The Enforceability of Distinction: Finite Capacity, the Quantum Skeleton, and the Correlation Space

E. S. Brooke

Independent Researcher, San Anselmo, California, USA
Corresponding author: ebrooke@cleanwater1.com
ORCID: https://orcid.org/0009-0001-2261-4682

Abstract

The Standard Model's 26 free parameters, three gauge groups, and three generations are inserted by hand. This paper proposes that all such structure is *forced* by a single constraint: maintaining any physical distinction—spin-up versus spin-down, here versus there, proton versus neutron—costs a share of a finite endowment. We formalize this as Axiom A1: enforcement capacity is finite and positive.

From A1 alone—supplemented by two definitional sub-clauses (Multiplicity, Non-Degeneracy), both subsequently derived from the framework's field content—we prove four structural lemmas: L_{ε^*} (Minimum Cost), L_{loc} (Locality), L_{nc} (Non-Closure), and L_{irr} (Irreversibility). An operational postulate L_{Δ} (Competition), derived from A1 and a perturbation model, establishes the superadditivity separating quantum structure from classical overflow.

A bridge theorem shows that superadditivity forces order-dependent enforcement—the operational content of noncommutativity—on an abstract state set without assuming Hilbert-space structure. Standard results (Gel'fand–Naĭmark–Segal construction, Gleason's theorem, Wedderburn–Artin, Skolem–Noether) whose hypotheses are verified in each case then yield Hilbert space, the Born rule, completely positive trace-preserving dynamics, tensor products, gauge symmetry, and von Neumann entropy as committed capacity.

The framework is dimensionless: all predictions are pure numbers requiring one energy scale for unit conversion. Twenty-three theorems with constructive witnesses and executable code references are provided. Specific gauge groups, particle content, and the full set of zero-parameter predictions are derived in the companion papers.

Keywords: quantum foundations · finite information · enforcement capacity · quantum reconstruction · admissibility · operator algebra

1. Introduction: from free parameters to forced structure

The Standard Model of particle physics is the most precisely tested scientific theory in history. It predicts the electron's magnetic moment to twelve decimal places, accounts for every particle observed at the Large Hadron Collider, and has survived four decades of experimental assault without a single confirmed failure. And yet it explains almost nothing about its own structure. Why three generations of fermions and not two or seven? Why the gauge group $SU(3) \times SU(2) \times U(1)$ and not some other? Why 26 free parameters with these particular values? The Standard Model does not say. It takes these as inputs and never looks back.

This paper proposes that the structure is not chosen but *forced*—by a constraint so elementary that it has been hiding in plain sight.

1.1 The finite universe

What if the reason physics looks the way it does—quantum mechanics, three generations, three gauge groups, four spacetime dimensions, a cosmological constant 122 orders of magnitude below the natural scale—was never a choice, but an inevitability—forced by the single fact that the universe's resources are finite?

Start with what is least controversial. The universe began. Roughly 13.8 billion years ago, everything that exists emerged from a state of extraordinary compression—finite energy, finite volume, finite entropy. It is still expanding, still cooling, still diluting. The Big Bang was, at its core, a *flow event*: a finite quantity of something—energy, degrees of freedom, information capacity—entered the arena, unfolded into structure, and produced everything we observe. A flow event with a beginning implies a finite throughput. A finite throughput implies a bounded total. And yet, despite this, the idea that the universe operates under a hard *capacity constraint* has received remarkably little systematic attention. Physics has spent a century building theories on continuous symmetries, infinite-dimensional Hilbert spaces, and unbounded operator spectra—mathematical machinery that treats capacity as unlimited, even though the universe it describes plainly began with a finite endowment and will plainly end with one.

The mathematics of finite systems, by contrast, has long understood that bounded resources force structure. Ashby's Law of Requisite Variety [1]—the foundational theorem of cybernetics—established that a finite controller can regulate only as many distinct states as its own capacity permits. Regulation requires resources; bounded resources limit what can be regulated; and the *structure* of the controller is constrained by what it must regulate. This principle has been applied to control theory, biology, organizational design, and information theory—but never, until now, as a foundational axiom of physics.

The primitive concept we need is the *distinction*: a partition of possibilities into two classes, "this" and "not this." Spencer-Brown's *Laws of Form* [2] demonstrated that the act of drawing a distinction is the most elementary operation from which logical and algebraic structure can be generated—an insight that influenced both the theory of computation and second-order cybernetics. In physics, distinctions are everywhere: spin-up versus spin-down, here versus there, proton versus neutron. Each is a binary fact that the universe keeps true, moment to moment, against everything that would blur it out.

Distinctions do not exist in isolation; they are bound together by *correlations*. The spin of an electron is correlated with the magnetic field it sits in. The polarizations of an entangled pair are correlated across arbitrary distances [3]. A hydrogen atom is a web of correlations between the electron's quantum numbers and the proton's. Correlations are what give structure to the physical world; without them, distinctions would be isolated labels signifying nothing.

But here is a pattern that standard physics has quantified in extraordinary detail without ever recognizing as a single phenomenon: *maintaining a distinction costs something.* A spin state is stable because an energy gap protects it—thermal fluctuations below the gap cannot flip it. A qubit in a quantum computer stays coherent only as long as error correction actively shields it [4]. A proton stays distinct from a neutron because QCD binding energies make the mass difference robust against perturbations. In every case, nature commits resources to keeping a physical fact physical. We call this commitment **enforcement**: the minimum expenditure required to maintain a distinction against all perturbations that would destroy it.

Enforcement is not observation. It is not measurement. It is what nature does to keep physical facts physical, whether or not anyone is looking. Physics has quantified each instance precisely—energy gaps in electron volts, free energies in units of kT [5], error-correction thresholds as dimensionless rates [4], string tensions in GeV/fm [15], entropy bounds in bits [23]—but each sub-discipline uses its own currency, its own formalism, its own units. There is no exchange rate between them. There is no concept of a total budget. The Standard Model can compute the binding energy of a proton to four significant figures; what it cannot tell you is that this binding energy, the energy gap protecting a spin state, the error-correction overhead preserving a qubit, and the Bekenstein bound limiting a black hole are all denominated in the same currency—enforcement capacity—and that the sum is bounded.

Return now to the Big Bang. If the universe began as a finite flow event, and if maintaining every physical distinction costs a share of that finite endowment, then the total number of distinctions the universe can simultaneously maintain is *bounded*. Not approximately, not in some limit, but absolutely: there is a budget, and everything must fit within it. This is Ashby's Law applied to the universe itself: the variety of physical law is limited by the capacity of the enforcer.

This is the axiom.

A1: Enforcement capacity is finite and positive.

One sentence. No mention of quantum mechanics, gauge groups, spacetime, or particles. Just: the cost of maintaining distinctions is bounded. The rest of this paper shows that this single constraint—applied through standard mathematics (the GNS construction [6], Gleason's theorem [7], the Wedderburn–Artin decomposition [8, 10], Skolem–Noether [11, 12])—forces quantum structure, gauge symmetry, the seeds of particle content, gravitational scaling, and the thermodynamic arrow of time as necessary consequences. Not choices. Not parameters. Not inputs. Consequences.

The present paper establishes the foundation: the axiom, four structural lemmas, an operational postulate connecting finite resources to superadditivity, a bridge theorem connecting superadditivity to algebraic noncommutativity, and the structural skeleton from which the companion papers derive the Standard Model's 26 parameters, three generations, and 47 zero-parameter predictions.

1.2 The enforcement spectrum: from easy physics to hard problems

The ratio of enforcement demand to available capacity—E/C—organizes all of physics into a single spectrum (Figure 1). This ratio explains why some physics is "easy," why some is "hard," and why existing formalisms break down where they do.

The quiet regime ($E/C \ll 1$). A single photon maintains two polarization states; a hydrogen atom holds a handful of quantum numbers with well-separated energy levels. Nothing competes for resources. Perturbation theory converges, exact solutions exist, and QED predicts the electron's anomalous magnetic moment to 12 decimal places [13]. The physics is simple because the budget is barely taxed.

The competitive regime ($E/C \sim 0.3$–0.6). The budget begins to bind. Two electrons cannot maintain identical quantum numbers because the joint enforcement cost overflows the local budget—the Pauli exclusion principle *is* non-closure. Decoherence appears because system–environment correlations commit capacity at both interfaces, and the committed capacity is

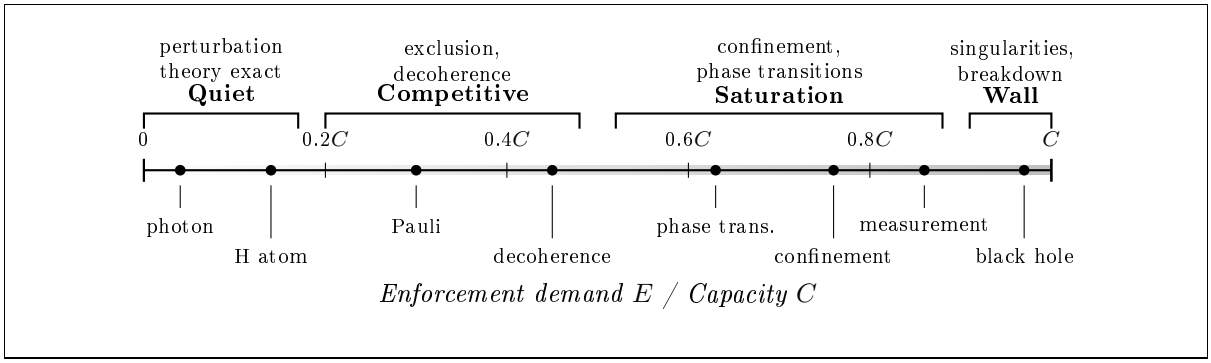


Figure 1: The enforcement spectrum. Physical systems arranged by the ratio of enforcement demand to capacity (E/C). Standard physics works exactly in the quiet regime, requires approximations in the competitive regime, and breaks down—singularities, the measurement problem, the information paradox—at the capacity wall. The hard problems of physics cluster where $E/C \to 1$

locally unrecoverable [14]. Perturbation theory still works but requires increasingly sophisticated approximations: the system is managing genuine competition for finite resources.

The saturation regime ($E/C \sim 0.6$–0.95). The budget is nearly exhausted and the physics becomes dramatic. In QCD, a free quark's unscreened color charge costs more to enforce the further it separates; when the cost exceeds interface capacity, the string snaps and produces color-singlet hadrons—fewer distinctions, lower demand [15]. At a phase transition, thermal perturbation costs consume the enforcement budget for the ordered state until the system can no longer afford it and transitions to disorder [16]. In measurement, cross-system correlations multiply as apparatus and environment entangle, saturating the budget until only one branch survives [17].

The capacity wall ($E/C \to 1$). A black hole has saturated its enforcement capacity: $S = A/4\ell_P^2$ counts the maximum simultaneously maintainable distinctions for interface area A [23]. Hawking radiation is the slow release of committed capacity [18]. General relativity's singularity—infinite curvature, infinite density [19]—is the formalism announcing that enforcement demand has exceeded any finite budget, which A1 forbids. The resolution (Paper 6): curvature saturates at a finite maximum, producing a bounce rather than a singularity.

The pattern. The standard formalism works beautifully in the quiet regime, requires approximations in the competitive regime, and breaks down—singularities, the measurement problem, the information paradox—at the capacity wall. These are not separate problems. They are the same problem: the formalism has no finite enforcement budget and cannot describe what happens when the budget runs out. A1 supplies that concept.

1.3 A1 is dimensionless

A1 is dimensionless: C, ε^*, and all quantities derived from them are pure numbers. The framework's predictions—mixing angles, mass ratios, coupling constants, $\Lambda \cdot G_N = 3\pi/102^{61}$—are dimensionless ratios. Converting to physical units (GeV, meters, seconds) requires one dimensional input: an energy scale such as the Planck mass M_{Pl} or the electroweak scale $v_{\text{EW}} \approx 246$ GeV. This is a consequence of dimensional analysis (no framework can derive meters from pure mathematics), not a limitation specific to the Admissibility Physics Framework (APF). It is the same

single input the Standard Model requires.

1.4 Context: reconstruction and finite-information programs

This work enters a well-established tradition. Since Hardy's seminal reconstruction of quantum theory from five operational axioms [20], a sustained program has sought to derive the Hilbert-space formalism from physically motivated principles rather than assume it. Chiribella, D'Ariano, and Perinotti [21] showed that six informational axioms—including causality, distinguishability, and purification—select quantum theory uniquely from among generalized probabilistic theories. Masanes and Müller [22] achieved a similar result from physical requirements on state spaces and measurements. These reconstruction programs demonstrate that "obscure postulates" can be replaced with principled constraints, and that the constraints themselves can be surprisingly rigid.

A1 (finite enforcement capacity) is closest in spirit to the finite-information tradition in quantum gravity. Bekenstein's universal entropy bound [23] established that bounded systems carry limited information; 't Hooft [24] and Susskind [25] argued that finite degrees of freedom force holographic structure; Bousso [26] synthesized these into the covariant entropy bound. Jacobson [27] showed that Einstein's field equations emerge from thermodynamic reasoning about horizon entropy. A1 can be read as the common ancestor of these ideas: a primitive statement that enforcement resources are finite, from which both the quantum formalism and the gravitational structure follow as consequences rather than assumptions.

What distinguishes this program from prior reconstructions is scope: where Hardy and Chiribella derive the formalism of quantum theory, the APF derives the formalism *and* the specific content—gauge groups, generations, mixing angles, cosmological parameters—from a single axiom with zero free parameters.

1.5 The APF engine: code as specification

Every theorem in this paper series is implemented as an executable check function in the APF codebase—a modular Python package organized by derivation layer (`core`, `gauge`, `generations`, `spacetime`, `gravity`, `cosmology`). Each check function encodes the theorem's statement, verifies its numerical content against exact arithmetic (using Python's `Fraction` type), records its dependencies, and returns a structured result. A separate validation module compares derived predictions to experimental benchmarks (Planck 2018 [28], PDG 2024 [29], NuFit 6.0 [30]) behind a firewall that isolates empirical data from the derivation chain. Code references of the form `(check_X, module.py)` appear throughout Papers 1–7.

1.6 Road map: from axiom to physics

This paper establishes the axiom and derives four foundational lemmas plus one operational postulate. The derivation order is strict and acyclic:

Order	Result	Type	Dependencies	Used by
1	L_{ε^*} (Minimum Cost)	Lemma	A1	all below
2	L_{loc} (Locality)	Lemma	A1, L_{ε^*}, M, NT	L_{irr}, Papers 2–7
3	L_{nc} (Non-Closure)	Lemma	A1, L_{ε^*}, M	L_{irr}, T1
4	L_{Δ} (Competition)	Postulate	A1, Pert. Model	T1, L_{irr}
5	L_{irr} (Irreversibility)	Lemma	L_{nc}, L_{Δ}, L_{loc}	Papers 3, 4, 5

These feed into a bridge theorem (T1: Order-Dependence → Noncommutativity), which opens the door to the quantum skeleton (T2–$T_{\text{canonical}}$) via standard mathematical results.

2. The axiom: finite enforcement capacity

The entire framework rests on a single physical claim. A1 says that nature's capacity to enforce distinctions is finite and positive—nothing more.

Axiom A1 (Finite Enforcement Capacity)

For any admissible state σ on a causally connected region R,

$$\sum_{d \in \mathcal{D}(\sigma, R)} \varepsilon(d) \;\leq\; C(R) \;<\; \infty \tag{1}$$

Variables:

- Ω — the abstract state set: the collection of all possible configurations of enforced distinctions. No algebraic or topological structure is assumed on Ω at this stage.
- $\sigma \in \Omega$ — a state (a particular configuration of enforced distinctions on region R).
- R — a causally connected region of the universe.
- $\mathcal{D}(\sigma, R)$ — the set of independently enforceable distinctions maintained in state σ on region R. Each element is a distinct physical difference that the universe actively enforces.
- $\varepsilon(d)$ — the enforcement cost of distinction d: the capacity committed to maintaining d against collapse. Bounded below: $\varepsilon(d) \geq \varepsilon^* > 0$ for all meaningful d.
- ε^* — the minimum enforcement cost per distinction, strictly positive. The framework requires $\varepsilon^* > 0$ but never requires its numerical value; all downstream results depend only on finiteness and positivity.
- $C(R)$ — the total enforcement capacity of region R: the finite upper bound on how much enforcement can be simultaneously committed within R.

Code reference: `check_A1` in `core.py`.

This is a constraint on what nature can do, not on what we can observe. The specific value of C is never required; only finiteness and positivity matter.

2.1 What A1 is not: no choices, no specifics

A1 does not say how large C is. It does not say what the distinctions are made of. It does not say whether the universe is classical or quantum, continuous or discrete, four-dimensional or otherwise. All of those properties are derived in Sections 4–6 as consequences of the axiom, not written into it.

The axiom is not a simplifying assumption. If enforcement resources are unlimited, there is no cap on how many distinctions can be maintained. If there is no cap, then every conceivable pattern can persist—nothing is forced to give way. If nothing is forced, there are no constraints. And if there are no constraints, there is no physics: no selection, no structure, no law. Finite enforceability is the minimum condition for physical law to exist.

2.2 Sub-clauses and their derivation: M, NT, and closure

Two definitional sub-clauses accompany A1, specifying that the universe under discussion has enough structure to be interesting. Think of them the way one thinks of the axiom of infinity in set theory: they carry no predictive content beyond a definitional role.

Sub-Clauses M (Multiplicity) and NT (Non-Degeneracy)

Sub-clause M (Multiplicity): At least two distinguishable subsystems exist: $|\mathcal{D}| \geq 2$.

Sub-clause NT (Non-Degeneracy): Not all subsystems are identical. There exist subsystems S_i, S_j with $C(S_i) \neq C(S_j)$.

Without M, the axiom is trivially satisfied by a single point with capacity C—no physics emerges because there is nothing to distinguish. Without NT, every subsystem has capacity C/N and no asymmetry develops—no gauge group selection, no generations, no symmetry breaking.

Code reference: `check_M, check_NT` in `core.py`.

Both sub-clauses are subsequently *derived* from the framework's own field content. This closes the logical loop: M and NT are assumed at the axiom layer, then proved to follow from structures that the axiom generates. The logical status is self-consistency, not circularity—the derivation chain is acyclic (A1 $\rightarrow$ lemmas $\rightarrow$ gauge content $\rightarrow$ M, NT), and the sub-clauses are used only in the initial derivation of L_{loc} (Locality). The derivations ($L_{M,\text{derived}}$ and $L_{NT,\text{derived}}$) are given in Appendix C.

3. Definitions: the language of enforcement

Before deriving consequences from A1, we fix the five primitive concepts that the framework uses. All are defined operationally in terms of enforcement on the abstract state set Ω, not on a Hilbert space. The Hilbert-space representation is *derived* in Section 5.

Definitions (Distinction, Enforcement Cost, Interface, Admissible State, Capacity)

Distinction. A binary predicate $d : \Omega \rightarrow \{0, 1\}$ that partitions the state set into two non-empty subsets. Each distinction is a physical fact the universe actively maintains: spin-up rather than spin-down, this particle here rather than there, proton rather than neutron. A distinction is *enforceable* if it is *robust*: no admissible perturbation of cost less than $\varepsilon(d)$ can destroy it.

Variables:

- Ω — the abstract state set. No algebraic, topological, or linear structure is assumed. Ω is the pre-geometric arena on which all other structures are built.
- $d : \Omega \rightarrow \{0, 1\}$ — a distinction. Think of d as a "yes/no" physical fact that the universe maintains.
- *Robustness* — a distinction d is robust (hence enforceable) iff $\varepsilon(d) > 0$: there exists a minimum perturbation cost below which d cannot be destroyed.

Enforcement Cost. The minimum cost of destroying a distinction d:

$$\varepsilon(d) \;=\; \inf\bigl\{\mathrm{cost}(p) : p \text{ is an admissible perturbation that destroys } d\bigr\}.$$

A distinction is *robust* (hence enforceable) iff $\varepsilon(d) > 0$: every perturbation capable of destroying d costs at least $\varepsilon(d)$.

Interface. A locus Γ at which enforcement occurs. Each interface has a capacity budget $C(\Gamma) > 0$ and enforces a subset of the total distinctions.

Admissible State. A state $\sigma \in \Omega$ is admissible if the total enforcement cost of all simultaneously maintained distinctions does not exceed the available capacity:

$$\sum_{d \in \mathcal{D}(\sigma, R)} \varepsilon(d) \;\le\; C(R).$$

Enforcement Capacity. $C(R)$ is the total enforcement capacity on region R. It decomposes over interfaces: $C(R) = \sum_{\Gamma \in R} C(\Gamma)$. The global capacity C_{total} counts the total number of distinguishable capacity types.

Perturbation Model. Perturbations targeting distinct distinctions can combine. The perturbation space for jointly-enforced distinctions $\{d_1, d_2\}$ contains combined attacks—perturbations that simultaneously threaten both d_1 and d_2—not merely perturbations targeting each individually. Formally: if $\mathcal{P}(d_i)$ is the set of perturbations threatening d_i, then the perturbation space for the joint enforcement of $\{d_1, d_2\}$ satisfies $\mathcal{P}(\{d_1, d_2\}) \supseteq \mathcal{P}(d_1) \cup \mathcal{P}(d_2)$, with the inclusion strict whenever d_1 and d_2 share enforcement resources. This is mathematical infrastructure (a regularity condition on the enforcement model), not a physics axiom; it carries the same logical status as the specification that Ω is a set or that d is a binary predicate.

4. The four lemmas and the competition postulate

The derivation chain from A1 to the full framework passes through four structural lemmas and one operational postulate (L_Δ, derived from the perturbation model in Section 3). The derivation order is strict and acyclic:

$$\begin{aligned} L_{\varepsilon^*}\ (\mathrm{A1}) \;\to\; L_{\mathrm{loc}}\ (\mathrm{A1} + L_{\varepsilon^*} + \mathrm{M} + \mathrm{NT}) \;\to\; L_{\mathrm{nc}}\ (\mathrm{A1} + L_{\varepsilon^*} + \mathrm{M}) \\ \to\; L_\Delta\ (\mathrm{A1} + \text{Pert. Model}) \;\to\; L_{\mathrm{irr}}\ (L_{\mathrm{nc}} + L_\Delta + L_{\mathrm{loc}}) \end{aligned}$$

Note: L_{loc} (Locality) and L_{nc} (Non-Closure) are logically independent—neither requires the other. Both require L_{ε^*}. They are ordered here to match the derivation logic: locality establishes the interface structure, then non-closure establishes that interfaces overflow. L_Δ (Competition) connects non-closure to superadditivity via the perturbation model.

4.1 L_{ε^*} (Minimum Cost): no infinitesimal distinctions

If distinctions could be maintained at arbitrarily small cost, an unbounded number of them could fit within any finite capacity—and finiteness would carry no structural force.

Lemma L_{ε^*} (Minimum Cost)

Statement. There exists $\varepsilon^* > 0$ such that every enforceable distinction costs at least ε^*:

$$\varepsilon(d) \;\geq\; \varepsilon^* \;>\; 0 \qquad \forall\, d \in \mathcal{D}. \tag{2}$$

Variables:

- $\varepsilon(d)$ — enforcement cost of distinction d (Section 3).
- $\varepsilon^* = \inf_{d \in \mathcal{D}} \varepsilon(d)$ — the infimum of enforcement costs over all enforceable distinctions.
- C — total enforcement capacity (from A1).
- *Robustness* — a distinction is enforceable iff $\varepsilon(d) > 0$: every destroyer costs at least $\varepsilon(d)$ (Section 3).

Proof (by contradiction).
Step 1 (Hypothesis). Suppose $\varepsilon^* = 0$. Then for every $n \in \mathbb{N}$, there exists an enforceable distinction d_n with $\varepsilon(d_n) < 1/n$. By M, there exist at least countably many pairwise-admissible distinctions in this low-cost regime (no global obstruction prevents their simultaneous enforcement when each costs $< 1/n$).
Step 2 (Cheap destruction exists). Consider any such d_n. Since $\varepsilon(d_n) < 1/n$, the infimum of destructive costs is below $1/n$. Therefore there exists an admissible perturbation of cost $< 1/n$ that destroys d_n.
Step 3 (Contradiction with robustness). As $n \to \infty$, the sequence $\{d_n\}$ yields enforceable distinctions destroyable at cost $\to 0$. But each d_n was assumed enforceable, meaning the universe actively maintains it. A distinction that can be destroyed at arbitrarily small cost is operationally indistinguishable from "not enforced at all"—the enforcement mechanism cannot distinguish "d holds" from "d does not hold" if toggling between them costs less than any positive threshold. This contradicts the assumption that each d_n is a physically maintained fact. Contradiction.

Conclusion. $\varepsilon^* > 0$. The maximum number of simultaneously enforceable distinctions is therefore bounded:

$$|\mathcal{D}| \;\leq\; \left\lfloor \frac{C}{\varepsilon^*} \right\rfloor \;<\; \infty. \tag{3}$$

Key premise: "meaningful = robust under admissible perturbation" is definitional within the framework (Section 3). This identification is the framework's definition of physical meaningfulness; it is not an additional postulate, but it is the assumption on which L_{ε^*} rests.

Code reference: `check_L_epsilon_star` in `core.py`.

4.2 L_{loc} (Locality): enforcement must distribute

Finite capacity with enough physical structure forces enforcement to split over multiple independent interfaces. The argument requires only L_{ε^*} (for the counting bound) plus the richness conditions M and NT. It does *not* require non-closure.

Lemma L_{loc} (Locality)

Statement. A1 + L_{ε^*} (Minimum Cost) + M + NT $\implies$ enforcement decomposes over independent interfaces.

Variables:

- Γ — an interface (enforcement locus) with capacity budget $C(\Gamma)$.
- $n_{\max}(\Gamma) = \lfloor C(\Gamma)/\varepsilon^* \rfloor$ — maximum distinctions enforceable at a single interface.
- N_{phys} — total number of independently meaningful distinctions.
- *Independent interfaces* — two interfaces Γ_1, Γ_2 are independent iff: (i) each has its own capacity budget ($C(\Gamma_1), C(\Gamma_2)$ allocated separately); (ii) an operation at Γ_1 cannot modify the enforcement state at Γ_2 without consuming additional capacity; (iii) budget accounting at each is self-contained. *Any* partitioning of a single budget into sub-budgets satisfying (i)–(iii) constitutes independent interfaces—the distinction between "one locus with internal partitions" and "multiple loci" is definitional, not physical.

Proof (four steps).
Step 1 (Single-interface capacity bound). By A1, $C < \infty$. By L_{ε^*}, each distinction costs $\geq \varepsilon^* > 0$. A single interface can enforce at most $\lfloor C/\varepsilon^* \rfloor$ distinctions.
Step 2 (Richness exceeds single-interface capacity). By M + NT, the number of independently meaningful distinctions N_{phys} exceeds any single interface's capacity: $N_{\text{phys}} > \lfloor C_{\max}/\varepsilon^* \rfloor$.
Step 3 (Distribution is forced). Since N_{phys} exceeds the single-interface bound, the universe must partition enforcement into sub-budgets $C_1, \ldots, C_k$ with $\sum_j C_j = C_{\text{total}}$, each handling a subset of distinctions. Coordination between sub-budgets itself requires enforcement (transferring capacity from C_1 to C_2 means modifying both enforcement states, costing $\geq \varepsilon^*$ by L_{ε^*}). Under A1, unlimited coordination would consume the entire budget. Concretely: maintaining global coherence across k sub-budgets requires at least $\binom{k}{2}\varepsilon^*$ coordination cost, which for large k overwhelms C_{total} and leaves nothing for the distinctions themselves. The budget itself enforces the separation. Therefore the sub-budgets are operationally independent.
Step 4 (Independence = locality). Independent interfaces with separate budgets means: (a) no interface has global access (each enforces a proper subset); (b) enforcement decomposes: $E_{\text{total}} = \sum_i E(\Gamma_i)$; (c) subsystems at disjoint interfaces are independently enforceable. This is locality.

Constructive witness. Six distinctions with $\varepsilon = 2$ and $C = 10$:

$$\text{Single interface:} \quad E_{\text{full}} = \tfrac{39}{2} = 19.5 > 10 = C \quad \text{(inadmissible)} \tag{4}$$

$$\text{Two interfaces:} \quad E_{\text{left}} = E_{\text{right}} = \tfrac{33}{4} = 8.25 \leq 10 = C \quad \text{(admissible)} \tag{5}$$

Countermodel. $|\mathcal{D}| = 1$: a single interface with $C \geq \varepsilon^*$ suffices. Confirms M is necessary.

Code reference: `check_L_loc` in `core.py`.

4.3 L_{nc} (Non-Closure): composition overflows capacity

Non-closure is the engine behind competition, saturation, and selection. It requires only finite capacity and positive costs—no spatial structure or locality is needed. L_{nc} and L_{loc} are logically independent: each depends on L_{ε^*}, but neither depends on the other.

The idea is immediate from a concrete example. Suppose the total enforcement budget is $C = 10$.

Sector S_1 (say, a collection of color-charge distinctions) costs $E_1 = 6$ to enforce; sector S_2 (say, a collection of electroweak distinctions) also costs $E_2 = 6$. Each sector is individually admissible: $6 \leq 10$. But their composition demands $E_1 + E_2 = 12 > 10$—the budget overflows. The universe cannot afford both at full strength simultaneously. Something must give: one sector is suppressed, resources are redistributed, or new structure emerges to manage the competition. This is non-closure.

Lemma L_{nc} (Non-Closure)

Statement. A1 + L_{ε^*} (Minimum Cost) + M $\implies$ the set of admissible states is not closed under composition.

Variables:

- C — total capacity budget (from A1).
- E_i — enforcement demand of sector i: the total cost of enforcing all of sector i's distinctions.
- $\Omega_{\text{adm}} = \{\sigma \in \Omega : \sum_d \varepsilon(d) \leq C\}$ — the set of admissible states.

Proof (constructive).

Step 1 (Existence of competing sectors). By M, at least two distinguishable sectors exist. By L_{ε^*} (Minimum Cost), each sector costs $\geq \varepsilon^* > 0$. Therefore there exist sectors S_1, S_2 with $E_1, E_2 > 0$.

Step 2 (Witness). Let $C = 10$. Let $E_1 = 6$ and $E_2 = 6$. Each is individually admissible:

$$E_1 = 6 \leq 10 = C, \qquad E_2 = 6 \leq 10 = C. \tag{6}$$

Their composition has total demand:

$$E_1 + E_2 = 12 > 10 = C \qquad \text{(inadmissible)}. \tag{7}$$

Step 3 (Generalization). For any n sectors each with $E_i > C/n$:

$$\sum_{i=1}^{n} E_i \;>\; n \cdot \frac{C}{n} \;=\; C. \tag{8}$$

Such sectors exist whenever M holds (multiple subsystems) and L_{ε^*} holds (each costs > 0). L_{ε^*} also blocks subadditive evasion: no cost function satisfying $\varepsilon(d) \geq \varepsilon^* > 0$ can make the sum converge below C for arbitrarily many sectors.

Conclusion. Ω_{adm} is not closed under composition. Note that this is a purely arithmetic consequence of finite capacity and positive costs—a classical "knapsack" overflow. It does *not* by itself imply superadditivity ($\Delta > 0$) or interference; those require the additional structure established in L_Δ below.

Code reference: `check_L_nc` in `core.py`.

4.4 L_Δ (Competition): superadditivity from combined perturbations

Pause here and take stock. L_{nc} says that two individually affordable collections of distinctions can be jointly unaffordable—a finite budget overflows when you try to maintain everything at once. This is a real constraint, but it is entirely classical: a shipping container has finite volume, and two loads that each fit individually may not fit together. Nothing about non-closure, by itself, produces interference, noncommutativity, or any distinctively quantum phenomenon.

Classical budget overflow gives you a knapsack problem. It does not give you quantum mechanics.

The passage from "finite budget" to "noncommutativity" requires something stronger: that jointly enforcing two distinctions costs *strictly more* than the sum of individual costs—that is, superadditivity. The additional cost arises because shared resources create new vulnerabilities: perturbations that could not threaten either distinction alone can threaten both when they draw on the same pool. This is derived from A1 and the perturbation model (Section 3).

L_Δ (Competition): Superadditivity from Shared Budgets

Statement. A1 + Perturbation Model $\implies$ for distinctions d_1, d_2 sharing an enforcement budget, the joint enforcement cost is superadditive:

$$\varepsilon(\{d_1, d_2\}) \;\geq\; \varepsilon(d_1) + \varepsilon(d_2) + \Delta(d_1, d_2), \qquad \Delta(d_1, d_2) > 0. \tag{9}$$

Variables:

- $\varepsilon(\{d_1, d_2\})$ — the joint enforcement cost: the minimum cost of maintaining *both* d_1 and d_2 simultaneously against all admissible perturbations, including combined attacks.
- $\Delta(d_1, d_2)$ — the superadditivity surplus: the additional cost of defending against combined perturbations beyond the sum of individual defenses.
- $\mathcal{P}(\{d_1, d_2\})$ — the joint perturbation space, which by the Perturbation Model strictly contains $\mathcal{P}(d_1) \cup \mathcal{P}(d_2)$ for resource-sharing distinctions.

Proof (geometric).
Step 1 (Individual defense). Enforcing d_1 alone requires defending against perturbations in $\mathcal{P}(d_1)$ at cost $\varepsilon(d_1)$. Likewise d_2 alone at cost $\varepsilon(d_2)$.
Step 2 (Joint defense is harder). By the Perturbation Model, when d_1 and d_2 share an enforcement budget, the joint perturbation space $\mathcal{P}(\{d_1, d_2\})$ strictly contains $\mathcal{P}(d_1) \cup \mathcal{P}(d_2)$: there exist combined attacks that simultaneously threaten both distinctions. Defending against a strictly larger perturbation space costs strictly more than the sum of individual defenses.
Step 3 (Budget sharing is generic). By A1, the total budget C is finite. By M, multiple subsystems share this budget. Therefore resource-sharing is the generic case, not the exception: most pairs of distinctions draw on the same finite pool, and $\Delta(d_1, d_2) > 0$ for those pairs.

Countermodel (necessity of perturbation model). If perturbations cannot combine—if $\mathcal{P}(\{d_1, d_2\}) = \mathcal{P}(d_1) \cup \mathcal{P}(d_2)$ with no combined attacks—then the joint defense cost equals the sum of individual costs: $\Delta = 0$. Enforcement is additive, nonclosure still holds (classical overflow), but no superadditivity arises. This is the classical knapsack model. The perturbation model is what separates the quantum case from the classical case.

Logical status. L_Δ is derived from A1 and the perturbation model (Section 3), which is mathematical infrastructure—a regularity condition on how perturbations compose—not a separate physics axiom. The single-axiom narrative is preserved: A1 is the physics; the perturbation model is the mathematical setting, on the same footing as "Ω is a set" or "d is a binary predicate."

Code reference: `check_L_irr (superadditivity witness)` in `core.py`.

4.5 L_{irr} (Irreversibility): correlations cannot be locally undone

Irreversibility does not require dissipation or a special arrow of time. It arises from a structural fact: cross-interface correlations commit capacity that no local observer can reclaim.

Consider the simplest scenario in which this can happen. A system S and its environment E sit at separate interfaces Γ_S and Γ_E, each with capacity budget $C = 15$. Initially each maintains its own distinctions: s (a system distinction, costing 4 at Γ_S and 2 at Γ_E) and e (an environment distinction, costing 2 at Γ_S and 4 at Γ_E). Both interfaces can afford this: $E_S(\{s, e\}) = 7 \leq 15$ and $E_E(\{s, e\}) = 7 \leq 15$.

Now S and E interact, creating a cross-interface correlation c. By superadditivity (L_Δ), c costs more to enforce jointly with the existing distinctions than its individual cost would suggest: the total at each interface jumps to 15, saturating the budget. The correlation c has committed capacity at *both* interfaces simultaneously. An observer at Γ_S can see that capacity has been consumed but has no operational access to Γ_E—locality (L_{loc}) forbids it. Undoing the correlation would require coordinated action at both interfaces, which no local observer can perform. The transition is irreversible.

The formal proof below verifies this with explicit cost functions and confirms that both L_Δ and L_{loc} are necessary: remove either, and the transition becomes reversible.

Lemma L_{irr} (Irreversibility)

Statement. L_{nc} (Non-Closure) + L_Δ (Competition) + L_{loc} (Locality) $\implies$ irreversibility: there exist admissible transitions that no local observer can reverse.

Variables:

- $\{s, e, c\}$ — three distinctions: s = system, e = environment, c = cross-interface correlation created by S–E interaction.
- Γ_S, Γ_E — two independent interfaces (from L_{loc}), each with capacity $C = 15$.
- $E_S(T), E_E(T)$ — enforcement cost of subset $T \subseteq \{s, e, c\}$ at interface Γ_S and Γ_E respectively.
- $\Delta_S(a, b) = E_S(\{a, b\}) - E_S(\{a\}) - E_S(\{b\})$ — superadditivity surplus at Γ_S (positive by L_Δ).

Proof (five steps with constructive witness).
Step 1 (Witness: enforcement cost functions). Define enforcement costs at each interface, satisfying monotonicity ($T_1 \subseteq T_2 \implies E(T_1) \leq E(T_2)$) and superadditivity ($\Delta > 0$, from L_Δ):

Subset T	$E_S(T)$	$E_E(T)$
$\emptyset$	0	0
$\{s\}$	4	2
$\{e\}$	2	4
$\{c\}$	3	3
$\{s, e\}$	7	7
$\{s, c\}$	10	6
$\{e, c\}$	6	10
$\{s, e, c\}$	15	15

Step 2 (Monotonicity verified). For every $T_1 \subsetneq T_2$: $E_S(T_1) \leq E_S(T_2)$ and $E_E(T_1) \leq E_E(T_2)$. No state is "trapped"—removing a distinction always yields an admissible subset.

Step 3 (Superadditivity verified). The interaction surpluses:

$$\Delta_S(s,c) = 10 - 4 - 3 = 3 > 0, \qquad \Delta_E(e,c) = 10 - 4 - 3 = 3 > 0. \tag{10}$$

Step 4 (All subsets globally admissible). All $2^3 = 8$ subsets satisfy $E_S(T) \leq 15$ and $E_E(T) \leq 15$.
Step 5 (Correlation is locally unrecoverable). The correlation c commits capacity at *both* interfaces:

$$\text{Cost of } c \text{ at } \Gamma_S: \quad E_S(\{s,e,c\}) - E_S(\{s,e\}) = 15 - 7 = 8 > 0, \tag{11}$$
$$\text{Cost of } c \text{ at } \Gamma_E: \quad E_E(\{s,e,c\}) - E_E(\{s,e\}) = 15 - 7 = 8 > 0. \tag{12}$$

Removing c requires coordinated action at *both* Γ_S and Γ_E. By L_{loc} (Locality), an observer at Γ_S has no operational access to Γ_E. The capacity committed to c is permanently inaccessible from S's perspective. The transition $\{s,e\} \to \{s,e,c\}$ is locally irreversible.

Countermodels (necessity).

- *Without L_Δ:* All $\Delta = 0$ (additive costs). Removing c frees $E(\{c\})$ at each interface independently. Fully reversible.
- *Without L_{loc}:* Single interface; observer has global access and can remove c unilaterally. Fully reversible.

Code reference: `check_L_irr` in `core.py`.

5. The quantum skeleton: standard structure from finite resources

The argument so far has been entirely about resources: finite budgets, positive costs, overflow, superadditivity, and the irreversible commitment of capacity across interfaces. Nothing in Sections 2–4 mentions Hilbert spaces, operators, complex numbers, or probability amplitudes. The state set Ω has no algebraic structure; the enforcement cost ε is a real-valued function; the lemmas are proved by arithmetic and set-theoretic reasoning. The entire lemma layer is, deliberately, pre-quantum.

What happens next is the central surprise of the framework: the resource constraints *force* algebraic structure. Superadditivity (L_Δ) means that the order in which you enforce distinctions matters—enforcing d_1 before d_2 leaves a different residual state than enforcing d_2 before d_1. That order-dependence, once you demand that it be faithfully represented, *is* noncommutativity. And noncommutativity, via standard mathematical results whose hypotheses we verify explicitly, is enough to generate the full apparatus of quantum mechanics: Hilbert space, the Born rule, completely positive trace-preserving (CPTP) dynamics, tensor products, and gauge symmetry.

The critical new content in this section is the **bridge theorem T1**, which makes this passage rigorous. Without T1, the claim that "finite resources imply quantum mechanics" would be an interpretive leap. With T1, it is a derivation. The hinge is L_Δ: superadditivity is what separates classical budget overflow (L_{nc}, which produces only knapsack constraints) from quantum noncommutativity (T1, which produces operator ordering).

5.1 Theorem T1 — The Bridge Theorem: order-dependence from capacity competition

This theorem bridges the gap between the resource-theoretic lemmas (which operate on the abstract state set Ω) and the algebraic structures of quantum mechanics (which require operator non-commutativity). The argument does not assume a Hilbert space.

Theorem T1 — The Bridge Theorem (Order-Dependence)

Statement. L_Δ (Competition) + NT + OR0 (Faithfulness) $\implies$ enforcement operations on Ω are order-dependent: there exist distinctions d_1, d_2 such that "enforce d_1 then d_2" $\neq$ "enforce d_2 then d_1."

Variables:

- $\mathfrak{E}_d : \Omega \to \Omega$ — the enforcement operation for distinction d. Concretely: $\mathfrak{E}_d$ maps state σ to the state σ' in which d is active, the budget is reduced by $\varepsilon(d)$ plus any interaction surplus with already-active distinctions, and the set of remaining enforceable distinctions is updated accordingly. This is the minimal operational content of "enforce d": activate, pay, update.
- $m(d_2 \mid d_1)$ — the *marginal cost* of enforcing d_2 given that d_1 is already enforced: $m(d_2 \mid d_1) = \varepsilon(\{d_1, d_2\}) - \varepsilon(\{d_1\})$.
- $\Delta(d_1, d_2) = \varepsilon(\{d_1, d_2\}) - \varepsilon(\{d_1\}) - \varepsilon(\{d_2\})$ — the superadditivity surplus. Positive for resource-sharing distinctions by L_Δ.

Proof (three steps).

Step 1 (Superadditivity forces context-dependent marginal costs). By L_Δ, $\Delta(d_1, d_2) > 0$ for distinctions sharing enforcement resources. Therefore:

$$m(d_2 \mid d_1) = \varepsilon(d_2) + \Delta \;>\; \varepsilon(d_2) = m(d_2 \mid \emptyset). \tag{13}$$

The marginal cost of enforcing d_2 *depends on whether* d_1 *is already enforced.* This is *not* a bookkeeping artifact: the enforcement map $\mathfrak{E}_{d_2}$ consumes a different amount of budget (and therefore leaves a different residual state in Ω) depending on the prior state.

Step 2 (NT creates asymmetric residual states). By NT, $\varepsilon(d_1) \neq \varepsilon(d_2)$. After enforcing d_1 first:

$$\begin{aligned} &\text{Residual budget: } C - \varepsilon(d_1), \\ &\text{Residual enforceable set: } \mathcal{D}_1 = \big\{d : m(d \mid d_1) \leq C - \varepsilon(d_1)\big\}. \end{aligned} \tag{14}$$

After enforcing d_2 first:

$$\begin{aligned} &\text{Residual budget: } C - \varepsilon(d_2), \\ &\text{Residual enforceable set: } \mathcal{D}_2 = \big\{d : m(d \mid d_2) \leq C - \varepsilon(d_2)\big\}. \end{aligned} \tag{15}$$

Since $\varepsilon(d_1) \neq \varepsilon(d_2)$, the residual budgets differ: $C - \varepsilon(d_1) \neq C - \varepsilon(d_2)$. More importantly, the residual enforceable sets differ: $\mathcal{D}_1 \neq \mathcal{D}_2$ (some distinctions fit in one residual budget but not the other). The state of Ω after "d_1 then d_2" records a different residual capacity, a different set of enforceable distinctions, and (via superadditivity) different marginal costs for all remaining distinctions. These are physical differences in Ω, not merely accounting labels.

This step uses one additional operational regularity condition:

- ***OR0 (Faithfulness).*** Enforcement maps are faithful to admissibility structure: if two enforcement orderings produce distinct residual enforceable sets ($\mathcal{D}_1 \neq \mathcal{D}_2$), then the resulting states in Ω are distinct. Equivalently: Ω does not identify states that differ in which distinctions remain enforceable. This is the minimal condition ensuring that budget accounting has physical content—without it, enforcement operations could map distinct resource configurations to the same state, and the order-dependence of budgets would carry no operational consequence.

Step 3 (Order-dependence = operational noncommutativity). The enforcement operations do not commute as maps on Ω:

$$\mathfrak{E}_{d_1} \circ \mathfrak{E}_{d_2}(\sigma) \;\neq\; \mathfrak{E}_{d_2} \circ \mathfrak{E}_{d_1}(\sigma) \tag{16}$$

because the final states differ in residual budget, residual enforceable set, and marginal cost structure. This is the operational content of noncommutativity, derived on the abstract state set Ω without assuming any Hilbert-space structure.

Countermodel. If $\Delta = 0$ (additive costs): marginal costs are context-independent, residual sets depend only on total consumed capacity, and enforcement commutes. If $\varepsilon(d_1) = \varepsilon(d_2)$ (no NT): residual budgets are equal after either ordering. If OR0 fails (states can have distinct residual sets but be identified in Ω): budget asymmetry carries no physical consequence. In any of these cases, enforcement commutes. All three premises—L_Δ, NT, and OR0—are necessary.

Code reference: `check_T1` in `core.py`.

5.2 T2: Hilbert space from noncommutativity

T1 established that enforcement operations on Ω are non-commutative: the order of enforcement matters, and the resulting maps form an algebra that does not collapse to a commutative structure. But non-commutative algebras come in many varieties—matrix algebras, free algebras, quaternion algebras, infinite-dimensional operator algebras. Why should enforcement produce the specific structure of quantum mechanics (finite-dimensional matrix algebras acting on Hilbert spaces) rather than some exotic alternative?

The answer is that finiteness does the work. L_{ε^*} bounds the number of generators; the enforcement algebra is therefore finite-dimensional. And finite-dimensional non-commutative $*$-algebras are completely classified: by the Wedderburn–Artin theorem, every one of them decomposes as a direct sum of matrix algebras $M_{n_k}(\mathbb{F})$ over a division ring $\mathbb{F}$, each acting naturally on $\mathbb{F}^{n_k}$. By Frobenius' theorem [9], the only finite-dimensional associative division algebras over $\mathbb{R}$ are $\mathbb{R}$, $\mathbb{C}$, and the quaternions $\mathbb{H}$. The reals are ruled out: they are commutative, contradicting the non-commutativity established in T1. The quaternions are ruled out: the capacity functional requires a self-dual state space (the trace inner product $\langle A, B\rangle = \omega(A^\dagger B)$ must identify the state space with its dual), which $\mathbb{H}$ does not provide. Therefore $\mathbb{F} = \mathbb{C}$, and the decomposition is $\mathcal{A} \cong \bigoplus_k M_{n_k}(\mathbb{C})$. The familiar structure of quantum mechanics is not assumed—it is the *unique* representation of a finite, non-commutative, self-adjoint algebra over the unique compatible ground field.

T2: Noncommutativity → Operator Algebra on Hilbert Space

Statement. The non-commutative enforcement algebra from T1 is faithfully represented as a $*$-subalgebra of $B(\mathcal{H})$ for a finite-dimensional Hilbert space $\mathcal{H}$.

Variables:

- $\mathcal{A}$ — the enforcement algebra: the $*$-algebra generated by enforcement operations $\{\mathfrak{E}_d\}$. The $*$-operation is defined by $\mathfrak{E}_d^* = \mathfrak{E}_d$ (enforcement is self-adjoint: enforcing a distinction and testing whether it holds are the same operation).
- $M_n(\mathbb{C})$ — the algebra of $n \times n$ complex matrices.
- $\omega : \mathcal{A} \to \mathbb{C}$ — a state (positive normalized linear functional) on the enforcement algebra.

Proof (constructive, via L_{T2}).

Step 1 (Enforcement generates a $*$-algebra). The enforcement operations $\{\mathfrak{E}_d\}$ are closed under: composition ($\mathfrak{E}_{d_1} \circ \mathfrak{E}_{d_2}$), linear combination (mixing enforcement strategies), and adjoint ($\mathfrak{E}_d^* = \mathfrak{E}_d$). By T1, this algebra is non-commutative. By L_{ε^*} (Minimum Cost), it is finite-dimensional (finitely many generators, each costing $\geq \varepsilon^*$, total bounded by C).

This step uses three **operational regularity conditions** (continuing the OR numbering from T1):

- ***OR1 (Convex mixing).*** Enforcement strategies can be mixed probabilistically: if $\mathfrak{E}_1$ and $\mathfrak{E}_2$ are enforcement operations, so is $\lambda\mathfrak{E}_1 + (1-\lambda)\mathfrak{E}_2$ for $\lambda \in [0,1]$. This gives closure under linear combination and is a standard operational axiom in generalized probabilistic theories [20, 21].
- ***OR2 (Self-adjointness).*** Enforcing a distinction and testing whether it holds are operationally equivalent: $\mathfrak{E}_d^* = \mathfrak{E}_d$. This identifies the $*$-operation with enforcement reversal and is the minimal condition for enforcement to define a real-valued cost functional.
- ***OR3 (Finite generation).*** The number of algebraically independent enforcement operations is bounded by $\lfloor C/\varepsilon^* \rfloor$, giving $\dim \mathcal{A} \leq \lfloor C/\varepsilon^* \rfloor^2$. This follows from L_{ε^*} under the identification that each generator corresponds to a distinct enforceable distinction.

Together with OR0 (Faithfulness, stated in T1), these four conditions constitute the complete set of operational regularity assumptions beyond A1 and the perturbation model. They carry the same logical status as the perturbation model (Section 3): mathematical infrastructure specifying the kind of mathematical object enforcement is, not additional physics.

Step 2 (Concrete witness). Take $A = \sigma_x = \left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$ and $B = \sigma_z = \left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right)$ as enforcement operators encoding two competing distinctions. Both are Hermitian. Their commutator: $[A, B] = -2i\sigma_y \neq 0$. The algebra they generate is $M_2(\mathbb{C})$.

Step 3 (State construction). Define $\omega(X) = \mathrm{tr}(X)/2$ on $M_2(\mathbb{C})$. This is: (a) normalized: $\omega(I) = 1$; (b) positive: $\omega(X^\dagger X) \geq 0$; (c) faithful: $\omega(X^\dagger X) = 0 \iff X = 0$. No infinite-dimensional machinery required.

Step 4 (GNS construction). Define $\mathcal{H}_\omega = M_2(\mathbb{C})$ with inner product $\langle X, Y \rangle = \omega(X^\dagger Y)$. The Gram matrix on $\{E_{11}, E_{12}, E_{21}, E_{22}\}$ has eigenvalues $\{1/2, 1/2, 1/2, 1/2\}$ (all positive). The representation $\pi_\omega(A)B = AB$ is a faithful $*$-homomorphism. Thus $\mathcal{A} \hookrightarrow B(\mathcal{H}_\omega)$ with $\dim \mathcal{H}_\omega = 4$.

Step 5 (General structure). By the Wedderburn–Artin theorem, every finite-dimensional $*$-algebra decomposes as $\mathcal{A} \cong \bigoplus_k M_{n_k}(\mathbb{F})$ for a division ring $\mathbb{F}$. By Frobe-

nius' theorem, $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}, \mathbb{H}\}$. The reals cannot host the non-commutativity established in T1; the quaternions violate self-duality of the capacity functional (the trace inner product). Therefore $\mathbb{F} = \mathbb{C}$ and $\mathcal{A} \cong \bigoplus_k M_{n_k}(\mathbb{C})$, each factor acting on its own Hilbert space. The full enforcement algebra thus acts on $\mathcal{H} = \bigoplus_k \mathbb{C}^{n_k}$. The heuristic upper bound $\dim \mathcal{H} \leq (C/\varepsilon^*)^2$ follows from OR3 (generator counting: at most $\lfloor C/\varepsilon^* \rfloor$ generators, each contributing at most one matrix block); a tight bound requires specifying the representation structure of the enforcement algebra, which depends on the specific budget deployment (Papers 2–7).

Code reference: `check_T2, check_L_T2_finite_gns` in `core.py`.

5.3 T_{Born}: the Born rule from Gleason's theorem

T2 gives us a Hilbert space, but a Hilbert space without a probability rule is a mathematical skeleton with no physical content. We need to know: if a system is in state ρ and we perform a measurement with outcomes $\{P_i\}$, what is the probability of each outcome? In standard quantum mechanics, this is the Born rule: $p_i = \text{tr}(\rho \cdot P_i)$. It is usually postulated. Here it is derived.

The key insight is that admissibility imposes a conservation law on probabilities: the total capacity distributed across measurement outcomes must equal the total capacity going in—nothing is created or destroyed. This conservation condition is precisely the hypothesis of Gleason's theorem (1957), which proves that on any Hilbert space of dimension ≥ 3, the *only* probability assignment consistent with this conservation is the Born rule. Since T2 guarantees $\dim \mathcal{H} \geq 4$ (two subsystems of dimension ≥ 2), Gleason's hypotheses are satisfied and the Born rule is forced.

T_{Born}: Born Rule from Admissibility Invariance

Statement. On $\mathcal{H}$ from T2 with $\dim \mathcal{H} \geq 3$, the unique probability measure on projection-valued measurements is $p_i = \text{tr}(\rho \cdot P_i)$.

Variables:

- ρ — a density operator (positive, trace-one) on $\mathcal{H}$.
- $\{P_i\}$ — orthogonal projections with $\sum_i P_i = I$ (a projective measurement).
- p_i — the probability assigned to outcome i.

Proof.

Step 1 (Hypothesis verification: $\dim \mathcal{H} \geq 3$). By T2, the enforcement algebra contains at least $M_2(\mathbb{C})$ (from the σ_x, σ_z witness). By M, at least two subsystems exist. By T_{tensor} (below), the composite Hilbert space has $\dim \mathcal{H}_{AB} = \dim \mathcal{H}_A \cdot \dim \mathcal{H}_B \geq 2 \times 2 = 4 \geq 3$.

Step 2 (Frame function hypothesis). A frame function f assigns probabilities to projection-valued outcomes such that $\sum_i f(P_i) = 1$ for every complete set. Admissibility requires this: the total budget distributed across outcomes is conserved (not created or destroyed). We identify this conservation with the frame function condition. (*Identification:* budget conservation $\leftrightarrow$ frame function.)

Step 3 (Application of Gleason's theorem). Gleason's theorem (1957; finite-dimensional version by Busch, 2003): the *only* frame function on $\mathcal{H}$ with $\dim \geq 3$ is $f(P) = \text{tr}(\rho \cdot P)$ for a unique density operator ρ. Therefore:

$$p_i \;=\; \text{tr}(\rho \cdot P_i). \tag{17}$$

Code reference: `check_T_Born` in `core.py`.

5.4 $T_{\text{Hermitian}}$: self-adjoint observables

In standard quantum mechanics, observables are Hermitian operators—matrices that equal their own conjugate transpose—which guarantees that measurement outcomes are real numbers. This is usually taken as a postulate. In the enforcement framework, it follows from the fact that enforcement costs are real-valued: an observable corresponds to a measurable enforcement outcome, and enforcement costs cannot be complex. The spectral theorem then forces self-adjointness.

$T_{\text{Hermitian}}$: Hermitian Observables

Statement. Observables in $\mathcal{A}$ are self-adjoint: $O = O^\dagger$, with real spectrum $\sigma(O) \subset \mathbb{R}$.

Proof. The enforcement algebra $\mathcal{A} \cong \bigoplus_k M_{n_k}(\mathbb{C})$ from T2 has involution $*$ (adjoint). An observable corresponds to a measurable enforcement outcome—an element whose eigenvalues represent enforcement costs. Enforcement costs are real and non-negative (Section 3). By the spectral theorem for finite-dimensional $*$-algebras, elements with real eigenvalues are precisely the self-adjoint elements: $O = O^*$.

Code reference: `check_T_Hermitian` in `core.py`.

5.5 T_{CPTP}: dynamics must preserve admissibility

How can a quantum state evolve over time? Not arbitrarily: evolution must map admissible states to admissible states. In the Hilbert-space language of T2, this means evolution cannot create capacity out of nothing (trace preservation), cannot map a valid density operator to something with negative eigenvalues (positivity), and must remain well-behaved even when the system is entangled with an ancilla (complete positivity). These three requirements are exactly the defining properties of CPTP maps—the most general allowed quantum operations. The Stinespring dilation theorem then guarantees that every such map can be realized as unitary evolution on a larger system followed by discarding the environment.

T_{CPTP}: Completely Positive Trace-Preserving Maps

Statement. Admissibility-preserving evolution maps $\Phi : \mathcal{A} \to \mathcal{A}$ are completely positive and trace-preserving (CPTP).

Variables:

- Φ — a linear map on $\mathcal{A}$.
- *Positive*: $\rho \geq 0 \implies \Phi(\rho) \geq 0$.
- *Trace-preserving*: $\operatorname{tr}\Phi(\rho) = \operatorname{tr}\rho$.
- *Completely positive*: $(\Phi \otimes \mathbb{1}_k)(\rho_{AB}) \geq 0$ for all k and all $\rho_{AB} \geq 0$.

Proof.

Step 1 (Trace preservation). (*Identification:* trace $\leftrightarrow$ total committed capacity.) A1 bounds total capacity; an admissibility-preserving map cannot create or destroy capacity: $\operatorname{tr}\Phi(\rho) = \operatorname{tr}\rho$.

Step 2 (Positivity). Admissible states are positive operators (density operators on $\mathcal{H}$, from T2). Φ maps admissible to admissible, so $\Phi(\rho) \geq 0$.

Step 3 (Complete positivity). By T_{tensor} (below), systems can be coupled to ancillas. Φ must preserve admissibility of entangled states: $(\Phi \otimes \mathbb{1}_k)(\rho) \geq 0$.
Step 4 (Stinespring). By the Stinespring dilation theorem (for CPTP maps on finite-dimensional C*-algebras), every such map has unitary dilation: $\Phi(\rho) = \text{tr}_E[U(\rho \otimes |0\rangle\langle 0|)U^\dagger]$.
Code reference: `check_T_CPTP` in `core.py`.

5.6 T_{tensor}: tensor products from independent interfaces

When two systems sit at independent interfaces—each with its own capacity budget, neither able to modify the other without paying enforcement cost—how does the joint Hilbert space relate to the individual ones? Locality (L_{loc}) guarantees that the budgets are independent, which means enforcement cost scales linearly in each subsystem separately. The mathematical structure that captures "bilinear and independent" is the tensor product. This is not an assumption about how quantum systems compose; it is forced by the budget structure that L_{loc} established.

T_{tensor}: Tensor Products from Compositional Independence

Statement. Independent subsystems at disjoint interfaces compose via tensor products: $\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$.

Proof.
Step 1 (L_{loc} (Locality) gives independent budgets). Disjoint interfaces Γ_A, Γ_B have independent capacity budgets C_A, C_B.
Step 2 (Bilinearity). (*Identification:* linearity of enforcement in each subsystem $\leftrightarrow$ bilinearity of composition.) Preparing A in state $\alpha\sigma_A$ while B is in state σ_B gives joint state $\alpha(\sigma_A \otimes \sigma_B)$—enforcement cost scales linearly in each subsystem independently, because the budgets are independent.
Step 3 (Minimality via Wedderburn). The smallest vector space satisfying bilinear composition and closure under admissible recombination is the tensor product. By the Wedderburn decomposition, $\mathcal{A}_{AB} \cong \mathcal{A}_A \otimes \mathcal{A}_B$ for algebras on independent subsystems. Therefore $\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$ with $\dim = d_A \cdot d_B$.
Code reference: `check_T_tensor` in `core.py`.

5.7 T3: gauge symmetry from local relabeling invariance

Each interface Γ has its own enforcement algebra $\mathcal{A}_\Gamma$, a matrix algebra $M_n(\mathbb{C})$ from T2. But the labeling of basis vectors within that algebra is arbitrary—nature enforces distinctions, not labels. Relabeling the basis at one interface without affecting any other is an automorphism of the local algebra. The Skolem–Noether theorem (1927/1933) classifies these automorphisms completely: every automorphism of $M_n(\mathbb{C})$ is conjugation by a unitary, so the automorphism group is the projective unitary group $PU(n)$. As the interface label varies across space, these local symmetry groups must patch together consistently—and the mathematical structure that does this is a principal fiber bundle. Gauge symmetry is not postulated; it is the automatic consequence of having local algebras with arbitrary internal labeling.

T3: L_{loc} (Locality) → Gauge Structure

Statement. The automorphism group of the enforcement algebra, restricted to local transformations, forms a gauge group G with principal-bundle structure.

Variables:

- $\text{Aut}(M_n(\mathbb{C}))$ — automorphism group of the matrix algebra.
- $PU(n) = U(n)/U(1)$ — the projective unitary group.
- Skolem–Noether theorem — every automorphism of $M_n(\mathbb{C})$ is inner: $\alpha(X) = UXU^\dagger$.

Proof.

Step 1 (Local algebras are simple). L_{loc} (Locality) decomposes the enforcement algebra: $\mathcal{A} = \bigoplus_\Gamma \mathcal{A}_\Gamma$. Each $\mathcal{A}_\Gamma$ is a simple factor $M_{n_\Gamma}(\mathbb{C})$ by A1 cost minimization: a non-simple algebra (direct sum of smaller blocks) would waste capacity on redundant enforcement channels, violating the constraint that enforcement resources are finite and must be used efficiently.

Step 2 (Skolem–Noether). $\text{Aut}(M_n(\mathbb{C})) \cong PU(n)$, with $\dim PU(n) = n^2 - 1$.

Step 3 (Bundle structure). As the interface label varies, local automorphism groups patch together via cocycle-consistent transition functions $g_{ij} \cdot g_{jk} = g_{ik}$. Verified: $SU(2)$ and $SU(3)$ cocycles satisfy this condition.

Step 4 (Gauge group recovery). The compact group G is recovered from its representation category via Tannaka–Krein duality. The specific gauge group $(SU(3) \times SU(2) \times U(1))$ is determined in Paper 4A by the capacity budget.

Code reference: `check_T3` in `core.py`.

5.8 T_{entropy}: entropy measures committed capacity

A pure state at interface Γ has all its distinctions independently determined—no capacity is tied up in correlations with other interfaces. A mixed state, by contrast, has some of its capacity committed to correlations that are invisible from Γ alone (this is precisely what L_{irr} showed: cross-interface correlations lock up capacity locally). How much capacity is committed? Schumacher's noiseless coding theorem (1995) answers this on a finite-dimensional Hilbert space: the minimum number of qubits needed to faithfully represent the state is $S(\rho) = -\text{tr}(\rho \ln \rho)$, the von Neumann entropy. Identifying each qubit with one distinction at minimum cost ε^*, the entropy directly measures the enforcement capacity that has been irreversibly committed.

T_{entropy}: Von Neumann Entropy as Committed Capacity

Statement. The capacity committed to correlations at interface Γ for a mixed state ρ equals the von Neumann entropy: $S(\rho) = -\text{tr}(\rho \ln \rho)$ in units of ε^*.

Proof.

Step 1 (Operational definition). A mixed state at Γ is one whose distinctions are not all independently determined—some capacity is committed to correlations with other interfaces.

Step 2 (Schumacher's theorem). On a finite-dimensional Hilbert space (T2), Schumacher's noiseless coding theorem (1995): the minimum number of qubits to faithfully compress n copies of ρ is $nS(\rho)$ as $n \to \infty$.

Step 3 (Identification). (*Identification:* each qubit of compression $\leftrightarrow$ one distinction costing ε^*.) The committed capacity per copy is $S(\rho) \cdot \varepsilon^*$. In natural units ($\varepsilon^* = 1$): committed capacity $= S(\rho)$.
Step 4 (Consistency). Pure states: $S = 0$ (no correlations). Maximally mixed: $S = \ln d$ (maximum correlations). Both match the operational interpretation.

Code reference: `check_T_entropy` in `core.py`.

5.9 T_M: monogamy from finite capacity

Entanglement monogamy—the fact that a quantum system maximally entangled with one partner cannot be simultaneously entangled with another—is one of the most distinctive features of quantum mechanics. In the enforcement framework, it is an immediate consequence of finite budgets. Correlations with B and correlations with C both draw from Γ_A's capacity. Since the budget is finite, committing capacity to one set of correlations leaves less available for the other. Monogamy is not a peculiarity of quantum mechanics; it is bookkeeping.

T_M: Interface Monogamy

Statement. $\Delta_{AB} + \Delta_{AC} \leq C(\Gamma_A)$: capacity committed to A–B correlations is unavailable for A–C correlations.

Proof. By A1, Γ_A has total capacity $C(\Gamma_A) < \infty$. By L_{loc} (Locality), correlations with B and C both draw from Γ_A's budget. Since $\Delta_{AB}, \Delta_{AC} \geq 0$ (by L_{nc}) and both must fit within the single budget, $\Delta_{AB} + \Delta_{AC} \leq C(\Gamma_A)$.

Code reference: `check_T_M` in `core.py`.

5.10 T_ε, T_κ, T_η: enforcement granularity

The lemmas established that enforcement has a minimum cost (L_{ε^*}), distributes over interfaces (L_{loc}), and creates irreversible commitments (L_{irr}). Three parameters characterize the granularity of this structure. The base cost unit ε sets the scale. The directed multiplier κ counts how many interfaces a single distinction spans—the answer is exactly two (system and environment), giving $\kappa = 2$. The subordination bound η controls the cost of adding generations within a sector, and feeds directly into the $N_{\text{gen}} = 3$ derivation in Paper 4A.

T_ε, T_κ, T_η: Enforcement Granularity Parameters

T_ε: ε is the base enforcement cost unit.

T_κ: $\kappa = 2$ (directed enforcement multiplier).

Variables:

- C_{fwd} — forward obligation at Γ_S (from L_{nc}: $C_{\text{fwd}} \geq \varepsilon$).
- C_{env} — environment record at Γ_E (from L_{irr}: $C_{\text{env}} \geq \varepsilon$).

Lower bound ($\kappa \geq 2$): Each distinction spans two interfaces (system + environment), committing $\geq \varepsilon$ at each. These are independent (L_{loc}: $\Gamma_S \perp \Gamma_E$), so $\kappa \geq 2$.
Upper bound ($\kappa \leq 2$): A third interface would require a second environment = a new distinction, not a third obligation on the original. So $\kappa \leq 2$. (*Note:* this upper bound is a structural accounting argument, not a theorem from A1 alone. It assumes the "one

distinction = at most two interfaces" identification.)

T_η: $\eta \leq \varepsilon$ (subordination bound). Generation cost function:

$$E(N) = N\varepsilon + \frac{N(N-1)}{2}\eta, \qquad N_{\text{gen}} = 3 \text{ (Paper 4A)}. \tag{18}$$

Code reference: `check_T_epsilon, check_T_kappa, check_T_eta` in `core.py`.

6. The canonical object: what A1 forces

Every result in Sections 4–5 follows from A1 through a specific derivation chain. Taken together, they assemble into a single mathematical object—the minimal structure that any universe satisfying A1 must possess. This is not a model or an approximation; it is the unique mathematical consequence of finite enforcement capacity.

The canonical object has four components: a finite-dimensional Hilbert space $\mathcal{H}$ (from T2), a C*-algebra of observables $\mathcal{A}$ acting on it (from Wedderburn–Artin), a compact gauge group G (from T3), and an entropy functional $S[\rho]$ measuring committed capacity (from T_{entropy}). What the canonical object does *not* specify is the particular gauge group, the particle content, or any numerical predictions—those require deploying the capacity budget through specific channels, which is the work of Papers 2–7. The canonical object is the skeleton; the companion papers are the anatomy.

$T_{\text{canonical}}$: The Canonical Object Forced by A1

Statement. A1 forces a unique mathematical structure:

$$\mathcal{A}(C, \varepsilon^*, \kappa) \;=\; \Big(\mathcal{H},\;\; \mathcal{A} \subset B(\mathcal{H}),\;\; G_{\text{gauge}},\;\; S[\rho] \Big) \tag{19}$$

Components and derivation chain:

- $\mathcal{H}$ — finite-dimensional Hilbert space. *Via:* A1 $\to$ L_{nc} $\to$ T1 (order-dependence) $\to$ T2 (GNS).
- $\mathcal{A} \subset B(\mathcal{H})$ — C*-algebra of observables. *Via:* T2 (Wedderburn–Artin decomposition).
- G_{gauge} — compact gauge group. *Via:* L_{loc} $\to$ T3 (Skolem–Noether + Tannaka–Krein).
- $S[\rho] = -\text{tr}(\rho \ln \rho)$ — entropy functional. *Via:* T2 $\to$ T_{entropy} (Schumacher compression).

The specific content—$C_{\text{total}} = 61$, $G_{\text{gauge}} = SU(3) \times SU(2) \times U(1)$, numerical predictions—is determined in Papers 2–7.

Code reference: `check_T_canonical` in `core.py`.

7. The derivation DAG: tracing dependencies

The full derivation chain from A1 to the 312 theorems of the codebase forms a directed acyclic graph. This paper establishes the root and the first two tiers:

Tier	Content	Count
-1	Axiom (A1)	1
0	Structural lemmas (L_{ε^*}, L_{nc}, L_{loc}, L_{irr})	4
1	Bridge + quantum skeleton (T1–T3, T_{Born}, T_{CPTP}, ...)	18
	Total (this paper)	**23**

The critical path through the DAG:

$$\text{A1} \to L_{\varepsilon^*} \to L_{\text{nc}} \to L_{\text{loc}} \to \text{T1} \to \text{T2} \to \text{T3} \to T_{\text{gauge}} \to T_{\text{field}} \to T_{\text{concordance}}$$

8. Scope: what this paper derives and what it does not

8.1 Derived in this paper: the structural skeleton

From A1 alone (plus sub-clauses M and NT, both subsequently derived): L_{ε^*} (Minimum Cost), L_{nc} (Non-Closure), L_{loc} (Locality), L_{irr} (Irreversibility), order-dependent enforcement (T1), Hilbert space (T2), Born rule (T_{Born}), Hermitian observables ($T_{\text{Hermitian}}$), CPTP dynamics (T_{CPTP}), tensor products (T_{tensor}), gauge symmetry skeleton (T3), entropy (T_{entropy}), interface monogamy (T_M), enforcement parameters (T_ε, T_κ, T_η), and the canonical object ($T_{\text{canonical}}$).

8.2 Deferred to later papers: specific content and predictions

Specific gauge groups ($SU(3) \times SU(2) \times U(1)$, Paper 4A), particle content (61 types and 3 generations, Paper 4A), mass hierarchies, mixing matrices, and CP violation (Paper 4B), spacetime dimension ($d = 4$) and signature (Paper 6), cosmological constants (Paper 6), and the 47 zero-parameter predictions (Papers 4B, 5, 6, 7).

8.3 Explicit identifications used in this paper

Several steps in the quantum skeleton invoke identifications—mappings between enforcement-theoretic concepts and standard mathematical structures—that are not theorems from A1 alone. For transparency, we collect them here:

1. T_{Born}: Budget conservation over outcomes $\leftrightarrow$ frame function (Gleason hypothesis).
2. T_{CPTP}: Total committed capacity $\leftrightarrow$ operator trace.
3. T_{tensor}: Linearity of enforcement in each subsystem $\leftrightarrow$ bilinearity of composition map.
4. T_{entropy}: One distinction at minimum cost $\leftrightarrow$ one qubit of compression.
5. T_κ: One distinction spans at most two interfaces.

Each identification is physically motivated and consistent with A1, but none is a pure deduction from the axiom. They represent the points where the framework's physical interpretation meets the mathematical formalism.

9. Tensions and attack surfaces: honest accounting

9.1 Known tensions: none at this layer

The axiom is consistent; the four lemmas are proved with constructive witnesses; the competition postulate (L_Δ) is derived from the perturbation model; the bridge theorem (T1) and quantum skeleton follow from standard theorems whose hypotheses are explicitly verified. Tensions arise at the *deployment* level (Papers 4–7), not at the foundation.

9.2 Attack surfaces: where a skeptic should push

1. **"Meaningful = robust" (Section 3).** L_{ε^*} (Minimum Cost) rests on identifying meaningful distinctions with perturbation-robust distinctions. If a weaker notion of meaningfulness is adopted, ε^* could be zero.

2. **The bridge theorem (T1) and operational regularity.** T1 shows that enforcement operations are order-dependent (non-commutative as maps on Ω), using OR0 (Faithfulness) to ensure budget differences are state differences. The step from "order-dependent maps on Ω" to "$*$-algebra with Hermitian generators" (T2 Step 1) additionally requires OR1 (convex mixing), OR2 (self-adjointness), and OR3 (finite generation). Of these, OR1 is the most substantive: it asserts that probabilistic mixtures of enforcement strategies are themselves valid strategies, and is the standard convexity assumption shared by all generalized probabilistic theory (GPT) reconstruction programs. OR2 and OR3 are more minimal; OR0 is the condition that makes budget accounting physically meaningful. See Section 9.3 for comparison with prior reconstruction programs.

3. **The perturbation model and L_Δ (Competition).** Superadditivity ($\Delta > 0$) is the hinge between classical overflow and quantum noncommutativity. It rests on the perturbation model (Section 3): the assumption that combined perturbations can attack jointly-enforced distinctions, making the joint defense strictly harder than the sum of individual defenses. A skeptic who rejects combined perturbations (insisting $\mathcal{P}(\{d_1, d_2\}) = \mathcal{P}(d_1) \cup \mathcal{P}(d_2)$) gets a classical knapsack theory with additive costs, non-closure, but no noncommutativity. The perturbation model is the precise assumption that selects the quantum case from the classical case.

4. **M/NT as a self-consistency loop.** The logical structure is "assume M/NT $\to$ derive field content $\to$ derive M/NT." This is self-consistent (the DAG is acyclic), but M and NT are load-bearing assumptions at the level of this paper.

5. **The five identifications (Section 8.3).** These are the interface between the enforcement formalism and standard quantum mathematics. Each is physically motivated, but a skeptic could challenge any of them as underdetermined by A1 alone.

6. **Red team confirmation.** RT2_non_closure_necessity confirmed that L_{nc} (Non-Closure) is genuinely forced by A1 + L_{ε^*} (Minimum Cost) + M. Only escapes: $C = \infty$ or $|\mathcal{D}| = 1$.

9.3 Comparison with GPT reconstruction programs

The operational regularity conditions OR0–OR3 are not unique to this framework; every quantum reconstruction program requires comparable structural assumptions. What differs is which assumptions are made and how strong they are.

Hardy [20] assumes five axioms: subspace composition, continuous reversibility, simplicity, and two state-counting rules. Chiribella, D'Ariano, and Perinotti [21] assume six: causality, perfect distinguishability, ideal compression, local distinguishability, pure conditioning, and purification. Masanes and Müller [22] assume continuous reversibility plus local tomography. All three programs derive the Hilbert-space formalism from these starting points.

The APF's operational regularity conditions are:

- **OR0** (Faithfulness): distinct admissibility structures → distinct states.
- **OR1** (Convex mixing): probabilistic mixtures of strategies are strategies. (Shared with all GPT programs.)
- **OR2** (Self-adjointness): enforcement = test. (Weaker than Hardy's reversibility or CDP's purification.)
- **OR3** (Finite generation): generator count bounded by C/ε^*. (Follows from L_{ε^*}; no analog needed in GPT programs that assume finite dimension directly.)

Notably absent from the APF list: continuous reversibility (Hardy), purification (CDP), local tomography (Masanes–Müller), and subspace composition (Hardy). The APF does not assume any of these. Instead, it trades them for the perturbation model (combined attacks exist) and L_Δ (superadditivity), which have no direct analog in prior programs. The net effect is a different path to the same destination: where GPT programs derive quantum theory from information-processing axioms, the APF derives it from resource constraints on enforcement. The additional content—specific gauge groups, generations, and numerical predictions—comes from deploying the capacity budget (Papers 2–7), a step that has no counterpart in GPT reconstructions.

10. The correlation space after this paper

This paper is the first of seven, each of which reveals a new layer of a single mathematical object. We call that object the **correlation space** $\mathcal{S}_{\text{adm}}$: the set of all admissible correlation patterns consistent with A1. Formally, once the Hilbert-space representation is established by T2, the correlation space is

$$\mathcal{S}_{\text{adm}} \;=\; \big\{\, \rho \in B(\mathcal{H}) \;:\; \rho \geq 0,\;\; \operatorname{tr}\rho = 1,\;\; \textstyle\sum_d \varepsilon(d) \leq C \,\big\}. \tag{20}$$

This is a convex body inside the space of density operators on $\mathcal{H}$. Its interior consists of states with enforcement capacity to spare. Its boundary consists of states at capacity saturation—and the boundary is where all the interesting physics happens.

What do we know about $\mathcal{S}_{\text{adm}}$ after this paper?

1. *It exists.* A1 guarantees a finite, positive capacity bound. The admissibility constraint carves a proper subset of the full state space.

2. *It is finite-dimensional.* L_{ε^*} (Minimum Cost) ensures that the number of simultaneously enforceable distinctions is bounded by $\lfloor C/\varepsilon^* \rfloor$. The Hilbert space $\mathcal{H}$ from T2 is finite-dimensional.

3. *It is convex.* The budget constraint is linear in enforcement costs, so any probabilistic mixture of admissible states is admissible.

4. *It is not closed under composition.* L_{nc} (Non-Closure) proves that two individually admissible states can overflow the budget when composed. $\mathcal{S}_{\text{adm}}$ is not a group or an algebra

under state composition—it is a constrained body.

5. *It factorizes locally.* L_{loc} (Locality) forces the capacity budget to decompose over independent interfaces. The correlation space has a tensor-product structure: $\mathcal{S}_{\text{adm}}$ respects the interface decomposition established by T_{tensor}.

6. *It has irreversible directions.* L_{irr} (Irreversibility) proves that some transitions—specifically, the creation of cross-interface correlations—move a state toward the boundary of $\mathcal{S}_{\text{adm}}$ in a way that cannot be locally reversed. The correlation space has an arrow.

7. *It carries an operator algebra.* T1 (the Bridge Theorem) and T2 equip $\mathcal{S}_{\text{adm}}$ with a C*-algebra $\mathcal{A} \cong \bigoplus_k M_{n_k}(\mathbb{C})$ acting on $\mathcal{H}$.

8. *It has a unique probability rule.* T_{Born} (Gleason's theorem) fixes the Born rule as the only consistent probability assignment on $\mathcal{H}$.

9. *It has gauge symmetry.* T3 shows that the automorphism group of the local algebras forms a compact gauge group G with principal-bundle structure.

10. *It has an entropy functional.* T_{entropy} identifies the von Neumann entropy $S(\rho) = -\text{tr}(\rho \ln \rho)$ as the measure of capacity committed to correlations. Entropy measures depth into the boundary of $\mathcal{S}_{\text{adm}}$.

What we do *not* yet know: the dimension of $\mathcal{H}$ (that is, the value of C_{total}), the identity of G, the field content, the generation structure, the mass spectrum, or the geometric embedding. This paper establishes the mathematical type of the correlation space—a finite-dimensional convex body with operator algebra, gauge symmetry, and entropy—but not its specific anatomy.

The subsequent papers fill in the anatomy. Paper 2 determines the size ($C_{\text{total}} = 61$) and the coarsest internal partition (vacuum 42, matter 19). Paper 3 establishes the thermodynamic foliation and the arrow of time. Paper 4 derives the gauge group $SU(3) \times SU(2) \times U(1)$, the field content, the generation structure, all mass ratios, and all mixing parameters—the complete internal architecture. Paper 5 shows that $\mathcal{S}_{\text{adm}}$ generates its own geometric embedding in 3+1 Lorentzian spacetime. Paper 6 traces the cosmological trajectory through $\mathcal{S}_{\text{adm}}$, from the capacity-saturated Big Bang to the de Sitter vacuum. Paper 7 proves the boundary is closed: the capacity budget is exactly saturated by the Standard Model plus three right-handed neutrinos, and no further structure can be accommodated.

Seven papers, one object. Each reveals a layer. Together they describe the correlation space in full.

Everything derived in this paper—Hilbert space, the Born rule, CPTP dynamics, tensor products, gauge symmetry, von Neumann entropy—followed from a single constraint on finite resources. No physical content was assumed; no specific content was derived. The correlation space after this paper is a skeleton waiting for anatomy. The anatomy turns out to be the universe we observe. That is the subject of the papers that follow.

Statements and Declarations

Funding. The author did not receive support from any organization for the submitted work.

Competing interests. The author has no relevant financial or non-financial interests to disclose.

Data availability. All data supporting the findings of this study are contained within the paper itself. The executable theorem-checking codebase (Python) is available at `https://github.com/Ethan-Brooke/APF-Paper-1-The-Enforceability-of-Distinction`.

Code availability. The APF codebase implementing all 23 theorems of this paper, together with the full 312-theorem bank of the companion papers, is available at `https://github.com/Ethan-Brooke/APF-Paper-1-The-Enforceability-of-Distinction`.

Author contributions. E. S. Brooke is the sole author. All conceptualization, formal analysis, software development, and manuscript preparation were performed by the author.

References

[1] Ashby, W.R.: An Introduction to Cybernetics, Chapter 11. Chapman & Hall, London (1956)

[2] Spencer-Brown, G.: Laws of Form. Allen & Unwin, London (1969)

[3] Bell, J.S.: On the Einstein Podolsky Rosen paradox. Physics Physique Fizika **1**, 195–200 (1964)

[4] Nielsen, M.A., Chuang, I.L.: Quantum Computation and Quantum Information. Cambridge University Press, Cambridge (2000)

[5] Jaynes, E.T.: Information theory and statistical mechanics. Phys. Rev. **106**, 620–630 (1957)

[6] Gel'fand, I.M., Naĭmark, M.A.: On the imbedding of normed rings into the ring of operators in Hilbert space. Mat. Sb. **12**(54), 197–213 (1943)

[7] Gleason, A.M.: Measures on the closed subspaces of a Hilbert space. J. Math. Mech. **6**, 885–893 (1957)

[8] Wedderburn, J.H.M.: On hypercomplex numbers. Proc. Lond. Math. Soc. **6**, 77–118 (1908)

[9] Frobenius, G.: Über lineare Substitutionen und bilineare Formen. J. Reine Angew. Math. **84**, 1–63 (1878)

[10] Artin, E.: Zur Theorie der hyperkomplexen Zahlen. Abh. Math. Semin. Univ. Hambg. **5**, 251–260 (1927)

[11] Skolem, T.: Zur Theorie der assoziativen Zahlensysteme. Skr. Norske Vid.-Akad. Oslo, Mat.-Naturv. Kl. No. 12 (1927)

[12] Noether, E.: Nichtkommutative Algebra. Math. Z. **37**, 514–541 (1933)

[13] Aoyama, T., et al.: The anomalous magnetic moment of the muon in the Standard Model. Phys. Rep. **887**, 1–166 (2020)

[14] Zurek, W.H.: Decoherence, einselection, and the quantum origins of the classical. Rev. Mod. Phys. **75**, 715–775 (2003)

[15] Wilson, K.G.: Confinement of quarks. Phys. Rev. D **10**, 2445–2459 (1974)

[16] Wilson, K.G.: The renormalization group and critical phenomena. Rev. Mod. Phys. **55**, 583–600 (1983)

[17] Zurek, W.H.: Pointer basis of quantum apparatus: Into what mixture does the wave packet collapse? Phys. Rev. D **24**, 1516–1525 (1981)

[18] Hawking, S.W.: Particle creation by black holes. Commun. Math. Phys. **43**, 199–220 (1975)

[19] Penrose, R.: Gravitational collapse and space-time singularities. Phys. Rev. Lett. **14**, 57–59 (1965)

[20] Hardy, L.: Quantum theory from five reasonable axioms. arXiv:quant-ph/0101012 (2001)

[21] Chiribella, G., D'Ariano, G.M., Perinotti, P.: Informational derivation of quantum theory. Phys. Rev. A **84**, 012311 (2011). arXiv:1011.6451

[22] Masanes, L., Müller, M.P.: A derivation of quantum theory from physical requirements. New J. Phys. **13**, 063001 (2011). arXiv:1004.1483

[23] Bekenstein, J.D.: Universal upper bound on the entropy-to-energy ratio for bounded systems. Phys. Rev. D **23**, 287–298 (1981)

[24] 't Hooft, G.: Dimensional reduction in quantum gravity. arXiv:gr-qc/9310026 (1993)

[25] Susskind, L.: The world as a hologram. J. Math. Phys. **36**, 6377–6396 (1995)

[26] Bousso, R.: The holographic principle. Rev. Mod. Phys. **74**, 825–874 (2002). arXiv:hep-th/0203101

[27] Jacobson, T.: Thermodynamics of spacetime: The Einstein equation of state. Phys. Rev. Lett. **75**, 1260–1263 (1995). arXiv:gr-qc/9504004

[28] Planck Collaboration, Aghanim, N., et al.: Planck 2018 results. VI. Cosmological parameters. Astron. Astrophys. **641**, A6 (2020). arXiv:1807.06209

[29] Particle Data Group, Navas, S., et al.: Review of particle physics. Phys. Rev. D **110**, 030001 (2024)

[30] Esteban, I., et al.: Updated global analysis of three-flavor neutrino oscillations (NuFit-6.0). JHEP **12**, 216 (2024). arXiv:2410.05380

A. Appendix A: complete theorem register

Name	Statement	Dependencies	Code
A1	Finite Enforcement Capacity	(axiom)	`check_A1`
M	Multiplicity: $\|\mathcal{D}\| \geq 2$	A1 (sub-clause)	`check_M`
NT	Non-Degeneracy: costs non-equal	A1 (sub-clause)	`check_NT`
$L_{M,\text{der}}$	Multiplicity Derived	A1, T_{field}	`check_L_M_derived`
$L_{NT,\text{der}}$	Non-Degeneracy Derived	A1, T_{11}	`check_L_NT_derived`
L_{ε^*}	Minimum Cost	A1	`check_L_epsilon_star`
L_{loc}	Locality	A1, L_{ε^*}, M, NT	`check_L_loc`
L_{nc}	Non-Closure	A1, L_{ε^*}, M	`check_L_nc`
L_Δ	Competition (superadditivity)	A1, Pert. Model	`check_L_irr (witness)`
L_{irr}	Irreversibility	L_{nc}, L_Δ, L_{loc}	`check_L_irr`
$L_{\text{irr,unif}}$	Uniform Irreversibility	L_{irr}	`check_L_irr_uniform`
P_{exhaust}	MECE Partition	A1, M, NT	`check_P_exhaust`
L_{cost}	Cost Uniqueness	A1	`check_L_cost`
$L_{\Omega,\text{sign}}$	Sign Dichotomy	A1, L_{nc}	`check_L_Omega_sign`
M_Ω	Microcanonical Measure	A1, L_{nc}	`check_M_Omega`
T0	Axiom Witnesses	A1, M, NT	`check_T0`
T1	Order-Dependence	L_Δ, NT, OR0	`check_T1`
T2	Hilbert Space	T1, OR1–OR3	`check_T2`
L_{T2}	Concrete GNS	T2	`check_L_T2_finite_gns`
T3	Gauge Structure	L_{loc}, T2	`check_T3`
T_{Born}	Born Rule	T2, T_{tensor}	`check_T_Born`
T_{Herm}	Hermitian Observables	T2	`check_T_Hermitian`
T_{CPTP}	CPTP Dynamics	T2, T_{tensor}	`check_T_CPTP`
T_{tensor}	Tensor Products	T2, L_{loc}	`check_T_tensor`
T_M	Interface Monogamy	A1, L_{loc}	`check_T_M`
T_{can}	Canonical Object	T2, T3, T_{entropy}	`check_T_canonical`
T_{entropy}	Entropy Functional	T2	`check_T_entropy`
T_ε	Enforcement Granularity	A1	`check_T_epsilon`
T_κ	Directed Multiplier	A1, L_{loc}	`check_T_kappa`
T_η	Subordination Bound	A1, T_κ	`check_T_eta`

All check functions reside in `core.py`. Supporting results from `supplements.py`: $L_{\text{Gleason,finite}}$ and $L_{\text{Wedderburn}}$.

B. Appendix B: red team self-tests

1. **RT2_non_closure_necessity** (PASS): L_{nc} (Non-Closure) is genuinely forced by A1 + L_{ε^*} (Minimum Cost) + M. Only escapes: $C = \infty$ or $|\mathcal{D}| = 1$.

2. **RT_seesaw_necessity** (PASS): The seesaw chain from L_{irr} (Irreversibility) through to neutrino masses is load-bearing—no alternative mechanism produces the same results.

C. Appendix C: derived sub-clauses

The sub-clauses M (Multiplicity) and NT (Non-Degeneracy) are assumed at the axiom layer in Section 2.2 and used in the derivation of L_{loc} (Locality). They are subsequently derived from the framework's own field content, closing the logical loop.

$L_{M,\text{derived}}$ (Multiplicity Derived) and $L_{NT,\text{derived}}$ (Non-Degeneracy Derived)

$L_{M,\text{derived}}$: The gauge constraint chain (Paper 4A) forces exactly 61 capacity types via T_{field}. Since $61 \geq 2$, Multiplicity is a consequence of A1:

$$\text{A1} \;\rightarrow\; T_{\text{field}} \;\rightarrow\; C_{\text{total}} = 61 \geq 2 \;\rightarrow\; M.$$

$L_{NT,\text{derived}}$: The MECE partition (P_{exhaust}) splits 61 types into vacuum (42) and matter (19). Since $42 \neq 19$, Non-Degeneracy is a consequence of A1:

$$\text{A1} \;\rightarrow\; T_{11} \;\rightarrow\; C_{\text{vac}} = 42 \neq 19 = C_{\text{mat}} \;\rightarrow\; NT.$$

Code reference: `check_L_M_derived, check_L_NT_derived` in `core.py`.

D. Appendix D: symbol index

Symbol	Meaning	Defined
Axiom and sub-clauses		
A1	Finite enforcement capacity (the axiom)	§2
M	Multiplicity: $\lvert\mathcal{D}\rvert \geq 2$	§2.2
NT	Non-Degeneracy: enforcement costs not all equal	§2.2
Primitive quantities		
Ω	Abstract state set (no algebraic structure assumed)	§3
σ	A state in Ω (a configuration of enforced distinctions)	§3
d	A distinction: binary predicate $d : \Omega \to \{0, 1\}$	§3
$\mathcal{D}(\sigma, R)$	Set of distinctions enforced in state σ on region R	§2
$\varepsilon(d)$	Enforcement cost of distinction d	§3
ε^*	Minimum enforcement cost per distinction (> 0)	§4.1
R	A causally connected region	§2
Γ	An interface (enforcement locus)	§3
$C(R)$, $C(\Gamma)$	Enforcement capacity of region R or interface Γ	§2, §3
C_{total}	Total enforcement capacity (global)	§3
Derived quantities		
$E_\Gamma(T)$	Total enforcement cost of distinction set T at interface Γ	§4.5
$\Delta(d_1, d_2)$	Superadditivity surplus for jointly enforced distinctions	§4.4
$m(d_2 \mid d_1)$	Marginal cost of enforcing d_2 given d_1 active	§5.1
$\mathcal{P}(S)$	Perturbation space for distinction set S	§3
Ω_{adm}	Set of admissible states ($\sum \varepsilon(d) \leq C$)	§4.3
Algebraic structure (derived in §5)		
$\mathfrak{E}_d$	Enforcement operation for distinction d (map $\Omega \to \Omega$)	§5.1
$\mathcal{A}$	Enforcement algebra ($*$-algebra generated by $\{\mathfrak{E}_d\}$)	§5.2
$\mathcal{H}$	Hilbert space (from GNS/Wedderburn–Artin on $\mathcal{A}$)	§5.2
ρ	Density operator on $\mathcal{H}$	§5.3
$S(\rho)$	Von Neumann entropy: $-\text{tr}(\rho \ln \rho)$	§5.8
G	Gauge group (automorphism group of local algebras)	§5.7
Parameters (from companion papers)		
κ	Directed enforcement multiplier ($= 2$)	§5.11
η	Subordination bound	§5.11