Binary Refinement from Minimal Distinction

Why Dyadic Loading Reduces to a Single Uniformity Hypothesis

Keith Taylor

VERSF Theoretical Physics Programme

General Reader Summary

The puzzle this paper takes on

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 (developed in the companion species paper) needs exactly those loads to make "three generations" come out of a fixed budget. So the powers of two are doing 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 asks whether the doubling is forced, or merely assumed.*

The answer this paper reaches is honest but partial, and it is worth stating the partiality up front because it is the whole point. There is a clean intuitive argument that refinement *must* double: built from the Fold — the smallest possible distinction, which has exactly two sides — refinement cannot grow slower than doubling without failing to actually distinguish anything new, and cannot grow faster than doubling without sneaking in more than one distinction at a time. Squeezed from both sides, the branching factor would be exactly two.

But when the two halves of that squeeze are examined closely, the work divides between *two* hidden assumptions, not one. The first is that the Fold really is binary — has exactly two sides — for this is what rules out a *three*-way branch (a "minimal distinction" with three outcomes would either not be minimal or not be a single Fold); call this **binarity**. The second is that each step of refinement is *one* clean Fold, the same operation once at every level — for without it a sequence could alternate "double, double, do-nothing" and dodge the lower argument, or bundle two Folds into one step and dodge the upper; call this **uniformity**. The two do different jobs: binarity decides that the branch is a *power of two* (two, four, eight — not three or five); uniformity decides that the power is the *first* one (two, not four or eight). Neither alone gives doubling. So the paper's real finding is not "doubling is proved" but something sharper: *the vague question "why powers of two?" splits cleanly into two precise questions — why binary, and why uniform — each with a clear home.* Binarity is inherited from the Fold; uniformity is the open hinge this paper isolates, with a clear statement of what proving it would require.

Abstract

Several branches of the VERSF programme repeatedly 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. At present dyadic growth is an **assumption**, not a derived result. This paper asks whether binary refinement can be forced from more primitive principles, and reaches a result that is conditional in a single, sharply-located way.

The natural argument is a two-sided squeeze. *Sub-binary* refinement appears to fail distinguishability preservation (a step that branches by fewer than two realises no new distinction); *super-binary* refinement appears to violate minimality (a step branching by more than two encodes more than one independent distinction, hence more than one Fold). Between them, the branching factor would be forced to **2**.

The central finding of this paper concerns *how* the squeeze depends on its premises, and it is sharper than "uniformity is the one hinge." The result rests on **two** hinges that divide the labour. The **Fold's binarity** (B, §2) answers *why a power of two at all* — it is what excludes b = 3 and every non-power-of-two, since a ternary "minimal distinction" either falsifies binarity or is inadmissible by it; uniformity contributes nothing to this, because a ternary branch is not a multi-Fold step. **Uniformity** (U) answers *why the first power rather than a higher one* — the lower bound excludes the sub-unit (vacuous) case and the upper bound's composite part excludes b = 4, 8, … (multi-Fold steps), both being uniformity conditions. Neither hinge alone gives b = 2: binarity without uniformity permits any power of two (via composite or vacuous steps); uniformity without binarity permits b = 3 (a single non-binary distinction). Only (B) ∧ (U) forces b = 2. So the squeeze adds no independent content beyond these two: *(B) selects the base, (U) selects the exponent.*

This does not derive dyadic loading. It reduces it to the conjunction **(U) ∧ (B)** — the **Uniform Refinement Hypothesis** together with the **Fold's inherited binarity** — with a substrate-level motivation for (U) (refinement-by-minimal-distinction is, by construction, the repeated application of one Fold) and the standing of (B) flagged as architecture-dependent (§2: inherited if the fold architecture fixes two-valuedness; this-paper-assumption if it fixes only minimality). Both conjuncts are open in the relevant sense — (U) as a derived result, (B) conditionally on the architecture — so the honest statement is *reduced to (U) ∧ (B)*, not "reduced without residue to (U) alone." The URH is moreover *physically expected though not logically necessary* (§7.1): violating it means denominating refinement in a non-Fold unit — a variable-length metre — a natural assumption with a structural reason, not an arbitrary one. The paper identifies the candidate result that would *dissolve* (U) entirely — a **Step Individuation Theorem** establishing the Fold as the uniquely admissible refinement unit (§7.2) — as the natural sequel, while noting it too presupposes (B). The dyadic-loading problem — premise P2 of the companion species paper's generation-count candidate — is thereby converted from a diffuse "why powers of two?" into two precisely-stated questions: *why binary* (B, §2) and *why uniform* (U, §§7–8).

*Epistemic markers: **(inherited)** for results imported from prior VERSF papers; **(proven)** for results established here; **(conditional)** for results holding under stated inputs; **(open)** for what remains undecided.*

Table of Contents

---

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 $\mathcal{D}$ = 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. The question is general, not specific to flavour: whenever a structure undergoes repeated refinement, one may ask how many distinguishable sectors emerge per step — the *branching factor* b. The candidate values 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 answer reached here is conditional, and the conditionality is the contribution. There is a clean squeeze argument (§§4–5) that appears to force $b = 2$ from two sides. But the squeeze, examined honestly, rests entirely on a single assumption about how refinement steps are individuated (§6), and the paper's real result is to *reduce* the dyadic-loading question to that one assumption (§7) and state precisely what closing it would require (§8). We do not prove dyadic loading; we show it has exactly one hinge.

---

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 this paper must 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, and it must be either inherited or marked here. If the fold architecture establishes the Fold as two-valued (a single cut, two sides), then 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 **(this paper)**, and it is itself an assumption — a weaker cousin of the very dyadic question under investigation. We flag this because the entire squeeze converts "minimal" into "binary" at this step: 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 not a harmless definition; it 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 **(inherited)** and mark the dependency rather than bury it.

---

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 with depth; 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, giving

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

The question is whether admissible refinement forces branching factor 2. The two following sections give the squeeze — the lower bound ($b \geq 2$) and the upper bound ($b \leq 2$) — and then §6 shows that both, examined carefully, rest on the same hypothesis.

---

4. The Lower Bound — and Its Dependence on Uniformity

The intuitive argument. Suppose a refinement step branches by fewer than two. Then it 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 $b \geq 2$.

Distinguishability bound (conditional)

A refinement step that introduces a genuinely new distinction has branching factor $b \geq 2$, since a new distinction separates at least two alternatives.

Where it leaks. This bound 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 merely 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. And that requirement is not a distinguishability fact; it is a statement about what counts as a step. Concretely, a non-uniform schedule

$b = 2, 2, 1, 2, 2, 1, \ldots$

never violates distinguishability at any *genuine* step, yet averages below 2 and does not produce $N_n = 2^{n-1}$. The lower bound therefore controls the *average* branching only if the schedule contains no vacuous steps — which is a uniformity condition, not a consequence of distinguishability. So the lower bound is **conditional on uniformity** (no vacuous steps), and we mark it so rather than presenting it as an independent half of the squeeze.

We attach no ■: the bound is conditional on an unstated premise (no vacuous steps) that §6 identifies with uniformity, so it is not proved even relative to stated premises until that premise is named.

5. The Upper Bound — and Its Dependence on Uniformity

The intuitive argument. Suppose a refinement step branches faster than binary — say $b = 3$, producing 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 supposed to be minimal, a step contributing more than one Fold violates minimality, so $b \leq 2$.

Minimality bound (conditional)

A refinement step branching by more than two encodes more than one independent binary distinction; if a step is constrained to a single Fold, $b \leq 2$.

Where it leaks — and the leak reveals that the upper bound does less than it appears. Read the Minimal-Distinction Principle (§2) exactly: it says *one Fold* contributes *one* binary distinction. It does **not** say *one step* contributes *one Fold*. Minimality, as defined, is a property of Folds — a Fold cannot be sub-divided — not a property of refinement steps. So a step that branches by 4 does not violate the Minimal-Distinction Principle; it is simply a step composed of *two Folds*. To forbid it, one needs a separate premise: **one refinement step = exactly one Fold**. That premise — one step is one Fold — is the uniformity hypothesis of §7 in another guise (every step implements *the* minimal-distinction operation, one Fold). So for composite **integer-multi-Fold** branching — $b = 4$ (two Folds), $b = 8$ (three Folds), and the rest of the powers of two above 2 — the upper bound is **conditional on uniformity**, exactly as the lower bound is.

But this is only half of what "$b \leq 2$" needs, and the other half is not uniformity's to give. Forbidding multi-Fold steps rules out $b = 4, 8, \ldots$ — the *powers of two* above 2. It does **not** rule out $b = 3$. A ternary branch is not a clean multi-Fold step (it is not two Folds; $\log_2 3$ is not an integer), so "one step = one Fold" does not obviously exclude it. An earlier draft waved ternary away as "one Fold plus a partial Fold, inadmissible for a different reason" — but *a partial Fold is not a notion the substrate supplies*: the Fold is indivisible (§2). So a ternary distinction is one of exactly two things:

- **(a)** an admissible *single non-binary minimal distinction* — in which case the Minimal-Distinction Principle's binarity clause (§2) is false, and the paper's seed is gone; or
- **(b)** inadmissible — in which case *what rules it out is §2's binarity*, that the minimal distinction has exactly two outcomes, **not** the minimality/uniformity bound of this section.

Either way, the exclusion of $b = 3$ is done **entirely by §2's binarity**, with uniformity contributing nothing. This is the honest finding, and it *strengthens* the consolidation rather than weakening it: the upper bound's sole work is forbidding composite integer-multi-Fold steps ($b = 4, 8, \ldots$) via uniformity; the exclusion of $b = 3$ — and of every non-power-of-2 branching factor — traces to the Fold being binary (§2), not to the bound. So the upper bound carries the (B) dependency at least as heavily as it carries uniformity: without (B), "$b \leq 2$" fails at $b = 3$ no matter how cleanly steps are individuated. We attach no ■: the bound is conditional on uniformity (for the composite exclusions) *and* on §2's binarity (for the ternary and non-power-of-2 exclusions), and the latter dependency is the one the earlier draft obscured.

6. The Consolidation: Two Hinges, Not One

The two preceding sections are usually presented — including in an earlier draft of this material — as two independent bounds ($b \geq 2$ and $b \leq 2$) squeezing the branching factor to 2, with a separate, secondary "uniformity" issue left open at the end. That presentation is misleading, and correcting it is this paper's main structural contribution — but the correction is subtler than "uniformity is the one hinge," and getting it exactly right matters because it identifies *two* hinges that divide the work between them.

Examined honestly, the squeeze decomposes not into "lower bound + upper bound" but into two questions with two different answers:

- **Why is the branching factor a power of two at all (rather than 3, 5, …)?** This is answered *entirely by the Fold's binarity* (B, §2). A non-power-of-two branch — paradigmatically $b = 3$ — is excluded because the minimal distinction has exactly two outcomes; a ternary "minimal distinction" would either falsify (B) or be inadmissible *by* (B). Uniformity contributes nothing here (§5): "one step = one Fold" does not exclude $b = 3$, because a ternary branch is not an integer-multi-Fold step.
- **Given that branching is a power of two, why is it 2^1 and not $2^2 = 4$ or $2^3 = 8$?** This is answered *entirely by uniformity* (U). The lower bound (§4) rules out the *sub*-unit case (vacuous steps dragging the average below 2); the upper bound's composite part (§5) rules out the *multi*-Fold cases ($b = 4$ as two Folds, $b = 8$ as three). Both are uniformity conditions — no vacuous steps, one step = one Fold — and (B) contributes nothing to *this* question, since powers of two are already on the table.

So the labour divides cleanly: **(B) selects the base — branching is a power of two; (U) selects the exponent — the power is the first.** Neither alone gives $b = 2$. (B) without (U) permits $b \in \{1, 2, 4, 8, \ldots\}$ (powers of two, but any of them, via composite or vacuous steps); (U) without (B) permits $b = 3$ (a single non-binary minimal distinction, uniformly applied). Only **(B) ∧ (U)** forces $b = 2$.

Consolidation

The squeeze is not "two independent bounds plus a loose uniformity end." It rests on exactly **two** hinges, which divide the work: the **Fold's binarity (B)** forces the branching factor to be a power of two (excluding $b = 3$ and all non-powers, §5), and **uniformity (U)** forces that power to be the first (excluding the vacuous case $b < 2$ and the composite cases $b = 4, 8, \ldots$, §§4–5). Neither hinge alone yields $b = 2$; together they force it. There is no version of the result resting on uniformity alone — an earlier framing that called uniformity "the sole hinge" understated (B)'s role, which the ternary analysis of §5 makes unavoidable.

This is sharper and more honest than "binary is established, uniformity is open," and sharper than the earlier "uniformity is the sole hinge." It says: *why a power of two* is settled by the Fold (B, inherited); *which power* is the open question (U). The squeeze adds no independent content beyond (B) and (U); given both, $b = 2$ is immediate, and the elaborate two-sided presentation collapses to *a uniform refinement built from binary Folds branches binarily, trivially.* The remaining work is establishing (U) — and, to the extent (B) carries the architecture-dependent caveat of §2, confirming (B) as well.

We state the conditional result accordingly, and at its true strength:

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 **two** hinges: (U) **(open**, §7–8**)** — selecting the exponent; and (B) **(inherited**, §2, with the attribution caveat there**)** — selecting the base, and carrying *more* of the load than an earlier draft credited (it alone excludes b = 3). No ■: the result is conditional on (U) and (B), neither fully closed here; it is an implication, not a theorem.

---

7. The Uniform Refinement Hypothesis

The consolidation locates the entire problem in one hypothesis.

Uniform Refinement Hypothesis (URH)

Every admissible refinement step implements exactly one genuine minimal-distinction operation — one Fold — applied to every distinguishable sector of the frontier, and the same operation at every depth. Equivalently: refinement steps are individuated such that each contributes exactly one binary distinction per sector, neither zero (no vacuous steps) nor more than one (no multi-Fold steps), and acts across the whole frontier (no partial-frontier steps), uniformly in depth.

Status: (open).

The substrate-level motivation, stated as motivation and not proof. The URH is not arbitrary; it is what "refinement *by minimal distinction*" most naturally means. If refinement is, by construction, the repeated introduction of the substrate's minimal distinction — the Fold — then each step *is* one Fold, because the Fold is the unit in which refinement is denominated, and there is no sub-Fold step (minimality, §2) and no reason for a step to bundle several Folds rather than apply them as successive steps. On this reading the URH is close to definitional: a refinement step that did nothing would not be a refinement step, and a step that did several distinctions at once would be several steps. So the URH is the *expected* structure of refinement-by-Fold — where a non-uniform alternative would require refinement to be denominated in some unit other than the Fold, which the substrate does not supply.

But — and this is the honest line — "expected because natural" is not "derived." The URH could fail if admissible refinement turns out to permit composite steps (several Folds committed together as one admissible event) or idle steps (depth advancing without a new distinction). Whether the commitment dynamics permit such steps is a question about the substrate's refinement structure that this paper does not settle. So the URH is motivated, not proved, and the Binary Refinement Result inherits that status precisely.

7.1 Why Uniformity Is Expected — Logical Necessity versus Physical Expectation

It would be a mistake to leave the reader thinking the URH is *arbitrary*. It is open — not proven — but very far from arbitrary, and the distinction matters enough to name directly. Ask the URH two questions and they have different answers:

- *Is the URH logically necessary?* No. One can coherently describe refinement schemes that violate it — composite steps, idle steps — without contradiction. Nothing in pure logic forces uniformity.
- *Is the URH physically expected, given the substrate?* Strongly yes — for the reason already given in §7's motivation: refinement is denominated in Folds, and a non-uniform step measures refinement in a unit that varies step to step (the variable-length metre), which is less an alternative dynamics than a refusal to use the substrate's own unit.

So the honest status of the URH is **physically expected but not logically necessary, and open as a derived result** — stronger than "possible assumption," weaker than "theorem." The reader should leave §7 thinking not "URH is arbitrary" but "URH is what refinement-by-Fold almost certainly is, pending a derivation that the Fold is forced as the refinement unit" — which is the candidate of §7.2.

7.1.1 Depth Is Not a Direction

A further motivation for uniformity — and specifically for one of its three conditions — comes from the status of refinement depth itself. The inherited result *Depth Is Not a Direction* **(inherited)** establishes that refinement depth is not a geometric coordinate analogous to a spatial dimension: it has no intrinsic metric, supports no reversible transport, and represents no physical location. Depth indexes successive stages of coarse-graining and descriptive resolution — a hierarchy of descriptions, not a place.

This bears on uniformity in a precise and limited way: it motivates the **no-vacuous-steps** condition (§8, condition 2), and that one specifically. If depth were a direction, one could imagine "empty," "skipped," or "partially occupied" levels — positions along an axis through which something propagates, some of which happen to be unoccupied. But depth is not a direction, so a refinement level has no existence independent of the distinction that constitutes it: a level exists *only* insofar as a new admissible distinction has been committed. On that reading a vacuous step — advancing the bookkeeping index while committing no new distinction — is not merely disallowed by fiat but physically unnatural, because there is no empty level for it to advance into. Depth-being-descriptive is therefore a substrate-level reason to expect *no vacuous steps*.

It is worth being exact about what this does and does not motivate. It supports condition 2 (no vacuous steps). It does **not** by itself support condition 1 (no multi-Fold steps) or condition 3 (no partial-frontier steps) — those concern *how much* a step commits and *across how much of the frontier*, not whether empty levels can exist, and depth-being-descriptive is silent on them. So this subsection strengthens the *no-vacuous-steps* leg of uniformity, not the whole URH; the other two legs rest on the considerations of §§5–6 and the step-individuation candidate of §7.2.

A note on the generation-load reading, to keep this paper within its scope. Once depth is descriptive rather than geometric, the sequence 1, 2, 4, 8, … is naturally read as *distinguishability capacity generated by successive refinement strata*, not as distance along an axis — and the companion species paper's closure-capacity condition 1 + 2 + 4 = 7 is correspondingly read there as exhaustion of admissible capacity within the K = 7 architecture rather than of geometric extent. But that capacity-exhaustion reading presupposes the species paper's *sum-rule* (its premise P3, that the strata loads add within one register), which that paper marks **conditional**. It is out of scope here: this paper concerns why loading is *dyadic* (premise P2), not whether the loads *sum to a budget* (P3). So the depth-is-descriptive point is adopted here only for what it motivates about *this* paper's subject — the unnaturalness of vacuous refinement steps — and the 1 + 2 + 4 = 7 gloss is noted as **(inherited from the species paper; conditional on its P3)**, not asserted.

This does not prove uniformity, and the URH remains open. It explains why *one* of uniformity's conditions — no vacuous steps — is the natural expectation of a substrate in which depth earns meaning only through the distinctions it commits.

7.1.2 From Motivation to Consistency Condition

§7.1.1 motivates no-vacuous-steps. There is a stronger move available, which upgrades two of the three uniformity conditions from separately-assumed regularities into consequences of a single deeper principle — and answers the referee's sharpest question along the way: *why should generation depth correspond to counting distinctions at all?* The answer runs through the inherited depth-is-descriptive result, and is exactly as strong as that inheritance, which we mark.

The argument has two principles. The first answers the referee:

Distinguishability Realization Principle (conditional on the inherited depth-is-descriptive result)

A refinement level exists if and only if it realises at least one new admissible distinction.

The reasoning: depth is not a location (inherited), so depth levels cannot be individuated *geometrically* — there is no axis-position to tell one level from the next. The only thing that can individuate one refinement level from another is therefore the appearance of new distinguishability structure. So a level just *is* the realisation of a new distinction; a putative "level" that realises none is not a level. This is why depth corresponds to counting distinctions: not by stipulation, but because — once depth is descriptive rather than geometric — there is nothing *else* for a level to be individuated by. The referee's question is answered by the inherited result, and inherits its status.

The second principle unifies the upper and lower bounds, which the paper has so far treated as separate pathologies:

Refinement Unit Consistency Principle (conditional on Fold-uniqueness-as-unit, §7.2)

If the Fold is the unique admissible unit of distinction, then a consistent refinement hierarchy must denominate every step in that one unit: exactly one Fold per step.

Here is the unification. The paper has treated $b = 1$ and $b = 4$ as two different failures — a vacuous step and a composite step. They are *the same failure*: a refinement hierarchy that sometimes commits zero Folds per step ($b = 1$) and sometimes two ($b = 4$) is **measuring refinement in a variable-length unit**. It is the metre stick whose length changes between measurements — and just as that makes "distance" ill-defined, a variable Fold-count per step makes "depth" ill-defined as a distinguishability hierarchy. Both $b = 1$ and $b = 4$ violate one principle: refinement must be denominated in a *fixed* unit. Given that the Fold is the unique admissible unit (the §7.2 premise), the fixed unit is one Fold, and both the vacuous (0-Fold) and composite (2-Fold) steps are ruled out *together*, as the single requirement that the refinement unit not vary.

This unification is genuine and physical, and it covers exactly conditions 1 (no multi-Fold) and 2 (no vacuous) — and *only* those two. It does **not** absorb the other exclusions, and saying so keeps the claim honest:

- **$b = 3$ is not a variable-unit pathology.** A ternary step is a *single, consistent, non-binary* unit — a fixed unit, just the wrong size. The Refinement Unit Consistency Principle, being about unit *constancy*, does not exclude it; $b = 3$ is excluded by the Fold being *binary* (B, §2), exactly as established in §5–6. The metre-stick argument unifies 1 and 2; (B) handles 3.
- **Partial-frontier steps (condition 3) are not a unit-size pathology** either — they concern *where* the unit acts, not its size or constancy — so they remain a separate condition, motivated by the frontier-uniformity considerations of §8, not by this principle.

What this upgrade achieves, and its honest ceiling. Together the two principles convert conditions 1 and 2 from assumed regularities into a *consistency condition on what a refinement level is*: levels are individuated only by realised distinctions (Realization Principle), and a hierarchy of such levels must use a fixed distinction-unit (Unit Consistency Principle), so each level commits exactly one Fold — no more (composite), no less (vacuous). The URH stops looking like an arbitrary "every step is the same" regularity and starts looking like *what a distinguishability hierarchy must be to be one*.

But the ceiling must be stated, because the upgrade relocates the open premise rather than removing it. The Realization Principle is conditional on the **inherited** depth-is-descriptive result. The Unit Consistency Principle is conditional on **Fold-uniqueness-as-unit** — that the Fold is the *unique* admissible unit, so there is no second unit to switch to and no determinate-unit-less option. And that uniqueness is precisely the **Step Individuation Theorem candidate of §7.2**, still open. So this route and the §7.2 route are not two independent strengthenings — *they are the same terminus reached from two directions*: both bottom out at Fold-uniqueness-as-the-refinement-unit. The §7.1.2 framing (consistency condition) and the §7.2 framing (uniqueness theorem) are one open question wearing two descriptions. What §7.1.2 adds is the *interpretation*: it shows that proving Fold-uniqueness would not merely "establish a regularity" but would establish that uniformity is constitutive of depth being a hierarchy at all — which is why the URH, though open, is far more than a convenient assumption. It is a consistency condition awaiting one inherited result (depth-is-descriptive) and one candidate theorem (Fold-uniqueness), neither proven here, both named.

7.2 The Step Individuation Theorem (candidate)

That last clause points at something sharper than a motivation: a candidate theorem that, if established, would not merely *support* the URH but *dissolve* it — turning binary refinement into an automatic consequence of the Fold's binarity with no uniformity assumption left standing. It is worth stating explicitly as the natural next target, marked as a candidate and not claimed here. As §7.1.2 showed, this is the same terminus that the consistency-condition route reaches from the other direction: *Fold-uniqueness-as-unit* is what both the "uniformity is a consistency condition" framing (§7.1.2) and the "uniformity is dissolved by a theorem" framing (here) bottom out at. They are one open question under two descriptions.

Step Individuation Theorem (candidate; not established here)

Once refinement is denominated in minimal distinctions, the admissible refinement unit is *uniquely* one Fold: an admissible refinement step neither composes several Folds nor advances vacuously, because the Fold is the unique indivisible unit of admissible distinction and refinement is the commitment of such units one at a time.

If this held, the chain would shorten dramatically. The URH would no longer be a separate open hypothesis — it would be a *consequence* of step individuation — and the whole result would collapse to:

```
Fold is binary  (§2, inherited)
   ↓  Step Individuation Theorem (refinement unit is uniquely one Fold)
binary refinement  (automatic)
```

with no uniformity premise to discharge. Binary refinement would become a direct consequence of the Fold's binarity, full stop.

But the candidate has its own open premise, and naming it is what keeps this from being circular. The seductive form of the argument — "refinement is counted in Folds, so each step is one Fold" — proves nothing on its own, because it relocates the question onto *whether refinement admits a unique denomination, and whether that denomination is the Fold*. That is precisely the thing the Step Individuation Theorem would have to establish, not assume. Concretely, the candidate closes only if the commitment dynamics establish (a) that admissible distinction has a *unique* indivisible unit (no finer denomination than the Fold, and no coarser admissible composite that fails to decompose), (b) that refinement *is* the one-at-a-time commitment of that unit (no admissible refinement event that is not a single such commitment), and (c) that the unit acts across the *whole frontier* at each step (no partial-frontier steps). These are exactly the three sub-claims §8 identifies as the closure conditions of the URH — *no-multi-Fold-steps ∧ no-vacuous-steps ∧ no-partial-frontier-steps* — which confirms the relationship: **the Step Individuation Theorem is what proving §8's three conditions would amount to.** So §7.2 and §8 are not two separate programmes; the Step Individuation Theorem is the *theorem* whose proof is the *task* of §8, viewed from the other side. (And both presuppose (B), §2: they fix that the unit is applied once everywhere, but that the unit is *binary* is inherited, not established by step-individuation.)

Status: (candidate; open). The Step Individuation Theorem is not established here and is flagged as the natural next paper. If proven, it discharges the URH entirely and makes binary refinement automatic given the Fold's binarity. Its own non-circularity hinge is that the uniqueness of the Fold as refinement unit must be *derived* from the commitment dynamics, not read off the assumption that refinement is "counted in Folds" — the same discipline that separates a derivation from a restatement throughout this programme.

---

8. What Would Close the Hypothesis

As with the premises of the companion species paper, the paper cannot perform the derivation (it requires the formal commitment/refinement dynamics) but can state exactly what closing the URH would consist of, so it is a defined task rather than a hope.

The URH closes if and only if the substrate's refinement dynamics establish **step individuation**: that an admissible refinement event is exactly one Fold-commitment — no more, no less. Concretely, two sub-claims, each with a named refuter:

1. **No multi-Fold steps (forces the upper bound).** Show that the commitment dynamics do not admit the simultaneous commitment of two or more independent Folds as a single refinement event — that any such composite necessarily decomposes into successive single-Fold steps. *Refuter:* if the dynamics permit a genuinely simultaneous multi-Fold commitment that does not decompose, super-binary branching is admissible and the upper bound fails.
2. **No vacuous steps (forces the lower bound).** Show that an admissible refinement event necessarily commits a *new* distinction — that depth cannot advance without distinguishability strictly increasing. *Refuter:* if the dynamics permit depth-advance

without a new committed distinction (an idle step), sub-average branching is admissible and the lower bound fails.

If both hold, every step is exactly one Fold, the URH is established, and — *given the Fold's binarity (B, §2)* — the Binary Refinement Result follows. The (B) dependency must be carried here as visibly as §6 carries it: the two conditions below force $b = 2$ *only* given (B). On their own they force b to be a power of two reachable as a single non-composite, non-vacuous step — but "single non-composite step" excludes $b = 4, 8, \ldots$ (composite) and $b < 2$ (vacuous), leaving $b \in \{2, 3, 5, \ldots, \text{any non-composite count}\}$; it is (B) that then removes 3 and every non-power-of-two, by the §5 analysis. So *no-multi-Fold-steps ∧ no-vacuous-steps* forces $b = 2$ **given (B)**, exactly as §6's Binary Refinement Result states with its (B) premise. The §8 conditions are the (U) hinge; they silently presuppose (B) to rule out ternary, and we name that presupposition rather than let the two conditions appear self-sufficient.

A third condition, easily missed, belongs alongside them: **frontier-uniformity.** The conclusion $N_{n+1} = 2\,N_n$ requires not only that each step be one Fold *in depth* (the two conditions above) but that the step act *uniformly across the frontier* — every distinguishable sector at depth n branches, not merely some. "No vacuous steps" forbids a step that does nothing; it does not obviously forbid a step that branches some sectors while idling others, which would also break the clean 2^{n-1} (giving N_{n+1} between N_n and $2\,N_n$). So the URH's "applied uniformly at every depth" must be read to carry frontier-uniformity as well as depth-uniformity:

3. **No partial-frontier steps (forces multiplicative compounding).** Show that an admissible refinement step acts on *every* distinguishable sector of the frontier, not a proper subset — so that the count compounds multiplicatively (×2 across all sectors) rather than additively (a few sectors splitting). *Refuter:* if the dynamics permit a step that branches some sectors and leaves others unrefined, N_n grows sub-exponentially and the dyadic law fails even with conditions 1 and 2 holding.

So the precise statement is: **conditions 1–3 (no multi-Fold, no vacuous, no partial-frontier steps), together with (B), force $N_n = 2^{n-1}$.** Conditions 1–3 are the full content of (U); (B) is inherited (§2). If all hold, dyadic loading is forced for arbitrary depth; if any fails, it is not, and the recurrence of powers of two needs another explanation. The URH is therefore not vaguely "plausible"; it is this precise conjunction, each conjunct a checkable claim about the commitment dynamics with a stated refuter — a task handed to the substrate-dynamics programme.

Proving conditions 1–3 is exactly the content of the **Step Individuation Theorem** (§7.2): condition 1 (no multi-Fold steps) is the uniqueness of the Fold as indivisible refinement unit from above, condition 2 (no vacuous steps) its necessity from below, and condition 3 (no partial-frontier steps) the requirement that the unit act everywhere on the frontier at once. Together they establish that the admissible refinement unit is uniquely one Fold applied across the whole frontier. So the §8 task and the §7.2 candidate theorem are the same result seen from two sides — the task is what one *does*; the theorem is what one *has* when the task is done. Discharging it (given (B)) would dissolve the URH rather than merely supporting it (§7.2).

---

9. Consequences for Generation Structure

If the URH holds *and the Fold is binary* (B, §2), refinement branches binarily, generation depth carries dyadic loading 1, 2, 4, 8, …, and premise P2 of the companion species paper's generation-count candidate (§9.1 there) is discharged — *derived for arbitrary depth* from refinement structure, rather than read off the three-term $\mathcal{D}$ = diag(1, 2, 4). Strictly, P2 reduces to **(U) ∧ (B)**, not to (U) alone: uniformity fixes that the branch is the first power of two, binarity fixes that it is a power of two at all (§6). To the extent (B) is architecturally closed (§2), the residue is just (U); to the extent (B) carries the §2 caveat, there is a second, smaller residue, and the honest phrase is *reduced to (U) ∧ (B)*.

This does not by itself yield three generations. The companion paper's candidate rests on three premises: P1 (a generation-blind closure budget K = 7), P2 (dyadic loading), and P3 (shared occupancy, so loads add). This paper addresses **only P2**, and only by reducing it to the URH. Even with P2 discharged, three generations would still require P1 and P3, each open or conditional in their own ways (the species paper states their closure conditions). So the contribution here is exactly one link in that chain: *why powers of two*, reduced to *why uniform single-Fold steps*. The generation count is not derived here and is not claimed to be.

---

10. 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) of the generation-count candidate 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 and reaches a more precise conclusion than the species paper stated in passing: the squeeze rests on **two** hinges that divide the labour — the Fold's binarity (B), which excludes non-powers-of-two including b = 3, and uniformity (U), which excludes the vacuous and composite cases — so that *(B) selects the base and (U) selects the exponent*. An intermediate framing that called uniformity "the sole hinge" understated (B), which alone carries the ternary exclusion (§5).

The species paper should be read as inheriting this result at the status given here: P2 is **conditional on (U) ∧ (B)**; (U) is **open** with the three closure conditions of §8 (no multi-Fold, no vacuous, no partial-frontier steps); and (B) — the binarity that makes the target a power of two, and specifically excludes 3 — is **inherited** from the Fold (§2), with the architecture-dependent caveat there. The chain across the two papers is therefore:

```
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)
   ↓  this paper §6 — two hinges
Why a power of two (not 3)?  →  (B) Fold binarity, §2, inherited
Why the first power (not 4, 8)?  →  (U) uniformity, §§7–8, open
   ↓  this paper §8 (the (U) task)
```

```
[ open — no-multi-Fold ∧ no-vacuous ∧ no-partial-frontier steps, a question
about the
  commitment dynamics; (B) supplies why the branch is a power of two and not
3 ]
```

Each step is a genuine reduction: the open question gets smaller and nearer the substrate. The terminus for (U) is a sharply-posed question about step individuation in the commitment dynamics — answerable in principle within the programme, with a clear expected answer (the Fold is the unit of refinement) and stated refuters (§8) — while (B) is inherited and carries its own §2 caveat.

---

11. What the Paper Establishes

Established (conditionally, with the conditions named):

- *Given* the Fold's binarity (B, §2, inherited) and uniformity (U, the URH, §7, open), binary branching is forced and dyadic loading 2^{n-1} follows (§6, the Binary Refinement Result). This is an implication, not a theorem; it carries no ■.
- The squeeze rests on **two hinges that divide the labour**, not on uniformity alone (§6): the **Fold's binarity (B)** selects the *base* — it is what excludes b = 3 and every non-power-of-two (§5), with uniformity contributing nothing there; **uniformity (U)** selects the *exponent* — the lower bound excludes the vacuous case and the upper bound's composite part excludes b = 4, 8, …. Neither alone forces b = 2. This is the paper's main structural finding, and it corrects an earlier framing that called uniformity "the sole hinge," which understated (B).
- The dyadic-loading problem (P2 of the companion paper) is **reduced to (U) ∧ (B)** (§9) — to the URH together with the Fold's inherited binarity, not to the URH alone — and (U) is given a precise closure condition: *no-multi-Fold-steps ∧ no-vacuous-steps ∧ no-partial-frontier-steps* (§8), each with a named refuter.
- The URH is **physically expected though not logically necessary** (§7.1): violating it requires denominating refinement in a non-Fold unit, analogous to a variable-length metre — so it is a *natural* assumption with a substrate-level reason, not an arbitrary one. (This is a clarification of (U)'s status, not a proof of it.)
- Two of (U)'s three conditions are **upgraded from assumed regularities to a consistency condition** (§7.1.2): given the inherited depth-is-descriptive result, refinement levels are individuated only by realised distinctions (Distinguishability Realization Principle), and a hierarchy of such levels must use a fixed distinction-unit (Refinement Unit Consistency Principle) — which unifies the b = 1 (zero-Fold) and b = 4 (two-Fold) pathologies as *one* variable-unit violation, covering conditions 1 and 2. (b = 3 remains excluded by (B); condition 3 remains separate.) This is conditional on the inherited result and on Fold-uniqueness (§7.2), not proven.
- The path to *dissolving* (U) entirely is identified as a candidate next result, the **Step Individuation Theorem** (§7.2): if the Fold is the uniquely admissible refinement unit acting across the whole frontier, uniformity is automatic and binary refinement follows

from (B) alone. Stated as a candidate with its own open hinge, and noting it too presupposes (B).

Not established (open or out of scope):

- the **URH / (U)** itself (§7) — physically expected (§7.1) but not derived; its closure is the three-condition §8 task (no multi-Fold, no vacuous, no partial-frontier steps);
- the **binarity of the Fold (B)** if the fold architecture fixes only minimality and non-degeneracy rather than two-valuedness (§2) — the seed-level dependency, and (per §5) carrying *more* load than an earlier draft credited, since (B) alone excludes ternary;
- the **Step Individuation Theorem** (§7.2) — the candidate that would discharge (U); flagged as the natural next paper, with its non-circularity hinge named (the Fold's uniqueness as refinement unit must be derived, not assumed) and its presupposition of (B) noted;
- the **generation count** — this paper addresses only P2 of the species paper's candidate, not P1 or P3 (§9);
- **closure capacity** and **shared occupancy** — inherited concerns of the species paper, out of scope here.

The honest summary: this paper does not derive dyadic loading. It shows the dyadic-loading question has *exactly two hinges* — (B) why a power of two, (U) why the first power — and that (B) is inherited (architecture-dependent, §2) while (U) reduces to a sharply-stated, substrate-near question about step individuation (§8) with a known proof-strategy and stated refuters. It converts "why does the programme keep producing powers of two?" into those two precisely-located questions. That reduction, not a derivation, is the contribution.

---

12. Conclusion

The recurrence of dyadic structure across the VERSF programme — generation loads 1, 2, 4; the closure-capacity saturation 1 + 2 + 4 = 7 — motivates asking whether binary refinement is fundamental or assumed. This paper argues it is *neither yet*, but that the question has a single, sharply-locatable hinge.

The natural argument is a two-sided squeeze: refinement slower than binary fails distinguishability, refinement faster than binary violates minimality, so the branching factor is two. But examined honestly the squeeze rests on *two* hinges that divide the labour, not one. The Fold's binarity (B) does the work of excluding b = 3 and every non-power-of-two — a ternary "minimal distinction" either is not minimal or is not admissible, and uniformity has nothing to say about it. Uniformity (U) does the work of excluding the vacuous case below and the composite cases (b = 4, 8, …) above, fixing the power at the first. Neither alone forces b = 2: binarity permits any power of two, uniformity permits 3. So the squeeze adds nothing independent — *(B) selects the base, (U) selects the exponent* — and given both, the result is *uniform refinement by a binary Fold branches binarily*, almost a tautology once both premises are granted, with the entire content in the premises.

This is the honest shape of the result, and it is more useful than an overstated theorem would be. It reduces the dyadic-loading problem — premise P2 of the companion species paper's generation-count candidate — to the conjunction of two hinges: the Fold's binarity (B, inherited, §2) and the Uniform Refinement Hypothesis (U, open). It states precisely what closing (U) requires: that the commitment dynamics admit no composite (multi-Fold), no idle (vacuous), and no partial-frontier refinement steps (§8). If those hold *and* (B) holds, dyadic loading follows for arbitrary depth, P2 is discharged, and the generation-count candidate rests only on its remaining premises. If either hinge fails, the powers of two need another explanation.

Binary refinement would then emerge not as an assumption, and not as a theorem proved here, but as a consequence of two things: the Fold being the minimal *binary* distinction (B), and refinement being the uniform single-Fold repetition of exactly that across the frontier (U). The first is inherited — though, per §2, only conditionally on the architecture fixing two-valuedness rather than mere minimality — and it does more work than it first appears, since (B) alone excludes ternary (§5). The second is the open hinge this paper most sharply isolates. *Why a power of two* is (B); *why the first power* is (U).

That hinge is open but not arbitrary. Uniformity is what refinement-by-Fold *almost certainly is*: to violate it is to count refinement in a unit other than the Fold — a variable-length metre — which is less an alternative dynamics than a refusal to use the substrate's own unit (§7.1). So the URH is physically expected though not logically necessary, and the natural next step is not to *assume* it more confidently but to *dissolve* it: the Step Individuation Theorem (§7.2), should it be established, would show the Fold is the uniquely admissible refinement unit, at which point uniformity is automatic and binary refinement follows from the Fold's binarity with nothing left to assume. That theorem is the proper sequel to this paper — and its non-circularity will turn, as everything here has, on deriving the Fold's uniqueness as a unit rather than reading it off the claim that refinement is "counted in Folds."

Refinement does not double because we observed three generations.

It would double if — and exactly if — each step is one binary Fold and no more; the recurrence of powers of two is the shadow of that single, still-open hypothesis — a hypothesis the substrate appears to want, and which the next paper must either force or surrender.