The Shared Phase Obligation — Why the Dyadic-Loading and Born-Rule Programmes Converge

Binary Refinement and Step Individuation: Dyadic Loading Reduced, by Equivariant Refinement and No-Pre-Individuation, to the Single Phase-Forcing Obligation It Shares with the Born Rule

Keith Taylor

VERSF Theoretical Physics Programme

General Reader Summary

The puzzle

Across several parts of the VERSF programme the same sequence keeps appearing — 1, 2, 4, 8, … *— powers of two. The flavour programme attaches the loads 1, 2, 4 to the three particle generations; the closure-capacity argument needs exactly those loads to make "three generations" come out of a fixed budget. So the powers of two do real work, and the obvious question is: why two? Why does each step of "refinement" double the count rather than tripling it or growing some other way?*

This paper chases that question to the bottom, in two movements, and the end-state is honest about its own shape: not a proof that the count doubles out of nothing, but a *conditional theorem* — the doubling is forced once one names the conditions, and the last of those conditions turns out to be something the framework already expects to hold, with the burden falling on anyone who would deny it.

Movement I — the reduction. There is a tempting argument that refinement *must* double. Build everything from the Fold — the smallest possible distinction, which has exactly two sides. Then refinement can't grow slower than doubling without failing to distinguish anything new, and can't grow faster without secretly doing more than one distinction at once. Pinched from both sides, the branching factor would have to be two. But look closely and *both* pincers rely on the same unspoken assumption: that each step is *one* clean minimal distinction — the same move, once, every time. Call that **uniformity**. So Movement I does not prove doubling; it proves something sharper and more useful — that the entire puzzle collapses onto this one assumption. Granted uniformity, doubling is immediate, because the Fold has two sides. The whole weight rests on uniformity.

Movement II — trying to remove the assumption. Could one *prove* that refinement must go one Fold at a time, so that uniformity stops being an assumption? Movement II tries. Part of it works outright: refinement can't *idle* — a step that distinguishes nothing isn't a step, because a step is an act of committing a new fact, and committing nothing is no act at all. Another part looked dangerous — that refinement might *bundle* several distinctions into one uncuttable lump, which seemed to clash with the programme's insistence that nature contains genuinely joint facts (entanglement). It turns out not to clash: a joint fact can be cut by ordinary two-way cuts while staying joint. The real difficulty lies deeper. The number that is supposed to double is the size of a *menu of possibilities* that exists before anything is decided, and that menu is carved out of a smooth, featureless "capacity" with no built-in pieces and no built-in schedule. To get doubling, the carving has to be even-handed across the whole menu at each stage — and the worry was that "even-handed across the menu at once" smuggles in a notion of *simultaneity*, hence of time, in a programme where time itself is supposed to be built out of the very acts of deciding. That would be a vicious circle.

The resolution — favouritism, not timing. The circle is broken by a symmetry argument, and the key is to see that the worry was about the wrong thing. Suppose two possibilities are genuine *twins* — nothing admissible tells them apart. A rule that refined one twin and not the other would have to be reading some label distinguishing them. But by assumption there is no such label. So the rule must treat twins alike: refine both or neither. This needs no notion of "at the same instant" at all — it is simply that the rule cannot play favourites between things it cannot tell apart. The supposed timing problem was never a timing problem; it was a fairness problem, and fairness is settled by the absence of labels, not by a clock.

The same move disposes of the last apparent loophole — that one twin might simply be *abandoned*, refined forever while the other never is. To sustain that, the rule would have to *stand* in a permanently different relation to the two: always act on this one, never on that one. But a standing difference in how the rule treats two indistinguishable things is *itself* a way of telling them apart — the very label that, by assumption, isn't there. So abandonment is not some extra danger needing a separate law to forbid it; it is the same forbidden favouritism, merely held steady over time. Where could a real difference between twins come from? From a genuine feature — but then they were never twins. From the capacity supply — but that is a bare quantity that doesn't know one possibility from another. From nothing — in which case it cannot happen.

What is left — and who carries the burden. That exhausts the possibilities but one, and pinning it down is the last turn. Two would-be twins might differ not in any feature of their own, but in how they *overlap everything else* on the menu. A difference in overlap is itself a real, admissible difference — so if two sectors differ that way they were never true twins, and the rule may treat them differently, and the doubling is not forced. One cannot wriggle out by simply *defining* twins to be alike in overlap too: the rule only ever has to be fair between sectors *it* cannot tell apart, so what matters is not how we choose to draw the word "twin" but whether the rule, in deciding what to refine, already takes overlap into account. So the result is a *conditional theorem*: given that the Fold has two sides, that the rule plays no favourites among indistinguishables, and that the overlap structure is built from the very same notion of distinguishability the rule already reads — rather than added as a separate ingredient — the powers of two are forced.

The decisive point is that this last condition is *not a hopeful guess but the framework's default*, and seeing why shifts the burden. The programme already holds, as a basic principle, that the world carries no structure beyond what distinguishability supplies — no hidden labels, no extra ingredients bolted on. An overlap that secretly told twins apart would be exactly such an extra ingredient, and so is forbidden — unless it is, after all, part of what distinguishes things, in which case the rule sees it too and there is no secret. Either way there is no hidden difference. So the natural reading is not "maybe overlap conceals something" but the reverse: overlap *should* be a derived shadow of the same distinguishability that already drives refinement, and anyone who would deny the doubling owes a concrete account — *what* is the extra something the overlap carries, and *why* can the rule not see it? Until such a thing is named, the presumption sits with the powers of two. What honestly remains is not to hope the condition holds but to *audit the construction* — to confirm that the overlap structure, as actually built, has not quietly smuggled some convention in and dressed it up as a real feature. The expectation, and the burden of proof, both rest on the side of the doubling.

There is one striking thing about where that audit ends. The overlap, as the theory actually builds it, comes down to a single ingredient — a kind of *phase* — and the whole question becomes whether that phase is fixed by distinguishability or floats free of it. That turns out to be the *very same* question another part of the programme must answer for an entirely different reason: the rule that says probabilities go as the square of an amplitude (the famous $P = |\psi|^2$) needs that same phase to be pinned down in the same way. So two of the deepest "why" questions — *why do the particle generations come in powers of two?* and *why is probability the square of an amplitude?* — turn out to be one question seen from two sides. Answer the one about the phase, and both close together.

Abstract

Several branches of the VERSF programme generate dyadic structure: generation depth carries loads 1, 2, 4, and the closure-capacity argument of the companion species paper requires that growth law to derive a finite generation count. Dyadic growth is at present an assumption. This paper reduces it to a single hypothesis, attempts to discharge that hypothesis, and reaches a terminus that is conditional in one sharply-located and ordinary way.

Movement I (the reduction). The natural argument for binary branching is a two-sided squeeze: *sub-binary* refinement fails distinguishability preservation, *super-binary* refinement violates minimality, so the branching factor is forced to 2. The central finding is that **both** bounds depend on the *same* premise — that each refinement step is exactly one minimal-distinction (Fold) operation, applied uniformly at every depth. Uniformity is therefore not a residue but the sole load-bearing hypothesis. This reduces dyadic loading without residue to the **Uniform Refinement Hypothesis** (URH), with closure conditions **no-vacuous-steps** ∧ **no-multi-Fold-steps**, plus a third condition (frontier-uniformity) that surfaces in Movement II.

Movement II (the attempt to discharge it). The **Step Individuation Theorem** (SIT) — that the Fold is the uniquely admissible refinement unit — would dissolve the URH. Pursued rigorously it splits. **No-vacuous-steps is [Proven]**, conditional on irreversible commitment and the emergence of order from commitment: a step committing nothing is incoherent, not merely forbidden. **No-multi-Fold-steps does not collide with non-separability**: that conflates decomposition into *local* parts (blocked by entanglement) with decomposition into *binary* cuts (all the condition needs); the Bell basis decomposes into joint commuting binary commitments (ZZ then XX) while remaining non-locally-decomposable, the enabling fact being that committed records are mutually exclusive (orthogonal) by irreversibility. Admissibility closure therefore does **not** force a non-decomposing joint primitive. The residue lies at the **candidate menu**: the doubling count N_{n+1}/N_n counts the pre-commitment menu (commitment *selects*, factor → 1), characterised by the corpus as internally individuated carvings of a scalar, additive, partition-free, resolution-dependent capacity Vol_op(M) with **no refinement-space dynamics**. Hence a Capacity Quantisation Theorem is likely refutable by resolution-sweep and would yield only *integer* loading; *dyadic* loading needs a **complete binary tree** — completeness + binarity + frontier-uniformity — of which only binarity is supplied. Frontier-uniformity appears order-laden (simultaneity across the frontier), threatening a circularity with emergent order.

The resolution and the terminus. That circularity is **defeated** by an **isomorphism-equivariance** argument: if live sectors at a stage are mutually isomorphic (no admissible invariant separates them), a refinement rule blind to inadmissible labels cannot refine one while leaving an isomorphic partner unrefined — selective refinement would require a symmetry-breaking selector excluded by no-pre-individuation. This forces **all-or-none across alike sectors** from *label-blindness, not simultaneity*, so the order-loop of frontier-uniformity is removed rather than relocated. Read on the *rule* rather than the limit, the same principle disposes of the one apparent residue: permanent starvation of an isomorphic sector is a *standing disposition-asymmetry* — "refine A, never B" as a policy of the rule — which is itself an equivariance violation, present at the first asymmetric disposition and not deferred to a limit. **No-starvation is therefore a theorem, not a posit**, and needs no capacity-conservation law (which is fortunate, as FP1 supplies none). A trichotomy closes it: a standing asymmetry must be grounded in geometry (then the sectors were never isomorphic), in capacity (but capacity is a sibling-blind additive scalar and cannot ground it), or in nothing (then starvation is impossible). The only place an admissible sibling-distinguishing datum could hide is the **candidate-layer support functional**, and the apparent "hidden overlap" loophole reduces to a clean biconditional. Since the equivariance σ acts on the whole candidate-layer structure including s, overlap cannot break symmetry *within* an isomorphism class (that would contradict σ being an isomorphism); but counting s as admissible structure merely *relocates* the content into the definition of isomorphism, and enriching that definition is no free lunch — shrinking the isomorphism classes weakens the all-or-none constraint by exactly as much as it forbids hidden overlap (a conservation of loopholes). The substantive question is therefore whether **refinement-siblinghood** (alike in the data the rule reads) and **support-isomorphism** (alike including s) *coincide*. The result is biconditional: **P2 closes if and only if the support profile supervenes on the refinement-relevant data** — iff s adds no admissible distinguishing power beyond what the refinement rule already reads. The entire dyadic result thus reduces to one checkable fact about how the corpus *constructs* s: does it introduce admissible distinctions the refinement dynamics does not already see? This is captured as a **Conditional Support-Closure Theorem**: if the refinement rule F and the support functional s are both functions of the one admissible distinguishability datum D, with no support degrees of freedom beyond D, then F-siblings are support-isomorphic trivially, supervenience holds, no-starvation follows by equivariance, and **Fold-binarity + equivariant refinement + common distinguishability origin of F and s ⟹ $N_n = 2^{n-1}$**. This does not prove dyadic loading absolutely; it proves it *relative to a no-extra-support-structure condition*, with a named refuter — dyadic loading fails iff the construction of s introduces a support distinction between D-identical sectors. Crucially the close is not bought by *defining* isomorphism to include s (free, and settling nothing, since equivariance binds only on pairs the rule itself treats alike); the substantive content is whether the rule's decision data already determine s — the definition is free, the coincidence of F's data with s's data is the theorem. The condition is not merely *plausibly* met: by no-pre-individuation (§9, inherited), the substrate admits no structure beyond admissible distinguishability data, so an independent support degree of freedom is prohibited unless itself admissible — in which case it is already part of D and supervenience holds anyway. Supervenience is therefore the programme's *default*, and the burden lies with any challenger. The formal audit of the constructed s can moreover be carried out, since the corpus builds s from one object — the geometric phase θ of the Double Square kernel. The audit's verdict: θ is *not* D-determined by inspection (the metric is symmetric; θ is oriented, carrying a holonomy the metric does not contain), so supervenience is not free; it holds iff θ is forced from D. What remains genuinely open is therefore one identified proposition, **shared with the Born-rule derivation**: that *finite distinguishability forces continuous U(1) phase* (the Layer-2 obligation of *Physical Necessity of Quantum Probability Structure*). The shared content is the *forcing*: "D forces continuous U(1)" gives the Born chain the continuity that makes probability quadratic *and* the generation chain the D-determination that makes s supervene on D — neither forced-discrete nor continuous-but-free would serve both. So "why powers of two?" and "why $P = |\psi|^2$?" are one open question wearing two hats; dyadic loading closes if and only if that one proposition closes, which the companion reports established for finite-holonomy impossibility and deferred for the continuity-and-U(1) step. The programme's minimum-distinguishability scale is no objection to that continuity: finite distinguishability limits phase *observability*, not the phase *catalogue*, and the phase is accumulated transport holonomy rather than a local distinguishability coordinate — so the open proposition is properly a global-consistency claim (does holonomy-consistency under unbounded composition force U(1)?), which is also the form that delivers continuity *forced from D*, as both chains require.

*Epistemic markers: **[Inherited]**, **[Proven]** (relative to stated inputs), **[Conditional]**, **[Conjectural]**, **[Open]**.*

Table of Contents

Movement I — The Reduction to Uniformity

1. Introduction

The same sequence recurs across the programme. The flavour programme associates generation depth with the loads 1, 2, 4. The closure-capacity argument of the companion species paper (*Particle Species as Representation Classes of Persistent Fold Defects*, §9.1) requires exactly that dyadic growth law to derive a finite generation count: three sectors of loads $1 + 2 + 4 = 7$ saturate a closure register of capacity $K = 7$, a fourth of load 8 overflows it. The dyadic law is therefore load-bearing — and, in that paper, explicitly **open**: it is read off the three-term operator $\mathcal{D} = \mathrm{diag}(1, 2, 4)$, which presupposes three generations, rather than derived for arbitrary depth.

This paper isolates that open premise and asks whether it can be forced. Whenever a structure undergoes repeated refinement, one may ask how many distinguishable sectors emerge per step — the *branching factor* b. The candidates are $b = 1$ (no growth), $b = 2$ (binary), $b = 3$ (ternary), …, and the question is whether $b = 2$ is *forced* by what refinement is, or merely observed.

The investigation runs in two movements. Movement I (§§1–8) shows there is a clean squeeze argument that appears to force $b = 2$ from two sides, but that the squeeze rests entirely on a single assumption about how refinement steps are individuated — *uniformity* — so that the real result is a reduction of dyadic loading to that one hinge, with its closure conditions stated precisely. Movement II (§§9–17) takes up the candidate theorem that would discharge the hinge, the Step Individuation Theorem, follows it down through the substrate primitives to an apparent circularity, and then removes that circularity by a symmetry argument, leaving a single ordinary capacity condition. We do not prove dyadic loading; we reduce it, without residue, to that condition.

2. The Fold as Minimal Binary Distinction

The primitive operation of the VERSF substrate is the Fold. A Fold is a minimal admissible distinction: it separates one possibility from another, and — being minimal — it separates into exactly two, A or ¬A. A distinction with three or more outcomes is not minimal; it already contains more than one independent binary distinction.

Minimal-Distinction Principle

One Fold contributes exactly one independent binary distinction: it has exactly two outcomes, A and ¬A.

Status: [Inherited], *with one clause to attribute carefully.* That the Fold is the *minimal* admissible distinction is inherited from the fold architecture. That a minimal distinction has *exactly two* outcomes — its binarity specifically, not merely its minimality — is the clause the whole paper leans on. If the fold architecture establishes the Fold as two-valued (a single cut, two sides), the binarity is [Inherited] and §2 stands as written. If the architecture establishes only minimality and non-degeneracy (more than one outcome) without fixing the count at two, then "exactly two" is an assumption — a weaker cousin of the very dyadic question under investigation: if a minimal distinction could be three-valued, the Fold would be ternary and the branching factor would be 3, not 2. So the binarity of the Fold is the seed of the result, and the paper's standing is exactly the standing of that one inherited clause. We proceed taking the Fold as binary and mark the dependency rather than bury it. This is the dependency labelled (B) throughout.

3. Refinement and the Branching Factor

Let N_n denote the number of distinguishable sectors after refinement to depth n. A refinement rule specifies how N_n evolves; the *branching factor* at a step is the ratio N_{n+1}/N_n. Dyadic loading is the claim that this ratio is 2 at every step:

$N_{n+1} = 2\,N_n$, $N_n = 2^{n-1} = 1, 2, 4, 8, 16, \dots$

The two following sections give the squeeze — the lower bound ($b \geq 2$) and the upper bound ($b \leq 2$) — and §6 shows that both rest on the same hypothesis. (The precise meaning of "N_n after refinement to depth n" returns in §§15–17; it is left at its intuitive reading until then.)

4. The Lower Bound and Its Dependence on Uniformity

The intuitive argument. A refinement step that branches by fewer than two produces fewer than two distinguishable descendants from a previously distinguishable sector — it introduces no genuinely new distinction. A step that introduces a real distinction must separate at least two alternatives, so a *genuine* refinement step has $b \geq 2$.

Where it leaks. This forbids a step that *claims to refine but branches below two*; it does **not** forbid $b = 1$ outright, because a step that introduces *no* new distinction does not violate distinguishability *preservation* — it preserves every existing distinction and adds nothing. The $b = 1$ vacuous step is excluded only if one rules out vacuous steps, i.e. only if every step is *required* to be a genuine refinement. That requirement is not a distinguishability fact; it is a statement about what counts as a step. A non-uniform schedule $b = 2, 2, 1, 2, 2, 1, \dots$ never violates distinguishability at any *genuine* step, yet averages below 2 and does not produce $N_n = 2^{n-1}$. So the lower bound controls the *average* branching only if the schedule contains no vacuous steps — a uniformity condition, not a consequence of distinguishability. The lower bound is **conditional on uniformity** (no vacuous steps).

5. The Upper Bound and Its Dependence on Uniformity

The intuitive argument. A step branching faster than binary — say $b = 3$ — produces three distinguishable descendants. Three outcomes encode more than one binary distinction ($\log_2 3 > 1$), so a ternary step carries more than one Fold's worth of distinction. If refinement is minimal, a step contributing more than one Fold violates minimality, so $b \leq 2$.

Where it leaks. Read the Minimal-Distinction Principle (§2) exactly: *one Fold* contributes *one* binary distinction. It does **not** say *one step* contributes *one Fold*. Minimality is a property of Folds — a Fold cannot be sub-divided — not of refinement steps. A step branching by 4 is cleanly two Folds and violates nothing about Folds; a 3-fold branch is one Fold plus a partial one, inadmissible for the separate reason that the Fold is binary (§2) and a partial Fold is not a substrate object — so the exclusion of *ternary* branching is the work of the Fold's binarity (B), not of the minimality bound. To forbid integer-multi-Fold steps ($b = 4, 8, \ldots$) one needs a separate premise: **one refinement step = exactly one Fold**. Without it, super-binary steps are admissible as multi-Fold steps and the upper bound fails. That premise is the same content as uniformity. So the upper bound, like the lower, is **conditional on uniformity**; its work is to forbid composite integer-multi-Fold steps, while the exclusion of non-powers-of-two is carried by (B).

6. The Consolidation: Uniformity Is the Sole Hinge

Examined honestly, **both bounds depend on the same premise**, the uniformity premise:

- the **lower** bound (§4) controls the branching only if there are *no vacuous steps* — otherwise an alternating schedule averages below 2 while never failing distinguishability;
- the **upper** bound (§5) holds only if *one step = one Fold* — otherwise a super-binary step is an admissible multi-Fold step, violating nothing about minimal distinctions.

"No vacuous steps" and "one step = one Fold" are two faces of one condition: **each refinement step is exactly one genuine minimal-distinction operation** — not zero, not more than one, the same at every depth. That is uniformity.

Consolidation

Uniformity is not a residual assumption left over after the squeeze; it is the *sole* load-bearing hypothesis. Given uniformity, both bounds hold and binary branching is forced. Absent uniformity, *neither* bound holds: the lower leaks through vacuous steps, the upper through multi-Fold steps.

The squeeze adds no independent content beyond the binarity of the Fold and the uniformity of steps; given both, $b = 2$ is immediate. The work is entirely in establishing that refinement *is* uniform in this sense.

Binary Refinement Result (conditional on uniformity and on the Fold's binarity)

If (U) each admissible refinement step is exactly one genuine minimal-distinction operation, applied uniformly at every depth, and (B) the minimal distinction is binary (§2), then refinement branches binarily: $N_{n+1} = 2\,N_n$, hence $N_n = 2^{n-1}$.

Status: [Conditional], on (U) ([Open], §§7–17) and (B) ([Inherited], §2, with its caveat).

7. The Uniform Refinement Hypothesis

Uniform Refinement Hypothesis (URH)

Every admissible refinement step implements exactly one genuine minimal-distinction operation — one Fold — and the same operation at every depth: each step contributes exactly one binary distinction, neither zero (no vacuous steps) nor more than one (no multi-Fold steps), uniformly across the live frontier.

Status: [Open].

The URH is not arbitrary; it is what "refinement *by minimal distinction*" most naturally means. There are two distinct questions, with different answers. *Is the URH logically necessary?* No — one can coherently describe schemes that violate it. *Is it physically expected, given the substrate?* Strongly yes: refinement is *denominated in Folds*, and to let some steps carry two Folds and others none is to measure refinement in a unit that varies step to step — a variable-length metre, a denial that the unit is a unit. So the honest status is **physically expected but not logically necessary, and open as a derived result**.

8. The Closure Conditions, and the Theorem That Would Discharge Them

The URH closes if and only if the refinement dynamics establish **step individuation**: that an admissible refinement event is exactly one Fold-commitment, applied uniformly. Three sub-claims, each with a named refuter:

1. **No multi-Fold steps** (upper bound). No single event commits two or more independent Folds without decomposing into successive single-Fold steps. *Refuter:* a genuinely simultaneous, non-decomposing multi-Fold commitment.
2. **No vacuous steps** (lower bound). An admissible refinement event necessarily commits a *new* distinction. *Refuter:* an idle step.
3. **Frontier-uniformity** (the doubling itself). When refinement advances at a depth, *every* live sector is carved, not some. *Refuter:* a selective (depth-first) schedule refining one sector while leaving others. *(This third condition is implicit in the companion species paper's sketch; Movement II makes it explicit, since the first two alone do not force the count to double — §14.)*

The candidate result whose proof is this conjunction is the **Step Individuation Theorem**:

Step Individuation Theorem (candidate)

Once refinement is denominated in minimal distinctions, the admissible refinement unit is *uniquely* one Fold, applied uniformly: an admissible refinement step neither composes several Folds nor advances vacuously, and when it advances it advances across the whole live frontier.

Its non-circularity hinge — the discipline separating a derivation from a restatement — is that the uniqueness and uniformity of the Fold-step must be *derived* from the substrate, not read off the slogan that refinement is "counted in Folds." Movement II takes up this candidate and follows it to its terminus.

Movement II — The Attempt to Discharge the Hinge

9. Inherited Setup for the Descent

Movement II draws on inherited structure beyond the Fold of §2.

Irreversible commitment. *[Inherited]* The substrate operation fixing one admissible possibility as a committed fact, irreversibly — one of the three primitives, alongside finite distinguishability and admissibility closure. A committed fact is a *new* fact by construction.

Emergent order. *[Inherited]* Physical temporal order is not a background parameter but a relational ordering induced by commitment: the ordering *is* the structure distinguishable committed facts induce (Distinguishability Principle paper §6; no background continuum is assumed, per house convention).

Substrate sequencing versus physical time. *[Inherited]* The programme distinguishes *substrate sequencing* from *physical temporal order*, and the two must not be conflated in any circularity argument. Tick accumulation provides a pre-commitment sequencing structure through which candidate evolution may proceed; it is not observable physical time. Physical time, by contrast, emerges only after commitment and record formation (the entry above). An argument that some pre-commitment construction is circular must therefore show dependence on *physical temporal order* — which is genuinely downstream of commitment — and not merely on substrate sequencing, which is available before commitment and carries no such circularity. This distinction is used to keep §15 honest: the apparent obstruction there is not that frontier-uniformity needs *any* sequencing (ticks supply that) but that it appears to need something stronger.

No-pre-individuation. *[Inherited]* The substrate supplies no external labels, sector indices, preferred carvings, or non-native conserved charges; admissible data are invariants, not presentations. Distinctions among sectors are physical only if carried by an admissible invariant, not by labelling or position in a diagram. This principle is the source the equivariance argument of §16 draws on.

The two layers. *[Inherited, QM reconstruction and NPI/Seal-Trichotomy]* The programme distinguishes a *candidate (pre-commitment) layer* and a *committed (record) layer*.

- *Candidate layer.* Pre-commitment dynamics are unitary; the structure is a *menu of candidate refinements* with overlapping support. The probability machinery rests on a positive-semidefinite support functional $s(A, B)$ with $s(A, B) \neq 0$ for distinct candidates in general — the candidate alternatives are **non-orthogonal**, the overlap carrying interference once complexified.
- *Record layer.* Commitment selects and freezes one admissible refinement; post-commitment states are orthogonal/projective. Committed records of a register are **mutually exclusive** — to commit to A is to exclude ¬A — hence orthogonal.

Commitment is the map from the non-orthogonal candidate field to a mutually-exclusive record.

Capacity. *[Inherited, NPI/Seal-Trichotomy]* The candidate menu is not indexed by a pre-given labelled outcome space; admissible refinements are *internally individuated carvings* of an unresolved region's capacity. Capacity Vol_op(M) is a **bare additive scalar**; the operational dimension d_op is a **capacity dimension**, not a labelled index set; Σ(M) is resolution-dependent capacity-counting; capacity is **resolution-dependent**; and FP1 is **silent on dynamics** — no conservation of a function on refinement-space, no rule forcing carvings into a complete tree. The Seal-Trichotomy papers state capacity forces neither seal nor flow.

10. The Lower Half — No Vacuous Steps Is Proven

The lower closure condition (§8.2) is provable from inherited primitives only, by two independent arguments.

From commitment-as-fact-fixing. A refinement step is a commitment event; refinement *is* the committing of a distinction. Commitment is by construction the fixing of a *new* admissible possibility (§9). A purported step fixing no new distinction performs no such transition. Such a "step" is not an instance of commitment; there is no idle step to forbid — "a commitment that commits nothing" is a category error, like a cut that separates nothing.

From emergent order. Depth is read off the commitment ordering, not off a background parameter. "Depth advanced, but nothing was distinguished" is incoherent: a position in the commitment ordering is induced by a commitment; an advance with no commitment has nothing to induce it — not an idle step but a non-step. In a theory with background time an idle step is coherent and must be excluded by hand; where order is emergent from commitment, there is no "passing" for nothing to happen during.

Status: [Proven], conditional on two inherited inputs — that refinement steps are commitment events, and that order is emergent from commitment. The one open input is the identification "refinement step = commitment event"; this is the natural reading and is not the disputed point (which is individuation, counting, and uniformity, not whether steps are commitments). On that reading the lower-bound dependency is genuinely dissolved.

11. The Upper Half I — The Decomposition Collision Dissolves

A natural first reading holds that no-multi-Fold-steps (§8.1) collides with the programme's non-separability commitment. That reading is mistaken.

Two decompositions, conflated. No-multi-Fold-steps requires joint distinctions to decompose **into successive single-Fold (binary) steps**. The apparent collision invokes the inherited fact (Distinguishability Principle §3.3) that a joint state's statistics are *not fixed by its marginals* — entanglement, contextuality. But these are two different decompositions: into **local/marginal** parts (factorising *across subsystems*), and into **binary cuts** (splitting the joint possibility space by successive two-way cuts). Non-separability blocks the first; no-multi-Fold-steps requires only the second.

The Bell construction — binary without local. Take the four Bell states — a genuinely entangled, non-factorisable four-way joint distinction. Discriminating them decomposes into single binary Folds: measure **ZZ** parity (splits $\{\Phi^+, \Phi^-\}$ from $\{\Psi^+, \Psi^-\}$), then **XX** parity (splits the ± within each). Since $[ZZ, XX] = 0$, these are two compatible binary commitments, each one Fold, committed one at a time, reaching the four-element entangled basis. ZZ and XX are *joint* parities — each requires both subsystems — so non-separability is intact: the Bell distinction remains non-locally-decomposable (LOCC cannot discriminate the Bell states) while being **binary-decomposable into joint Folds**.

Orthogonality forced; closure forbids rather than forces the joint primitive. The Bell tree is a tree over *mutually exclusive committed outcomes*. Mutual exclusivity of committed records is **orthogonality**, *forced* by irreversibility (§9), not assumed. Any set of mutually orthogonal admissible outcomes admits a binary-cut decomposition (orthogonal projectors give a binary tree "P_1 vs rest, P_2 vs rest, …").

Admissibility closure does not force a non-decomposing joint primitive. Mutual exclusivity of records yields binary-decomposability of any single commitment's outcome set; closure, if anything, forbids a non-decomposing primitive rather than forcing one.

Status: [Proven], relative to inherited irreversibility and non-separability — the collision is a conflation; binary-decomposability holds at the orthogonal record layer; closure does not host the residue.

12. The Upper Half II — The Count Lives at the Candidate Layer

§11 clears the record layer. But the upper bound is about a *count* — N_{n+1}/N_n — and the decisive question is which layer it inhabits.

It inhabits the **candidate** layer, by what it counts. N_n is the number of distinguishable *alternatives* after refinement to depth n — the size of the menu, where the count of alternatives *grows*. Commitment does the opposite: it *selects* one alternative and freezes it, taking the count to one. A commitment event has no branching factor in the relevant sense; its factor is → 1, not

2. The four Bell records exist as a menu of four *before* the ZZ/XX commitment selects among them.

So the orthogonality of §11 lives at exactly the layer where the dyadic *count* is **not** defined (the record layer). The layer where the count *is* defined (the candidate menu) is, by §9, **non-orthogonal**: overlapping support, $s(A, B) \neq 0$. The clean binary-decomposition argument is unavailable there. This is confirmed by the corpus: the menu is *internally individuated carvings of capacity*, with **no external record basis** and no pre-given exclusive labels. The count cannot borrow record-orthogonality, because there is no record basis indexing the menu.

Consequence. The URH's conditions do not all live in the same primitive. No-vacuous-steps is settled at commitment (§10). No-multi-Fold-steps and frontier-uniformity are facts about the *candidate menu* — the NPI candidate-refinement / capacity layer. The descent continues there.

13. Capacity Is Not Atomic — Quantisation Buys Integers, Not Powers of Two

At the candidate layer the menu is carved from capacity Vol_op(M). The obvious target — a **Capacity Quantisation Theorem**: that Vol_op is allocated in indivisible Fold-quanta — does not hold up.

Likely refutable. The inherited characterisation (§9) describes Vol_op as **scalar, additive, partition-free, resolution-dependent** — the properties of a *continuous* capacity. With d_op = $\log_2$(Vol_op) a bit-count, the count is integer only when Vol_op is a power of two; for generic regions and generic ε it is not (a capacity of 3 gives d_op = $\log_2 3 \approx 1.585$ — the "one Fold plus a partial one" of §5, here generic). Resolution-dependence makes it worse: as ε varies continuously, Vol_op sweeps continuously through every non-power-of-two value, which an atomic quantity cannot do. The strong claim "Vol_op is intrinsically Fold-quantised" most likely ends in **refutation by resolution-sweep**.

Integerisation versus doubling. A weaker route survives: *committable* capacities are integer-bit-valued, because a committed record is discrete — one cannot freeze $\log_2 3$ of a distinction. The Fold-atom is then in the **commitment interface**, not in Vol_op. But integerisation gives committable capacities 1, 2, 3, 4, … — **integer** loading, not **dyadic** 1, 2, 4, 8. A 3-bit committable capacity is a three-Fold structure, but whether its menu presents $2^3 = 8$ leaves or merely 3 sectors depends on **completeness of the carving tree**, which quantisation does not supply.

Integerisation ⟹ integer loading; at the commitment interface. **Doubling** ⟹ a **complete binary tree** (completeness + binarity + frontier-uniformity). Not a quantisation fact.

Even granting integerisation in full, dyadic loading is untouched. The descent continues to the tree.

14. Doubling Requires a Complete Binary Tree

Doubling needs the carving tree complete and binary at every level. Decompose:

- **(G) binary generators** — every carving is a two-way Fold-cut. *Near-free* from (B): if the only primitive carving is a Fold, every carving composes binary carvings.
- **(C) complete and uniform schedule** — completeness + frontier-uniformity: at each depth every live sector is carved, by exactly one Fold, none left uncarved. Only then does the leaf count double.

(G) is the binarity of the *generator*; (C) is dyadic loading itself in candidate-layer vocabulary. The seductive move "carvings are binary, so the menu is 2^n" smuggles (C) into (G). Naming (C) as the unsolved part separates derivation from relabelling.

What sources (C)? Not FP1: by §9 it carries no conservation on refinement-space and no dynamics — it cannot force any schedule. So (C) needs structure beyond FP1. The first two URH conditions do not supply it: no-vacuous-steps (§10) forbids an *idle* step but not a *selective* one; no-multi-Fold-steps (§11) forbids *bundling* but not *partiality*. The unsupplied piece is precisely **frontier-uniformity** — that when refinement advances, it advances across the whole frontier rather than depth-first into one sector. §§15–16 confront it.

15. The Apparent Obstruction — Frontier-Uniformity Looks Order-Laden

Frontier-uniformity, read naively, requires that the live sectors advance *together* — a depth level n being a stage at which the whole frontier is carved at once. The worry is that this needs a notion of **co-level completion across the frontier**: not merely that carvings are sequenced, but that a whole level is closed off before the next begins. Here the distinction of §9 must be applied carefully. Mere *substrate sequencing* — tick accumulation — is available pre-commitment and supplies an order in which candidate carvings may proceed; it carries no circularity, since it does not depend on physical time. So the obstruction is *not* that frontier-uniformity needs sequencing at all (ticks supply that). The worry is the stronger one: that *co-level completion* — every sector at depth n carved before any reaches depth n+1 — looks like it needs something beyond bare sequencing, a notion of simultaneity or frontier-wide synchrony, and *that* stronger structure is the one that would, if it had to be imported as a primitive, presuppose physical temporal order rather than mere ticks.

On the naive reading the missing slot may loop — but only via the stronger notion. Substrate sequencing (ticks, pre-commitment) is not circular and is freely available. What looks order-laden is *frontier-wide completion*: closing a level as a level. If that required physical temporal order — order emergent from commitment — it would be circular, since the carvings to be completed are pre-commitment candidate structure. Whether tick accumulation alone supplies co-level completion, as opposed to merely *some* refinement schedule, is not obvious and is taken up in §15A. Admissibility closure does not supply it either: closure constrains which *configurations* are admissible, not the *completeness* of their generation (and §11 showed it does not even force the joint primitive).

This would be the terminus — dyadic loading obstructed because co-level completion appears to need more than substrate sequencing can give — *if* frontier-uniformity genuinely required that stronger synchrony. The next section shows it does not: the same all-or-none can be obtained from a symmetry that invokes no notion of completion-as-a-level at all, neither physical time nor co-level closure, so the question of whether ticks supply the stronger structure need never be answered.

15A. Does Tick-Accumulation Supply the Missing Refinement Dynamics?

Before turning to the symmetry route, it is worth asking the direct question, because the answer locates exactly what substrate sequencing can and cannot do — and shows that the *sequencing* route, pursued on its own, lands on the same hinge the *symmetry* route reaches in §17.

The natural hope is that tick-accumulation (TPB), the pre-commitment sequencing the substrate already supplies (§9), might *itself* generate refinement: ticks defining refinement depth, candidate carvings proceeding tick by tick, a tree growing as ticks accumulate. Three sub-questions sharpen this.

- *Can tick-accumulation generate refinement at all?* Plausibly yes: sequencing is precisely what a step-by-step carving needs, and it is available pre-commitment without circularity. This is the part the §9 distinction rescues — refinement *may* proceed through substrate sequencing, and a critic who says "you forgot ticks" is answered here.
- *Can ticks define refinement depth?* Plausibly yes: depth as tick-count is a coherent pre-commitment notion, again carrying no dependence on physical time.
- *Can a* **complete binary tree** *emerge from tick dynamics — every live sector carved at each level — as opposed to* **some** *refinement schedule?* This is the one that does not follow. Sequencing orders carvings; it does not, by itself, force *which* sectors are carved at a given stage, nor that *all* of them are. A tick-driven schedule could perfectly well carve one sector repeatedly while neglecting others — a depth-first walk is a valid use of sequencing — and that is exactly the non-uniform schedule that fails dyadic loading.

So the honest verdict on TPB: tick-accumulation may well supply sequencing sufficient for candidate-layer evolution, and so dissolves the worry that refinement needs physical time to occur. But sequencing alone does not imply *complete* refinement. The outstanding question is therefore not whether refinement may occur pre-commitment — it may — but whether tick dynamics *force complete binary refinement* rather than permitting arbitrary schedules.

And here the sequencing route rejoins the main line. "Do tick dynamics force the whole frontier to advance, rather than some sub-set?" is asking whether the dynamics can treat one live sector differently from another at a given stage. If the sectors are genuine siblings — alike in the admissible data the dynamics reads — then forcing some and neglecting others is exactly the disposition-asymmetry §17 shows equivariance forbids; and whether the dynamics *can* so distinguish them is exactly the supervenience question of §17.1–17.2 (does the rule's data determine the overlap profile that would separate them?). So TPB does not open a second, independent fork. Asking whether tick-accumulation forces *complete* rather than *arbitrary* refinement is the sequencing-flavoured statement of the same hinge the symmetry route reaches:

frontier-completeness and sibling-supervenience hold or fail together. The §16 symmetry argument is the cleaner route to it, because it bypasses the question of what ticks supply entirely; §17 then locates the residue precisely.

15B. Why No-Pre-Individuation Implies Equivariant Refinement

The equivariance argument of §16 may appear to introduce a new principle — that admissible refinement rules commute with isomorphisms, $F(\sigma x) = \sigma(F(x))$. This section shows that equivariance is not an additional physical postulate but the operational form of an already-inherited commitment: no-pre-individuation (§9).

No-pre-individuation as a minimality principle. The programme holds that admissible physical structure is exhausted by admissible distinguishability relations: no external labels, preferred indices, hidden identifiers, or presentation-dependent distinctions are admitted as physical data, and a distinction is physically meaningful only if carried by an admissible invariant. As a familiar illustration, not a premise, modern physics routinely identifies physical content with invariant structure — gauge-related descriptions are not distinct physical states; manifold points carry no individuality independent of relational structure; coordinate labels are not observables. The common thread is that physical laws may depend only on physically meaningful distinctions. No-pre-individuation is the candidate-layer instance of this, and the argument below rests on it alone, not on the analogy.

Why selective refinement requires additional structure. Let two candidate sectors A and B be isomorphic, $A \cong B$ — no admissible invariant distinguishes them — and suppose a rule acts selectively: $A \mapsto (A_1, A_2)$ while $B \mapsto B$. The rule has treated A and B differently. What grounds the difference? There are exactly three possibilities.

1. **An admissible invariant distinguishes them.** Then A and B were never isomorphic and the premise fails — harmless.
2. **An inadmissible label distinguishes them.** Then the rule depends on structure the theory explicitly excludes — a direct violation of no-pre-individuation.
3. **Nothing admissible distinguishes them.** Then the rule is not a function of admissible data, contradicting that it is an admissible rule. The reply "then refinement is merely stochastic" does not reopen this: an admissible stochastic rule must have σ-invariant *statistics* — a distribution over outcomes not keyed to any admissible distinction is symmetric between A and B — so a selective *bias* toward A over B still requires a distinguishing datum, returning us to (1) or (2).

No fourth possibility exists.

Equivariance as the operational form of the principle. The conclusion is immediate: if an admissible rule may depend only on admissible distinctions, it cannot distinguish sectors related by an admissible isomorphism, which is exactly $F(\sigma x) = \sigma(F(x))$ for every admissible σ. Equivariance is therefore not an independent axiom layered on top of no-pre-individuation; it is that principle applied to refinement dynamics. The all-or-none result of §16 does not rest on a separate symmetry postulate — it follows from the requirement that refinement depend only on admissible distinctions.

The standing of the conclusion — calibrated to the s-audit of §17.4. What the trichotomy establishes is that equivariance is what no-pre-individuation *requires* of any admissible refinement rule: a non-equivariant F would, by the argument, either read surplus structure (forbidden) or fail to be a function of admissible data (not a physical law). So equivariance is the **default**, and the burden falls on any challenger to *exhibit the admissible distinction* a selective rule would need — absent which, selective refinement is not a physical possibility but a relabelling of identical structure. This does not yet certify that the refinement rule *as actually constructed* imports no presentation-dependent choice; that is a construction-audit, exactly parallel to the audit of the support functional s in §17.4. No-pre-individuation thus discharges two audits of identical epistemic structure — one for F (this section), one for s (§17.4) — each converting "could surplus structure enter?" from an open worry into a default-plus-burden, with the only residue the inspection of the built object. The requirement is inherited; the construction-audit for F is the parallel residue.

16. Equivariance Removes the Order-Loop

Frontier-uniformity has two possible sources, and they are not equivalent. One is *co-level completion* — closing each level as a level, which looked to need more than substrate sequencing can give (§15, the apparent obstruction). The other is *label-blindness* — the rule cannot prefer one alike sector over another. The second needs no notion of completion-as-a-level at all, neither physical time nor frontier-wide closure, and it is the correct one.

The equivariance argument. Suppose that at a stage the live sectors are mutually isomorphic: no admissible invariant distinguishes one from another (§9, no-pre-individuation), so they differ only by presentation, labelling, or position in a refinement diagram — none physically admissible as a distinguishing datum. Let F be the refinement operation and σ an isomorphism exchanging two live sectors. A physically admissible refinement rule is **equivariant** under such isomorphisms — not as a fresh posit but as the operational form of no-pre-individuation, derived in §15B:

$$F(\sigma(x)) = \sigma(F(x)).$$

This expresses that refinement may not depend on an inadmissible label: exchanging two isomorphic sectors before refining and refining before exchanging must give the same physical result. Now consider a *selective* step refining one isomorphic sector while leaving its partner unrefined,

$$A \mapsto (A_1, A_2),\ B \mapsto B, \text{ with } A \cong B \text{ via } \sigma.$$

Applying σ before refinement gives the opposite outcome, $B \mapsto (B_1, B_2)$, $A \mapsto A$. The two are not related by a mere relabelling of the rule; the rule has *physically* selected one isomorphic sector over another, violating equivariance — unless some symmetry-breaking datum distinguishes A from B. By no-pre-individuation no such datum is present; introducing one (a hidden sector label, preferred carving, boundary condition, non-native charge) is exactly the surplus structure the admissibility arguments exclude. Hence selective refinement is inadmissible among isomorphic sectors, and the admissible alternatives at a stage are: **all alike live sectors refine, or none do.**

Why this removes the loop, rather than relocating it. The all-or-none conclusion is derived from *label-blindness*, not from any notion of completion-as-a-level. At no point does the argument invoke "at the same depth," "at the same instant," or even co-level closure; it invokes only that the rule cannot read an inadmissible label. Frontier-uniformity — the third and last unsupplied piece of (C) — is therefore a **symmetry fact, not a completion fact**, and the §15 worry (that co-level completion would presuppose physical temporal order) does not arise, nor need one settle whether tick-accumulation supplies co-level completion (§15A). This is the materially new step: the obstruction §15 identified is *defeated*, not merely moved to another primitive, and it is defeated on an axis — fairness between indistinguishables — orthogonal to sequencing altogether.

Equivariance is a constraint on the rule, at the level of disposition. One subtlety must be fixed before the argument is pressed further, because a careless reading destroys it. Equivariance constrains F — the refinement *rule* as a function of admissible data — not the bare instant of action. F(A) = F(B) for isomorphic A, B means *the rule assigns them the same fate as a function of their admissible data*, not that it acts on them at the same instant. The distinction is load-bearing:

- **Transient asymmetry** — A carved before B within a single sweep, the ordering set by a non-admissible presentation detail (labelling, position in a diagram) — does *not* violate equivariance. It is a difference in presentation order, not in disposition, and label-blindness already quotients it out. A principle forbidding *any* momentary depth difference would forbid refinement from ever occurring, since a sweep necessarily carves its sectors in *some* order.
- **Standing asymmetry** — A's disposition is "refine," B's is "never refine," as a standing policy of the rule — *does* violate equivariance, because it is a difference in F as a function of admissible data, with no admissible invariant to ground it.

The quantifier sits on the disposition, not the instant. This is what makes "transient ordering fine, standing exclusion forbidden" a principled line rather than a convenient one, and it is the form in which the argument now does all its remaining work.

The status of the equivariance principle. "An admissible refinement rule is equivariant under sector-isomorphisms" is the refinement-space form of no-pre-individuation — label-blindness applied to the refinement operation. §15B derives it: by the three-possibility argument there, a non-equivariant rule must either read surplus structure (forbidden) or fail to be a function of admissible data (not a physical law), so equivariance is **[Inherited]** as a *requirement* of no-pre-individuation, not a fresh posit. The only residue is the construction-audit — whether the rule *as actually built* imports a presentation-dependent choice — which is **[Open]**, exactly parallel to the audit of s in §17.4. Either way it is not a synchrony assumption, which is the whole point.

The two apparent gaps collapse into one. The argument earns "all alike sectors refine **or none do**," and two gaps seemed to separate that from doubling: a ragged-frontier induction (the frontier must *stay* level), and a layer check (equivariance must hold where the count lives). The first is not independent of equivariance at all, once equivariance is read at the disposition level — which is the substance of §17. Permanent starvation of an isomorphic sector is *itself* a standing disposition-asymmetry, hence a direct equivariance violation, not a separate capacity condition to be discharged downstream. What survives is the *single* second gap: whether the disposition-level symmetry between siblings genuinely holds at the non-orthogonal candidate layer, where the count lives. §17 collapses the first gap and isolates the second as the lone terminus.

17. Starvation Is Disposition-Asymmetry — The Single Residue

The equivariance argument changes the terminus decisively. The earlier reading left a fork — a *synchronic* count obstructed by an order-loop against a *diachronic* count reducing to a capacity question — and then, in a first correction, left two conditions: equivariance-at-the-candidate-layer and no-permanently-starved-sector. The disposition-level reading of §16 collapses these further. The fork is gone (the order-loop is removed, §16), and the two conditions are *not independent*: no-starvation is equivariance itself, evaluated on the rule rather than on the limit.

Starvation is not a passive outcome; it is an active standing asymmetry. "Permanently starved" sounds innocent, but unpacked physically it is a *policy* of the refinement rule:

sector A continues to refine; sector B never refines.

To sustain that, the substrate must, at every relevant step, *act on A and not on B* — allocate carving to A, deny it to B. That is not one delayed step; it is a standing difference in how the rule treats A versus B. If A ≅ B with no admissible invariant separating them, this standing difference is precisely a difference in F as a function of admissible data — a disposition-asymmetry — which equivariance forbids directly (§16). The violation does not wait for "depth 10" to accumulate; it is present in the *standing policy itself*, at the level equivariance governs. One need not integrate up to a contradiction: the contradiction is the policy.

This is stronger than the previous framing, which treated no-starvation as a fact about the limit ($F^n(A) \neq F^n(B)$ for large n) needing its own argument. The limit is only the accumulation of per-rule decisions, and each is already in equivariance's jurisdiction. So:

No-Starvation as a Theorem, not a Posit

Permanent starvation of an isomorphic sector is itself a violation of equivariant refinement. A standing policy "refine A, never B" with A ≅ B and no admissible invariant separating them is a disposition-asymmetry, which equivariance forbids. No-starvation is therefore not a fourth substrate condition but a *consequence* of equivariance — provided the siblings are genuinely isomorphic at the layer where refinement decisions are made.

The transient/standing distinction of §16 is what keeps this from proving too much: a sweep that carves A before B is *transient* (presentation-order, label-blind) and permitted; a rule that carves A forever and B never is *standing* and forbidden. Starvation is standing by definition ("never refines"), so it falls on the forbidden side without dragging admissible transient asymmetry with it.

Where could the asymmetry come from? — a trichotomy that closes the case. A standing "refine A, never B" must be grounded in *something*. There are exactly three possibilities:

- **Geometry** — a genuine admissible invariant distinguishes A from B. Then $A \ncong B$: the sectors were *never* isomorphic, the premise simply fails, and selective refinement is admissible and not a violation. (This does not threaten dyadic loading; it means the two were different sectors all along.)
- **Capacity** — capacity is allocated to A and withheld from B. But capacity is a *bare additive scalar* with no sibling-indexing (§13): the scalar does not know which sibling is which, so it carries no datum on which to ground an asymmetric allocation. Asymmetric allocation between isomorphic siblings would require a *label on top of* the scalar — exactly the surplus structure no-pre-individuation excludes. Capacity therefore cannot be the selector, *unless* the candidate-layer structure secretly distinguishes the siblings — which is the third branch.
- **Nothing** — no admissible datum grounds the asymmetry. Then the asymmetry cannot stand: starvation is impossible.

The middle branch is where the whole question now concentrates, and it has a single sharp form. Capacity itself is sibling-blind; the only place an admissible sibling-distinguishing datum could hide is the **candidate-layer support functional** $s(\cdot,\cdot)$ (§9, §12), the overlap structure on the non-orthogonal menu. If $s(A, C) \neq s(B, C)$ for some third sector C, then A and B carry distinct overlap profiles — an admissible invariant separating them — they were never isomorphic *at the candidate layer*, and selective (hence non-dyadic) refinement is admissible. If instead the support functional respects sibling isomorphism — record/geometric isomorphism implying isomorphism of the overlap structure — then no admissible datum anywhere grounds the asymmetry, the middle branch is empty, and starvation is impossible by the theorem above.

The single residue. The entire dyadic-loading result now rests on one well-posed question about one object:

Does the candidate-layer support functional $s(\cdot,\cdot)$ ever assign isomorphic siblings distinct overlap profiles?

If **no** — if isomorphism at the record/geometric level implies isomorphism of the overlap structure — then no admissible datum can ground asymmetric refinement, the No-Starvation Principle is **proven** (not posited), and the chain closes: **Fold-binarity (B) + equivariant refinement ⟹ uniform refinement ⟹ dyadic loading**, for arbitrary depth. If **yes** — if overlap can break sibling symmetry — then that overlap asymmetry *is* the admissible selector, the siblings were never truly isomorphic at the candidate layer, and selective non-dyadic refinement is admissible after all, so dyadic loading fails (or holds only over overlap-symmetric sub-families, with a separate account of why those are the relevant ones).

17.1 The Residue Is a Supervenience Question, Not a Hidden-Overlap One

The "single residue" as just stated has a tempting quick answer that is correct but does *less* than it appears, and seeing why pins the genuinely final hinge.

The equivariance σ acts on the *whole* candidate-layer structure, the support functional included. If A ≅ B via σ, then σ is by definition an automorphism preserving all admissible structure, hence preserving s: s(A, C) = s(B, σC) for every corresponding C. A sibling-distinguishing difference in overlap profile is therefore *impossible from inside an isomorphism class* — such a difference would mean σ failed to preserve s, i.e. σ was not an isomorphism after all. So the dichotomy is exhaustive and the answer to the face-value question is a clean *no*: equal-in-all-admissible-respects sectors have equal overlaps trivially.

But that "no" is bought by counting s as admissible structure — by *including the support functional in the definition of isomorphism*. The substantive content has not vanished; it has migrated from the theorem into the definition, and with it the risk. The honest question is no longer "can overlap break symmetry within a class?" (no, by construction) but:

Are there sectors the refinement dynamics treats as genuine siblings — alike in whatever admissible data the rule F actually reads — that nonetheless differ in support profile?

Two notions of sameness have come apart and must be named:

- **refinement-siblinghood** — alike in the admissible data F reads to decide refinement;
- **support-isomorphism** — alike including the overlap functional s.

The face-value theorem forces all-or-none only *within a support-isomorphism class*. If support-isomorphism is *strictly finer* than refinement-siblinghood — if s distinguishes sectors F does not — then the equivariance constraint binds only on the finer classes, selective refinement *between* support-classes that are refinement-siblings is fully admissible (they are not isomorphic on the rich definition), and that selective refinement is exactly the non-uniform schedule that breaks dyadic loading. Enriching the definition did not forbid the counterexample; it renamed it from "starvation of siblings" to "differential refinement of non-siblings," which nothing forbids.

Why a rich definition cannot be a free lunch (conservation of loopholes). One cannot simply *decree* that isomorphism includes s and collect dyadic loading, because the decree has a cost that exactly offsets the gain. Including s shrinks the isomorphism classes; the smaller the classes, the *fewer* pairs equivariance forces to refine together, hence the *weaker* the all-or-none constraint and the *more* admissible selective refinement becomes. Enriching the definition to close the starvation loophole opens the non-sibling loophole by the very same act. The two are conserved against each other, so no choice of definition settles the matter; a *fact* is required, not a stipulation.

The biconditional. That fact is supervenience:

P2 (dyadic loading) closes if and only if the support profile supervenes on the refinement-relevant data — i.e. iff refinement-siblinghood and support-isomorphism *coincide*, iff s adds no admissible distinguishing power beyond what F already reads.

If s supervenes on F's data, the two sibling-relations coincide: equivariance binds on exactly the refinement-siblings, no gap remains, no-starvation is proven, and dyadic loading follows. If s carries independent distinguishing power, the relations come apart, the gap between them is admissible non-uniform refinement, dyadic loading fails — and no enrichment of "isomorphism" repairs it, since enrichment only trades the starvation loophole for the non-sibling loophole.

This is non-circular and biconditional: it does not bank supervenience, and it does not pretend a definitional choice settles a substantive matter. The genuinely final hinge is therefore a checkable fact about how the corpus *constructs* the support functional:

Does the candidate-layer support functional s introduce admissible distinctions that the refinement dynamics F does not already see? If s is built entirely from refinement-relevant structure — read off the same data F is — supervenience holds and P2 closes. If s is an independent primitive with its own degrees of freedom, supervenience fails and that independence is the last obstacle, now precisely located.

So the terminus is not "is the definition of isomorphism rich enough" (a choice) but "does support supervene on refinement-relevant data" (a fact about the construction of s) — and that is the single question on which the entire dyadic result now turns.

17.2 The Definition Is Free; the Coincidence of Data Is the Theorem

One reading of §17.1 must be headed off, because it is tempting and it concedes nothing of substance while appearing to close the result. The reading is: *define candidate-layer isomorphism to include the support functional — then $A \cong B$ already means support-equal, support cannot break symmetry by construction, and P2 closes.* The first half is correct (it is §17.1's opening move); the conclusion does not follow, and seeing why fixes the terminus in its honest form.

The equivariance constraint forbids "refine A, not B" *only for pairs the refinement rule treats as isomorphic* — pairs alike in the data the rule **F** actually reads when it decides whether to carve. The decisive question is therefore whether F has s in front of it at the moment of decision.

- **If F reads s** — if the support profile is among the data the refinement rule responds to — then enriching isomorphism to include s is the correct definition, but it changes nothing, because refinement-siblinghood already *means* s-equal: the two notions were never apart. P2 closes — but the operative fact is "F reads s," a claim about the dynamics, not the definition.
- **If F does not read s** — if the refinement decision is made on data that underdetermines the overlap profile — then there exist sectors F treats identically (equivariance over F's

data forces all-or-none on them) that nonetheless differ in s. The rich definition calls these "not isomorphic," but that relabelling does not save dyadic loading: F, blind to s, refines them together regardless, so either the count proceeds dyadically *for a reason the rich definition denies*, or some downstream use of s later treats them differently and the count and the dynamics come apart. The definition has renamed the sectors without changing what the rule does.

In neither case does the definitional choice decide the matter. What decides it is the relationship between **what F reads** and **what s encodes** — and that is exactly the supervenience fact of §17.1, now seen from its operative side. Adopting the full-candidate-layer definition of isomorphism is *free* and settles nothing on its own; the substantive content is whether the refinement rule's decision data already determine s. Equivalently:

Defining isomorphism to include s closes P2 **if and only if** the refinement dynamics already sees enough to determine s. The definition is free; the coincidence of F's data with s's data is the theorem.

This is the honest statement, and it points the next paper at a construction to inspect rather than a definition to adopt: *do F and s draw on the same substrate data?* There is reason to expect they do, and to expect the answer favourable — in the programme s is built from the distinguishability structure, and the refinement rule is *also* a creature of that same distinguishability structure; they are not independent inputs but both downstream of the one substrate relation. If that is borne out, F's data and s's data are the same data, supervenience holds, and P2 closes for real. But that is a sentence to be *earned* by checking the construction of s against the inputs of F — not installed by defining isomorphism to include s.

This much can be stated as a theorem, with the supervenience condition named as its hypothesis rather than assumed silently.

Conditional Support-Closure Theorem

Let **D** be the full admissible distinguishability data at the candidate layer. Suppose:

1. **F**, the refinement rule, is a function only of D (it reads no inadmissible label — mild, just admissibility);
2. **s**, the candidate-layer support functional, is also fully determined by D;
3. no independent support degrees of freedom are introduced beyond D.

Then any two sectors indistinguishable to F have identical support profiles — both being functions of the same D — so refinement-siblinghood and support-isomorphism coincide. The supervenience condition therefore holds, no-starvation follows by equivariance (§17), and dyadic loading closes:

Fold-binarity (B) + equivariant refinement + common distinguishability origin of F and s $\Longrightarrow N_n = 2^{n-1}$.

What the theorem does and does not do. The deduction from premises 1–3 is airtight: two quantities that are both functions of the same argument cannot disagree on inputs where that argument agrees, so F-siblings are support-isomorphic trivially. But premises 2–3 *are* the supervenience condition, stated as hypothesis — the theorem does not derive supervenience, it shows supervenience *suffices* and names the exact form it must take: **s carries no degree of freedom beyond D.** So this does not prove dyadic loading absolutely; it proves it **relative to a no-extra-support-structure condition**. The theorem's value is to reduce the entire dyadic result to one inspectable, falsifiable property of the construction of s.

The refuter, named. Dyadic loading fails *if and only if* the construction of s introduces a support distinction between D-identical sectors — an admissible overlap difference not fixed by the distinguishability data F reads. That is exactly what a counterexample must exhibit: two sectors alike in all of D, differing in s. If the construction of s never produces one — if s is read off D and adds nothing — the theorem fires and the chain closes. A conditional theorem whose condition has a named refuter is the programme's standard form, and this is its instance for P2.

[Proven] the theorem (premises 1–3 $\Longrightarrow$ dyadic loading). *[Open]* premises 2–3, the no-extra-support-structure condition — but §17.3 shows the burden on this lies the other way from how a bare "open condition" suggests.

17.3 The Burden Lies with the Challenger — No-Pre-Individuation and the Support Functional

It would understate the position to leave the no-extra-support-structure condition marked merely "plausibly met." The asymmetry between its two failure modes is not a matter of expectation; it is fixed by an inherited primitive.

The support functional s is introduced to measure relationships *between candidate alternatives*. It must be built from *something*, and there are exactly two possibilities. **Option 1: s is derived from the distinguishability data**, s = s(D); then the support profile carries no information not already in D, supervenience is automatic, and the Conditional Support-Closure Theorem fires. **Option 2: s is an additional primitive**, s = s(D, E) for some extra structure E; then a new physical degree of freedom has been introduced, requiring a definition, an interpretation, a reason it exists, and an account of why refinement cannot read it.

These options are *not symmetric*, and the asymmetry is supplied by **no-pre-individuation** (§9), already inherited. That principle states that the substrate admits no structure beyond admissible distinguishability data — no external labels, no surplus primitives, nothing the distinguishability relation does not carry. An independent support degree of freedom E is precisely such surplus structure, and is therefore *prohibited* unless E is itself admissible distinguishability data — in which case $E \subseteq D$ and Option 2 collapses into Option 1. Option 2 is not a co-equal branch; it is a branch that, to be taken, must either violate an inherited principle or dissolve into Option 1.

So the no-extra-support-structure condition is **the default reading of the programme, not a hopeful one**, and the burden lies with whoever would deny it. The residue is best stated as a challenge:

If s does not supervene on D, exhibit the additional admissible datum it carries — a concrete degree of freedom present in s and not in D, together with an account of why refinement cannot read it.

No-pre-individuation makes this demand hard to meet: any candidate E is either inadmissible (forbidden outright) or admissible (hence already part of D, hence no counterexample). The abstract worry "perhaps s contains more information" — which a referee can always raise — is forced, by this challenge, into the concrete "*what* information?", and the dichotomy closes against the objector in either branch.

The one guard, so the burden-shift does not become an overclaim. No-pre-individuation forbids an independent admissible *primitive* E; it does not, by itself, certify that the support functional *as actually constructed* in the Born/Double-Square machinery is in fact s(D) and nothing more. A construction can smuggle surplus structure in unintentionally — a basis choice, a representation-dependent phase, an ordering convention promoted to a physical datum — exactly the covert dependence the §6 relabelling-trap and the §17.1 conservation-of-loopholes warnings guard against. So two claims must be kept apart: *no admissible E exists* (no-pre-individuation: true, inherited) and *the constructed s honours this* (a property of the machinery, not yet formally audited). The principle flips the **presumption** — supervenience holds unless a covert dependence is exhibited — but it does not substitute for the **audit**: confirming that the s the corpus actually builds carries no convention-masquerading-as-datum.

The honest final status is therefore stronger than "expected" and short of "proven." Dyadic loading is a theorem under the programme's default reading, in which no-pre-individuation is honoured and s = s(D); its remaining openness is not "whether an extra degree of freedom might exist" (the principle says not) but "whether the constructed support functional covertly imports one" — a formal audit of the construction, with the presumption, and the burden of producing a counterexample, both resting on the challenger.

17.4 The Audit Resolved — The Terminus Coincides with the Born-Rule Layer-2 Obligation

§17.3 leaves one task: a formal audit of the constructed support functional s, to confirm it imports no convention beyond the distinguishability data D. That audit can now be carried out, because the corpus constructs s explicitly. In the Double Square derivation the candidate-layer overlap is the non-orthogonal kernel

W(P, P′) = e^{i(θ(P) − θ(P′))},

and the support profile s is built entirely from one object: the geometric phase θ. The supervenience question s = s(D) therefore reduces, without residue, to a single question about θ:

Is θ read off the admissible distinguishability data D, or is it independent structure?

The audit's verdict: supervenience is not free. The distinguishability datum is a symmetric metric d; the phase θ is oriented and antisymmetric, and a closed loop carries a holonomy θ(L) that d does not contain. Isometry-consistency constrains θ to be invariant under the reversible group, but invariance alone does not fix θ: distinct admissible phase assignments — distinct holonomy classes — are compatible with one and the same metric. Read at the level of the constructed object, θ is thus not a function of d by inspection, and so s carries, on its face, a datum beyond d. This is exactly the failure mode the §17.3 guard isolates: not a forbidden admissible primitive (no-pre-individuation excludes those), but a convention-carrying construction. The audit therefore *confirms* the §17.3 stance rather than undercutting it — it has located the one candidate covert dependence, named it (the phase), and shown it is the kind the principle does not auto-forbid.

What would close it, and the proposition that must do the work. The programme already holds the only defence that can close this. The companion *Physical Necessity of Quantum Probability Structure* argues that the phase is not free structure but is **forced from D**: finite distinguishability, tick-denominated commitment, and temporal extensibility together require the holonomy set to be infinite, then continuous, then minimally U(1). The operative word is *forced*: the shared proposition is not "θ happens to be continuous U(1)" but "**D forces continuous U(1) phase**," and the distinction is load-bearing. Two near-misses show why neither weaker reading suffices:

- *Forced-but-discrete.* If D forced θ onto a merely finite holonomy set, θ would be D-determined — supervenience would hold and the generation chain would close — but the Born chain would *not* obtain the continuous U(1) it needs to make probability quadratic.
- *Continuous-but-free.* If θ were continuous U(1) yet still unfixed by the metric (distinct holonomy classes, one metric, now in the continuous setting), the Born chain would have its U(1) but supervenience would *fail* — s would carry a continuous datum beyond d, and the generation chain would not close.

"θ is D-determined" (what supervenience needs) and "θ is continuous U(1)" (what quadratic probability needs) are thus independent axes; neither entails the other. It is precisely *forcing-from-D* that discharges both at once — "D forces continuous U(1)" is simultaneously "θ is continuous U(1)" (the Born requirement) and "θ is determined by D" (the supervenience requirement), because forcing-from-D *is* D-determination. The single proposition that does double duty is therefore the forcing claim, not the bare continuity claim and not bare D-determination.

The coincidence, stated at its true strength. With that precision, two arcs of the programme bottom out at one proposition. The Born-rule chain needs D to force continuous U(1) phase, to make probability quadratic. The generation-count chain needs D to force θ — and the forcing makes θ D-determined, so that s supervenes on D and no-starvation follows. "Why powers of two?" and "why $P = |\psi|^2$?" are not independent open questions; they are one open question — *does finite distinguishability force continuous U(1) phase?* — wearing two hats, the shared content being the forcing, which delivers continuity to the one chain and D-determination to the other.

Where that defence stands — reported at the companion's own assessment, not endorsed in this paper's voice. The companion marks its load-bearing step — the upgrade from infinite holonomy to *continuous* U(1) holonomy — as incomplete: *in the companion*, the impossibility of purely finite holonomy is established, while the continuity-and-U(1) step is reduced to four technical items and deferred to a dedicated treatment (*Why Finite Distinguishability Forces Continuous Phase*). The companion characterises this as the sole remaining foundational obligation of its Born-rule chain. This paper inherits that step at the companion's stated status; if the companion's status changes, the claim here tracks it.

The open proposition is a global-consistency claim, not a local-density one — and the minimum-distinguishability scale is no objection to it. A natural challenge runs: the programme posits finite distinguishability (distinction bottoms out at one Fold), so why should the phase set be continuous rather than finite? The challenge conflates two claims that must be separated. Finite distinguishability limits phase *observability* — two states whose phase differs by less than the minimum distinguishable amount cannot be locally separated by a single comparison — but it does not limit the phase *catalogue*, the set of values the phase variable may take. The two come apart exactly as instrument resolution comes apart from the quantity measured: a ruler marked in millimetres does not force the line it measures to have millimetre length; finite pixels do not pixelate the world. So "finite distinguishability ⟹ finite phase set" is a non-sequitur the moment one notices that phase need not be a *local distinguishability coordinate*.

And in the programme it is not. The companion's result is named "finite holonomy impossible → continuous U(1)" — the structure of a *closure-under-composition* argument, not a fineness-of-resolution one: local distinctions are finite, but transport compositions accumulate without bound (arbitrarily long commitment histories, arbitrarily many loops), and the requirement that phase assignment remain *consistent under unbounded composition* is what forecloses a finite holonomy group. Continuity, on this reading, enters through the group law governing accumulated transport, not through arbitrarily fine local distinction. The minimum-distinguishability scale therefore bears on the wrong axis to be an obstruction: phase here is accumulated holonomy structure, not a local label, and the scale constrains observability, not the catalogue. The obligation, correctly stated, is thus: *does consistency of holonomy under unbounded composition force the phase group to be continuous U(1)?* — a global-transport question, with the minimum-distinguishability scale explicitly set aside as a category error rather than answered. (This is the form the companion's argument must take given the name of its result; whether that argument closes is the deferred step reported above, not settled here.)

This form is, moreover, the *favourable* one for the generation chain specifically. The chain needs not merely continuous U(1) but continuous U(1) **forced from D** (§17.4, the forcing that does double duty). A composition-consistency route delivers continuity in exactly that form: if continuity follows from consistency requirements on transport built out of admissible commitments, then it is forced *by the admissible structure itself*, which is precisely D-determination. A local-density route, by contrast, would have made continuity a fact about resolution and might have yielded continuity-without-D-determination — one of the two near-misses ruled out above. So the global-consistency reading does not merely rescue continuity from the minimum-scale objection; it is the reading under which closing the continuity step closes *both* hats at once.

Status. *[Proven]* that the audit reduces to a single question about θ — the corpus builds s from θ alone (the Double Square kernel). *[Proven]* that θ is not D-determined by inspection of the construction, so supervenience is not free (holonomy is the witness: distinct holonomy classes, one metric). *[Proven]* that the proposition discharging both chains is specifically *D forces continuous U(1) phase*, and that neither forced-discrete nor continuous-free would suffice for both. *[Proven]* that the minimum-distinguishability scale is no objection to continuous phase: finite distinguishability limits phase *observability*, not the phase *catalogue*, and phase here is accumulated transport holonomy, not a local distinguishability coordinate — so the open proposition is a global-consistency claim (does holonomy-consistency under unbounded composition force U(1)?), on which the scale does not bear. *[Conditional]* that supervenience holds — conditional on the forcing proposition. *[Inherited, at the companion's stated status]* the partial standing of that proposition: finite-holonomy impossibility established in the companion, continuity-and-U(1) deferred there. *[Open]* that one proposition, and nothing else: the dyadic terminus and the Born-rule terminus are the same open proposition, best read as the global-transport question above. The earlier "audit of s" is therefore not a separate obligation — discharging the companion's continuity-and-U(1) step discharges it.

Everything else in the chain is discharged. The squeeze reduces to uniformity (§6); no-vacuous-steps is proven (§10); the multi-Fold collision is a conflation and does not live in closure (§§11–12); the count lives at the candidate menu (§12); capacity is non-atomic and sibling-blind (§13); frontier-uniformity is a symmetry fact, not a completion fact, so the order-loop is removed (§§15–16); no-starvation is equivariance read on the rule, not a separate posit (§17); the apparent hidden-overlap loophole reduces, biconditionally, to a single supervenience fact (§17.1); and that fact, named as the no-extra-support-structure condition, makes dyadic loading a *conditional theorem* (§17.2). What remains is one inspectable, falsifiable property of the construction of s — whether it carries any degree of freedom beyond the distinguishability data the refinement rule reads — and the whole of "why powers of two?" closes the moment that property is confirmed.

Status. *[Proven]* that starvation of genuinely-isomorphic siblings is a disposition-level equivariance violation (§17), that capacity is sibling-blind and cannot ground the asymmetry (§13), that overlap cannot break symmetry *within* an isomorphism class (§17.1), and that enriching the definition of isomorphism cannot close P2 for free (conservation of loopholes, §17.1). *[Proven, biconditional]* that P2 closes **iff** the support profile supervenes on the refinement-relevant data (§17.1), and *[Proven, conditional]* the Conditional Support-Closure Theorem: premises 1–3 $\Longrightarrow$ dyadic loading (§17.2). *[Default of the programme, burden on the challenger]* the no-extra-support-structure condition: by no-pre-individuation an independent support degree of freedom is prohibited unless admissible (hence already in D), so supervenience is the presumption and the residue is a challenge — exhibit a concrete admissible datum carried by s and not by D (§17.3). *[Open]* only the formal audit: whether the *constructed* s covertly imports a convention beyond D — not whether such a datum may exist in principle (the principle says not). *[Inherited as a requirement of no-pre-individuation; F-construction-audit open, parallel to s]* the equivariance principle itself, per §§15B, 16.

Movement III — Position, Consequences, Status

18. Consequences for Generation Structure and the Dyadic Chain

If dyadic loading holds — and under the programme's default reading it does, since by no-pre-individuation the no-extra-support-structure condition is the presumption and the Conditional Support-Closure Theorem fires (§§17.2–17.3) — refinement branches binarily, generation depth carries loads 1, 2, 4, 8, …, and premise **P2** of the companion species paper's generation-count candidate is discharged, *derived for arbitrary depth* rather than read off the three-term $\mathcal{D} = \mathrm{diag}(1, 2, 4)$. With cumulative load $G_n = \sum_{k=1}^{n} 2^{k-1} = 2^n - 1$ and a closure register of capacity K = 7, the admissibility condition $G_n \leq K$ gives $G_1 = 1$, $G_2 = 3$, $G_3 = 7$, $G_4 = 15 > 7$ — exactly three generation-depth sectors admissible. This addresses **only P2**; three generations still require P1 (a generation-blind closure budget K = 7) and P3 (shared occupancy, so loads add). The full chain:

```
Why exactly three generations?           (species paper §9.1)
   ↓  given P1 (K = 7, flavour-blind) + P3 (shared occupancy)
Why dyadic loading 1, 2, 4, 8, … ?         (P2)
   ↓  §6  — the squeeze's two bounds each reduce to uniformity
Why uniform single-Fold refinement?        (the URH, §7)
   ↓  §8  — closure conditions: no-vacuous ∧ no-multi-Fold ∧ frontier-uniformity
   ├─ no-vacuous-steps:    [Proven] (§10) — commitment + emergent order
   ├─ no-multi-Fold:       not a closure collision (§11); count at candidate menu (§12);
   │                       capacity non-atomic, quantisation → integers only (§13)
   └─ frontier-uniformity: looked order-laden (§15) — BUT
                           equivariance derives all-or-none from label-blindness,
                           not simultaneity → order-loop removed (§16)
   ↓
RESIDUE (§17):  starvation = standing disposition-asymmetry ⟹ equivariance violation
                (NOT a separate condition; no-starvation is a theorem, not a posit)
   trichotomy on what could ground "refine A, never B":
     geometry → admissible invariant → A ≇ B (premise fails, not equivariance)
     capacity → sibling-blind scalar (§13) → cannot ground it alone
     nothing  → starvation impossible
CONDITIONAL SUPPORT-CLOSURE THEOREM (§17.2):
   if F = F(D) and s = s(D) with no support DOF beyond D,
   then F-siblings are support-isomorphic, supervenience holds, and
     Fold-binarity (B) + equivariant refinement + common origin of F and s
        ⟹  N_n = 2^{n-1}
```

```
   — proves dyadic loading RELATIVE TO a no-extra-support-structure
condition;
     refuter: a construction of s giving D-identical sectors distinct
profiles
   premises 2–3 (s = s(D), no extra DOF) are the condition — and by
     no-pre-individuation (§9) they are the DEFAULT, not a hope:
       an independent support DOF is surplus structure → prohibited
         unless admissible → then already in D → supervenience anyway
   burden on the challenger (§17.3): exhibit an admissible datum in s but not
D
       — any candidate is inadmissible (forbidden) or admissible (already in D)
   genuinely open: only the audit of the CONSTRUCTED s (covert convention?),
     not whether an extra DOF may exist in principle (principle says not)
```

Two features. First, the binarity dependency (B) — "why 2 and not 3" — is untouched throughout; it remains the Fold's, at the §2 status, and even the exclusion of ternary branching in §5 is (B)'s work, not the upper bound's. Second, the terminus is a *conditional theorem* (§17.2): Fold-binarity + equivariant refinement + the common distinguishability origin of F and s ⟹ dyadic loading. The third input is a named condition — that s carries no degree of freedom beyond the data F reads — and by no-pre-individuation (§17.3) it is the programme's *default*, not a hope: an independent support degree of freedom is surplus structure the substrate prohibits unless it is itself admissible, in which case it is already in D. The burden thus falls on any challenger to exhibit a concrete admissible datum carried by s and not by D; what stays open is only a formal audit of the *constructed* support functional. So the result is not a fork, nor an open question with a balanced prior, but a theorem holding under the programme's default reading, awaiting a construction-audit whose presumption favours it.

19. Relation to the Companion Species Paper

This paper is the expansion of §9.2 of *Particle Species as Representation Classes of Persistent Fold Defects*. There, the dyadic-loading premise (P2) was noted to reduce to a binary-branching question, with a two-sided squeeze sketched and a uniform-branching lemma flagged as the gap the squeeze hides. This paper develops that sketch to a stronger conclusion: both squeeze bounds reduce to uniformity (§6); uniformity's lower half is proven (§10); its multi-Fold half is shown not to collide with non-separability and not to live in closure (§§11–12); and its frontier-uniformity half — the genuine gap the species paper flagged — is supplied by isomorphism-equivariance (§16), which — read on the rule rather than the limit — also makes no-starvation a theorem rather than a posit, turning P2 into a conditional theorem: Fold-binarity + equivariant refinement + the common distinguishability origin of F and s ⟹ dyadic loading (§17.2).

The species paper should be read as inheriting this result at the status given here: P2 is a **conditional theorem** (Conditional Support-Closure Theorem, §17.2) — Fold-binarity + equivariant refinement + the common distinguishability origin of F and s ⟹ dyadic loading — with the binarity that makes the target 2 rather than 3 **inherited** from the Fold (§2), no-vacuous-steps **proven** (§10), no-starvation reduced from a separate posit to a **consequence of equivariance** (§17), and the one remaining hypothesis named, and by no-pre-individuation (§17.3) the *default* of the programme rather than a balanced open question — an independent support degree of freedom would be surplus structure the substrate prohibits, so the burden lies with any challenger to exhibit it. Each step of the chain is a genuine reduction; the terminus is not an open question but a theorem holding under the programme's default reading, its one remaining task the same proposition the Born-rule chain awaits — that finite distinguishability forces continuous U(1) phase — on which the supervenience of s on D, and with it dyadic loading, jointly turn (§17.4).

20. What the Paper Establishes

Established (with conditions named):

- *Movement I.* The squeeze's **two bounds each depend on uniformity** (§§4–5); uniformity is the sole hinge (§6); dyadic loading (P2) is **reduced without residue to the URH together with the Fold's inherited binarity**, with closure conditions *no-vacuous ∧ no-multi-Fold ∧ frontier-uniformity* (§§6–8). The exclusion of *ternary* branching is (B)'s work, not the upper bound's (§5).
- *Movement II.* **No-vacuous-steps is [Proven]** (§10). **The decomposition collision is a conflation** (§11): binary-decomposability ≠ local-decomposability; the Bell basis decomposes into joint commuting binary Folds while remaining non-locally-decomposable; record-orthogonality, forced by irreversibility, yields binary-decomposability generally; **admissibility closure does not force a non-decomposing joint primitive**. **The dyadic count lives at the candidate layer** (§12): commitment selects (→ 1) rather than branching, and the menu is non-orthogonal capacity carvings with no record basis. **Capacity is non-atomic** (§13): a Capacity Quantisation Theorem is likely refutable by resolution-sweep and would yield only integer loading; doubling requires a complete binary tree. **Doubling needs frontier-uniformity, unsupplied by FP1** (§14). **The apparent order-loop in frontier-uniformity is [defeated]** (§§15–16): the worry was never that refinement needs *sequencing* (substrate ticks supply that pre-commitment, without circularity, §9, §15A) but that *co-level completion* looked to need physical temporal order; isomorphism-equivariance derives all-or-none from label-blindness, invoking neither physical time nor co-level closure, so the loop is removed rather than relocated, and the question of whether tick-accumulation forces *complete* (vs arbitrary) refinement need never be answered — it coincides, in any case, with the supervenience hinge below (§15A). The equivariance principle is plausibly inherited from no-pre-individuation. **No-starvation is a theorem, not a posit** (§17): permanent starvation of an isomorphic sector is a *standing disposition-asymmetry* — a difference in the refinement rule as a function of admissible data — which equivariance forbids directly, at the level of the rule rather than in the limit. A trichotomy closes the case: a standing "refine A, never B" must be grounded in geometry (then A ≇ B, the premise fails honestly), in capacity (but capacity is a sibling-blind scalar, §13, and cannot ground it), or in nothing (then starvation is impossible). The only place an admissible sibling-distinguishing datum could hide is the candidate-layer support functional. **The hidden-overlap loophole reduces to a biconditional, and thence to a conditional theorem** (§17.1–17.2): since the equivariance σ preserves s, overlap cannot break symmetry *within* an isomorphism class; counting s as admissible structure only relocates the content into the definition of isomorphism; and enriching that definition is no free lunch (conservation of loopholes). The substantive content is whether refinement-siblinghood and support-

isomorphism coincide, and the **Conditional Support-Closure Theorem** states it: if F and s are both functions of the one distinguishability datum D, with no support degrees of freedom beyond D, then F-siblings are support-isomorphic trivially, no-starvation follows by equivariance, and **Fold-binarity + equivariant refinement + common distinguishability origin of F and s ⟹ dyadic loading**. This proves dyadic loading not absolutely but *relative to a no-extra-support-structure condition*, with a named refuter.

Not established (open or out of scope):

- **the no-extra-support-structure condition** (§§17.2–17.4, premises 2–3 of the Conditional Support-Closure Theorem) — *the one remaining hypothesis, and by no-pre-individuation the programme's default rather than a balanced open question*. An independent support degree of freedom E with s = s(D, E) is surplus structure the substrate prohibits (§9) unless E is itself admissible — in which case E ⊆ D and supervenience holds anyway; the burden lies with the challenger. The audit of the *constructed* s, moreover, can be carried out (§17.4): the corpus builds s from the geometric phase θ, and the audit's verdict is that θ is not D-determined by inspection (the metric is symmetric, θ carries a holonomy it does not contain), so supervenience holds iff θ is *forced* from D. This is the **same proposition the Born-rule derivation awaits** — that finite distinguishability forces continuous U(1) phase — the shared content being the forcing, which gives the Born chain its quadratic-probability continuity and the generation chain its D-determination at once (neither forced-discrete nor continuous-free serving both). So the lone open item is one identified proposition shared across two arcs of the programme, reported by the companion as established for finite-holonomy impossibility and deferred for the continuity-and-U(1) step. The natural next target — and not a *separate* obligation: discharging the companion's step discharges the audit;
- the **equivariance principle** — *derived*, not posited: §15B shows it is the operational form of no-pre-individuation (a non-equivariant rule must read surplus structure or fail to be a function of admissible data), so it is [Inherited] as a requirement, with the only residue the F-construction-audit (whether the rule as built imports a presentation-dependent choice) — [Open], exactly parallel to the s-audit of §17.4;
- the **binarity of the Fold** (B) if the architecture fixes only minimality and non-degeneracy (§2) — the seed-level dependency;
- the **generation count**, **P1**, **P3**, **closure capacity**, **shared occupancy** — out of scope.

The honest summary: this paper does not derive dyadic loading absolutely. It reduces it without residue to a single uniformity hinge, proves the half of that hinge which proves, dissolves a collision, traces the surviving half down through the candidate menu and the non-atomic, sibling-blind scalar capacity to the frontier-uniformity requirement — where an isomorphism-equivariance argument removes the order-loop (on an axis orthogonal to sequencing, so the role of substrate ticks need never be settled), the same argument (read on the rule) makes no-starvation a *theorem* rather than a posit, and the apparent hidden-overlap escape reduces, biconditionally, to a single condition stated as the **Conditional Support-Closure Theorem**. The upshot is a conditional theorem — *Fold-binarity + equivariant refinement + the common distinguishability origin of F and s ⟹ the powers of two* — whose one condition is, by no-pre-individuation, the programme's *default reading*: supervenience holds unless a challenger can exhibit surplus support structure the substrate prohibits. The audit of that condition, carried out on the explicitly-constructed s (§17.4), resolves it to a single proposition the programme already needs elsewhere: that finite distinguishability *forces* continuous U(1) phase. "Why powers of two?" and "why $P = |\psi|^2$?" are thereby revealed as one open question — the forcing of the phase from D — wearing two hats. That conversion — from a diffuse "why doubling?" to a conditional theorem holding by default, whose sole remaining obligation coincides with the Born-rule chain's own — is the contribution.

21. Conclusion

The recurrence of dyadic structure across the programme — generation loads 1, 2, 4; the closure saturation 1 + 2 + 4 = 7 — motivates asking whether binary refinement is fundamental or assumed. The answer reached here is *neither yet*, but the question has a single, sharply-locatable hinge, and that hinge has now been traced to one ordinary capacity condition.

The natural argument is a two-sided squeeze, and both halves rest on the same premise — that each step is exactly one genuine minimal-distinction operation, applied uniformly. The Step Individuation Theorem would have made that premise a consequence. Taken to the bottom it does not, but it shows where the work is. The lower half is settled: refinement cannot idle, because a step is a commitment, a commitment fixes a new fact, and order — being emergent from commitment — leaves no background along which depth could advance while nothing is fixed. The multi-Fold half does not collide with non-separability: a joint distinction cuts binarily even when it does not factor across subsystems, the Bell basis shows it, and the mutual exclusivity of committed records makes binary-decomposability generic. Closure does not host the residue.

The real work was the third condition — that when refinement advances, it advances across the whole frontier rather than depth-first into one sector. Read as a demand for simultaneity, this looked order-laden, and order is built from commitment: a loop. But the demand is not for simultaneity. If the live sectors are genuinely alike — no admissible invariant separating them — then a rule that refined one and spared its twin would be reading a label that is not there. Refinement is all-or-none across alike sectors because it is *blind to inadmissible labels*, not because the carvings happen at one instant. The loop is removed, not relocated: frontier-uniformity is a symmetry fact.

What seemed to remain — whether any sector falls permanently behind — is not a separate capacity question at all, and the way it dissolves is the sharpest step in the argument. Starvation sounds passive, but to keep one sector refining while an isomorphic sibling never refines, the rule must *stand* in a different relation to the two: act on A, deny B, as a policy. If A and B are genuinely alike, that standing policy is a difference in the rule as a function of admissible data — exactly what equivariance forbids. The violation is not in some far-off limit; it is in the policy itself, at the first asymmetric disposition. So permanent starvation of an isomorphic sector is not an additional condition to be discharged; it is an equivariance violation, and no-starvation is a *theorem*, not a posit. Where could the asymmetry come from? From geometry — but then the sectors were never isomorphic. From capacity — but capacity is a bare scalar that does not know which sibling is which, and so cannot ground an asymmetric allocation without a label the substrate forbids. From nothing — and then it cannot stand. The trichotomy is exhaustive, and it leaves exactly one hiding place for an admissible distinguishing datum: the overlap structure of the candidate menu.

So the chain closes on itself, twice over. Equivariance removes the order-loop that threatened the whole descent; and equivariance, read on the rule rather than the limit, forbids the starvation that looked like the last open condition. One hiding place remains — that the would-be twins differ in how they overlap the rest of the menu — and here the final turn is to see what kind of question it is. Overlap cannot break the symmetry from *inside* an isomorphism class, because the symmetry preserves the overlap structure too; so the tempting move is to fold overlap into the definition of "alike" and be done. But that is a free lunch that isn't free: the richer the notion of alike, the fewer sectors count as twins, and the weaker the fairness rule becomes — tightening the definition to forbid hidden overlap permits selective refinement of the newly-non-twins by exactly the same stroke. The loopholes are conserved against each other, so no definition settles it. What settles it is a fact about how the overlap structure is *built*: if it is only an echo of the features refinement already responds to, it adds no new way to tell sectors apart and the doubling is forced; if it is an independent ingredient the theory tracks on its own, that independence is the last gap. This is the Conditional Support-Closure Theorem: if the refinement rule and the overlap structure are both functions of the one distinguishability datum, with no overlap degrees of freedom beyond it, then sectors alike to the rule are alike in overlap trivially, no-starvation follows, and the powers of two follow with them. One cannot reach this by *defining* "alike" to include overlap, because the fairness rule binds only between sectors the refinement rule cannot tell apart; what matters is whether that rule already takes overlap into account, not how we elect to draw the word "twin." The definition is free; the coincidence of what refinement reads with what overlap encodes is the theorem. So the result is not a derivation of dyadic loading absolutely, but a derivation *under the programme's default reading* — and that reading is not a hopeful guess but what an inherited principle delivers. By no-pre-individuation, the substrate carries no structure beyond admissible distinguishability data; an independent overlap degree of freedom would be exactly the surplus structure the principle prohibits, admissible only if already part of the distinguishability data — in which case the overlap supervenes after all. Supervenience is therefore the default, and the burden falls on anyone who would deny it: to do so is to name a concrete admissible feature the overlap carries that refinement does not already read, and the principle makes that hard, since any candidate is either forbidden or already counted. What stays genuinely open is not whether such a feature might exist — the principle says not — but a formal audit of the overlap structure as actually constructed, checking that no convention has been quietly promoted to a physical datum; and the presumption going into that audit favours the theorem.

That audit can in fact be carried out, and it ends somewhere telling. The corpus builds the overlap from a single object, the geometric phase of the Double Square kernel; so the whole question is whether that phase is read off the distinguishability data or stands free of it. By inspection it stands free — the distinguishability datum is a symmetric distance, while the phase is oriented and carries a holonomy the distance does not contain — so supervenience is not had for nothing. It holds exactly if the phase is *forced* from the distinguishability data. And that is the very proposition another arc of the programme already needs: the Born-rule derivation requires the same phase to be forced into continuous U(1) form to make probability come out quadratic. The two needs are one need, because *forcing the phase from the data* delivers both at once —

continuity to the probability rule, determinacy to the overlap. "Why powers of two?" and "why $P = |\psi|^2$?" are therefore not two open questions but one, wearing two hats; and that one question — does finite distinguishability force continuous U(1) phase? — is, by the companion programme's own account, established for the impossibility of a merely finite phase and still owed for the step to continuity.

Refinement does not double because we observed three generations.

Refinement cannot idle: a commitment that commits nothing is not a commitment.

Refinement cannot bundle in any way that matters: joint distinctions cut binarily, and committed facts are exclusive.

Refinement cannot play favourites among alike sectors: a rule blind to inadmissible labels refines all or none — fairness between indistinguishables, an axis orthogonal to sequencing, so the question of whether ticks supply frontier-completion (let alone physical time) need never be answered.

Refinement cannot starve a sector forever: a standing policy "refine A, never B" between siblings is a difference in the rule with no admissible datum to ground it — equivariance read on the rule itself.

Overlap cannot break the symmetry from inside, cannot be closed by redefinition, and — by no-pre-individuation — cannot carry an independent degree of freedom without naming surplus structure the substrate forbids. So the burden falls on anyone who would deny that overlap is built from the same distinguishability the rule already reads.

The whole of "why powers of two?" closes as a conditional theorem holding by default: Fold-binarity, equivariant refinement, and the common distinguishability origin of the refinement rule and the overlap structure together force the powers of two — and that common origin is what no-pre-individuation already entails, surplus support structure being prohibited unless admissible and so already counted. The one place the default must still be earned is the construction of the overlap, which the corpus builds from the geometric phase; the audit there resolves to a single proposition — that finite distinguishability forces continuous U(1) phase — and that proposition is not ours alone: it is exactly what the Born-rule derivation awaits. "Why powers of two?" and "why $P = |\psi|^2$?" are one open question wearing two hats. The recurrence of powers of two is the shadow of that single forcing, and the work that closes it closes the Born rule with it.

22. Status of Claims

Claim	Status
Fold as minimal admissible distinction	Inherited
Fold binarity (exactly two outcomes), dependency (B)	Inherited (§2) if architecture fixes two-valuedness; otherwise

Claim	Status
	an assumption — seed-level dependency, flagged
Irreversible commitment; emergent order; no-pre-individuation; two layers; scalar/additive/resolution-dependent capacity; FP1 dynamics-silent	Inherited (QM reconstruction, NPI / Seal-Trichotomy)
Lower bound ($b \geq 2$) depends on no-vacuous-steps	Conditional on uniformity (§4)
Upper bound ($b \leq 2$) depends on one-step-one-Fold; ternary exclusion is (B)'s work	Conditional on uniformity (§5); ternary exclusion Inherited via (B)
Consolidation: uniformity is the sole hinge of the squeeze	Proven (structural, §6)
Dyadic loading (P2) reduced without residue to URH ∧ (B); closure conditions no-vacuous ∧ no-multi-Fold ∧ frontier-uniformity	Proven as a reduction (§§6–8)
No-vacuous-steps (lower half)	Proven, conditional on commitment-as-step and emergent order (§10); idle step incoherent
"Refinement step = commitment event" as the load-bearing identification	Inherited reading; flagged
Decomposition collision is a conflation (binary ≠ local decomposability)	Proven (§11)
Bell basis decomposes into joint commuting binary Folds (ZZ, XX) while non-locally-decomposable	Proven (worked construction, §11)
Record orthogonality forced by irreversibility ⟹ binary-decomposability generic	Proven, relative to inherited irreversibility
Admissibility closure does NOT force a non-decomposing joint primitive	Proven (§11)
Dyadic count N_n lives at candidate menu, not record layer; commitment selects (→ 1)	Proven, relative to corpus characterisation (§12)
Capacity Quantisation Theorem (Vol_op intrinsically Fold-atomic)	Likely refutable by resolution-sweep (§13)
Integerisation (committable capacity whole-Fold-valued) ⟹ integer loading	Conditional; yields integers, not powers of two (§13)
Doubling requires complete binary tree (completeness + binarity + frontier-uniformity)	Proven (decomposition of the requirement, §14)
Frontier-uniformity not supplied by FP1; first two URH conditions do not force it	Proven (§14)
Substrate sequencing (ticks, pre-commitment) vs physical temporal order (post-commitment); circularity needs the latter	Inherited (§9); distinguishes the two notions of "order"

Claim	Status
Frontier-uniformity needs not mere sequencing (ticks supply that) but *co-level completion*; only that stronger notion looked order-laden	Conditional on the co-level-completion reading (§15)
Tick-accumulation (TPB) may supply sequencing sufficient for pre-commitment refinement, but sequencing alone does not force *complete* (vs arbitrary) refinement	Proven that sequencing ≠ completeness (§15A)
"Do ticks force complete refinement?" coincides with the supervenience hinge (frontier-completeness and sibling-supervenience hold or fail together)	Proven coincidence (§15A); not an independent fork
Isomorphism-equivariance ⟹ all-or-none across alike sectors (selective refinement inadmissible)	Proven, relative to: isomorphic siblings, equivariant rule, no symmetry-breaking selector (§16)
All-or-none from label-blindness, NOT simultaneity ⟹ order-loop removed, not relocated	Proven (§16) — the materially new step
Equivariance derived from no-pre-individuation: a selective rule on isomorphic sectors must read an admissible invariant (then not isomorphic), an inadmissible label (forbidden), or nothing admissible (not a function of admissible data; stochastic escape closed — σ-invariant statistics still need a distinguishing datum to bias)	Proven, three-possibility argument (§15B)
Equivariance principle (label-blind refinement rule)	Derived from no-pre-individuation as a *requirement* (§15B, three-possibility argument); residue is the F-construction-audit, Open and parallel to the s-audit (§17.4)
Ragged-frontier induction: it collapses into equivariance — a frontier going ragged via starvation is the standing asymmetry §17 forbids, not an independent induction	Proven, relative to §17
Layer check now revealed as the *whole* residue, not one risk among conditions	See lone open question below (§17)
Synchronic/diachronic fork dissolved; order-loop removed	Proven (§§16–17)
Starvation = standing disposition-asymmetry (a difference in the rule on admissible data), present at the first asymmetric disposition, not in a limit	Proven, relative to equivariance read on the rule (§17)
No-starvation is a theorem, not a posit: starvation of isomorphic siblings violates equivariance directly	Proven, conditional on genuine sibling isomorphism at the decision layer (§17)
Transient (presentation-order) vs standing (policy) asymmetry: only standing forbidden	Proven; necessary or refinement could not occur (§16–17)

Claim	Status
Trichotomy on the source of a standing asymmetry: geometry ($\Longrightarrow$ not isomorphic) / capacity (sibling-blind scalar, cannot ground it) / nothing ($\Longrightarrow$ impossible)	Proven, exhaustive (§17)
Capacity cannot ground asymmetric sibling allocation (bare additive scalar, no sibling index)	Proven, relative to §13 characterisation
Overlap cannot break sibling symmetry from inside an isomorphism class (σ preserves s)	Proven (§17.1)
Counting s in the definition of isomorphism only relocates the content; enriching the definition is no free lunch (conservation of loopholes)	Proven (§17.1)
Biconditional terminus: P2 closes $\Longleftrightarrow$ support profile supervenes on refinement-relevant data (refinement-siblinghood = support-isomorphism)	Proven, biconditional (§17.1)
Conditional Support-Closure Theorem: if F = F(D) and s = s(D) with no support DOF beyond D, then F-siblings are support-isomorphic, no-starvation follows, and Fold-binarity + equivariant refinement + common origin of F,s $\Longrightarrow N_n = 2^{n-1}$	Proven, given premises 1–3 (§17.2)
Premises 2–3 (s = s(D); no extra support DOF) **are** the supervenience condition stated as hypothesis, not derived	Noted explicitly (§17.2)
Dyadic loading proven *relative to a no-extra-support-structure condition*, not absolutely	Conditional theorem (§17.2)
Refuter, named: dyadic loading fails iff the construction of s gives two D-identical sectors distinct support profiles	Proven as the falsifier (§17.2)
No-pre-individuation makes the no-extra-support-structure condition the programme's *default*: an independent support DOF is surplus structure, prohibited unless admissible (hence already in D)	Proven from inherited no-pre-individuation (§17.3)
Burden on the challenger: to deny supervenience, exhibit a concrete admissible datum carried by s and not by D, with why F cannot read it	Stated as the residual challenge (§17.3)
Audit of the constructed s carried out (§17.4): corpus builds s from the geometric phase θ alone (Double Square kernel)	Proven (§17.4)
θ not D-determined by inspection (symmetric metric vs oriented phase carrying holonomy) — supervenience not free	Proven (§17.4)
Proposition discharging both chains is *D forces continuous U(1) phase* (the forcing); forced-discrete serves only generation, continuous-free serves only Born	Proven that forcing does double duty; neither near-miss suffices (§17.4)
Lone open proposition: finite distinguishability forces continuous U(1) phase — shared with the Born-rule chain; dyadic terminus = Born-rule terminus	Open; *[Inherited, at companion's status]* finite-holonomy impossibility established there, continuity-and-U(1) deferred (§17.4)

Claim	Status
Minimum-distinguishability scale is no objection to continuous phase: it limits phase *observability*, not the phase *catalogue* (ruler/pixel: instrument resolution ≠ quantity granularity)	Proven (§17.4)
Open proposition correctly read as global-consistency, not local-density: phase is accumulated transport holonomy, so continuity (if forced) comes from composition-consistency under unbounded commitment — which delivers the *forced-from-D* form both chains need	Proven that this is the right form; the forcing itself is the deferred step (§17.4)
Generation count; P1, P3; closure capacity; shared occupancy	Out of scope