No Pre-Individuation and the Seal-Trichotomy

Why FP1 Has No Ledger to Keep, and Why the Only Surviving Classical Reading Is a Purely Dynamical Seal

Keith Taylor VERSF Theoretical Physics Programme versf-eos.com

---

For the General Reader

A rolling die has its six faces *before* it lands. The faces are there, labelled, with probabilities attached, while the die is still tumbling; the landing only picks which of the pre-existing faces comes up. Call that the *ledger* picture — separate accounts, one per outcome, all present in advance, with the commitment (the landing) merely *revealing* which account was already going to win.

The companion papers showed that whether the deep substrate is a die-like *ledger* or a *bath* — one shared pool of possibility that gets carved up only when reality commits — is exactly what decides whether quantum probability comes out as the square of an amplitude. Bath gives the square; ledger gives something classical and non-quantum. The previous paper proved the two readings are genuinely different and left which one holds as an open choice.

This paper removes most of that choice, and it does so by looking at what the substrate's founding rule actually provides. The founding rule (FP1) says: a region of reality has a certain *capacity* to hold distinguishable things. That is all it says. It gives you a region and a number. It does **not** hand you a list of labelled outcomes the way the die comes pre-printed with 1 through 6. The labels — the separate outcomes — are not in the rule. They appear only once reality carves the region a particular way.

That single observation disproves the strongest version of the ledger. The die's picture requires the labelled accounts to exist *first*; the founding rule supplies no labels at all, only a region and its capacity. And we can say something sharper than "the labels are absent": the rule counts *how many* distinguishable things fit in a region, and that count *changes when you look more finely*. A fixed set of pre-named outcomes does not grow when you resolve more sharply — but the founding rule's count does. So the count cannot be a fixed list of outcomes; it is a measure of capacity, realized into specific outcomes only by a choice of how finely and in which directions to carve. The die-like picture is not merely unsupported by the rule — it *contradicts* the rule's own form.

A more careful opponent retreats one step. Grant, he says, that there are no pre-given outcomes; still, perhaps the *menu of allowed carvings* is fixed in advance — a settled list of ways the region may be cut, handed in from outside. We dispose of this too, and the way we do it is the heart of the paper. The founding rule supplies only the region and its capacity; so any allowed carving must be defined *from* that capacity, not from some external list handed in alongside. A carving-menu pinned down by anything other than the region's own capacity would be extra structure the rule never provided — the same smuggling the die committed, now one level up. So the menu of carvings is not handed in from outside; it is generated from within, by the region's capacity itself. We call this the **Internality Axiom** — and, crucially, it is not something we assume but something the founding rule forces.

What remains — and this is now the *entire* remaining mystery of why probability is a square — is one last question about that internally generated menu of carvings. Is it *alive* (the carvings flow into one another, capacity moving between them, until commitment freezes one) or *frozen* (a fixed set of sealed, separate carvings, one of which commitment simply picks)? Alive is the bath, and gives the square. Frozen is the last surviving classical reading, and does not. We call the assertion that it is alive **Reversible Connectedness (RC)**, and we do not settle it; settling it by preference would be the very move the programme has refused throughout. But we show it is *all* that is left — that the die is disproved, the externally-handed-in menu is disproved, and the only classical survivor is this single strange object: a menu generated from the region's own capacity, yet sealed rather than flowing. And we note that whether it flows is recognisably the same "is reality's reversible motion continuous?" question the programme already had to answer elsewhere. So the square reduces to one mild bookkeeping principle and one already-familiar continuity question — and nothing more.

Abstract

The Squaring Residue, the linearity companion, and *Bath or Ledger* reduced the Born exponent to **$\ell^2 \Longleftarrow$ PC $\wedge$ (off-diagonal mixing) $\wedge$ (connectivity)**, then showed mixing and connectivity reduce to the *bath reading* of pre-factual conserved weight, which — modulo a flagged identification (Claim 7.1.1) — is the dynamical-partition reading of the *Packing* paper's refinement obstruction (Obstruction B). *Bath or Ledger* left the bath/ledger fork **open** and **adopted** the bath reading without forcing it, the bare packing axiom permitting the rigid alternative.

This paper isolates two admissibility conditions on the family of admissible refinements of a region M, parallel in form and asymmetric in status:

- **IA (Internality Axiom):** admissible refinements are internally individuated by M's distinguishability-capacity structure, not by an external label set or pre-fixed menu.
- **RC (Reversible Connectedness):** admissible (internally individuated) refinements are connected under admissible capacity-preserving reversible motion.

The central results are: **FP1 $\Longrightarrow$ No External Individuation (proven)**, whence **IA (internality) follows** by exhaustive dichotomy — capacity is the only thing FP1 gives M, so $\neg$external $\Longrightarrow$ internal — with only the *uniqueness* of internal individuation left unclaimed and unneeded; **B $\Longleftrightarrow$**

RC given IA (proven), with **RC open**; and a **seal classification (proven)** that narrows any RC-failure to a single enforcer: the only data-level seal (an adjoined external component-charge) is excluded dynamics-free, and a native separating function cannot seal — the seal-making property is conservation, conservation is dynamical, and FP1 entails the conservation of no function — so every surviving seal is the irreducibly *dynamical* one, which is RC itself.

The **No-Pre-Individuation theorem (NPI)** establishes that FP1, as stated, attaches distinguishability capacity to an operational *region* M — its primitive datum is the scalar Vol_op(M), not an outcome-indexed tuple (C_1, …, C_d) over a pre-given label set Ω. The operational dimension d_op is a *capacity dimension*, not a labelled basis; and the distinguishability set Σ(M) is a *resolution-dependent cardinality* (|Σ(M)| grows as Δ_op shrinks), not a fixed labelled outcome space. The decisive supporting result is the **Coarse-Graining Non-Primitivity Theorem**: FP1 supplies no commitment-independent grouping of distinguishable content into a fixed number of outcome classes — any candidate fails on one of three exhaustive grounds (it adjoins external structure; or it is itself a refinement, hence commitment-selected; or it imposes a resolution-independent class-count FP1 does not contain). This theorem closes the microstructure objection (that a fixed Ω might survive by reading Σ(M) as growing microstructure within fixed outcomes) and, recurring at three levels — states, menu, refinement-components — does load-bearing work throughout. Hence the **primitive ledger PL** (per-outcome capacities existing and separately conserved before commitment — the classical die) requires structure FP1 supplies at no level and whose strongest form contradicts the axiom's resolution-dependence. **Proven (given FP1 as stated): PL is disproved.** The die is correctly *not* a counterexample: it carries its outcome space as external, commitment-independent, resolution-independent input — exactly what FP1 omits.

With PL disproved, an **Exhaustion Lemma** establishes that any FP1-admissible pre-factual conservation structure is either the **emergent ledger EL** (generated sub-capacities frozen) or the **bath B** (generated sub-capacities reallocable). EL is the *frozen-refinement* reading (Obstruction B negated); B is the *dynamical-refinement* reading (Obstruction B affirmed). So **bath-vs-ledger collapses, with no residual primitive-ledger branch, onto Obstruction B itself**, and **the Born residue compresses to PC ∧ Obstruction B** — proven, exhaustive, independent of whether the bath is ultimately forced.

We then force the *individuation* axis and isolate the *topology* axis. IA is **proven from NPI** (Lemma 7.1): external individuation re-imports the index structure NPI excludes, one level up, and is excluded by the adjunction (extra-data) argument alone (we do not lean on a resolution-independence claim — see §7.2). IA deliberately does **not** decide whether the internal menu is connected or sealed — it fixes the *source* of individuation, not the *topology*. The remaining axis is RC, which, given IA, is equivalent to the bath: **B ⟺ RC**. We do not assume RC; asserting it would be the affirmative fiat. The surviving classical reading is therefore neither the die nor an externally-retrieved menu — both disproved — but the single narrow object: an *IA-satisfying yet RC-failing* refinement family, internally individuated from M's own capacity yet sealed rather than flowing.

Net, stated at its honest two tiers. **Proven:** IA — the menu of carvings is internally individuated; the primitive ledger (the die) and the externally-retrieved menu are both disproved from the primitive; and the Born residue compresses, exhaustively, to **PC ∧ Obstruction B**. **Open frontier:** given IA, the bath — hence the squared norm — follows **if and only if RC**, which we name and decline to assume. So this paper does not deliver the square; it delivers IA and isolates RC as the single gate. RC and the phase-axis continuity upgrade the companions already carry are the same *species* of condition — connectedness of admissible reversible motion — so RC is plausibly not a new debt at all, in which case the programme's only genuinely new commitment beyond a debt already incurred is PC.

Contents

Scope and Conditional Status

This is the fourth companion to *The Squaring Residue*, after the linearity paper and *Bath or Ledger*, and inherits their framework and conventions. It does not assert the bath reading by preference. It isolates two admissibility conditions on the admissible-refinement family — the Internality Axiom (IA) and Reversible Connectedness (RC) — proves IA from the primitive, proves that the bath is equivalent to RC given IA, and declines to assume RC.

The two axioms, stated up front:

IA (Internality Axiom). Every admissible refinement of M is internally individuated by M's distinguishability-capacity structure (Vol_op(M), Δ_op, d_op): not by an external label set, and not by a pre-fixed menu whose members are individuated independently of M's capacity relations.

RC (Reversible Connectedness). Any two admissible (internally individuated) refinements of M are joined by a continuous path of admissible capacity-preserving reversible transformations of the substrate.

They are **parallel in form** (each constrains the admissible-refinement family) but **asymmetric in status**. What is *proven* is the exclusion **FP1 ⟹ No External Individuation** (§7.2): admissible refinements cannot be individuated by structure adjoined independently of M's capacity. IA — *every admissible refinement is internally individuated* — then **follows**: capacity is the only thing FP1 gives M, so ¬external ⟹ internal is an exhaustive dichotomy and internality is proven. What is *not* claimed is the *uniqueness* of the internal individuation (a referee may grant several internally individuated refinement structures); nothing downstream needs uniqueness, only internality. RC is an *open premise* (§7.3). We do not force it: because FP1 is silent on dynamics, asserting that reversible motion connects the carvings (flow) is as adjoined as the seal it would deny, so forcing RC would be the bath preference made asymmetric — the fiat the programme refuses. What §7.5 does instead is *narrow* RC: a seal classification (Theorem 7.5.2) shows the only data-level seal is the adjoined external charge — a native separating function cannot seal, since the seal-making property is conservation and FP1 supplies the conservation of no function — so every surviving seal is dynamical, which is RC (§7.5). IA constrains the family's *individuation*; RC constrains its *connectedness*. The phase-axis continuity upgrade the companions carry constrains the connectedness of admissible phase motion, so RC and that upgrade are the same species of condition, while IA is a different species (§7.4, §10.6).

Epistemic labelling is maintained throughout:

- **proven** — follows from the stated inputs (FP1 as stated, the inherited companion results, the declared elevation of §2.3) by an argument given here;
- **conditional** — follows given a named, separately-open premise;
- **conjectural** — asserted as plausible with structural support but no argument approaching proof.

NPI and the disproof of PL are **proven** given FP1 as stated (§4–5), defended across all candidate sources of a covert index set (outcomes, d_op, Σ(M), and a fixed coarse-graining of Σ(M)). The Exhaustion Lemma and the collapse of the residue to PC ∧ Obstruction B are **proven** (§6), independently of the bath-forcing attempt. (All inherited reductions carry the bath ⟺ Obstruction B identification, Claim 7.1.1, flagged once here; throughout this paper "proven" is relative to that inherited input, and we do not repeat the modulo on each theorem.) IA is **proven from NPI** (Lemma 7.1, §7.2). RC is **open**; given IA, B ⟺ RC (§7.3). Time is emergent throughout; "pre-factual" / "pre-commitment" denote the substrate configuration before an irreversible commitment event. A commitment event is here *declared* to be the selection of one admissible refinement of the unresolved region (§2.3) — a new but natural elevation, stated explicitly and held strictly weaker than both IA and RC.

A methodological note, because it is the discipline of the paper. The previous papers refused to let a contested conclusion ride in on a word — "reversible = isometry," "respects superposition," "genuine possibility." The trap this paper avoids is the symmetric one on the affirmative side: proving NPI and then announcing "therefore bath." NPI does *not* entail bath; it entails IA — internality, the individuation axis (uniqueness of the internal individuation is neither claimed nor needed) — and leaves the topology axis (RC) open. A second discipline guards IA itself: IA is held to *individuation*, not *topology*, and does not pretend to exclude a sealed internal menu, which is precisely the residue RC governs. The honest yield is the disproof of PL and external individuation, the exhaustive collapse onto Obstruction B, the derivation of IA, and the reduction of the remainder to RC.

---

1. Introduction: the trap, and the move that avoids it

Bath or Ledger reduced the Born exponent's mixing leg to a fork in the conservation structure of pre-factual weight: a single shared **bath** (capacity reallocable among alternatives, ⟹ continuous off-diagonal mixing ⟹ ℓ^2) versus separate **ledgers** (per-outcome accounts each conserved, ⟹ diagonal torus ⟹ classical, non-ℓ^2). It proved the equivalence mixing ⟺ bath, identified the fork with the unmined "Balance" content of BCB and — pursued to the packing primitive — with the *Packing* paper's refinement obstruction (Obstruction B), and then, scrupulously, declined to assert the bath: the bare axiom permits the rigid ledger, and adopting bath was flagged as a *reading*, not a theorem.

The temptation, having localized the question to the substrate's founding rule, is to argue from the commitment ontology directly: pre-commitment there are no separate facts, capacity belongs to the unresolved whole, *therefore* the conserved quantity is one shared pool — bath. This is the trap. It proves too quickly, and the thing it skips is exactly what classical probability supplies as a counterexample.

It is true that the ontology forces **reversibility** of pre-commitment dynamics — almost by definition, since commitment *is* the irreversible event and everything prior to it has not committed. But reversibility is *shared by the ledger*: rephasing each account and permuting accounts are reversible, bijective, total-conserving maps (Theorem 4.1 of *Bath or Ledger* — ledger ⟹ torus — is a theorem about a fully reversible dynamics with no mixing). So "everything is reversible until commitment" cannot select the bath; the die's pre-landing dynamics is reversible too. Reversibility is true and cheap; it is not the discriminator.

What the ontology *appears* to force, and what this paper makes precise, is **no pre-commitment individuation** — no labelled outcomes, hence no per-outcome accounts to conserve. But here too we must be exact, because there is a fallback the strong version overruns. "No pre-commitment accounts" disposes of the ledger *only if* the ledger needs its accounts to exist before commitment. The classical die does: its faces are labelled before it rolls. But a subtler ledger lets the labels appear *at* the carving and *then* freezes them — and conflating that with the die is how the fast argument cheats.

The contribution of this paper is threefold, and it is cleanest stated through the two axioms. First, disprove the primitive ledger (the die — PL) from the primitive: NPI (§4). Second, derive the Internality Axiom — the menu of carvings is individuated from M's own capacity, not handed in from outside — so that the externally-supplied refuge is disproved alongside the die: NPI ⟹ IA (§7.2). Third, show that what then remains lives entirely on the Connectedness Axiom —

whether the internally individuated carvings flow or seal — and that this is exactly the refinement-axis of a continuity question the programme already faced: B ⟺ RC (§7.3).

---

2. Inherited setup, FP1 as stated, and the two axioms

2.1 The primitive (FP1)

The foundational finite-packing axiom of the *Packing* paper bounds the distinguishability content of an operational region:

$|\Sigma(M)| \leq Vol\_op(M) / \Delta\_op^{d\_op}$,

where M is an operational region, Σ(M) the distinguishability set realizable within it, Vol_op(M) the operational volume (the distinguishability *capacity*) of M, Δ_op the minimal operational resolution, and d_op the operational dimension.

We record the structural facts this paper turns on, and nothing more:

The primitive datum of FP1 is the scalar capacity Vol_op(M) of a region M. The bound is on the *cardinality* |Σ(M)|, a function of the resolution Δ_op and dimension d_op. FP1 does not provide, and does not presuppose: an outcome-indexed tuple (C_1, …, C_d) over a pre-given label set; a pre-fixed menu of carvings individuated independently of M's capacity; nor a commitment-independent grouping of Σ(M) into a fixed number of outcome classes. There is no Ω in the axiom.

These are textual observations; §4 and §7 turn them into theorems and defend them against the covert-structure attacks.

2.2 Inherited results

From the companions we take as given: the carrier H0 (pre-factual $\psi = (c_1, \ldots, c\_d) \in \mathbb{C}^d$ as a path-sum); the separable conserved normalization $N(\psi) = \Sigma_i h(|c_i|)$ (BCB-additivity N1, diagonal phase C2, faithfulness N4, monotonicity H-mono); the Diagonal-Torus selection (continuous off-diagonal mixing ⟺ ℓ^2); the Possibility-Connectivity Theorem (mixing ⟺ bath); and the reduction — modulo Claim 7.1.1 — of bath to the dynamical-partition reading of Obstruction B. The Born exponent's standing form on entry is **ℓ^2 ⟸ PC ∧ Obstruction B**, with Obstruction B's affirmative resolution *adopted*, not forced.

2.3 One declared elevation: commitment as refinement-selection

We make explicit one statement strongly implied by the programme but not previously elevated, and we are exact about its strength.

Declaration (commitment = refinement-selection). Before commitment there is an unresolved region M with capacity Vol_op(M). The *admissible refinements* of M are the ways that region may be carved into resolution-stable sub-region capacities $\text{Vol\_op}(M) = \Sigma_i \text{Vol\_op}(M_i)$. A commitment event selects one admissible refinement, freezing it into the realized outcome structure.

This is *new but natural*: it names what commitment does to the region — it selects a carving — in the *Packing* paper's refinement vocabulary. It is held **strictly weaker** than both axioms, and the boundary is the paper's central anti-smuggling guard:

Anti-smuggling (the Declaration asserts neither axiom). The Declaration says commitment selects *a* refinement from the admissible menu $\mathscr{R}$. It says **nothing** about (i) how the members of $\mathscr{R}$ are individuated — by M's own capacity structure (internal, IA) or by an external label set / pre-fixed list (external); nor (ii) whether the members, however individuated, are sealed or connected under reversible motion (RC). IA addresses (i); RC addresses (ii). The Declaration is therefore compatible with the emergent ledger and the bath alike, and the results of §7 are **not** true by definition. A reader should verify that nothing in the Declaration prefers internal to external, nor connected to sealed.

2.4 The two axioms (stated; status established later)

We state the two admissibility conditions here so the whole argument can be read as discharging one and isolating the other.

IA (Internality Axiom). Every admissible refinement of M is internally individuated by M's distinguishability-capacity structure: not by an external label set, and not by a pre-fixed menu individuated independently of M's capacity relations.

RC (Reversible Connectedness). Any two admissible (internally individuated) refinements of M are joined by a continuous path of admissible capacity-preserving reversible transformations of the substrate.

Their statuses, established below: **IA (internality) is proven** — FP1 ⟹ No External Individuation (§7.2), and since capacity is the only thing FP1 gives M, ¬external ⟹ internal is an exhaustive dichotomy, so internality itself follows; only the *uniqueness* of the internal individuation is not claimed, and is not needed. **B ⟺ RC given IA** (§7.3, proven). **RC is open**, with the seal classification (§7.5, proven) showing the only data-level barrier is the adjoined (external) one — every other seal is dynamical (Type 1 = RC), one continuity question on the refinement axis, declined here because FP1 is silent on dynamics.

---

3. Three readings, not two

The previous paper's fork was stated as bath vs. ledger. Resolving it requires seeing that "ledger" was two readings wearing one name, distinguished by *when* the per-outcome accounts come into being.

Definition 3.1 (Primitive ledger, PL). There exists, *before* any commitment, an outcome index set $\Omega = \{1, \ldots, d\}$ and per-outcome capacities $(C_1, \ldots, C_d)$ with each C_i separately conserved by pre-factual reversible dynamics; the total $\Sigma_i C_i$ is fixed; commitment *reveals* which i obtains. The accounts precede commitment. *This is the classical die*: Ω is the face-set, C_i the face-weights, commitment the landing.

Definition 3.2 (Emergent ledger, EL). No pre-commitment index set exists. A commitment selects a refinement (Declaration §2.3), generating a decomposition $\mathrm{Vol\_op}(M) = \Sigma_i \mathrm{Vol\_op}(M_i)$ and with it the indices i; thereafter the sub-capacities $\mathrm{Vol\_op}(M_i)$ **freeze** and are separately conserved. The accounts are *born at* the carving and rigid from then. No reallocation across refinements.

Definition 3.3 (Bath, B). No pre-commitment index set exists. Capacity is borne by the unresolved Vol_op(M) and is **reallocable across candidate refinements** until commitment freezes one. The carving is alive: dividing lines move, pooling and redistributing capacity, until commitment localizes a single refinement. Dynamical refinement.

Two facts about this trichotomy, used below:

- PL is distinguished from EL/B by *pre-commitment individuation*: PL has it, EL and B do not. NPI (§4) targets this axis.
- EL is distinguished from B by *frozen vs. flowing refinement*: EL freezes the carving, B keeps it alive. This axis is Obstruction B's reverse inclusion (§6); RC (§7) targets it.

The previous paper's "ledger" was PL and EL conflated. The die — the most vivid and threatening ledger — is PL. Separating them is what lets the substrate disprove the die without yet deciding the bath.

---

4. The No-Pre-Individuation theorem

4.1 The theorem and its defence against a covert index set

Theorem 4.1 (No Pre-Individuation — proven, given FP1 as stated). FP1's primitive datum is the scalar capacity Vol_op(M) of an unresolved region M. It does not supply a pre-commitment outcome index set. Consequently no pre-commitment per-outcome capacity assignment ($C_1, \ldots, C_d$) exists *as a datum of the primitive*: the indices i over which such an assignment would range are not present in FP1.

Proof. FP1 constrains a single scalar functional $Vol\_op(\cdot)$ of a region together with the resolution scale Δ_op and dimension d_op. The bound $|\Sigma(M)| \leq Vol\_op(M) / \Delta\_op^{d\_op}$ references M as an undivided region; the right-hand side is a number computed from M's capacity, the resolution, and the dimension. Nowhere does the axiom range over, quantify, or presuppose a set of labelled sub-outcomes $\{1, \ldots, d\}$: there is no symbol in FP1 denoting an outcome index, and no term decomposing Vol_op(M) into pre-assigned summands. A per-outcome assignment ($C_1, \ldots, C_d$) presupposes the index set $\{1, \ldots, d\}$ as its domain; that domain is absent. Hence ($C_1, \ldots, C_d$) is not extractable from FP1 without adjoining the index set as additional structure.

A hostile referee will object that FP1 *does* carry index-like structure, at one of three places: the dimension d_op, the cardinality of Σ(M), or the elements of Σ(M) themselves. We close all three, and then close the subtler microstructure variant.

Proposition 4.2 (d_op is a capacity dimension, not an index set — proven). The operational dimension d_op does not constitute a pre-commitment outcome index set, and in particular does not supply the Ω required by PL.

Proof. d_op enters FP1 as an exponent on the resolution scale Δ_op, fixing how capacity scales with resolution — it is the *dimension* of the region's capacity, a count of independent directions of distinguishability, not a labelled enumeration of realized outcomes. A dimension is a cardinal attached to capacity; an index set is a labelled domain over which per-outcome data is defined. The distinction is the standard one between the dimension of a vector space and a distinguished basis: knowing dim V = n does not single out n labelled vectors, and a choice of which n directions realize a decomposition is exactly a refinement, performed at commitment (Declaration §2.3). Thus d_op bounds *how much* and *in how many independent directions* capacity may be carved, while leaving *which* carving — which labelled outcomes — to refinement. It is partition-free.

Proposition 4.3 (Σ(M) is a resolution-dependent cardinality, not a labelled outcome space — proven, given FP1's form). The distinguishability set Σ(M) appearing in FP1 does not constitute a commitment-independent labelled outcome space, and so supplies no pre-commitment index set Ω.

Proof. There are two readings of Σ(M): *(A, capacity-counting)* Σ(M) is the set of distinguishable states M *can support*, realized only relative to a chosen refinement, with |Σ(M)| a measure of capacity at a given resolution; *(B, pre-labelled)* Σ(M) is a distinguished, labelled outcome space attached to M prior to and independently of any refinement — a sample space Ω.

FP1's own form forces (A). The axiom bounds the *cardinality* |Σ(M)| by $Vol\_op(M) / \Delta\_op^{d\_op}$, an explicit function of the resolution: decreasing Δ_op increases the bound by the factor $(\Delta\_op^{old} / \Delta\_op^{new})^{d\_op}$, so the number of distinguishable states M supports *grows as one resolves more finely*. A fixed, pre-labelled outcome space has no such dependence — a sample space does not acquire new elements when one looks more sharply. Hence a resolution-dependent Σ(M) cannot be a commitment-independent labelled outcome space without contradicting FP1's stated dependence on Δ_op. Reading (B) is inconsistent with the axiom; (A) holds. Under (A), Σ(M) carries a refinement-relative cardinality but no canonical labelled enumeration prior to a carving: "how many distinguishable states at resolution Δ_op" is well-defined, "which specific labelled states" is not, until a refinement is selected. So Σ(M) supplies no pre-commitment index set.

Theorem 4.4 (Coarse-Graining Non-Primitivity — proven, given FP1's form). FP1 supplies no commitment-independent coarse-graining of Σ(M) into a fixed number d of outcome classes. Precisely, any candidate coarse-graining Σ(M) → Ω into a fixed class-count d fails to be primitive on one of three exhaustive grounds:

- *(external)* it adjoins a class-structure individuated independently of M's capacity data — the (M, Ω) adjunction, structure FP1 does not contain (Theorem 4.1);
- *(internal-but-a-refinement)* it is definable from M's capacity data, in which case it is a *refinement* in the programme's sense, hence by the Declaration (§2.3) selected at commitment — not commitment-independent;
- *(fixed-count)* it imposes a commitment-independent class-count d which is *not derivable from M's capacity data* — neither |Σ(M)| (a resolution-relative cardinality, Proposition 4.3) nor d_op (a capacity dimension, not a label-count, Proposition 4.2); FP1 contains no such number, so a fixed d must be adjoined. (We do not claim a fixed d *contradicts* resolution-dependence — fixed macrostates over a growing microstate set is the consistent statistical-mechanics picture; the point is solely that d is not FP1-derivable and so is extra structure.)

Proof. A coarse-graining is either individuated independently of M's capacity data or from it. If independently, it is external structure absent from FP1 (Theorem 4.1) — the first ground. If from M's data, it is a partition of the distinguishable content into resolution-stable classes, which is precisely a refinement; by the Declaration such a refinement is the object commitment selects, so it is not commitment-independent — the second ground. Independently of which, a coarse-graining into a *fixed* commitment-independent number d of classes posits a count d that is not derivable from M's capacity data: FP1's only outcome-relevant numbers are |Σ(M)|, a resolution-relative cardinality (Proposition 4.3), and d_op, a capacity dimension rather than a count of realized classes (Proposition 4.2); neither is a fixed class-count, so d must be adjoined — the third ground. (This is an adjunction point, not a consistency point: a fixed d over growing |Σ(M)| is internally consistent, as stat-mech shows; what fails is its *derivability* from FP1.) Hence no commitment-independent fixed-count coarse-graining is primitive.

Remark 4.4.0 (Scope: groupings exist, none is primitive). The theorem does *not* deny that coarse-grainings of Σ(M) exist — they do, and they are exactly the refinements the Declaration speaks of. It denies only that any is *primitive*, i.e. commitment-independent and fixed in class-count. And it bears on the *individuation/grouping* question alone: it says nothing about connectedness, so a sealed internal refinement family remains perfectly consistent with it (the topology axis is §7.3's business, not this theorem's).

Remark 4.3.1 (The microstructure rebuttal — an application of Theorem 4.4). A subtler referee concedes that $|\Sigma(M)|$ grows with resolution and yet tries to keep Ω: "let $\Omega = \{1, \ldots, d\}$ be a *fixed* outcome space, and let finer resolution merely reveal more microstructure *within* each fixed outcome; $\Sigma(M)$ is then the growing set of microstates, while the d labels stay fixed." This is the classical statistical-mechanics picture and the last place Ω can hide. It requires exactly a commitment-independent coarse-graining $\Sigma(M) \to \Omega$ into a fixed class-count d — which the Coarse-Graining Non-Primitivity Theorem (4.4) excludes: such a grouping is either external (adjoined data) or a refinement (commitment-dependent), and in either case its fixed class-count d is not derivable from M's capacity (neither $|\Sigma(M)|$ nor d_op), so it must be adjoined. (We do not claim a fixed d *contradicts* resolution-dependence — fixed macrostates over growing microstates is the consistent stat-mech picture; the point is solely non-derivability, exactly as in Theorem 4.4.) We do not deny microstructure may exist; we deny that FP1 bundles it into a fixed labelled outcome space. The bundling *is* the carving, and the carving is commitment's job. So the microstructure move does not preserve Ω against NPI; it is the very thing NPI excludes, relabelled.

Remark 4.3.2 (What survives without a refinement). The *bound* on $|\Sigma(M)|$ holds for M as given, with no refinement chosen — the capacity to distinguish is a pre-commitment fact. But the *determinate elements* of $\Sigma(M)$, and any grouping of them into outcomes, exist only relative to a refinement. Capacity precedes commitment; individuated states, and their bundling into outcomes, are generated by it.

Corollary 4.5 (The primitive ledger is disproved — proven, given FP1). PL is not a reading of FP1: it requires the pair (M, Ω) where FP1 supplies only M, and the Ω it requires is excluded at every candidate source — outcomes (Theorem 4.1), the dimension d_op (Proposition 4.2), the distinguishability set $\Sigma(M)$ (Proposition 4.3), and a fixed coarse-graining of $\Sigma(M)$ (Theorem 4.4). Reading $\Sigma(M)$ itself as Ω contradicts FP1's resolution-dependence (Proposition 4.3); the fixed coarse-graining is excluded instead by non-derivability (Theorem 4.4), not by a resolution-dependence claim.

Proof. By Definition 3.1, PL requires pre-commitment per-outcome capacities $(C_1, \ldots, C\_d)$ over an index set Ω. Theorem 4.1, Propositions 4.2–4.3, and Theorem 4.4 jointly show FP1 supplies no Ω at any candidate source, and that reading $\Sigma(M)$ as Ω is inconsistent with the axiom's Δ_op-dependence (Proposition 4.3), while the fixed coarse-graining is excluded by non-derivability (Theorem 4.4). So PL's conserved data is absent from the primitive; to instantiate PL one must adjoin (M, Ω), strictly more than — and in its Σ-forms inconsistent with — FP1's M. PL is therefore disproved as a reading of the primitive.

4.2 The die is correctly not a counterexample

Remark 4.4.1 (The die is PL, and PL is not FP1). Classical probability is the standing witness that a "shared possibility space with commitment" need not have mixing — and it is genuine: the die is fully reversible pre-landing, jointly available, committed at the landing, and rigid. NPI does *not* deny the die exists; it denies the die is a reading of *FP1*. The die carries $\Omega = \{1, \ldots, 6\}$ as **external, commitment-independent input** with a *fixed* cardinality that does not grow under finer inspection — precisely the (M, Ω) structure Corollary 4.5 identifies as extra and, in its fixed resolution-independent cardinality, the structure Proposition 4.3 shows FP1 cannot carry. So the die is PL, and PL is not a reading of the primitive: the die is a coherent possibility space, but it is *not* what FP1 describes, because FP1 has no pre-printed faces and its "how many faces" count is resolution-relative. This is the exact sense in which the substrate "has no ledger to keep" — there are no pre-commitment accounts, because there is no pre-commitment, resolution-independent label set to index them.

---

5. The burden flip

Corollary 4.5 inverts the dialectical situation the previous paper was in.

Before: bath was *adopted*; the bare axiom *permitted* the rigid ledger; the burden lay on the affirmative side to exclude the classical alternative, and it could not, so it flagged bath as a reading.

After: the rigid alternative the bare axiom was thought to permit is **PL**, and FP1 does **not** permit it — PL requires $(M, \Omega) \notin$ FP1, and its Σ-forms contradict the axiom's resolution-dependence. The bath reading B, by contrast, uses only the primitive. So **B adds nothing; PL smuggles in Ω.** The burden has flipped: it is now the *ledger* side that must justify adjoining an external index structure FP1 never supplied — and cannot, for the primitive form.

This is a genuine advance in standing, stated at exactly its strength: it disproves PL, not EL. It disposes of the *primitive* ledger — the die and everything like it — and leaves standing only the emergent ledger EL, which respects NPI and adjoins no external outcome set. EL is the real survivor, and §6 shows what it is while §7 disposes of its externally-supplied refuge via IA.

---

6. The collapse: bath-vs-ledger is Obstruction B, exhaustively

With PL disproved, the trichotomy of §3 reduces to the dichotomy B vs. EL — and that dichotomy is *exactly* Obstruction B, with no residual third branch.

Lemma 6.1 (Exhaustion — proven). Given (i) no pre-commitment index set (NPI, §4) and (ii) commitment generates indices through refinement (Declaration §2.3), any FP1-admissible pre-factual conservation structure is either EL or B; no third option exists.

Proof. By (i) the conserved data carries no pre-commitment per-index component, so any conservation acts on capacity individuated only via a refinement (ii). Fix an admissible refinement with generated sub-capacities $\{Vol\_op(M_i)\}$. A conservation structure either holds each Vol_op(M_i) invariant under all admissible reversible maps — *freezing*, EL (Definition 3.2) — or admits some admissible reversible map altering some Vol_op(M_k) at fixed total — *reallocation*, B (Definition 3.3). For each generated sub-capacity, "invariant under all admissible maps" and "moved by some admissible map" are complementary and exhaustive. PL is excluded (Corollary 4.5), so no structure with pre-commitment per-index conservation remains. Hence {EL, B} is exhaustive.

Theorem 6.2 (The fork is Obstruction B — proven). Under NPI and Lemma 6.1, the surviving readings are exactly B and EL, and:

B ⇔ dynamical refinement ⇔ Obstruction B's reverse inclusion holds, EL ⇔ frozen refinement ⇔ Obstruction B's reverse inclusion fails.

Hence bath-vs-ledger contains no branch outside Obstruction B: the fork *is* Obstruction B.

Proof. Exhaustiveness of {EL, B} is Lemma 6.1. The *Packing* paper's Obstruction B is the reverse inclusion that *every admissible decomposition is a resolution-stable refinement* of the unresolved region — admissible decompositions are not independently fixed but are exactly the (movable) refinements of a shared capacity. EL asserts the decomposition is fixed at first carving and rigid thereafter — a frozen list, the negation of the reverse inclusion. B asserts the decomposition is the live refinement of the shared Vol_op(M), reallocable until commitment — the reverse inclusion affirmed. So B ⇔ Obstruction B affirmative and EL ⇔ Obstruction B negated.

Theorem 6.3 (Compression of the Born residue — proven, independent of §7). The Born exponent's residual substrate dependence is

$\ell^2 \Longleftarrow \mathrm{PC} \wedge$ Obstruction B,

with the bath reading *not* a separate condition but the affirmative reading of Obstruction B, and the primitive-ledger escape route eliminated.

Proof. By the inherited reduction (§2.2), $\ell^2 \Longleftarrow \mathrm{PC} \wedge$ (mixing) ∧ (connectivity), and mixing ⇔ B. By Theorem 6.2, B ⇔ Obstruction B affirmative, EL (= Obstruction B negated) the only alternative, PL disproved. Connectivity is delivered by a single shared bath (Proposition 6.1 of *Bath or Ledger*), the affirmative reading of Obstruction B applied to one unresolved region (single-bath ⇔ the refinement of *one* region is *one* live carving; §10.3). Substituting, $\ell^2 \Longleftarrow \mathrm{PC} \wedge$ Obstruction B. The compression invokes no premise about *which* way Obstruction B resolves; it holds whether B or EL obtains.

This is the paper's first earned result, and it stands whether or not §7 succeeds: the Born residue is **PC ∧ Obstruction B**, the bath subsumed into Obstruction B, the die-like primitive ledger disproved. *Bath or Ledger* reached "PC ∧ Obstruction B" only modulo the live possibility that the ledger was an independent third thing; NPI plus the Exhaustion Lemma closes that, making the two-condition compression exhaustive rather than provisional.

7. IA derived, RC isolated, the Seal-Trichotomy

The compression of §6 leaves Obstruction B's *resolution* — B vs. EL — open. Forcing B means excluding EL. EL is the emergent ledger: no pre-commitment outcome indices, but a menu $\mathscr{R}$ of admissible refinements from which commitment selects, the selected carving frozen. EL can shelter on two independent axes, which must not be conflated:

- **Individuation:** are the members of $\mathscr{R}$ pinned down by M's own capacity structure (internal — IA), or by an external label set / pre-fixed list (external)?
- **Topology:** are the members of $\mathscr{R}$, however individuated, sealed or connected under admissible reversible motion (RC)?

A menu can be internal yet sealed; in principle external yet connected. The individuation axis is closed by NPI (§7.2: NPI $\Longrightarrow$ IA); the topology axis is the genuine residue (§7.3: B $\Leftrightarrow$ RC given IA). Keeping them apart is what makes the first a theorem and the second an honestly-named premise.

7.1 The forbidden thing, stated precisely

Earlier drafts used the word "retrieval," which did too much: it suggested a *statically existing* menu is forbidden, which is false and not what FP1 excludes. A menu may be generated once from M and then exist statically; selecting from such a static menu is not by itself the importing of an external index set. We forbid not *static existence* but *external individuation*:

External individuation: the members of $\mathscr{R}$ are individuated independently of the live capacity relations of M — by a label set, enumeration, or pre-fixed list whose identity does not derive from Vol_op(M), Δ_op, d_op.

Internal individuation: the members of $\mathscr{R}$ are individuated *by* the distinguishability-capacity structure of M — each admissible refinement is a carving *of* Vol_op(M), defined from M's own capacity data and nothing else.

IA is the assertion that all admissible refinements are internally individuated. This cut lets §7.2 prove what is provable without claiming what is not.

7.2 NPI $\Longrightarrow$ IA (the individuation axis — proven)

Lemma 7.1 (No External Individuation — proven; IA/internality follows). Given FP1 and NPI, no admissible refinement of M is externally individuated: none is fixed by a label set or pre-fixed menu individuated independently of M's capacity relations. What is *proven* is this exclusion, and **IA (internality) follows from it**: capacity is the only thing FP1 gives M, so once external individuation is excluded, internal individuation is the only source left — ¬external ⟹ internal is exhaustive. What is *not* claimed is that the internal individuation is unique; the downstream results use internality alone, never uniqueness.

Proof. FP1 supplies only M, its capacity Vol_op(M), the resolution Δ_op, and the dimension d_op. A refinement admissible under FP1 is one definable from these data; anything else adjoins structure FP1 does not contain. A menu $\mathscr{R}$ whose members are individuated by an external label set $\mathbb{I}$ — an enumeration of "available refinements" fixed independently of Vol_op(M) — adjoins exactly such structure: it is (M, $\mathbb{I}$) where FP1 gives M, in precise parallel to PL's (M, Ω). This is the adjunction argument alone, and we do not lean on a resolution-independence claim beside it — an external rule ("at each Δ_op, take the finest admissible cells") can produce resolution-*dependent* membership while being fixed from outside, so "external ⟹ resolution-independent" is false in general. The correct cut is cleaner: any resolution-dependence *deriving from M's capacity* is internal by definition, so the only genuinely external individuation is fixed-from-outside, excluded by the adjunction (extra-data) argument. External menu individuation is thus the Coarse-Graining Non-Primitivity Theorem applied one level up — grouping *carvings* into a fixed externally-fixed list — and dies by adjunction, not by a resolution-dependence gloss. So external individuation is inadmissible at every level, by the same argument that disproved PL (Corollary 4.5). The only menu compatible with FP1's bare datum is one whose members are carvings *of* Vol_op(M), individuated from M's own capacity structure — i.e. internally individuated. Hence IA.

Remark 7.1.1 (What IA does and does not establish). The proven content is No External Individuation — NPI applied at the menu level: external individuation is the primitive ledger's index set re-entering one level up, as a pre-fixed list of carvings, and the partition-free, resolution-dependent reading of FP1 excludes it there too. IA (internality) follows by the exhaustive ¬external ⟹ internal dichotomy; only uniqueness is unclaimed. The externally-retrieved emergent ledger is therefore disproved, alongside the primitive ledger. But IA is held strictly to the *individuation* axis. It establishes that the menu's members derive their identity from M's capacity; it establishes **nothing** about whether that internally individuated menu is sealed or connected. In particular, IA does **not** exclude a menu that is internally individuated yet a discrete, sealed family — for instance, the finitely many maximal resolution-stable partitions definable from M at scale Δ_op, each holding its sub-capacities fixed. Such a menu is internal (members are carvings of Vol_op(M), definable from M's data) and yet sealed (no reversible capacity-flow between them). It satisfies IA and is handed to RC. So IA closes external individuation without pretending to close sealing — the latter is the topology axis, governed by §7.3. This is the anti-overclaim guard, mirror to the §2.3 Declaration guard: just as the Declaration must not smuggle IA, IA must not smuggle RC.

7.3 B ⟺ RC given IA (the topology axis — open)

With external individuation excluded (IA, Lemma 7.1), EL survives only as an *internally individuated yet sealed* refinement family: carvings genuinely cut from M's capacity, but discrete and sealed, with no reversible capacity-flow between them. To exclude this last refuge we require an *operational* connectedness — physical accessibility under the admissible reversible group, not mere abstract path-connectedness of a refinement space (a connected space can harbour dynamically inaccessible regions; the premise must be about accessibility, not geometry). This is RC.

RC (Reversible Connectedness, operational). Any two admissible (internally individuated) refinements of M are joined by a continuous path of admissible *capacity-preserving reversible transformations* of the substrate — the connection realised by physical reversible motion along which capacity flows across the cuts. Equivalently: no internally individuated candidate carving seals, because capacity is demonstrably carried across any cut by admissible reversible dynamics.

Theorem 7.2 (B ⇔ RC, given IA — proven). With external individuation excluded (Lemma 7.1), the bath reading holds if and only if RC holds.

Proof. By IA the admissible menu is internally individuated, so the only surviving EL refuge is internal-but-sealed. (⇐) If RC holds, candidate carvings do not seal: admissible capacity-preserving reversible transformations carry capacity across the cuts along paths joining any two refinements. So no admissible refinement freezes its sub-capacities pre-commitment; every decomposition is a live carving of the shared pool, the internal-but-sealed refuge falls, and B obtains (Definition 3.3). (⇒) If B holds, capacity is reallocable across candidate refinements by admissible reversible dynamics; the reallocating maps furnish the continuous paths of capacity-preserving reversible transformations between refinements, which is RC. Hence B ⇔ RC.

We do not assume RC. RC concerns the *dynamical accessibility of refinements under the admissible reversible group* — whether physical reversible motion carries capacity across cuts. FP1 constrains the capacity of a region; it does not constrain the reversible accessibility of its carvings, and NPI's partition-free, resolution-dependent reading — which delivered IA — says nothing about whether the *internally individuated* carvings flow into one another. Asserting RC would be the affirmative fiat: "the carvings flow because a genuine substrate's carvings flow," the symmetric image of the word-fiats the programme has refused. We decline it. RC is the residue, in its sharpest form.

Status. With IA proven, the bath-forcing question reduces, with no remaining slack, to RC: a single operational premise about a single object — *is the internally individuated family of admissible carvings connected under capacity-preserving reversible motion, or are its candidates sealed?*

7.4 The surviving classical reading, named precisely

The narrowness of the survivor is the measure of the advance.

PL is disproved: no pre-commitment outcomes (§4). The externally-retrieved emergent ledger is disproved: IA (§7.2). What survives as the sole classical (non-ℓ^2) reading is neither — it is a substrate whose refinements are *all genuinely internally individuated*, cut from M's own distinguishability-capacity structure, yet which form a *discrete, sealed* family across which no reversible capacity-flow exists. It is the **IA-satisfying yet RC-failing** object. Commitment selects one such sealed internal carving; nothing reallocates.

This is far stranger and more constrained than "emergent ledger" suggested. It cannot import its menu from outside (IA); its carvings must be M's own; and it differs from the bath *solely* in that its internal carvings are sealed rather than flowing. The entire bath-vs-classical question now lives on that single property — sealed vs. flowing — of an otherwise fully internal, fully FP1-native structure. RC is precisely the assertion that the property is *flowing*. §7.5 sharpens this survivor further: a seal's enforcer is data or dynamics, the only data-level seal (the adjoined external charge) is excluded, and a native separating function cannot seal because capacity cannot supply conservation — so only a purely dynamical seal remains.

Remark 7.4.1 (RC is the continuity upgrade, refinement-axis — and IA is not). The two-axiom framing makes the decisive structural observation transparent. IA and RC are parallel admissibility conditions on the same object (the admissible-refinement family), but they constrain *different features*: IA constrains *individuation* (the source of the carvings), RC constrains *connectedness* (the reversible accessibility among them). The phase-axis continuity upgrade the companions carry (temporal-extensibility / holonomy-continuity, *Squaring* §11.3) constrains the *connectedness* of admissible phase motion. So **RC and the phase upgrade are the same species of condition — connectedness of admissible reversible motion — applied on two axes (refinement, phase); IA is a different species (individuation).** This is the structural reason RC, and not IA, is the candidate for unification with the phase upgrade. A single substrate principle asserting continuity of *all* admissible reversible motion would discharge RC and the phase-axis upgrade together. Since RC now carries the *entire* remaining bath-vs-classical discrimination (§7.4), this potential identity is the single most consequential open claim across the four-paper arc — conjectural, but decisive if it holds.

7.5 The cost of sealing: the seal-enforcer is data or dynamics, and data cannot enforce it

RC was isolated in §7.3 as the condition on which the bath turns, and §7.4 named the survivor it leaves. The natural next move is to *force* RC by forbidding the sealing barrier as adjoined structure — the discipline that disproved the die. We must not make that move carelessly, and the reason is a fact this paper itself states (§7.3): **FP1 is silent on dynamics.** The data-axis discipline (NPI, IA, Theorem 4.4) forbids adjoining *primitive data* not in FP1; it earned the burden-flip of §5 because there the question was what the axiom *contains*. RC is a question about *dynamics* — whether reversible motion connects the carvings — and on the dynamics axis FP1 entails neither flow nor seal. So a genuinely symmetric no-adjoined-structure discipline cannot forbid the seal while permitting the flow: relative to a dynamics-silent axiom, *both* are adjoined, and forbidding only the seal would be the bath preference in a discipline's clothing — the very fiat the programme refuses. We therefore do **not** assert a "no extra barriers" axiom, and we do **not** claim RC is forced. We do something narrower and honest: classify a seal by its *enforcer*, show the only data-level enforcer is the adjoined (external) charge — because the seal-making property is conservation, and capacity cannot supply conservation — and concede the irreducibly dynamical seal as RC.

Lemma 7.5.1 (Component-Separation — proven, given C_ref). If internally individuated refinement space is disconnected and admissible reversible motion is continuous (C_ref, the refinement-axis instance of the phase-axis continuity upgrade), then no continuous admissible path crosses between components: component membership is invariant under the identity component of the admissible reversible group, hence a quantity conserved under all continuous admissible motion.

Proof. A continuous admissible path is the image of a connected parameter interval under a continuous map; the continuous image of a connected set is connected; a connected subset of a disconnected space lies in one component. So the identity component of the admissible group maps each refinement-component into itself, and any function constant on components (distinct values on distinct components) is invariant under all continuously-reachable admissible motion.

This says only that *if* there is a seal under continuous dynamics, *then* there is a separating invariant — it converts a topological seal into a conserved charge. It does not, by itself, exclude anything: whether that invariant can be *enforced* is the question, and the enforcer is what we now classify.

Theorem 7.5.2 (Seal classification — proven). Any seal of internally individuated refinement space is enforced, in its seal-making (conservation) content, by either *data* or *dynamics*:

- **Type 3 (external charge):** the seal is enforced by a component-label *not* derivable from M's capacity (Vol_op(M), Δ_op, d_op) — adjoined data.
- **Type 1 (dynamical):** the seal is enforced by a restriction on the admissible reversible maps — the carvings are all M-definable and (in the bath ontology) entertained pre-commitment, but the admissible dynamics does not connect them.

A third descriptive case — **Type 2 (native charge):** the *separating function* is derivable from M's capacity — is not a third enforcer. Because the seal-making property is conservation, and conservation is dynamical (Lemma 7.5.2a), a native-charge seal's enforcer is dynamical: Type 2 collapses into Type 1, its separating function merely happening to be M-derived. So the genuine partition is binary — data (Type 3) or dynamics (Type 1) — and we treat the two enforcers in turn.

Type 3 is excluded, dynamics-free. An external component-charge is data adjoined to FP1 — structure individuated independently of M's capacity, the (M, Ω) adjunction at the refinement-component level. It is excluded by exactly the adjunction argument that disproves PL (Corollary 4.5) and external menus (Lemma 7.1): FP1 supplies only M's capacity; a component-label not derivable from it is extra primitive data the axiom does not contain. No dynamics premise is used; this is a clean data-axis exclusion. (for Type 3)

Type 2 collapses into Type 1: a native charge can separate, but cannot seal. The exclusion of the native-charge seal does *not* run through the claim that M-derived separating functions fail to exist — they exist, and conceding this is what makes the argument airtight. Take any candidate carving r ∈ $\mathscr{R}$ and form v(r) = f({Vol_op(M_i(r))}), a functional of r's own sub-capacities. This varies across $\mathscr{R}$ and is derived from M's capacity (the carvings are M-internal, by IA). Nor may we call v "post-commitment": the bath ontology itself entertains candidate carvings *before* commitment (Definition 3.3 has capacity reallocable across candidate refinements pre-commitment), so candidate-space is available pre-commitment, and v is a perfectly good pre-commitment, M-derived, refinement-*varying* function. So a separating function can be native. The error would be to think that settles anything.

Lemma 7.5.2a (No Native Seal — proven). A native (M-derived) function on refinement-space may separate candidate carvings, but cannot by itself *seal* them: the seal-enforcing property is conservation under the admissible dynamics, and FP1 entails the conservation of no function whatsoever.

Proof. A separating function v partitions $\mathscr{R}$ into its level sets, but the components are *sealed* only if the admissible reversible dynamics preserves that partition — i.e. only if v is conserved under every admissible map (this is exactly what Lemma 7.5.1 extracts from a continuous seal: the separating invariant must be dynamically preserved). Conservation is a dynamical predicate: whether a given function is invariant under the admissible maps is a fact about *which maps are admissible*, not about M's capacity. FP1 is silent on dynamics (the keystone of §7.3). Therefore FP1 cannot entail that v — or any M-derived function — is conserved. So the *separating* content of a native charge is real but inert; its *sealing* content (conservation) is supplied not by capacity but only by a restriction on the admissible dynamics. A native-charge seal is therefore, in its operative (seal-enforcing) content, a dynamical seal — Type 1 — wearing native-function clothing.

This collapses the apparent trichotomy into an honest **binary**. A seal's enforcer is either **data** — an external component-charge adjoined to FP1 (Type 3, excluded by the adjunction argument, dynamics-free) — or **dynamics** — a restriction on the admissible reversible maps (Type 1 = RC, open). The "native charge" (Type 2) is not a third box: because conservation is never capacity-forced, a native-charge seal's seal-enforcing content is dynamical, so it *is* a Type-1 seal whose separating function happens to be M-derived. The result is sharper than "two of three excluded": **there is no data-level barrier other than the adjoined one, because data cannot enforce conservation.** And it is symmetric — it does not smuggle the bath, because the bath's free reallocation is equally a dynamical fact FP1 does not force, which is exactly why RC stays open. Capacity forces neither seal nor flow; both are dynamical; the residue is dynamical.

Remark 7.5.2b (why the conservation fulcrum, not the post-commitment one). An earlier version excluded the native charge by calling the sub-capacity functional "post-commitment." That is not clean: the bath's own ontology entertains candidate sub-capacities pre-commitment, so the EL/seal objector may borrow the same candidate-space and define the separating function on it pointwise, with no commitment. The conservation fulcrum avoids the wedge entirely — it concedes the pre-commitment separating function and denies only that capacity can make it *conserved*. It also removes any dependence on M being structureless: even a richly graded M yields separating functions, never data-level conservation, so the argument does not lean on a property of the *Packing* construction (§10.7). It is the paper's own thesis — FP1 silent on dynamics — applied once more, which is why it reads as inevitable rather than imported.

Type 1 is the residue, and is not excludable from FP1. A dynamical seal adjoins no data: the carvings are M-definable and (in the bath ontology) pre-commitment-available, with at most a native separating function — the dynamics simply fails to connect them. Excluding it would require asserting that admissible reversible motion *does* connect them, i.e. asserting flow — which, against a dynamics-silent axiom, is as adjoined as the seal. **We do not assert it.** Type 1 is precisely Reversible Connectedness (RC): whether the internally individuated carvings are dynamically connected. It is the same species as the phase-axis continuity upgrade (both ask whether substrate reversibility is continuous, here on the refinement axis), and it is **open**.

Honest status, and the cost. The only *data-level* seal is the external charge (Type 3), excluded dynamics-free by the adjunction argument; the native charge does not add a second data-level seal, because data cannot supply the conservation a seal needs (Lemma 7.5.2a) — its seal-making content is dynamical, Type 1. So the sole surviving classical (non-ℓ^2) reading is the Type-1 dynamical seal: an internally individuated family (with, at most, a native separating function) that the admissible reversible dynamics contingently fails to connect. RC is exactly the assertion that this does not occur. We record a **lean**, and flag it as a lean, not a proof: a substrate continuous in phase motion but discontinuous in refinement motion would have its two reversible axes disagree on continuity, which would itself want justification — so Type 1 is *unnatural* relative to the continuity the programme already requires on the phase axis. But because FP1 is silent on dynamics, asserting flow is as uncosted-against as asserting seal, so this is a lean and not an exclusion. **RC remains open**; what §7.5 establishes is that the only data-level barrier is the adjoined one — there is no native data-level seal, because capacity cannot enforce conservation — so any RC-failure is a restriction on the dynamics, one continuity question on one axis rather than a family of barriers.

Remark 7.5.3 (Why this is the honest form). The earlier impulse — a "No Extra Barriers" axiom forbidding the seal as adjoined structure — was the bath preference made asymmetric: it forbade the barrier while permitting the flow, on a dynamics axis where FP1 licenses neither. The seal classification replaces that asymmetric premise with a symmetric one: the only seal FP1 excludes is the one that adjoins *data* (the external charge, Type 3, dynamics-free, by the adjunction argument); every other seal — including one with a native separating function — has its seal-making content in the *dynamics*, which FP1 does not legislate, and is conceded as Type 1 = RC. The bath is forced iff Type 1 is excluded, and Type 1 is the continuity question we decline to prejudge. The symmetry is exact: capacity forces neither seal nor flow, both are dynamical, so the residue is dynamical — which is why RC stays open and the bath is not smuggled. This keeps the genuine content — the only data-level barrier is the adjoined one, and data cannot enforce conservation — without claiming the closure FP1's dynamical silence forbids.

---

8. Why this is not the fiat

The paper makes claims of very different strength, and their separation is the discipline.

What is proven (no fiat): PL is disproved by FP1 (NPI, the resolution-relativity of Σ(M), and the Coarse-Graining Non-Primitivity Theorem, §4); IA — external individuation of the refinement menu disproved (§7.2); the residue collapses exhaustively to PC ∧ Obstruction B (Lemma 6.1, Theorem 6.3). None of these asserts the bath. NPI is a textual-structural fact about the primitive — capacity on a region, a resolution-dependent cardinality, no index set and no fixed coarse-graining — defended across all candidate sources of a covert Ω. IA is the same fact at the menu level. The die is handled honestly: identified as PL and shown not to be a reading of FP1.

What is held to its proper scope (no overclaim): IA is held to *individuation*, not *topology*. It excludes the externally-supplied menu; it does **not** exclude the internally individuated yet sealed menu, and we say so explicitly (Remark 7.1.1, §7.4). Pretending IA closed sealing would be the symmetric overreach to the fiat — claiming the topology axis was settled when only the individuation axis was. We refuse it.

What is conditional (fiat refused, premise named): the bath is forced only under RC (§7.3), which we do not assume. We resisted the temptation to *force* RC by forbidding the sealing barrier as adjoined structure — because FP1 is silent on dynamics, and on a dynamics-silent axiom forbidding the seal while permitting the flow is the bath preference made asymmetric, not a discipline (§7.5). Instead the seal classification (Theorem 7.5.2) sorts seals by their *enforcer*. The only data-level enforcer is the external charge (Type 3, excluded dynamics-free by adjunction). A native separating function (the apparent Type 2) does not add a data-level seal: the seal-making property is conservation, conservation is dynamical, and FP1 entails the conservation of no function (Lemma 7.5.2a), so a native-charge seal is dynamical — Type 1 — with an M-derived separating function. The surviving seal is therefore Type 1, purely dynamical, which is precisely RC; excluding it would require asserting flow, as adjoined against a dynamics-silent axiom as the seal. **We do not assert it. RC remains open.** This mirrors the linearity companion (which stopped at PC) and *Bath or Ledger* (which stopped at the open fork): the affirmative resolution is reached only by an open condition, flagged as open — here narrowed to a single continuity question on one axis, the only data-level barrier (the external charge) eliminated and the native charge shown unable to seal.

The asymmetry is the point. The *negative* results (PL disproved, IA, the exclusion of data-level seals) are forced by the primitive — they are about what the axiom *contains*. The *positive* result (bath) is reached only via RC, which is about *dynamics*, on which FP1 is silent. We do not let the strength of the former leak into the latter — that leakage ("NPI, therefore bath," or "no-extra-barriers, therefore flow") is the trap, the symmetric image of the word-fiats the programme refuses. What §7.5 adds is not closure but *structure*: the classical survivor is narrowed to the single Type-1 dynamical seal, its data-level forms (Types 2–3) eliminated, and a lean — explicitly not a proof — recorded against the survivor's requiring the refinement axis to be discontinuous where the phase axis is continuous.

---

9. Assembly: the consolidated Born residue

Across the four companion papers the Born exponent now stands as follows:

$\ell^2 \Leftarrow$ PC ∧ Obstruction B. (Theorem 6.3 — proven, exhaustive; all inherited reductions modulo Claim 7.1.1, see Scope)

with:

1. **PC** — the mild decomposition-independence principle (linearity companion); the one genuinely new bookkeeping commitment across all companion papers. *To be stated and checked.*
2. **Obstruction B** — the refinement obstruction, now wearing *five* hats: ODG's refinement structure (*Packing*); the linearity of reversible transport (linearity companion); the bath reading and hence mixing (*Bath or Ledger*); the dynamical-partition reading of the packing primitive (this paper, §6); and the disproof of the primitive ledger as a separate branch, so that Obstruction B's two readings — B (dynamical) and EL (frozen) — are now *exhaustive*, the die-like PL removed (§4–6).
3. **Obstruction B affirmative ⟺ RC** — its affirmative resolution is equivalent (via IA, §7) to Reversible Connectedness: the internally individuated admissible refinements are connected under capacity-preserving reversible motion (§7.3). RC is **open**. §7.5 does not force it but narrows it: the seal classification (Theorem 7.5.2) shows the only *data-level* seal is the adjoined external charge (Type 3, excluded dynamics-free), while a native separating function cannot seal because the seal-making property is conservation and FP1 supplies the conservation of no function (Lemma 7.5.2a). So every surviving seal is dynamical — Type 1 = RC — and the residue beyond PC is a single continuity question on the refinement axis, the same species as the phase-axis upgrade the programme already owes; FP1's silence on dynamics is exactly why it is not closed here.

So the proven headline is IA together with the exhaustive compression **PC ∧ Obstruction B**, the primitive ledger and the externally-retrieved emergent ledger — between them, every externally-individuated classical reading — disproved from the primitive. The frontier is RC, and it is **open**: given IA, the square follows from PC together with RC. What §7.5 adds is not closure but a sharper map of RC: the only data-level seal is the adjoined external charge (Type 3, excluded dynamics-free), and a native separating function cannot seal (conservation is dynamical, and FP1 supplies it for no function — Lemma 7.5.2a), so every surviving seal is dynamical — Type 1 = RC — and RC reduces to a single continuity question on the refinement axis, declined here because FP1 is silent on dynamics. We prove IA; we narrow RC to one continuity question; we do not assert it.

Remark 9.1 (What changed, precisely). *Bath or Ledger* delivered PC ∧ Obstruction B *modulo* the live possibility that the ledger was a third, independent condition and the lingering force of the die. This paper removes both: PL is disproved (NPI + Coarse-Graining Non-Primitivity, Theorem 4.4), the ledger collapses into EL = ¬Obstruction B (Exhaustion Lemma), external individuation is disproved (IA), and the compression becomes exhaustive. It does not close

Obstruction B, and it does not force RC — to do so would assert flow on a dynamics-silent axiom, the fiat the programme refuses. What it does is narrow the surviving classical reading by the seal classification (Theorem 7.5.2): the externally-charged seal (Type 3) is excluded by the data-axis adjunction argument, and the natively-charged one cannot seal — conservation is dynamical and FP1 supplies it for no function (Lemma 7.5.2a) — so it collapses into the dynamical case, leaving only the Type-1 dynamical seal, which is RC. The survivor is no longer "any sealed internal family" but specifically "a family the admissible reversible dynamics contingently fails to connect," against which we record a lean (it would split continuity across the two reversible axes) but not a proof.

---

10. Limitations and Open Problems

10.1 RC — open, narrowed to the Type-1 dynamical seal

The bath is forced iff RC holds (Theorem 7.2), and RC is **open**. We do not force it: forcing it would assert that admissible reversible motion connects the carvings (flow), which against a dynamics-silent axiom is as adjoined as the seal it would deny (§7.5). What §7.5 establishes instead is a narrowing. The seal classification (Theorem 7.5.2) sorts any RC-failure by its enforcer. The only *data-level* enforcer is the external component-charge (Type 3), excluded dynamics-free by the adjunction argument. A native separating function (the apparent Type 2) cannot enforce a seal: the seal-making property is conservation, conservation is dynamical, and FP1 entails the conservation of no function (Lemma 7.5.2a) — so its seal-making content is dynamical, Type 1. The sole survivor is therefore the Type-1 *dynamical* seal — carvings the admissible reversible dynamics contingently fails to connect, possibly carrying a native separating function — which is RC itself. So the controlling open object is one continuity question on the refinement axis, the same species as the phase-axis upgrade (§10.6). We record a weak prior toward continuity, and a lean against Type 1 (it would split continuity across the two reversible axes, §7.5), but decline to let either force RC.

10.2 The Declaration's strength (§2.3)

We elevated "commitment = refinement-selection" as new-but-natural, held strictly weaker than both axioms. A referee should verify the Declaration is genuinely compatible with internal and external individuation, and with sealed and connected menus (it is, by construction), so §7's results are not true by definition. If the Declaration were strengthened to assert internality or connectedness, §7 would be circular; we have kept it weak precisely to avoid this.

10.3 Single bath vs. sub-baths, on the refinement reading

Bath or Ledger §9.4 flagged single-bath vs. sub-baths (global isotropic ℓ^2 vs. block-diagonal sectors). On the refinement reading: is one unresolved region's refinement one connected internal family (single bath, RC on all admissible refinements) or several mutually sealed internal subfamilies (sub-baths, RC failing across blocks)? Folded into RC; the block-structured failure mode corresponds to super-selection sectors and should be tracked as a refinement of RC, not a separate condition.

10.4 FP1 as read

NPI and IA rest on FP1's primitive being the scalar $Vol_{op}(M)$, partition-free, with d_{op} a dimension (Proposition 4.2), $\Sigma(M)$ a resolution-relative cardinality (Proposition 4.3), and no commitment-independent fixed-count coarse-graining of $\Sigma(M)$ into outcome classes (Theorem 4.4). Two distinct facts about FP1 carry the load, and we keep them separate. (i) The *resolution-relativity* of $|\Sigma(M)|$: $|\Sigma(M)|$ grows as Δ_{op} shrinks, which is what makes $\Sigma(M)$ capacity-counting rather than a fixed sample space (Proposition 4.3). (ii) The *non-derivability* of a fixed class-count from M's capacity, which is what excludes fixed coarse-grainings and external menus — by the adjunction (extra-data) argument, *not* by a resolution-dependence claim (a fixed count over growing $|\Sigma(M)|$ is internally consistent, as stat-mech shows; what fails is its derivability from FP1). A referee should confirm against the *Packing* paper that Δ_{op} is the minimal resolution entering the bound as $|\Sigma(M)| \leq Vol_{op}(M)/\Delta_{op}^{d_{op}}$, so halving Δ_{op} multiplies the bound by $2^{d_{op}}$. Fact (i) is the inheritance behind Proposition 4.3; fact (ii) is what the disproof of PL and of external individuation lean on; the exclusion of a *native-charge* seal does not lean on it at all but on a separate, dynamical point (a native separating function cannot be conserved by capacity, Lemma 7.5.2a, §7.5/§10.7).

10.5 IA and the formal topologisation of refinement space

Lemma 7.1 (IA) is **proven given NPI**: external individuation re-imports an excluded structure. What remains formally open is the precise topologisation of the *internally individuated* admissible-refinement family under capacity-preserving reversible motion, needed to state RC with full rigor and to prove Theorem 7.2's equivalence as a lemma about that dynamical structure rather than a structural identification. We judge the identification correct and have given its content (capacity-reallocation as the reversible paths between carvings); the formal construction of internal refinement space with its admissible-reversible-group action is the outstanding technical task. The two-axis split (individuation / topology) localises the residue cleanly: individuation closed (IA/internality proven via No External Individuation, uniqueness unneeded); topology open (RC), with §7.5 showing the only data-level seal is the adjoined one and every other seal dynamical.

10.6 Continuity-upgrade unification

Remark 7.4.1's observation — that RC and the phase-axis continuity upgrade are the same *species* of condition, connectedness of admissible reversible motion — places the surviving Type-1 residue (§7.5, §10.1) on the same axis-family as a debt the programme already carries on the phase axis. We do not claim they are the *same* debt: that would require a single substrate principle establishing continuity of *all* admissible reversible motion, phase and refinement alike, which the programme has not proven (its continuity upgrade has been used most clearly for phase / holonomy motion). The unification — one continuity principle, all axes — is the natural capstone: if it held, the Born exponent would rest, beyond PC, on no premise the programme did not already owe. **Conjectural.** What is established here is only the kinship (same species, different axis), and the narrowing of RC's failure to the Type-1 dynamical seal that this kinship would address.

10.7 The native-grading residual — a dynamical question outside FP1

The exclusion of the native-charge seal (§7.5, Lemma 7.5.2a) does *not* depend on M being structureless, and we do not claim any "No-Native-Charge" question is closed by a property of the *Packing* construction. The argument is purely that conservation is dynamical and FP1 is silent on dynamics: even a richly graded M yields only *separating functions* on refinement-space, never the *conservation* that would seal. So a native grading on M — a $\mathbb{Z}_7$ closure structure, say — yields by itself only a Type-1 dynamical seal.

What would make such a grading actually seal refinement space is a *conservation law on the grading*, and that is a dynamical fact about the substrate's axioms *beyond* FP1. If the programme's further axioms impose a conserved refinement-charge on the closure structure, then that is an adjoined dynamical conservation law, RC genuinely fails in the graded sectors, and one has superselection sectors (§10.3) — block-structured rather than global ℓ^2. If they do not, the grading is a Type-1 separating structure with no seal, and RC is unaffected. So the honest statement is: **whether the closure architecture seals refinement space turns on whether it carries a *conserved* refinement-charge — i.e. on the dynamically-conserved-or-not status of the closure structure — which is a question about dynamical axioms beyond FP1.** It is a real open question the programme can adjudicate elsewhere, and it is cleanly *outside this paper*: this paper shows only that FP1 itself supplies no such conservation, so within FP1 the native grading is Type-1, not a data-level seal. We flag this rather than claim closure, because a reader who knows the $\mathbb{Z}_7$ result is entitled to ask exactly this question — and the answer lives in the substrate's dynamical axioms, not in FP1's capacity.

11. Conclusion

The mixing question, chased to the substrate's founding rule, was in danger of being settled by a fast argument: pre-commitment everything is reversible, capacity belongs to the unresolved whole, therefore the conserved quantity is one shared pool. We have shown why that argument proves too quickly — reversibility is shared by the ledger, and the die is the standing witness — and replaced it with structural results that earn most of the conclusion outright, organised around two admissibility conditions: the Internality Axiom (IA) and Reversible Connectedness (RC).

FP1 attaches distinguishability capacity to an unresolved *region*: its primitive is the scalar Vol_op(M), with d_op a capacity dimension and Σ(M) a resolution-dependent cardinality, not a labelled outcome index set — and it supplies no fixed coarse-graining of Σ(M) into outcomes, since any such grouping is a refinement, selected at commitment. The primitive ledger — the classical die — requires an external index set Ω that FP1 never supplies, and that in its strongest forms contradicts FP1's dependence on resolution: the founding rule's count of distinguishable states *grows when one resolves more finely*, and a fixed list of named outcomes does not. So the primitive ledger is **disproved** as a reading of the primitive; the die, though a coherent possibility space, is simply not what the founding rule describes.

With PL disproved, an Exhaustion Lemma completes the dichotomy: any admissible pre-factual conservation structure either freezes the generated sub-capacities (emergent ledger, frozen refinement) or reallocates among them (bath, dynamical refinement), with no third option. Bath-vs-ledger collapses, exhaustively, onto Obstruction B itself, and the Born exponent stands on **PC ∧ Obstruction B** — exhaustive now, not provisional.

We then forced the individuation axis and isolated the topology axis. IA is *derived*: FP1 supplies only the region and its capacity, so any admissible carving must be individuated from within that capacity, and a menu handed in from outside is the primitive ledger's index set re-entering one level up — NPI ⟹ IA, external individuation disproved alongside the die. We were careful not to let this prove more than it does: IA fixes the *source* of the carvings, not their *topology*, and a menu genuinely cut from M's capacity may still be sealed. That second axis — connectedness — falls only under RC, which we name and decline to assume. Given IA, the bath is forced if and only if RC holds.

So the survivor is named precisely, and §7.5 makes it narrower still — but by classification, not by fiat. We resisted forcing RC through a "no extra barriers" axiom, because FP1 is silent on dynamics: on a dynamics-silent axiom, forbidding the seal while permitting the flow is the bath preference in a discipline's clothing, the very move the programme refuses. Instead the seal classification sorts any RC-failure by its *enforcer*. A seal enforced by an external component-charge is adjoined data, excluded dynamics-free by the same argument that disproved the die. A seal whose separating function is *native* — built from M's own capacity — does not escape this: a separating function can indeed be native, but a seal requires that function to be *conserved* by the admissible dynamics, and conservation is a dynamical fact FP1 never supplies (it is silent on dynamics). So a native-charge seal's seal-making content is dynamical; it is not a second, data-level barrier but a dynamical seal wearing native-function clothing. What survives, then, is a single thing: a seal that adjoins no data at all — the carvings are all M's own, perhaps carrying a native separating function — where the admissible reversible dynamics simply, contingently, fails to connect them. That is RC, and FP1 cannot adjudicate it: capacity forces neither seal nor flow. We record a lean against it — a substrate continuous in phase motion but discontinuous in refinement motion would have its two reversible axes disagree, which would itself want justification — but a lean is not a proof, and we do not promote it to one.

The role of this paper, stated plainly. It is not to prove the bath, and it does not. What it accomplishes is the elimination of every classical alternative *except one narrow type*: the primitive ledger (the die) is disproved (§4); the externally-individuated menu is disproved (§7.2); and of the seals that could sustain an emergent ledger, the two that adjoin data — an external component-charge and a native one — are excluded (§7.5) — the external charge by the data-axis adjunction argument, the native one because a separating function, however M-derived, cannot *seal* without a conservation law, and conservation is dynamical, never supplied by capacity (Lemma 7.5.2a). The sole surviving non-quantum reading is the internally individuated family sealed by *dynamics alone* — carvings the admissible reversible motion contingently fails to connect — which FP1 cannot adjudicate because it is silent on dynamics. Every classical route that turns on adjoined *data* has been closed from the primitive; what remains is a single question about the *dynamics* of refinement, on which the axiom is silent and which we decline to settle by preference. That elimination, not a proof of the bath, is the result — and it is already a significant narrowing of where the Born exponent can fail to hold.

The squaring is no longer waiting on a beam-splitter, nor on a choice between bath and ledger, nor even on whether the menu of possibilities is handed in from outside. What this paper *proves* is IA: the die-like alternative is disproved, the externally-supplied menu is disproved, and reality's carvings are internally individuated from its own capacity. What it does *not* prove, and does not assume, is the bath — and it is careful not to, because the bath is a claim about dynamics and FP1 is silent on dynamics. Given IA, the bath, and with it the squared norm, follows if and only if RC: whether reality's own internally generated carvings are connected by reversible motion or sealed. The seal classification narrows that question as far as the primitive allows — the only data-level seal, the adjoined external charge, is eliminated, and a native separating function cannot seal because capacity cannot supply conservation — but the surviving seal is purely dynamical, and there the axiom has nothing to say. So the honest residue, beyond the mild PC, is one continuity question on one axis: is the refinement dynamics, like the phase dynamics the programme already studies, continuous? We prove IA; we narrow RC to that single question; we decline to answer it, because answering it by preference would be the fiat. That continuity question — recognisably the refinement-axis kin of one the programme already faces on the phase axis, and not assumed — is now, with the mild PC, the whole of what stands between finite distinguishability and the Born rule.