \exp(-iHt/\hbar)$

for some self-adjoint operator H (the Hamiltonian).

Differentiating $|\psi(t)\rangle = U(t)|\psi(0)\rangle$:

$i\hbar\, \partial|\psi(t)\rangle/\partial t = H|\psi(t)\rangle$

This is the **Schrödinger evolution form**.

For the general reader: The Schrödinger equation is the fundamental equation of quantum mechanics, governing how quantum states change over time. Physics students learn it as a postulate—"this is just how quantum mechanics works."

But here we've derived its *form* from admissibility constraints:

1. FD motivates finite operational distinguishability
2. IC distinguishes reversible from irreversible processes
3. Reversible processes must preserve distinguishability (Lemma 6.1)
4. Distinguishability-preserving processes must be unitary (Wigner/Kadison)
5. Continuous unitary processes must have Schrödinger form (Stone)

The logic is airtight—given the constraints and operational quantum axioms, the Schrödinger evolution form is *inevitable*.

Critical clarification on scope: We have derived the **Schrödinger evolution form for continuous reversible dynamics**, not derived quantum theory from scratch.

The logical structure:

1. **Admissibility (FD + IC)** → Contractivity of D; reversibility = D-isometry
2. **Operational quantum axioms** → Hilbert space representation
3. **Together** → Schrödinger evolution form

Key statement: Admissibility does not derive quantum theory; it explains why any admissible theory with a quantum-like operational state space must evolve unitarily between irreversible record formations.

What is NOT determined: The specific Hamiltonian H requires additional physical input.

For the general reader: The theorem tells us that quantum systems *must* evolve according to *some* Schrödinger equation. It doesn't tell us *which* Schrödinger equation—that depends on the specific Hamiltonian H, which encodes the particular physics of the system (what particles are present, how they interact, etc.).

Chapter 7: Operational Bounds on Soft Excitations

7.1 The Problem with Naive Mass Gap Claims

For the general reader: This section addresses a potential objection: "Doesn't physics have particles with zero mass (like photons)? Doesn't that violate finite distinguishability?"

The answer is subtle: zero-mass particles exist, but the number of *distinguishable* soft (low-energy) configurations is still finite under finite resources.

A naive formulation—"FD requires a mass gap"—is **false as stated**. Standard physics contains massless particles (photons, gravitons) and gapless excitations in many-body systems.

The correct statement concerns **operational distinguishability**, not energy gaps per se.

7.2 Correct Statement

Theorem 7.1 (No Unlimited Soft Distinctions): For fixed finite detector size R, finite integration time τ, finite total energy budget E, and fixed discrimination error δ, the number of distinguishable soft-excitation configurations is finite.

Scope clarification: The theorem concerns fixed finite operational resources; asymptotic symmetry charges become sharply defined only in limits ($R, \tau \to \infty$) outside the theorem's scope.

For the general reader: In plain terms: even with massless particles, you can't have infinitely many distinguishable low-energy states. Something must limit the number—either the energy cost of detection, gauge redundancy (some states are actually the same), or other physical bounds.

Proof: If arbitrarily many configurations with arbitrarily low incremental energy were all distinguishable within $\mathcal{O}(R, E, \tau, M, \delta)$, this would violate FD. ■

7.3 Compatibility with Known Physics

Why massless particles don't violate Theorem 7.1:

For the general reader: Here's why photons (which have zero mass) don't create infinite distinguishable states:

1. **Finite detection time/volume:** Distinguishing soft photons with wavelength λ requires detector size $\geq \lambda$ and integration time $\geq \lambda/c$. As $\lambda \to \infty$, required resources $\to \infty$.

Analogy: Imagine trying to distinguish radio waves of slightly different frequencies. The longer the wavelength (lower the frequency), the bigger your antenna needs to be and the longer you need to listen. For infinitely long wavelengths, you'd need an infinitely big antenna listening for infinite time.

2. **Gauge redundancy:** In QED, gauge-equivalent configurations represent the same physical state. The number of gauge-inequivalent soft configurations is bounded.

For the general reader: "Gauge redundancy" means that some configurations that look different mathematically are actually physically identical. Once you account for this redundancy, the number of genuinely distinct configurations is smaller than it first appears.

3. **Infrared dressing and superselection:** Physical charged states are dressed by soft photon clouds forming superselection sectors with finite internal distinguishability.
4. **Gravitational bounds:** Attempting to pack too many distinguishable soft quanta eventually forms a black hole.

Footnote (Infrared subtlety): The infrared structure of gauge theories is subtle: physically correct asymptotic states in QED involve soft dressing (e.g., Faddeev-Kulish-type constructions), and the infrared sector is tied to asymptotic symmetries and memory effects. These refinements do not contradict FD, because operational distinguishability remains bounded for fixed detector size, integration time, and error tolerance; rather, they clarify which degrees of freedom label physically meaningful asymptotic sectors and how they become accessible only in appropriate long-time/large-radius limits.

Why gapless many-body modes don't violate Theorem 7.1:

Gapless excitations exist in the thermodynamic limit. For finite systems: spectrum is discrete, mode number is bounded, finite temperature provides IR cutoff.

For the general reader: "Gapless excitations" occur in idealized infinite systems. Any real, finite system has a minimum excitation energy (even if very small). The "gapless" limit is a mathematical convenience, not physical reality.

Chapter 8: Admissibility Stress Test

For the general reader: Any good scientific theory should be falsifiable—there should be possible observations that would prove it wrong. Here we list what would falsify the admissibility framework.

8.1 Falsification Criteria

The framework is falsified by physical existence of any system exhibiting:

Criterion 1 (FD Violation): A bounded system with unbounded distinguishable states under uniform prepare/read within fixed finite resources.

What this would look like: A device of fixed size that can store any amount of data and retrieve it perfectly, with no increase in energy, time, or equipment needed.

Criterion 2 (IC Violation): Information erasure without any thermodynamic cost or trace.

What this would look like: A memory that can be reset without generating any heat whatsoever—perfect erasure with zero energy cost.

Criterion 3 (Soft Distinction Violation): Unbounded distinguishable stable excitations at arbitrarily low energy, all accessible with fixed finite resources.

What this would look like: Infinitely many different stable particle types, each requiring arbitrarily small energy to distinguish.

8.2 Status

No system satisfying any violation criterion has been observed or constructed. Black hole debates continue but proposed resolutions preserve FD. Quantum computers use reversible branching, not cost-free erasure.

Part III: Canonical Equations and Spine Compliance

For the general reader: Now we examine the major equations of physics to see which are *enforced* by admissibility (no choice—they must be true), which are *constrained* (narrowed down but not uniquely determined), and which are merely *compatible* (consistent with admissibility but requiring additional input).

Chapter 9: The Schrödinger Evolution Form

For the general reader: The Schrödinger equation governs how quantum systems evolve in time. It's the quantum analog of Newton's F=ma.

$i\hbar \, \partial|\psi(t)\rangle/\partial t = H|\psi(t)\rangle$

We derived above that this *form* is enforced by admissibility (given operational quantum axioms). The specific Hamiltonian H (which determines the particular physics) requires additional input.

9.1 Admissibility Status

Aspect	Status
Linear evolution	**Enforced** (given operational quantum axioms)
Unitarity	**Enforced**
Self-adjoint generator	**Enforced**
Specific Hamiltonian	Not determined
Hilbert space choice	Constrained, not uniquely determined

Clarification: "Enforced" means: given FD + IC + operational quantum axioms + continuity, the Schrödinger *form* is necessary. The specific H requires physics input.

Chapter 10: The Heisenberg Uncertainty Principle

For the general reader: The uncertainty principle says you can't simultaneously know both the position and momentum of a particle with perfect precision. The more precisely you know one, the less precisely you can know the other:

$\Delta x \cdot \Delta p \geq \hbar/2$

This isn't about measurement disturbing the system—it's a fundamental limit on how much information exists to be known.

10.1 Mathematical Statement

For observables x, p with $[x,p] = i\hbar$:

$\Delta x \cdot \Delta p \geq \hbar/2$

For the general reader: Why does this follow from FD? Because FD says you can only distinguish finitely many states in a finite region of "phase space" (the space of all possible positions and momenta). The uncertainty principle is a mathematical expression of this finite resolution—you can't pinpoint a state more precisely than about one "Planck cell" of phase space.

10.2 Admissibility Status

Aspect	Status
Existence of uncertainty tradeoffs	**Enforced** by FD (finite phase-space resolution)
Robertson-Schrödinger form	**Enforced** given Hilbert representation
Value of $\hbar$	Not determined (contingent)
Which observables are conjugate	Requires physical input

For the general reader: FD implies you can't resolve phase space infinitely finely—there must be *some* uncertainty tradeoff. The specific mathematical form (Robertson-Schrödinger) requires Hilbert space structure. The actual value of $\hbar$ is empirical.

Chapter 11: Maxwell's Equations

For the general reader: Maxwell's equations govern electricity and magnetism—how electric and magnetic fields are created by charges and currents, and how they change in time. They unify electricity, magnetism, and light into a single framework.

Unlike the Schrödinger equation, Maxwell's equations are NOT enforced by admissibility alone. They require additional assumptions about symmetry (specifically, U(1) gauge symmetry and Lorentz invariance).

11.1 Admissibility Status

Aspect	Status
Existence of gauge field	Not determined
U(1) structure	Not determined
Maxwell form given U(1) + locality + Lorentz	Strongly **constrained**
Vacuum reversibility	Compatible with IC
Finite distinguishable modes	Consistent with FD

For the general reader: Admissibility doesn't tell us that electromagnetic fields must exist or that they must have the structure they have. But it does tell us that *given* certain symmetry assumptions, Maxwell's equations are essentially the only possibility.

Chapter 12: The Einstein Field Equations

For the general reader: Einstein's field equations govern gravity—how matter and energy curve spacetime, and how curved spacetime affects the motion of matter:

$G_{\mu\nu} + \Lambda g_{\mu\nu} = (8\pi G/c^4)\, T_{\mu\nu}$

Like Maxwell's equations, these are not enforced by admissibility alone but are strongly constrained given certain symmetry assumptions (diffeomorphism invariance, locality, second-order derivatives).

12.1 Admissibility Status

Aspect	Status
Dynamical spacetime	Not determined
Diffeomorphism invariance	Not determined
Einstein form given Lovelock inputs	Uniquely **constrained**
Entropy bounds	Consistent with FD
Horizon thermodynamics	Consistent with IC

Chapter 13: The Second Law of Thermodynamics

For the general reader: The second law says that entropy (roughly, disorder or missing information) never decreases in an isolated system:

$S\_later \geq S\_earlier$

This is the most direct expression of IC. The second law isn't just compatible with admissibility—it's *enforced* by it.

13.1 Information-Theoretic Formulation

The data-processing inequality: for any quantum channel Φ,

$D(\rho \| \sigma) \geq D(\Phi(\rho) \| \Phi(\sigma))$

implies entropy monotonicity under irreversible processes.

For the general reader: The data-processing inequality says that no physical process can make two states *more* distinguishable than they started. You can only lose information, never create it. This is mathematically equivalent to the second law—and it follows directly from IC.

13.2 Admissibility Status

Aspect	Status
Entropy monotonicity	**Enforced** by IC
Data-processing inequality	**Enforced**
Landauer bound	**Enforced**
Specific entropy formulas	Requires statistical mechanics

Chapter 14: Summary — Classification Table

Feature	Enforced	Constrained	Compatible	Not Det.
Unitary structure (given quantum axioms)	✓			
Uncertainty tradeoffs	✓			
Robertson-Schrödinger (given Hilbert)	✓			
Entropy monotonicity	✓			
Landauer bound	✓			
Maxwell equations		✓		
Einstein equations		✓		
Bekenstein bounds			✓	
Holographic principle			✓	
Specific Hamiltonians				✓
Value of $\hbar$				✓
Coupling constants				✓

Note on classical theories: The "(given quantum axioms)" qualifier raises the question of whether FD + IC have analogous structure-preservation consequences in non-quantum theories. In classical stochastic theories with FD + IC, reversible dynamics would preserve distinguishability (e.g., total variation distance), yielding measure-preserving (symplectic/volume-preserving) flows rather than unitary groups. The admissibility constraints thus have representation-dependent structural consequences—a direction for future investigation.

For the general reader: This table is the key result of this section. It shows that:

- Quantum mechanics' basic structure (unitarity) and thermodynamics' second law are *inevitable*—any physically admissible theory must have them
- Electromagnetism and gravity are *constrained* but require additional symmetry assumptions
- Specific details (particular Hamiltonians, coupling constants) are not determined at all

Part IV: Gödel Reflection and the Tick Layer

Important disclaimer: The tick layer does not "escape" Gödel's incompleteness. Any sufficiently strong formalization of it is subject to incompleteness theorems. The connection is about architecture, not metaphysical escape.

For the general reader: This section connects our physics framework to a famous result in mathematical logic: Gödel's incompleteness theorems. The connection is suggestive rather than definitive—we're not claiming to "solve" incompleteness, just to illuminate interesting parallels.

Chapter 15: Gödel's Incompleteness

For the general reader: In 1931, Kurt Gödel proved one of the most profound results in mathematics: any consistent mathematical system capable of expressing basic arithmetic contains true statements that cannot be proven within that system.

More precisely:

- **Gödel sentence G:** "This statement is not provable in system T"
- If T is consistent, G is true but unprovable in T
- To prove G is true, you must step "outside" T to a meta-level

For consistent, sufficiently strong theory T:

- **Gödel sentence G_T:** $T \vdash G_T \leftrightarrow \neg Prov_T(\ulcorner G_T \urcorner)$ ("I am unprovable in T")
- **First Incompleteness:** If T consistent, $T \nvdash G_T$
- **Meta-level requirement:** Truth of G_T established from outside T

Chapter 16: Tick Layer as Process Substrate

For the general reader: Here's the connection to physics: the "outside" perspective that Gödel's theorem requires can be thought of as a *process*—systematically checking all possible proofs. This process can be modeled as a sequence of ticks.

Proof enumeration (systematically checking candidate proofs) can be implemented as a tick sequence—each computational step is a tick. The tick layer provides a **minimal process substrate for meta-computation**.

For the general reader: Imagine a computer program that:

1. Lists all possible proofs in order: proof #1, proof #2, proof #3, ...
2. For each proof, checks if it's a valid proof of G_T
3. If it finds one, stops; otherwise continues forever

Each step of this program is a tick. The tick layer provides the minimal "stuff" needed to run this program. The program operates *on* system T (checking proofs in T) but is not *within* T—it's at the meta-level.

Chapter 17: Precise Structural Parallel

The intended analogy is not merely "both have hierarchies," but that both domains require a principled separation between object-level descriptions and admissibility criteria that are not fully captured by the object level alone.

In formal logic, provability is an internal notion, while consistency/truth often require meta-level reasoning.

In our physical architecture, the tick layer provides an object-level process substrate, while physical implementability constraints (finite distinguishability and irreversible commitment relative to $\mathcal{O}$) act as admissibility criteria on which mathematical structures can be realized.

The parallel is one of admissibility stratification: validity conditions sit one level above the generative substrate. Both Gödelian meta-theory and physical admissibility layer serve as constraint levels that are not fully internal to the object-level description.

Chapter 18: Appropriate Interpretation

What we claim:

- Tick layer provides process substrate for meta-computation
- Gödel's "outside" can be operationally instantiated
- Structural parallel: admissibility stratification in both logic and physics

What we do NOT claim:

- Incompleteness is "solved" or "grounded"
- Physical processes escape formal limitations
- Minds transcend formal systems

The correspondence illuminates **architecture**, not metaphysics.

For the general reader: We're NOT saying that the tick layer "solves" Gödel's theorem or lets us escape its limitations. Any sufficiently powerful formal system—including a formal description of the tick layer—is subject to incompleteness.

What we ARE saying is that there's an interesting structural parallel: both physics and logic involve layered hierarchies where some things can only be "seen" from outside a given level. This parallel might be deep or might be superficial—we don't claim to know.

Part V: Conclusion

Summary

Layer 0 (Tick Layer): Minimal process ontology—events of change forming a locally finite causal poset. The raw "something happens" before physics.

Layer 1 (Admissibility Layer): Two universal constraints relative to $\mathcal{O}$(R, E, τ, M, δ):

- **FD:** Finite operational distinguishability (capacity constraint)
- **IC:** Existence of operationally non-invertible processes

Layer 2 (Effective Physics): Laws emerging as coarse-graining fixed points:

- **Enforced:** Unitarity (given quantum axioms), entropy monotonicity, Landauer bound, uncertainty tradeoffs
- **Constrained:** Maxwell, Einstein forms given symmetry inputs
- **Not determined:** Specific Hamiltonians, constants, matter content

For the General Reader: What Have We Achieved?

What this framework does:

1. Identifies two fundamental constraints (FD and IC) that any physical theory must satisfy
2. Derives several major features of physics from these constraints (unitarity, second law, uncertainty)
3. Clarifies which features are universal necessities vs. which require additional assumptions
4. Provides a clear hierarchy: tick layer → admissibility → effective physics

What this framework does NOT do:

1. Uniquely determine which physical theory describes our universe
2. Explain why the fundamental constants have the values they do
3. Resolve interpretational debates (measurement problem, etc.)
4. Eliminate the need for experimental physics

The bottom line: FD and IC are like "building codes" for the universe. Any possible physics must follow them. But just as building codes don't determine what your house looks like, FD and IC don't uniquely determine physics. They constrain without fully determining.

What We Have and Haven't Achieved

Achieved:

- Explicit operational definitions (resource-relative, not metaphysical)
- Clear assumption ledger with minimal bridges
- Derivation of Schrödinger *evolution form* for continuous reversible dynamics
- Completeness theorem for constraints

- Clean separation: necessity constraints vs dynamical selection

Not achieved:

- Derivation of quantum theory from scratch
- Unique determination of physics
- Explanation of coupling constants
- Resolution of interpretational questions

The bottom line: The result is a completeness theorem for constraints, not a completeness theorem for dynamics.

Open Questions

1. **Why these constraints?** Can FD and IC be derived from something deeper, or are they fundamental?
2. **Quantum gravity:** How do FD and IC manifest at Planck scale?
3. **Interpretation:** Which interpretation of quantum mechanics correctly implements IC?
4. **Selection:** What additional principles select our physics among admissible possibilities?
5. **Constants:** Is there a deeper explanation for specific values ($\alpha \approx 1/137$, etc.)?

These questions define ongoing research directions. The admissibility framework provides a structural skeleton; much remains to be filled in.

Appendices

Appendix A: Executive Spine Summary

```
                    THE ADMISSIBILITY SPINE

    The core logic of the entire framework in one box

  FINITE DISTINGUISHABILITY (FD)
  ───────────────────────────────
  "You can't distinguish infinitely many things with finite
   resources."
  (Capacity constraint; uniformity is the hinge)

  Consequences:
  • No unbounded information in bounded region
  • Uncertainty principle (finite phase-space resolution)
  • Bekenstein/holographic bounds
```

```
IRREVERSIBLE COMMITMENT (IC)
────────────────────────────
"Some processes can't be undone with finite resources."
(Operational; contractivity of D)

Consequences:
• Entropy monotonicity (Second Law)
• Landauer bound (erasure costs k_B T ln 2 per bit)
• Arrow of time

STRUCTURE PRESERVATION
──────────────────────
"Reversible = D-isometry" (from contractivity)
+ Quantum axioms → Hilbert

Consequences:
• Unitarity of quantum evolution (Wigner → Stone)
• Schrödinger evolution form
• Symplectic structure in classical limit

LAWS AS FIXED POINTS
────────────────────
"Laws = stable under {C_λ}" (RG-structural sense)

Consequences:
• RG-like emergence of effective theories
• Universality of admissibility-compliant laws
• Specific content requires symmetry + initial data
```

Appendix B: Notation

Symbol	Meaning	Plain English
$\mathcal{O}(R,E,\tau,M,\delta)$	Operational process class	All operations doable with stated resources
$D(\cdot,\cdot)$	Contractive distinguishability measure	How different two states are
$\hbar$	Reduced Planck constant	Fundamental quantum of action ($\sim 10^{-34}$ J·s)
k_B	Boltzmann constant	Converts temperature to energy ($\sim 10^{-23}$ J/K)
c	Speed of light	$\sim 3\times 10^{8}$ m/s
G	Newton's gravitational constant	Strength of gravity
$\mathcal{H}$	Hilbert space	Mathematical space of quantum states

Symbol	Meaning	Plain English
ρ, σ	Density matrices	Quantum states (including mixed states)
Φ	Quantum channel	Physical process on quantum states
$\{C_\lambda\}$	Coarse-graining family	Maps from fine to coarse descriptions
L_λ	Effective dynamics at scale λ	Laws at resolution λ
$(P, \preccurlyeq)$	Causal poset	Events with causal ordering
$T \vdash \varphi$	φ is provable in T	System T can derive statement φ
$\ulcorner \varphi \urcorner$	Gödel number of φ	Numerical encoding of formula φ

Appendix C: Key Theorems Used

Theorem	What It Says	How We Use It
Wigner's theorem	Transition-probability-preserving bijections are unitary/antiunitary	Links D-isometry to unitarity
Stone's theorem	Continuous unitary groups have exponential form	Yields Schrödinger evolution form
Kadison's theorem	Affine bijections preserving trace norm are unitary	Alternative route to unitarity
Lovelock's theorem	Einstein's equations are unique given symmetry assumptions	Explains uniqueness of gravity
Data-processing inequality	Physical processes can't increase distinguishability	Contractivity of D; second law
Diagonal lemma	Self-referential sentences exist	Constructs Gödel sentences
Landauer's principle	Erasing one bit costs $\geq k_B T \ln 2$	Thermodynamic cost of irreversibility

---

Appendix D: Classification Summary

Feature	Enforced	Constrained	Compatible	Not Determined
Unitary structure (given quantum axioms)	✓			
Entropy monotonicity	✓			
Landauer bound	✓			
Uncertainty tradeoffs	✓			
Robertson-Schrödinger (given Hilbert)	✓			
Maxwell equations		✓		

Feature	Enforced	Constrained	Compatible	Not Determined
Einstein equations		✓		
Bekenstein bounds			✓	
Holographic principle			✓	
Specific Hamiltonians				✓
Value of $\hbar$				✓
Coupling constants				✓
Matter content				✓
Initial conditions				✓

Appendix E: Glossary for General Readers

Admissibility: Whether something is physically possible; whether it could actually exist or happen in the physical world.

Arrow of time: Why time seems to flow in one direction; why we remember the past but not the future.

Capacity constraint: A bound on what's operationally achievable, not a metaphysical claim about reality.

Causal poset: A mathematical structure representing events with "before/after" relationships, but not necessarily a full time ordering.

Coarse-graining: Ignoring fine details; looking at the big picture rather than the microscopic level.

Contractivity: $D(\Phi(\rho), \Phi(\sigma)) \leq D(\rho, \sigma)$ for physical channels—processes can't increase distinguishability.

D-isometry: Transformation preserving D exactly; characterizes operational reversibility.

Entropy: A measure of disorder, uncertainty, or missing information. Roughly, how many different microscopic arrangements correspond to the same macroscopic state.

Finite Distinguishability (FD): The principle that with finite resources, you can only distinguish finitely many different states. A capacity constraint, not a metaphysical claim.

Fixed point: Something that doesn't change when you apply some operation to it. Physical laws are fixed points of coarse-graining.

Gödel's theorem: The mathematical result that consistent formal systems have true but unprovable statements.

Hamiltonian: The mathematical object encoding a system's total energy and how it evolves in time.

Hilbert space: The mathematical space in which quantum states live; an infinite-dimensional generalization of ordinary 3D space.

Holographic principle: The idea that the information in a 3D region is bounded by its 2D surface area.

Irreversible Commitment (IC): The principle that some physical processes cannot be undone with finite resources. Defined relative to $\mathcal{O}$, not to all conceivable operations.

Landauer's principle: Erasing one bit of information requires at least k_B T ln 2 of energy.

Meta-level: A perspective from "outside" a system, able to talk about the system rather than just within it.

Process ontology: A view that reality consists fundamentally of happenings/changes rather than static objects.

Second law of thermodynamics: Entropy never decreases in isolated systems; disorder tends to increase.

Tick: A primitive, minimal event of change—the most basic "something happened."

Uniform Readout: Discrimination cost independent of state index—the "interesting assumption" whose failure manifests FD.

Unitary: A transformation that preserves the structure of quantum states; reversible quantum evolution.
