Deriving the Commitment-Capacity Density from Primitive Commitment Structure

VERSF Theoretical Physics Programme

For the General Reader

Physics rests on the idea that the universe produces definite facts — that events happen, records form, and some things become irreversibly true. But what is the minimum cost, in space and energy, of one such fact being formed? And can two irreducible facts share the same minimal region of space at the same moment?

This paper answers both questions. It does so within the VERSF framework, which treats reality as built from primitive irreversible events — the smallest possible acts of physical commitment, each depositing exactly one irreducible unit of information into the fabric of the universe. These events are the bedrock of time, space, and all physical structure in the programme.

The first question — what is the minimum spatial cost of one primitive fact? — has a clean answer: there is a smallest region of space, called a coherence cell, whose size is set by the balance between energy and information at the threshold of one primitive event. This scale, ξ, is the programme's fundamental length.

The second question — can two such events share one coherence cell? — is harder, and is the central result of this paper. We prove that the answer is no. The proof works by exhaustion: we classify every logically possible way two primitive events could coexist in a single minimal region using the cell's own observable algebra — the complete set of measurable quantities defined within the cell. This classification is derived, not assumed. It yields three cases (fully separable, fully identified, or partially resolved by the algebra), and each case leads to a contradiction with the framework's own foundational rules. No configuration survives. A minimal region supports exactly one primitive fact.

This single-occupancy result — proved here as a theorem for the first time rather than assumed as a principle — has immediate consequences. It tells us what the energy density of the universe's fact-forming capacity must be: one primitive event per coherence cell area, giving $\rho c = \hbar c/\xi^3$. The coherence cell is a two-dimensional surface patch rather than a three-dimensional volume because, as established in a companion paper in the programme, "depth" — the third spatial-seeming parameter — is not a genuine spatial direction but a bookkeeping index for resolution and scale. The two genuinely spatial directions define an area; depth is not among them. From this, the key scaling law of the programme's gravity theory follows as a derived result rather than a postulate.

But the result is not merely about gravity. It is a structural statement about any universe that forms stable, irreversible facts under finite distinguishability: such a universe is subject to a packing limit on primitive events. The exclusion is not statistical — it is logical and geometric, forced by the very meaning of what it is for an event to be primitive and irreversible.

Table of Contents

Abstract

The effective density ρ entering the commitment-capacity invariant

$$\chi(L) = \rho L^3 / \hbar c$$

has until now been treated as an externally supplied vacuum energy density, estimated on dimensional grounds rather than derived from the primitive structure of the programme. This paper closes that gap. We show that, once primitive commitment, binary irreducibility, finite distinguishability, and local capacity competition are taken jointly, ρ is identified structurally as the *primitive commitment-capacity density*: the energy per unit two-area required to support primitive irreversible commitments at threshold.

The key step is the **No Multi-Primitive Occupancy theorem** (Theorem 4.1). Its proof proceeds by contradiction, using the cell's local observable algebra to derive — not assume — a complete partition of all pairwise event configurations into three cases: fully separable, fully identified, and partially resolved. Each case violates one of: primitive minimality, operational resolvability, or binary irreducibility. The trichotomy is an induced structure of the observable algebra, not an independent ontological classification. A separate Threshold Activation Lemma (Lemma 3.2) establishes that $\chi(\xi) = 1$ is saturating rather than merely permissive, securing the lower bound. Together these results establish $\mathrm{Nprim}(\xi^2) = 1$ on purely structural grounds.

A consequence is that the cubic scaling $\chi(L) = (L/\xi)^3$ is no longer an action-budget ansatz. It becomes a derived packing law, and the primitive triad of entropic, energetic, and geometric minimality is shown to be jointly rigid by a Consistency Theorem (Theorem 8.1).

---

1. Introduction and Motivation

The coherence scale relation

$$\chi(\xi) = 1 \Longrightarrow \xi = (\hbar c/\rho)^{(1/3)}$$

plays a central role in the VERSF gravity programme. It identifies ξ as the smallest bounded causal scale at which a single primitive irreversible commitment event becomes admissible. The cubic form of the invariant $\chi(L) = \rho L^3/\hbar c$ was established in the commitment-barrier paper on action-budget grounds, but the physical status of ρ was left at the level of an effective energy density — dimensionally reasonable, but not derived from the programme's primitive structure.

This paper addresses that deficiency. The payoff is substantial: once ρ is identified as the density of one primitive commitment event per coherence cell, the cubic invariant ceases to be an ansatz and becomes a theorem. The action-budget derivation and the present microscopic derivation then stand as mutually independent routes to the same result, each strengthening the other.

The central improvement over prior versions of this argument is threefold. First, the trichotomy used in the proof of Theorem 4.1 is now derived from the observable algebra rather than assumed as an intuitive classification. Second, the lower bound $\mathrm{Nprim}(\xi^2) \geq 1$ is established by a proper Lemma showing that the threshold is saturating. Third, the energy objection to Case 2 — that two indistinguishable events would still carry energy $2Ec$ — is answered by showing that the energy counting operator presupposes independent attribution, which is precisely what Case 2 denies.

The derivation proceeds in six main steps:

1. Establish the primitive observable algebra and derive the trichotomy of event configurations.
2. Identify the coherence cell and prove the Threshold Activation Lemma.
3. Prove by contradiction that no threshold cell supports two or more independently attributable primitive events (Theorem 4.1).
4. Derive the primitive capacity density ρc from single-cell occupancy.
5. Recover the cubic invariant as a consequence.
6. Establish the joint rigidity of the primitive triad (Theorem 8.1).

1.1 Scope and Consequences

The results of this paper bear on a question wider than gravity: what kind of universe can produce stable, irreversible facts at all?

Gravity enters as one downstream consequence. But the constraint established here is more fundamental: it is a structural limit on fact formation itself, applicable to any physical theory operating under finite distinguishability and primitive irreversibility. The following consequences follow directly.

Fact formation has a packing limit. Just as fermionic matter cannot be compressed without limit — not because of a force, but because of a logical constraint on quantum states — primitive facts cannot be packed without limit. The exclusion here is not statistical: it follows from the observable algebra structure, the entropy spectrum, and the operational definition of what it means for an event to be primitive and independently attributable. It holds regardless of the dynamics, the specific physical model, or the interaction terms. Any universe that forms stable irreversible facts under finite distinguishability is subject to this constraint.

Entropy is not a continuous fluid. By establishing that each primitive event contributes exactly one entropy quantum ln 2, and that no two such events can coexist in a single threshold cell, the result implies that entropy density is structurally constrained rather than freely divisible. This strengthens the interpretation of entropy as a count of irreducible commitments — a discrete quantity — rather than a smooth thermodynamic variable. The continuity of thermodynamic entropy emerges from large numbers of primitive events; at the threshold scale it is quantised.

Quantum non-classicality has a structural origin. The exclusion of indistinguishable multi-event configurations from independent attribution — Case 2 of Theorem 4.1 — parallels the structure of quantum mechanics at a deeper level. In quantum theory, indistinguishable contributions to amplitudes interfere rather than adding classically because they cannot be independently attributed. The present result, formulated at the pre-quantum level of primitive irreversible commitment, shares precisely this operational constraint: independent attribution is the condition under which multiplicity becomes physically meaningful. This structural parallel runs deeper than analogy: both cases are governed by the same operational constraint on independent attribution, suggesting a common origin for non-classical composition in quantum mechanics and the exclusion result here.

Information has an irreducible grain size. A minimal causal patch cannot encode more than one irreducible bit of committed information. This provides a microscopic realisation of holographic-style information bounds — bounds that in other approaches must be introduced as independent postulates or derived from gravitational considerations. Here they follow from the primitive structure of fact formation. The information capacity of any region is bounded by the number of threshold cells it contains, each carrying exactly one bit.

Two spatial dimensions are primitive; three-dimensional geometry is reconstructed. Theorem 2.3a establishes that the primitive substrate on which facts form is two-dimensional: one spatial dimension is insufficient for stable localised commitment, three or more are excluded by distinguishability decay under causal propagation, and two are the unique admissible primitive spatial structure. The apparent three-dimensional geometry of the world is reconstructed from correlations across successive causal updates. This is not a speculation — it is a consequence of the same finite-distinguishability constraints that force the packing limit.

Gravity is one consequence. Within the VERSF framework, identifying ρ as the primitive commitment-capacity density rather than an external vacuum energy density, and establishing the cubic invariant $\chi(L) = (L/\xi)^3$ as a derived packing law rather than an action-budget ansatz, completes a step in the gravitational derivation. But gravity is downstream. The upstream result is the constraint on fact formation, and that constraint applies to any physically meaningful theory — not to gravitational physics specifically.

The paper should therefore be read as establishing a foundational structural constraint, with gravity as one of its applications. The technical derivations that follow are in service of that constraint.

---

2. Prior Structural Inputs

2.1 Primitive Entropy Quantum

The admissible closure entropy is forced to take logarithmic form

$\tilde{S}(N) = \ln N$ [Theorem A]

under monotonicity, additivity, and the null-singleton condition. Binary minimality (Lemma B) then implies that the smallest non-trivial irreversible refinement is 1 → 2, giving the primitive entropy quantum in closure units:

$\tilde{\Theta}_0 = \ln 2$

Theorem A and Lemma B together entail that the admissible entropy spectrum is discrete at exactly the integer multiples of ln 2: the set $\{0, \ln 2, 2 \ln 2, 3 \ln 2, \ldots\}$. No value between consecutive multiples is admissible within the closure-entropy algebra. This is unconditional within the closure-entropy framework and is the operative constraint in Case 3 of Theorem 4.1.

2.2 Primitive Commitment Energy

The commitment-barrier chain establishes the primitive commitment energy scale

$Ec = r \cdot \hbar c/\xi$

with r a dimensionless coefficient that takes the value $r = 1$ under the physical entropy identification and natural primitive-unit normalisation (the primitive consistency theorem). Throughout this paper we work at $r = 1$:

$Ec = \hbar c/\xi$

2.3 Finite Distinguishability and Local Capacity Competition

The finite-distinguishability paper establishes that bounded physical systems must be described in terms of locally attributable capacity allocation. Cross-channel terms are excluded from the per-channel observable algebra because they are not operationally resolvable under finite distinguishability: shared constraints act through local capacity competition, not through jointly attributed observables.

This result is the operative constraint in Case 2 of Theorem 4.1.

2.3a On the Two-Dimensional Geometry of the Coherence Patch

The derivation of ρc requires identifying the geometric measure of the minimal commitment-supporting region. This paper adopts a two-dimensional area ξ^2 rather than a three-dimensional volume ξ^3. The justification is provided here as a self-contained theorem with proof sketch.

Theorem 2.3a (Dimensional Minimality of Fact-Producing Substrates). *Let a physical substrate satisfy: (i) finite distinguishability, (ii) localised irreversible commitment (primitive facts), (iii) operational resolvability within bounded causal regions, and (iv) path-independent causal propagation under finite capacity. Then the number of primitive spatial dimensions is exactly two.*

Proof sketch. The proof proceeds by elimination and sufficiency across four steps.

Step 1 — One spatial dimension is insufficient.

In a one-dimensional substrate, all relations reduce to total ordering. There is no notion of enclosure or interior–exterior separation: any bounded interval shares both endpoints with its complement, and a causal disturbance cannot be localised without propagating globally in at least one direction. As a consequence, no region can be isolated for stable irreversible commitment: any primitive commitment event either propagates throughout the substrate or fails to stabilise as a bounded local fact. One dimension therefore cannot support the localised, stable, irreversible commitment events required by condition (ii).

Step 2 — Two spatial dimensions are sufficient.

In two spatial dimensions, local cycles exist and closed curves enclose well-defined interior regions. Boundary size grows linearly with region size (perimeter ∝ radius), providing sufficient causal capacity to absorb the thermodynamic cost of irreversible commitment while preserving locality: the cost of maintaining a boundary grows no faster than the boundary itself. This permits scalable, stable fact storage: primitive commitment events can be localised within bounded regions that are causally isolated from their exterior at manageable boundary cost. Conditions (i)–(iv) are simultaneously satisfiable in two spatial dimensions.

Step 3 — Three or more spatial dimensions are excluded at the primitive level.

Under finite distinguishability and bounded per-event capacity (condition i), causal propagation must distribute commitment resources across an expanding wavefront. In D spatial dimensions, the number of addressable causal registers at radius r from a source grows as:

$$N_D(r) \propto r^{(D-1)}$$

For $D = 2$, N grows linearly: each successive shell carries one additional unit of boundary, and commitment capacity distributes without dilution. For $D \geq 3$, N grows superlinearly: the wavefront size grows as $r^{(D-1)}$, while commitment capacity supplied from the source grows as r^1 (linearly with propagation distance under finite-capacity constraints, since each causal step contributes one unit of capacity budget). The per-register capacity therefore scales as $r^1 / r^{(D-1)} = r^{(2-D)}$. For $D = 2$ this is $r^0 = 1$: constant, admissible. For $D \geq 3$ this is $r^{(2-D)} \rightarrow 0$ as r increases: per-register capacity decays to zero, falling below the threshold $\tilde{\Theta}_0 = \ln 2$ required to commit a fact at radius r. Causal propagation of primitive commitment is therefore non-admissible at large radius in $D \geq 3$ dimensions. Maintaining resolvability for primitive commitment at $D \geq 3$ would require violating at least one of locality (concentrating capacity non-locally), path-independence (allowing route-dependent capacity accumulation), or isotropy. Each of these violates condition (iv). Three or more primitive spatial dimensions are therefore structurally excluded by the conjunction of finite distinguishability and admissible causal propagation.

Step 4 — The third nominal parameter is a descriptive index, not a spatial direction.

A potential escape is to treat the third coordinate as a distinct spatial direction unrelated to coarse-graining. However, any parameter associated with scale, resolution, or renormalization flow fails all four necessary criteria for spatiality established in the companion paper (Taylor, "Depth Is Not a Direction," 2025, Theorem A.4): it lacks intrinsic metric structure (distances are scheme-dependent), has no locality structure (descriptions at different scales relate by functional dependence, not by independent interacting subsystems), supports no propagation (coarse-graining maps are non-injective, not automorphisms), and is irreversible by construction (a non-injective map R_z cannot admit an inverse on physical states). Such a parameter is a descriptive index rather than a geometric direction. This is not merely a failure of representation but a structural obstruction: any parameter governed by non-injective maps cannot support the group structure required for spatial translation, and therefore cannot contribute to spatial dimensionality. Any putative additional primitive direction would either reduce to a combination of the two already identified (not independent) or reintroduce this non-injective structure (excluded by the reversibility criterion). The surviving parameter count under the four-criteria filter is therefore exactly two — not "three minus one" in some abstract sense, but two as the cardinality of the set of parameters satisfying all four criteria simultaneously. This is not a reduction from an assumed higher-dimensional space, but a direct determination: the admissible set of parameters satisfying the defining criteria of spatial dimensionality has cardinality exactly two. No additional independent parameter satisfying all four criteria exists; any further parameter either fails at least one criterion or is functionally dependent on the existing two, and therefore does not increase dimensionality.

Conclusion. The only dimensionality that supports localised irreversible commitment ($D \geq 2$ from Step 1), preserves finite distinguishability under causal propagation ($D \leq 2$ from Step 3), and admits only genuinely spatial parameters (Step 4), is exactly $D = 2$. □

Corollary 2.3b (Geometry of the Coherence Cell). *The coherence cell — the smallest bounded causal patch supporting a primitive commitment event — has area* ξ^2*, not volume* ξ^3*. The primitive commitment-capacity density is* $\rho c = Ec/\xi^2 = \hbar c/\xi^3$*, and the commitment-capacity invariant takes the cubic form* $\chi(L) = \rho L^3/\hbar c = (L/\xi)^3$*.*

Proof. By Theorem 2.3a, the primitive spatial manifold is two-dimensional. The coherence cell (Definition 3.1) is the smallest bounded patch in that manifold, which is an area element ξ^2. Substituting into $\rho c = Ec/\xi^2$ with $Ec = \hbar c/\xi$ gives $\rho c = \hbar c/\xi^3$ and $\chi(L) = (L/\xi)^3$. □

Remark on operational versus emergent dimensionality. Theorem 2.3a concerns the dimensionality of the primitive substrate on which facts are formed. It does not assert that the effective reconstructed description of the world is two-dimensional everywhere. The apparent three-dimensional structure arises from correlations across successive causal updates; depth behaves as a derived coordinate encoding structured inter-slice differences rather than as an independent spatial direction. Fact formation is two-dimensional; three-dimensional geometry is reconstructed.

Compatibility with the commitment-barrier chain. The commitment-barrier chain establishes $Ec = \hbar c/\xi$ as a dimensional relationship involving only the linear scale ξ, not the cell area or volume. It therefore holds regardless of cell dimensionality. The coefficient $r = 1$ follows from the primitive consistency theorem, which depends on the entropy quantum ln 2 and the energy scale Ec, not on the dimensionality of the cell.

Consistency with the Bekenstein–Hawking companion result. The companion BH paper finds one independent binary commitment per boundary area $4\ell_p^2$ on a null horizon — consistent with 2D commitment geometry, and unexpected if the fundamental geometry were volumetric. That derivation is explicitly on a null surface and does not by itself establish 2D geometry in bulk; Theorem 2.3a provides the independent positive case.

2.4 Primitive Observable Algebra

We now make explicit a structural requirement on what counts as a physical observable at the threshold scale. This is needed both to define the trichotomy precisely and to answer the energy objection in Case 2.

Definition 2.4 (Primitive observable algebra). The *primitive observable algebra* $\mathcal{A}(\xi^2)$ at the coherence threshold is the algebra of quantities that are:

1. Definable from a single primitive commitment event within the threshold cell;
2. Jointly resolvable within the action budget of the cell;
3. Attributed per-channel in the sense of the finite-distinguishability paper.

In particular, any counting-type observable — such as a Hamiltonian that assigns energy nEc to n primitive events — belongs to $\mathcal{A}(\xi^2)$ only when the configuration satisfies the observable's *measurement preconditions*. This reflects the following foundational commitment of the framework, stated here explicitly:

Operationalist Precondition Principle (OPP). An observable is defined within $\mathcal{A}(\xi^2)$ only relative to a configuration in which its measurement preconditions are satisfiable. For a counting operator that assigns energy nEc to n primitive events, the measurement precondition is the existence of n independently attributable events in the sense of Definition 2.5. This principle is not derived in the present paper; it is a foundational constraint of the framework, grounded in the finite-distinguishability paper's requirement that observables admit per-channel attribution.

The OPP is invoked in Case 2 of Theorem 4.1 to establish that the 2Ec energy operator is not a member of $\mathcal{A}(\xi^2)$ for a Class I configuration — not as a consequence of the proof, but as the minimal condition under which multiplicity corresponds to physically distinct facts. Any framework that permits multiplicity without independent attribution must either abandon the identification of facts with observable distinctions or admit physically meaningless degrees of freedom. OPP is therefore not a choice but a constraint: it is what distinguishes a theory of facts from a theory of labels.

The deeper justification for OPP — why it is not an arbitrary framework choice — is the following lemma, which shows that OPP is not an additional assumption but a restatement of what it means for something to be a physical fact.

Lemma 2.4a (Attribution Condition for Physical Facts). *A physical fact is a distinction that is stably attributable within a bounded causal region. Any quantity that cannot be attributed to a distinct physical event within the local observable algebra does not correspond to a distinct fact but to an ungrounded multiplicity.*

Proof sketch. A physical fact, within the operationalist framework established by the finite-distinguishability paper, is defined by the existence of an attributable observable difference: something is a fact if and only if it can be stably registered as a difference by some $\hat{O} \in \mathcal{A}(\mathscr{R})$. A purported multiplicity that leaves every observable unchanged — that produces no attributable difference anywhere in $\mathcal{A}(\mathscr{R})$ — is not a fact but a nominal label with no physical correlate. The attribution condition is therefore not a restriction imposed on physics from outside; it is constitutive of what distinguishes facts from mere labels. □

OPP follows immediately: counting operators presuppose independently attributable events because, without independent attribution, multiplicity has no physical meaning — it is exactly the kind of ungrounded nominal label that Lemma 2.4a excludes. Case 2 of Theorem 4.1 therefore does not define configurations out of existence by fiat; it shows that configurations lacking independent attribution cannot correspond to multiple physical facts, because the very notion of multiple facts requires the attribution condition to be satisfied. The OPP is therefore not an additional postulate but the minimal condition under which multiplicity corresponds to physically distinct facts; without it, the notion of fact itself loses operational meaning.

2.5 Formal Definition of Independent Attribution

Definition 2.5 (Independent attribution). Two events e_1, e_2 occurring within a bounded causal region $\mathscr{R}$ are *independently attributable* if and only if there exist observables $\hat{O}_1$, $\hat{O}_2$ in the local per-channel observable algebra $\mathcal{A}(\mathscr{R})$ such that:

1. $\hat{O}_i$ certifies the occurrence of e_i for each $i \in \{1, 2\}$;
2. $\hat{O}_1$ and $\hat{O}_2$ are *jointly resolvable*: they can be simultaneously assigned definite values without invoking resources beyond the action budget of $\mathscr{R}$;
3. Neither $\hat{O}_1$ nor $\hat{O}_2$ requires reference to the other event for its definition or measurement — that is, they admit per-channel attribution in the sense of the finite-distinguishability paper.

If no such pair $(\hat{O}_1, \hat{O}_2)$ exists, the events are *not independently attributable*, regardless of whether they are numerically distinct or not.

Two events may be numerically distinct yet fail to be independently attributable if their joint attribution would exceed the local capacity budget or require substructure below the primitive threshold.

2.6 The Derived Trichotomy

A central requirement of the proof of Theorem 4.1 is that the case analysis is exhaustive. We derive the trichotomy here, before the proof, from the structure of the observable algebra. It is not an ontological classification imposed from outside; it is an induced partition on event pairs under $\mathcal{A}(\xi^2)$.

Definition 2.6 (Observable-algebra trichotomy). Let e_1, e_2 be two events within a coherence cell. Say that an observable $\hat{O} \in \mathcal{A}(\xi^2)$ *fully separates* (e_1, e_2) if $\hat{O}(e_1) \neq \hat{O}(e_2)$ and $\hat{O}$ independently certifies both events in the sense of Definition 2.5 — that is, $\hat{O}$ alone, or together with a jointly resolvable partner, satisfies all three conditions of Definition 2.5 for the pair. With this criterion in place, the pair falls into exactly one of the following three classes:

- **Class I (Fully identified):** $\forall\, \hat{O} \in \mathcal{A}(\xi^2)$, $\hat{O}(e_1) = \hat{O}(e_2)$. No observable assigns different values to the two events.
- **Class D (Fully separated):** $\exists\, \hat{O} \in \mathcal{A}(\xi^2)$ that fully separates (e_1, e_2) in the sense above: $\hat{O}(e_1) \neq \hat{O}(e_2)$ and $\hat{O}$ independently certifies both events per Definition 2.5.
- **Class P (Partially resolved):** Some observables agree on (e_1, e_2) and some assign different values, but no observable fully separates them in the sense of Class D — no observable in $\mathcal{A}(\xi^2)$ simultaneously assigns different values to the pair and satisfies the independent-certification conditions of Definition 2.5.

Lemma 2.7 (Exhaustiveness of the trichotomy). Classes D, I, and P are mutually exclusive and jointly exhaustive.

Proof. Mutual exclusivity follows directly from the definitions. Class I requires every observable to agree; Class D requires the existence of a fully separating observable (which therefore disagrees on the pair); these are incompatible, so $I \cap D = \emptyset$. Class D requires a fully separating observable; Class P explicitly denies the existence of any such observable; so $D \cap P = \emptyset$. Class I requires all observables to agree; Class P requires some to disagree; so $I \cap P = \emptyset$.

Joint exhaustiveness: take any pair (e_1, e_2). Either every observable agrees on the pair (Class I), or at least one observable disagrees. If at least one disagrees, either some disagreeing observable fully separates the pair per Definition 2.5 (Class D), or no disagreeing observable does (Class P). These three outcomes are the only logical possibilities. □

Remark 2.8. Class P is the residual class: event pairs where the observable algebra partially resolves the pair — some observables differentiate them, none fully separates and independently certifies them. It is not an independent ontological posit; it is defined precisely as the complement of $I \cup D$ within the space of event pairs.

Remark 2.9. No additional class is possible. Any proposed "exotic" configuration must be assessed via $\mathcal{A}(\xi^2)$. Any such configuration either has every observable agreeing (Class I), has a fully separating and independently certifying observable (Class D), or has neither (Class P). There is no fourth logical option consistent with a well-defined observable algebra, because Classes I, D, and P together cover all configurations by Lemma 2.7.

---

3. The Coherence Cell and Threshold Activation

Definition 3.1 (Coherence cell). A *coherence cell* is the smallest bounded causal patch whose action budget is sufficient to support one primitive irreversible commitment event. Its spatial scale is ξ, determined by $\chi(\xi) = 1$.

Lemma 3.2 (Threshold Activation). *At $\chi(\xi) = 1$, the coherence cell realises at least one primitive commitment event. The threshold is saturating: a configuration with zero events at $\chi = 1$ is structurally inconsistent with the definition of ξ.*

Proof. The action budget of a region of size L is $\chi(L) = \rho L^3/\hbar c$. At $L = \xi$, $\chi = 1$ by definition: this is the minimal action budget sufficient to support one primitive commitment. For $\chi < 1$, no primitive event is admissible — the action budget is sub-threshold. For $\chi = 1$, exactly the threshold budget is present. A configuration with zero primitive events at $\chi = 1$ would assert that the threshold budget is present but no event occurs — that is, the region has the capacity for one primitive commitment but realises none. This contradicts the definition of ξ as the *minimal commitment-supporting scale*: if zero events occur at $\chi = 1$, then ξ is not the minimal commitment-supporting scale (it supports no commitment), and the threshold condition $\chi(\xi) = 1$ is vacuous. Since ξ is defined as the scale at which the first primitive commitment becomes admissible and realisable, the threshold is not merely permissive but saturating. At $\chi(\xi) = 1$, at least one primitive event occurs. □

This establishes the lower bound $N_{prim}(\xi^2) \geq 1$ on structural grounds, independent of any probabilistic or dynamical argument.

---

4. No Multi-Primitive Occupancy at Threshold

We now prove the central result using the derived trichotomy of Definition 2.6 and Lemma 2.7.

Theorem 4.1 (No Multi-Primitive Occupancy). *A bounded causal region at the coherence threshold $\chi(\xi) = 1$ cannot support two or more independently attributable primitive commitment events.*

Proof. Suppose for contradiction that a coherence cell of scale ξ supports at least two independently attributable primitive commitment events e_1 and e_2, in the sense of Definition 2.5. By Lemma 2.7, the pair (e_1, e_2) falls into exactly one of Classes D, I, or P. We show each leads to a contradiction.

Case 1 — Class D: Fully separable events.

Suppose $(e_1, e_2) \in$ Class D: there exists $\hat{O} \in \mathcal{A}(\xi^2)$ with $\hat{O}(e_1) \neq \hat{O}(e_2)$.

By Definition 2.4, every observable in $\mathcal{A}(\xi^2)$ is definable from a single primitive commitment event within the threshold cell. An observable that separates e_1 from e_2 must resolve a physical distinction between them within the cell. Any observable that separates e_1 and e_2 induces a partition of the cell into distinguishable sub-configurations; such a partition corresponds to a refinement of the closure structure and therefore to a smaller admissible coherence scale — contradicting the minimality of ξ. Any observable that separates e_1 and e_2 therefore induces a refinement of the closure partition itself, not merely a descriptive distinction, and thus contradicts the definition of the primitive threshold. The distinction resolved is either spatial (the events occupy distinct sub-regions of the cell) or intrinsic (the events carry different internal properties).

Spatial separation: If $\hat{O}$ resolves distinct spatial sub-regions within the cell, those sub-regions constitute operationally distinct commitment-supporting structures below the threshold scale ξ. But ξ is defined as the *minimal* commitment-supporting scale (Definition 3.1): the existence of two independently resolvable sub-regions within the cell implies a refinement of the commitment structure below ξ, contradicting minimality. If such substructure existed, the coherence scale would be smaller than ξ.

Intrinsic separation: If $\hat{O}$ resolves a difference in internal properties — labels, charges, or quantum numbers — then e_1 and e_2 carry distinct internal degrees of freedom beyond the entropy quantum $\ln 2$. But a primitive commitment event is defined as carrying no internal structure beyond $\ln 2$; additional degrees of freedom would constitute a richer event type, contradicting binary irreducibility (Lemma B).

Mode structure: A referee may object that e_1 and e_2 could occupy distinct pre-existing modes within the cell — independent degrees of freedom that do not require sub-cell spatial resolution. This escape is closed as follows. The existence of multiple independent modes within the coherence cell would require a decomposition of the cell's degrees of freedom into distinguishable channels. By Definition 2.5, such channels would admit independent attribution and therefore constitute resolvable substructure within the cell. But resolvable substructure within the cell is precisely what Case 1 (spatial separation) already rules out. The primitive event does not occupy a pre-existing mode structure; it *defines* the minimal mode structure. Any mode decomposition below ξ contradicts the minimality of the coherence scale.

Case 1 leads to contradiction with primitive minimality and the definition of ξ. $\square_1$

Case 2 — Class I: Fully identified events.

Suppose $(e_1, e_2) \in$ Class I: $\forall\, \hat{O} \in \mathcal{A}(\xi^2)$, $\hat{O}(e_1) = \hat{O}(e_2)$.

Since no observable in $\mathcal{A}(\xi^2)$ distinguishes e_1 from e_2, there are no distinct certifying observables $\hat{O}_1 \neq \hat{O}_2$ in $\mathcal{A}(\xi^2)$ satisfying Definition 2.5. Any observable that certifies e_1 certifies e_2 equally; no observable separates them. The two events therefore share a single certifying channel. Their joint attribution reduces to a single attributed observable, not two independent ones, violating condition (1) of Definition 2.5.

Indistinguishable events correspond to the same equivalence class under the admissibility quotient on $\mathcal{A}(\xi^2)$. Within the admissibility quotient, equivalence classes are the physical carriers of facthood: distinct representatives of the same class do not correspond to distinct physical facts — they correspond to the same fact described redundantly. Since physical facts correspond to equivalence classes under the admissibility quotient, multiplicity within a class is not physically meaningful: it does not increase the cardinality of the fact space. Multiplicity within an equivalence class is therefore not merely unobservable but undefined — it does not correspond to a distinct element of the physical state space. Claiming that two fully identified events are nonetheless "two" is a statement with no physical content: it cannot be confirmed, refuted, or used in any derivation within the observable algebra. Under the framework's operationalist commitment, such a claim is inadmissible.

The energy objection: One might object that two fully identified events still carry combined energy 2Ec, so the Hamiltonian H with eigenvalue 2Ec provides a separating observable — distinguishing the two-event configuration from the one-event configuration and thereby placing the pair in Class D rather than Class I. This objection is closed by the Operationalist Precondition Principle (OPP, Definition 2.4). The OPP establishes, as a foundational commitment of the framework rather than a consequence of this proof, that a counting Hamiltonian assigning energy nEc belongs to $\mathcal{A}(\xi^2)$ only when n independently attributable events satisfy its measurement preconditions. This is stated prior to and independent of the proof of Theorem 4.1. Given OPP, the 2Ec operator's measurement precondition — two independently attributable events — is not satisfied in Class I, because Class I is defined as the case where no observable fully separates and independently certifies the pair. The 2Ec operator therefore does not belong to $\mathcal{A}(\xi^2)$ for a Class I configuration, not as a conclusion of the proof but as a consequence of the framework's foundational operationalist commitment. The only admissible energy assignment is Ec.

The bosonic objection: Bosonic multiple occupancy is operationally meaningful because distinct occupation numbers produce distinct observable signatures (photon counts, field amplitudes). At the threshold cell in Class I, no such signatures exist: every observable assigns the same value to both events. They are not two physical facts; they are one equivalence class with nominal multiplicity two, which is operationally identical to multiplicity one within $\mathcal{A}(\xi^2)$.

Case 2 leads to contradiction with operational resolvability and the admissibility quotient. $\square_2$

Case 3 — Class P: Partially resolved events.

Suppose $(e_1, e_2) \in$ Class P: some observables in $\mathcal{A}(\xi^2)$ agree on (e_1, e_2) and some assign different values, but no observable fully separates them in the sense of Class D.

Class P is not an independent ontological category. It is the residual class: the set of event pairs that are neither fully identified nor fully separated by the observable algebra (Remark 2.8). Every pair not in Class I and not in Class D is in Class P by construction.

To close Case 3 formally, we require the following bridge lemma, which connects observable partial resolution to the entropy spectrum.

Lemma 4.1a (Entropy–Resolution Correspondence). *Let two events be partially resolved within $\mathcal{A}(\xi^2)$ — that is, $(e_1, e_2) \in$ Class P. Then the distinguishability structure induced on the pair by $\mathcal{A}(\xi^2)$ corresponds to a partition refinement strictly between the trivial single-class partition (full identification) and a full binary refinement (full separation). Under Theorem A and Lemma B, any such intermediate refinement corresponds to an entropy increment strictly between 0 and ln 2 per event contribution — that is, the joint entropy increment satisfies:*

$\ln 2 < \Delta S_{joint} < 2 \ln 2$

Proof sketch. The *refinement lattice* referred to here is the partial order on partitions of the closure set induced by successive binary refinements under Lemma B: the bottom element is the trivial single-class partition (no distinction made), the first non-trivial level consists of binary partitions arising from one 1 → 2 splitting, subsequent levels from iterated splittings, and the ordering is refinement (finer partitions are higher). Each primitive commitment event corresponds to one step up this lattice — one complete binary refinement, a 1 → 2 splitting contributing exactly ln 2. Class I corresponds to the two events sharing a single step — one splitting, $\Delta S = \ln 2$, lattice position one level above the trivial partition. Class D corresponds to two fully independent steps — two independent splittings, $\Delta S = 2 \ln 2$, two levels above. Class P is defined as the residual: neither fully identified (one shared level) nor fully separated (two independent levels). By the definition of Class P, the pair occupies a lattice position strictly between these two endpoints — some observables agree (not at full independence) and some disagree (not at full identification), but no single observable achieves the full binary independence required by Class D.

One step in this correspondence requires making explicit: the move from Class I membership at the observable-algebra level to a single shared binary refinement at the substrate level assumes that observable equivalence entails substrate-level entropy sharing — that two events indistinguishable in $\mathcal{A}(\xi^2)$ correspond to one physical substrate transition rather than two. This is not independently derived here; it follows from the operationalist framework established by Lemma 2.4a. Under that framework, the observable algebra exhausts the physical content of the cell: there is no substrate reality beyond what $\mathcal{A}(\xi^2)$ registers. Class I membership therefore just is one physical fact, and one physical fact corresponds to one binary refinement and one entropy quantum. The correspondence is not an additional assumption but a consequence of the operationalist commitment already in place.

This intermediate lattice position corresponds, under the entropy–refinement correspondence of Theorem A and Lemma B, to an entropy increment strictly between ln 2 and 2 ln 2. □

A referee might object that two partially resolved events could still yield a joint entropy of 2 ln 2 — that is, that partial resolution at the observable level is compatible with full entropy production. This is ruled out by the correspondence: for the joint entropy to equal 2 ln 2, the events must correspond to two fully independent binary refinements, which requires a fully separating observable in $\mathcal{A}(\xi^2)$ — placing the pair in Class D, not Class P. Partial resolution explicitly violates the independence required for two separate binary refinements, and therefore cannot produce the entropy of two independent primitive events.

By Theorem A and Lemma B, admissible entropy increments are restricted to integer multiples of ln 2: the set $\{0, \ln 2, 2\ln 2, \ldots\}$. No value in the open interval $(\ln 2, 2\ln 2)$ is admissible. Lemma 4.1a establishes that Class P configurations require precisely such an inadmissible intermediate value. Two independent binary refinements correspond to a direct-product structure in the refinement lattice; Class P configurations, by definition, do not admit such a factorisation and therefore cannot produce the entropy of two independent primitive events. Class P is therefore structurally excluded by the entropy algebra.

Case 3 leads to contradiction with binary irreducibility and the Entropy–Resolution Correspondence. $\square_3$

Conclusion. By Lemma 2.7, Classes D, I, and P are exhaustive and mutually exclusive: every pair of events in $\mathcal{A}(\xi^2)$ belongs to exactly one class. Each class leads to a contradiction. Therefore, the assumption that two or more independently attributable primitive commitment events can occur within a single threshold coherence cell is false.

The upper bound Nprim(ξ^2) ≤ 1 is established. Combined with the lower bound Nprim(ξ^2) ≥ 1 from Lemma 3.2 (Threshold Activation), we obtain:

Nprim(ξ^2) $= 1$ □

The proof establishes not merely that multi-primitive occupancy is disfavoured, but that it is *structurally impossible*: every logically admissible configuration of two primitive events within a threshold cell violates a foundational constraint of the framework. Case 1 violates primitive minimality; Case 2 violates operational resolvability and the admissibility quotient; Case 3 violates the discreteness of the admissible entropy spectrum. Since the three cases are derived from and exhaustive within the observable algebra (Lemma 2.7), there is no configuration — however constructed — that escapes all three constraints simultaneously.

The result is independent of dynamical law, statistical interpretation, and specific field content: it follows solely from the structural requirements of distinguishability, irreversibility, and minimal commitment. Changing the physical model — its dynamics, its fields, its specific interaction terms — leaves the conclusion intact, because none of those elements enter the proof.

The theorem therefore establishes a constraint not merely within a chosen representation, but on any representation that supports finite distinguishability, primitive irreversibility, and physically attributable facts. Any such representation must either reproduce the structure of $\mathcal{A}(\xi^2)$ at the threshold scale or violate one of these three conditions — in which case it does not qualify as a framework for fact formation in the sense required by the programme.

Remark 4.2 (PCO as corollary). The Primitive Cell Occupancy claim — *a coherence cell at threshold supports exactly one independently attributable primitive commitment event* — follows immediately from Theorem 4.1 and Lemma 3.2. It is a corollary, not a principle.

---

5. Primitive Capacity Density

Theorem 2.3a is not a geometric aside but a load-bearing result: it fixes the dimensionality of the coherence cell as ξ^2. The identification of the primitive density ρc depends directly on dimensional minimality — without the 2D result, the cell geometry would be undetermined and the density derivation would not close. What follows is therefore conditional on Theorem 2.3a.

Once Theorem 4.1 is established, the density follows directly.

Definition 5.1 (Primitive commitment-capacity density). The *primitive commitment-capacity density* ρc is the energy density associated with primitive fact-formation capacity: the energy per unit two-area required to support primitive irreversible commitments at threshold. ρc is therefore not a vacuum property of space but a capacity property of the fact-forming substrate: it measures the maximum rate at which irreversible commitments can be supported per unit area per unit causal depth.

Derivation. By Remark 4.2, one coherence cell of area ξ^2 supports exactly one primitive commitment event of energy $Ec = \hbar c/\xi$. Therefore:

$$\rho c = Ec/\xi^2 = (\hbar c/\xi)/\xi^2 = \hbar c/\xi^3$$

More generally, retaining the dimensionless coefficient r:

$$\boldsymbol{\rho c = r\hbar c / \xi^3}$$

which reduces to $\hbar c/\xi^3$ under primitive consistency $r = 1$.

---

6. Main Theorem

Theorem 6.1 (Primitive Density Theorem). *Assume:*

1. *The primitive commitment energy is* $Ec = r\hbar c/\xi$ *with r dimensionless.*
2. *The coherence scale* ξ *is the minimal bounded causal scale supporting primitive commitment.*
3. *Binary irreducibility (Lemma B): the primitive entropy quantum is* $\tilde{\Theta}_0 = \ln 2$, *with admissible increments at integer multiples only.*

4. *Operational resolvability: independent attribution requires jointly resolvable per-channel observables in $\mathscr{A}(\xi^2)$ (Definition 2.5).*
5. *Primitive minimality: primitive events carry no internal substructure beyond ln 2.*

Then:

$$\rho c = r\hbar c / \xi^3$$

and under primitive consistency $r = 1$:

$$\rho c = \hbar c / \xi^3$$

Proof. By Theorem 4.1 and Lemma 3.2, assumptions (2)–(5) jointly imply that a coherence cell of area ξ^2 supports exactly one independently attributable primitive commitment event. By assumption (1), that event carries energy $Ec = r\hbar c/\xi$. The energy density of primitive commitment capacity is therefore $\rho c = Ec/\xi^2 = r\hbar c/\xi^3$. Setting $r = 1$ gives the stated result. □

---

7. Recovery of the Cubic Invariant

Corollary 7.1 (Derived Cubic Invariant). *For a bounded causal region of size L, the dimensionless primitive commitment capacity is*

$$\chi(L) = \rho c L^3 / \hbar c = r(L/\xi)^3$$

Under primitive consistency $r = 1$:

$$\chi(L) = (L/\xi)^3$$

and in particular $\chi(\xi) = 1$.

Proof. Substitute $\rho c = r\hbar c/\xi^3$ from Theorem 6.1 directly into the definition of $\chi(L)$. □

The threshold condition is recovered as a consequence of the primitive density theorem rather than imposed independently. The cubic invariant is no longer an action-budget ansatz; it is the microscopic derivation confirming that $\rho c = \hbar c/\xi^3$ is the correct primitive commitment-capacity density. The cubic form $\chi(L) = (L/\xi)^3$ was established independently in the commitment-barrier paper on action-budget grounds under 2D cell geometry; the present paper provides the microscopic derivation from single-cell occupancy under Theorem 2.3a and Theorem 4.1, and the two derivations agree exactly. Their agreement — from opposite directions, action-budget and microscopic — is what justifies confidence in the cubic invariant as the correct form.

The action-budget derivation (commitment-capacity paper) and this microscopic derivation now provide mutually independent routes to the same result. They agree exactly, which strengthens both.

8. Consistency Theorem: The Primitive Triad

A primitive commitment event is minimal in three simultaneous and mutually consistent senses:

Perspective	Minimal quantity	Value
Entropic	Entropy increment per event	$\tilde{\Theta}_0 = \ln 2$
Energetic	Energy per event at threshold	$E_c = \hbar c/\xi$
Geometric	Area per event at threshold	ξ^2

These are not three independent postulates. They are the same primitive event viewed in counting, energetic, and geometric language respectively. The primitive entropy quantum, the commitment energy scale, and the coherence cell geometry are not independent inputs but mutually constraining aspects of a single structure: the exclusion theorem (Theorem 4.1) ensures that these three minimalities coincide, yielding a unique density and a unique scaling law. The following theorem establishes their joint rigidity: no value of ρ_c other than $\hbar c/\xi^3$ is consistent with all three simultaneously.

Theorem 8.1 (Primitive Triad Consistency). *Given the primitive entropy quantum $\tilde{\Theta}_0 = \ln 2$, the primitive commitment energy $E_c = \hbar c/\xi$, and single-cell occupancy $N_{prim}(\xi^2) = 1$, the unique density consistent with all three minimality conditions is*

$$\rho_c = \hbar c/\xi^3.$$

Any other value of ρ_c either violates entropic minimality (assigning more or less than ln 2 per event), energetic minimality (assigning energy other than $\hbar c/\xi$ per event), or geometric consistency (placing more or fewer than one event per cell).

Proof. The density is defined as $\rho_c = E_c/\xi^2$. Substituting $E_c = \hbar c/\xi$ gives $\rho_c = \hbar c/\xi^3$. Any deviation from this value requires either $E_c \neq \hbar c/\xi$ (violating energetic minimality), $\xi^2 \neq$ (cell area) (violating geometric consistency), or $N_{prim} \neq 1$ (contradicting Theorem 4.1 and Lemma 3.2). Entropic minimality enters through the constraint on E_c: the energy $\hbar c/\xi$ is fixed by the commitment-barrier chain, which in turn is grounded in the entropy quantum ln 2. Deviating from $E_c = \hbar c/\xi$ would require a different entropy quantum, contradicting Theorem A and Lemma B. Therefore $\rho_c = \hbar c/\xi^3$ is the unique consistent value. □

Theorem 4.1 connects back to this triad directly: any attempt to place a second primitive event in the threshold cell produces a contradiction in at least one of the three registers — geometric (Case 1, substructure below ξ), operational (Case 2, no independent attribution), or entropic (Case 3, inadmissible intermediate increment). The triad is jointly rigid precisely because each register defends the others.

The primitive entropy quantum, commitment energy scale, and coherence cell geometry are therefore not independent inputs but mutually constraining aspects of a single structure. The exclusion theorem ensures their coincidence, and dimensional minimality fixes their geometric realisation.

---

9. Logical Status

Structurally closed by this paper

Claim	Status
$\tilde{S} = \ln N$, binary minimality, $\tilde{\Theta}_0 = \ln 2$	Unconditional from Theorem A and Lemma B
$Ec = \hbar c/\xi$	Established in commitment-barrier chain
Two-dimensionality of coherence patch (Theorem 2.3a)	**Proved here by elimination and sufficiency: 1D insufficient, D≥3 excluded, D=2 sufficient**
Geometry of coherence cell (Corollary 2.3b)	**Follows immediately from Theorem 2.3a**
Primitive observable algebra $\mathcal{A}(\xi^2)$ (Def. 2.4)	Stated here; grounded in finite-distinguishability paper
Operationalist Precondition Principle (OPP)	Foundational commitment stated here; not derived
Attribution Condition for Physical Facts (Lemma 2.4a)	**Grounded here; establishes OPP as necessary condition for facthood given the operationalist framework**
Independent attribution (Def. 2.5)	Stated here; operationally grounded
Observable-algebra trichotomy (Def. 2.6, Lemma 2.7)	**Derived here from $\mathcal{A}(\xi^2)$; mutually exclusive and exhaustive by construction**
Threshold Activation — lower bound $Nprim \geq 1$ (Lemma 3.2)	**Proved here from definition of ξ**
No Multi-Primitive Occupancy — upper bound $Nprim \leq 1$ (Theorem 4.1)	**Proved here by contradiction over derived trichotomy**
Entropy–Resolution Correspondence (Lemma 4.1a)	**Proved here; bridges Class P to inadmissible entropy increments**
PCO as corollary (Remark 4.2)	Follows from Theorem 4.1 + Lemma 3.2
$\rho c = \hbar c/\xi^3$ (Theorem 6.1)	Derived from Theorem 4.1 + Lemma 3.2 + prior results
$\chi(L) = (L/\xi)^3$ (Corollary 7.1)	Derived here; independently established in commitment-capacity paper
Primitive triad is jointly rigid (Theorem 8.1)	**Proved here**
Observable algebra completeness	Argued: extensions either refine ξ or violate attribution condition; formal proof conditional on Lemma 2.4a

Claim	Status
ρ is not an external input	Established here

On the Completeness of the Local Observable Algebra

The proof operates within $\mathcal{A}(\xi^2)$, the algebra of all physically attributable observables within a bounded causal region at the threshold scale. A natural objection is: could additional observables outside $\mathcal{A}(\xi^2)$ provide distinctions that the case analysis misses?

Any extension of the observable structure beyond $\mathcal{A}(\xi^2)$ must fall into one of two categories:

(i) Extensions that introduce additional distinguishability. If a proposed extended algebra $\mathcal{A}' \supset \mathcal{A}(\xi^2)$ resolves distinctions not resolvable within $\mathcal{A}(\xi^2)$, it resolves substructure or internal degrees of freedom below the threshold scale. By Case 1, such resolution either implies substructure below ξ (contradicting the minimality of the coherence scale) or internal degrees of freedom beyond ln 2 (contradicting binary irreducibility). An extension of type (i) does not extend the physics — it changes the threshold scale. The coherence scale for a theory with $\mathcal{A}'$ would be smaller than ξ, and $\mathcal{A}(\xi^2)$ would remain the correct algebra at that theory's threshold.

(ii) Extensions that introduce unresolvable distinctions. If $\mathcal{A}'$ posits distinctions that are in principle unresolvable by any measurement within the bounded causal region, those distinctions fail the attribution condition of Lemma 2.4a. They correspond to ungrounded multiplicities rather than physical facts, and are inadmissible under the framework's operationalist commitment.

Every physically meaningful extension therefore either (i) refines the coherence scale and is captured by $\mathcal{A}(\xi'^2)$ at the new threshold $\xi' < \xi$, or (ii) introduces unattributable distinctions excluded by Lemma 2.4a. In neither case does the extension invalidate Theorem 4.1 at threshold scale ξ. The conclusion applies to any representation supporting finite distinguishability, primitive irreversibility, and physically attributable facts — not merely to the specific representation of $\mathcal{A}(\xi^2)$. In particular, any purported extension that preserves physical meaning while escaping these two categories would require a notion of distinguishability that is neither locally attributable nor reducible to finer-scale structure, which is incompatible with the finite-distinguishability framework itself.

What remains conditional

The argument above reduces the completeness question to the attribution condition of Lemma 2.4a and the minimality of ξ, both already operative in the proof. The remaining open question is whether the attribution condition itself could be weakened or replaced by a different foundational operationalist commitment while still supporting the programme's goals. This is a question for the foundational theory of the observable algebra, not for the present paper.

The coefficient $r = 1$ enters through the commitment-barrier derivation and is treated as established from that chain.

Remark on geometric dimensionality. The invariant form $\chi(L) = (L/\xi)^3$ and the density $\rho c = \hbar c/\xi^3$ are derived here under the two-dimensional geometry established in Theorem 2.3a (coherence cell area ξ^2). That theorem proves dimensional minimality by elimination and sufficiency: one spatial dimension is insufficient for stable localised commitment, three or more are excluded by distinguishability decay under causal propagation, two are sufficient and constitute the unique admissible primitive spatial structure. The single-occupancy proof of Theorem 4.1 is independent of cell dimensionality — it depends only on the observable algebra structure — and holds in either geometric case. Theorem 2.3a is what determines the cubic form as operative.

---

10. Summary

The effective density ρ appearing in the commitment-capacity invariant is identified structurally as the primitive commitment-capacity density $\rho c = \hbar c/\xi^3$: the energy per unit two-area required to support primitive irreversible commitments at threshold.

The single-occupancy result is established by two independent structural results whose conjunction gives $N_{prim}(\xi^2) = 1$:

- **Lemma 3.2 (Threshold Activation):** the threshold $\chi(\xi) = 1$ is saturating — a cell at threshold realises at least one primitive event by definition of the threshold scale. This establishes the lower bound $N_{prim}(\xi^2) \geq 1$.
- **Theorem 4.1 (No Multi-Primitive Occupancy):** proved by contradiction over the derived observable-algebra trichotomy. Each case refutes a different aspect of multi-occupancy:
    - **Case 1** (Class D, fully separable): any separating observable in $\mathcal{A}(\xi^2)$ resolves substructure below ξ or internal labels beyond ln 2 — both contradictions.
    - **Case 2** (Class I, fully identified): events share the same equivalence class in the admissibility quotient; the $2E_c$ energy operator is inadmissible in $\mathcal{A}(\xi^2)$ for this configuration.
    - **Case 3** (Class P, partially resolved): requires an entropy increment in the open interval $(\ln 2, 2 \ln 2)$, which is excluded by the discrete spectrum of Theorem A and Lemma B.

No configuration survives. The occupancy is exactly one.

The cubic invariant $\chi(L) = (L/\xi)^3$ follows as Corollary 7.1, and Theorem 8.1 establishes that the primitive triad of entropic, energetic, and geometric minimality uniquely determines $\rho c = \hbar c/\xi^3$.

---

11. Positioning and Scope

The broader context and consequences of this paper's results are set out in the opening Scope and Consequences section. The core point bears restating here in one sentence: the constraint established in Theorem 4.1 — that a minimal causal patch supports exactly one independently attributable primitive commitment event — is a structural limit on fact formation in any finitely distinguishable universe, not a technical result specific to gravitational theory.

Gravity as one consequence

Within the VERSF framework, the result provides a crucial closure. It identifies ρ as the primitive commitment-capacity density rather than an externally supplied vacuum energy density, and converts the cubic invariant

$\chi(L) = \rho L^3/\hbar c$

from an action-budget ansatz into a derived packing law:

$\chi(L) = (L/\xi)^3$

This completes a specific step in the gravitational derivation. But as the opening section establishes, the upstream constraint governs fact formation in general — entropy quantisation, quantum non-classicality, information grain size, and the two-dimensionality of the primitive substrate all follow from the same structural requirements. Gravity inherits these constraints along with everything else.

Scope and Limitations

The argument assumes that $\mathcal{A}(\xi^2)$ is complete — that it captures all physically meaningful distinctions at the threshold scale. This is consistent with the finite-distinguishability results but would be further strengthened by an axiomatic derivation of the algebra's completeness. The coefficient $r = 1$ is inherited from the commitment-barrier derivation and is not re-derived here.

The result belongs at the level of foundational constraints, alongside locality, causality, and finite distinguishability — not as a model-dependent result specific to any particular physical domain.

The two-dimensionality established in Theorem 2.3a concerns the primitive substrate of fact formation — the level at which irreversible commitments occur. The apparent three-dimensional structure of physical space arises from correlations across successive causal updates, as established in the companion depth-reconstruction framework. The emergence of a third spatial dimension does not contradict the present result; it reflects the organisation of correlations across successive updates rather than an additional primitive degree of freedom. The third dimension is therefore not eliminated but reinterpreted as emergent: it is reconstructed from structured inter-slice differences rather than being primitive. This distinction matters for how the result should be read. It does not predict that the world looks two-dimensional at human scales; it predicts that the underlying fact-forming substrate is two-dimensional, and that three-dimensional geometry is a derived description of how those facts are organised across time.