The Individuation–Isolation Characterization

Why the Spectrally Isolated Modes Are the Generations — an Operational Derivation of the Species Reading

Keith Taylor

VERSF Theoretical Physics Programme

General Reader Summary

A companion paper built an instrument. It defined a property a matter mode can have — being *spectrally isolated*, meaning its separation from the rest of the substrate's modes stays open as the description is refined — and showed that the count of generations reduces to a single computable question: is the disputed fourth mode isolated, or does its separation close? That paper was scrupulous about one thing it did *not* prove. It assumed, without establishing, that the isolated modes *are* the generations. It called this a "reading," carried it openly, and flagged that the reading is what gives the whole reduction its meaning. Without it, computing that a mode is isolated tells you nothing about whether it is a generation.

This paper discharges that reading — or rather, it discharges it down to the smallest honest residue. The strategy is the one a careful sceptic would demand. It refuses to define a generation as "an isolated mode," because that would make the claim true by relabelling and prove nothing. Instead it defines a generation the way physics already understands the word: a generation is a *distinct species* — something the substrate can durably tell apart from the others. Made precise, this means a mode that carries its own record, witnessable above the threshold at which distinctions become real, and that stays separately witnessable as the description is refined rather than washing into its neighbours. This definition mentions no spectra, no resolvents, no isolating moats. It is built only from two principles the programme already holds: that a distinction is real only if it can be witnessed, and that a species must persist.

The paper then proves that this operationally defined notion — durable distinguishability — coincides exactly with spectral isolation. The bridge is a classical instrument of analysis: around an isolated mode one can build a clean extractor that reads that mode's record and nothing else, and such a clean extractor exists precisely when the isolating moat exists. So "you can durably tell it apart" and "it is spectrally isolated" turn out to be the same fact viewed from two sides.

Honesty requires naming the price. The coincidence is not unconditional. It rests on two stated premises: that the clean extractor is a genuinely realizable measurement and not merely a mathematical fiction, and that the substrate's discreteness keeps a persistent separation from shrinking to nothing. Both are plausible and both are argued; neither is proven here from first principles. So the achievement is precise rather than total: the companion paper's free-floating reading is replaced by a proven equivalence resting on two named premises. A by-product falls out for free — the companion paper had to *choose* to read its limit one way rather than another, and called the choice a stipulation; here that choice is shown to be forced by what durability means, not chosen. The number that decides the census is still uncomputed, and the test can still come out either way. What changes is that, when the instrument finally returns its verdict, the verdict will *mean* what the programme always wanted it to mean — and the paper says exactly which two premises that meaning hangs on.

Contents

Status Table — Read First

The strongest claim a speculative programme can make is an accurate one. This table states, before any argument, what the paper establishes and at what grade. Four grades, ordered **Proven ≻ Conditional ≻ Conjectural ≻ Open** (strongest to weakest); **[Inherited]** marks an imported result carrying its source grade; **[Def]** marks a definitional choice, not truth-apt, off the grade axis. Gate tags name an undischarged hypothesis a result waits on. Tags and [Def] are presentation, not grades.

Claim	Grade	Gating dependency
Operational individuation Gen(m) is well-defined from I3 + R2 + I5, with no reference to $\mathfrak{G}$	**[Proven]**	floor-quantization of the witness margin (§3)
The witness margin is floor-quantized: $\mu_n(m) \in \{0\} \cup [\delta^\star, \infty)$, floor applied before the supremum	**[Proven]**	distinguishability floor I3, via Definition 3.1 (§3)
B1-forward — the Riesz extractor is an admissible witness	**[Conditional · near-inherited]**	functional calculus $\subseteq \mathcal{O}$ (§4)
B1-reverse — no admissible observable out-separates the resolvent (the substantive premise)	**[Conditional]**	carried; candidate route from the ℓ-induced metric, not closed (§4)
The competing spectrum is entirely persistent: transient modes contribute none of it	**[Proven]**	subspace = ℓ-closure of Rec; R2 (§9)
A persistent mode's spectral *position* is eventually stationary	**[Conditional]**	I2 + R2 + E-stab (the [Conditional] grade is E-stab's) (§9)
E-stab — the spectral embedding is depth-stable on fixed integer data	**[Conditional · expected inherited]**	carried; to be confirmed from the $\mathscr{R}_n$ construction (§9.2a)
B2 reduces to a single static residue ¬ACC: persistent positions do not accumulate at an isolated value	**[Proven that B2 ⇔ ¬ACC]**	§9 — narrows the carried part of B2
¬ACC (the residue of B2)	**[Conditional]**	reduced from global assertion to one no-accumulation condition (§9)

Claim	Grade	Gating dependency
Sign-agreement: $\mu_n(m) > 0 \Leftrightarrow \Delta_n(m) > 0$ at each depth	**[Conditional]**	B1-reverse
Characterization: $\mathrm{Gen}(m) \Leftrightarrow \mathfrak{G}(m) > 0$ — isolated modes are exactly the species	**[Conditional]**	B1-reverse, ¬ACC
Species-census 4e_stab: companion's reading discharged **to B1-reverse + ¬ACC** (relocated, narrowed, not closed)	**[Conditional]**	B1-reverse, ¬ACC
The liminf reading of the gap functional is forced, not stipulated	**[Proven]**	on the operational definition of durability (§7)
Individuation and isolation are a priori separable — the equivalence has content	**[Conditional · realizability]**	separability specification, §8.2 (logical independence pending an instantiating operator)
B1-reverse and ¬ACC from first principles — possibly a shared discreteness root	**[Open]**	candidate routes only (§9.4)
The sign of $\mathfrak{G}(m_4)$ — whether the fourth mode is a species	**[Open · Gate: D5(depth)]**	operator construction; unchanged
Capacity catalogue; world-occupancy; gauge, chirality, mixing, spectrum	**[Open]**	not addressed here

---

1. The Carried Reading, and What Discharging It Honestly Requires

The companion gap-functional paper reduces the species-census to the sign of a single operator quantity, $\mathfrak{G}(m_4)$. Its reduction theorem is, by its own admission, near-immediate once the definitions are in place; the substance lives in a claim it carries but does not prove — that the spectrally isolated modes are the generations, so that the stable-level census 4e_stab is the count physics cares about. That paper marks the claim a *reading*, applies it where needed, and defers its establishment to a "dedicated characterization staged separately." This is that characterization.

The discharge has a trap, and naming it fixes the method. If a generation is *defined* to be a spectrally isolated mode, then "the isolated modes are the generations" is true by stipulation and proves nothing; the companion paper's reduction would remain a relabelling of its own definition. A genuine discharge must therefore define a generation from content that does **not** mention the gap functional, the resolvent, or isolating moats — and then *prove* that the independently defined notion coincides with spectral isolation. The coincidence, not the definition, is the result.

The programme already supplies the independent content. A generation is a distinct *species*: not merely something the substrate admits and that recurs, but something that maintains its own persistent, separately witnessable record. Two inherited principles make this precise without any spectral vocabulary. By the operational ontology (I3), a distinction counts as substrate structure only if it is witnessable at or above the distinguishability floor; a sub-floor distinction is not faint but absent. By the persistence condition built into the recurrent set (R2), a species is a structure that survives refinement; a distinction that lapses as the description deepens has not survived. Compose them: a mode is a generation iff it carries a record that is witnessably separate from the others, durably — from some refinement depth onward. This is the operational definition the paper will use, and §3 states it exactly. It refers to nothing the gap functional refers to.

The work of the paper is then to prove that this operational notion and spectral isolation are the same property, to state precisely the premises under which the proof goes through, and to demonstrate that the two notions are genuinely distinct objects whose coincidence carries content rather than synonyms whose coincidence is empty. The last task — separability — is what certifies that the trap has been avoided, and it is discharged in §8.

---

2. Inherited Inputs and Marker Calculus

Every truth-apt claim carries one of four grades — [Proven], [Conditional] (contingent on a stated, carried hypothesis), [Conjectural] (a candidate discharge), [Open] (no current discharge) — ordered Proven ≻ Conditional ≻ Conjectural ≻ Open, with a conclusion inheriting the **meet** (weakest) of the grades on its path. A definitional choice is not true or false; it carries no grade, is marked **[Def]**, and is justified by motivation. What follows from a definition is graded normally.

The following are imported and used, not re-proven. They are the companion paper's inputs, restated, together with two clauses (the floor and the observable algebra) that the companion left implicit and this paper makes explicit because it leans on them.

- **I1 — Refinement architecture.** [Inherited] Two binary exchangeable marks generate the three mark-levels R_0, R_1, R_2; silent on whether a non-mark admissible level exists.
- **I2 — Recurrent structure.** [Inherited] The closure evolution admits a recurrent structure (Orbit Count Theorem, integer return count K_c and commitment count p_v). The three mark-levels are recurrent and spectrally isolated.
- **I3 — Operational ontology, with floor.** [Inherited] A distinction is substrate structure only if witnessable at or above a distinguishability floor $\delta^\star > 0$; individuation is internal, not externally labelled. The floor is a fixed positive quantity supplied by finite distinguishability (H1): with finitely many distinguishable states at bounded refinement, separations below $\delta^\star$ are not resolved and count as absent.
- **R2 — Persistence (from the recurrent-set definition).** [Inherited] A generation-relevant mode survives admissible refinement: it does not decay into, or merge with, another mode's orbit as depth n increases. Persistence is an *eventual* condition — survival means holding from some depth onward.
- **I5 — Commitment functional and record.** [Inherited] The commitment functional ℓ on the admissible closure state space $\mathcal{A}$ supplies an intrinsic positive form $\langle\cdot,\cdot\rangle_\ell$. For a

recurrent mode m, its *commitment record* is the ℓ-content registered on m's closed orbit — a basis-independent quantity induced by the dynamics.

- **I6 — Operational observable algebra.** [Inherited, made explicit] $\mathcal{O}(\mathscr{R}, \ell)$ is the algebra of admissible internal observables: the closure, under composition and ℓ-pairing, of the bounded functional calculus of the evolution operator $\mathscr{R}$. This closure is in general *strictly larger* than the bounded functional calculus of $\mathscr{R}$ alone — composition and ℓ-pairing can build observables not expressible as a single function of $\mathscr{R}$. This is a deliberate choice, and it sets the stage for B1: because $\mathcal{O}$ is richer than the functional calculus, it is a substantive claim (not a near-tautology) that no observable in $\mathcal{O}$ can separate spectrally inseparable modes. An observable is *witnessable* in the sense of I3 iff it lies in $\mathcal{O}(\mathscr{R}, \ell)$ and resolves its target above δ⋆. The mark decomposition of I1 is *not* part of $\mathcal{O}(\mathscr{R}, \ell)$: the observables are built from dynamics and commitment, not from the mark sectoring.

Write $\mathscr{R}_n$ for the action of $\mathscr{R}$ at refinement depth n, Rec_ℓ($\mathscr{R}$) for the recurrent admissible set of the companion paper (modes satisfying recurrence, refinement-survival, and nontrivial ℓ-coupling), $\Delta_n(m)$ for the companion's finite-depth isolating-annulus radius, and $\mathfrak{G}(m) = \liminf_{n\to\infty} \Delta_n(m)$ for the gap functional. The companion's witness model **M_census** — in which a non-mark admissible recurrent mode m_4 exists, refuting recurrence-level exhaustiveness 4e_rec while leaving the mark count intact — is carried throughout.

---

3. Operational Individuation — A Species Defined Without Spectra

The definition below is the load-bearing independent object. It is stated entirely in the vocabulary of I3, R2, I5, and I6; it names no spectral quantity.

Definition 3.1 (witness margin). Call a separation between the commitment record of m and the joint record of Rec_ℓ($\mathscr{R}$) \ {m} *witnessed* by an observable W ∈ $\mathcal{O}(\mathscr{R}, \ell)$ iff W resolves it at or above the distinguishability floor δ⋆ in the sense of I3 — a separation W produces below δ⋆ is, by I3, not resolved and counts as no separation at all. The *witness margin* $\mu_n(m)$ is the supremum of the witnessed separations, in the ℓ-record metric, over admissible observables at depth n:

$\mu_n(m)$ = sup { sep_ℓ(W·m , W·[Rec \ {m}]) : W ∈ $\mathcal{O}(\mathscr{R}, \ell)$ and W witnesses the separation, i.e. sep_ℓ ≥ δ⋆ },

with the convention $\mu_n(m) = 0$ when no admissible observable witnesses any separation. The sup ranges only over witnessed separations; sub-floor metric separations are excluded from it by I3 before the supremum is taken, not discarded after. A witness achieving a positive margin is a contamination-free record for m: it reads m's content while suppressing every other recurrent mode's content.

Lemma 3.2 (floor-quantization of the margin). $\mu_n(m) \in \{0\} \cup [\delta\star, \infty)$. **[Proven, on I3 + Definition 3.1]**

Proof. By construction the supremum in Definition 3.1 is taken over a set whose every element is a witnessed separation, hence $\geq \delta\star$ by I3; the set is a subset of $[\delta\star, \infty)$. If that set is non-empty, its supremum lies in $[\delta\star, \infty)$ — in particular $\mu_n(m) \geq \delta\star$ even when the supremum is not attained, because the set is bounded below by $\delta\star$. If the set is empty — no admissible observable witnesses any separation above the floor — then $\mu_n(m) = 0$ by the convention. No value strictly between 0 and $\delta\star$ can arise: such a value would be the supremum of a non-empty subset of $[\delta\star, \infty)$, which is impossible, or the supremum of the empty set, which is 0. The quantization holds for the supremum, not merely per witness, because the floor is applied to each candidate separation before it enters the sup.

Remark 3.2a (why the floor is applied before the sup). The quantization is a property of the margin, not only of individual witnesses, precisely because Definition 3.1 admits a separation into the supremum only once it is witnessed above $\delta\star$. Were the sup taken over raw metric separations with the floor applied afterward, witnesses achieving separations approaching $\delta\star$ from below — each individually sub-floor, each resolving nothing — could drive the supremum to a value in $(0, \delta\star)$, and the trichotomy would fail. Definition 3.1 forecloses this by construction; the downstream arguments of §7 and §8 that depend on the sharp "exactly 0 or $\geq \delta\star$" rely on this reading and on no stronger one.

Lemma 3.2 is the quiet engine of the whole paper. It means the operational separation is not a continuous dial but a switch: a mode is either contamination-free witnessable (margin $\geq \delta\star$) or operationally indistinct (margin 0). There is no operational meaning to "slightly distinguishable."

Definition 3.3 (generation / durable species). A recurrent mode m is a *generation*, written Gen(m), iff its witness margin is eventually positive:

$$\text{Gen}(m) \Leftrightarrow \exists\, N \text{ such that } \mu_n(m) > 0 \text{ for all } n \geq N.$$

By Lemma 3.2 the "> 0" is equivalently "$\geq \delta\star$", so Gen(m) says the contamination-free record for m exists and persists from some depth onward — durable distinguishability. This is R2's persistence applied to the distinction itself: a species is a mode you can keep telling apart.

Remark 3.4 (what Gen(m) does not mention). Definition 3.3 is evaluable on any closure system $(\mathcal{A}, \ell, \mathscr{R})$ the moment the observable algebra $\mathcal{O}(\mathscr{R}, \ell)$ and the record metric are fixed. It does not invoke the resolvent, isolating annuli, the gap functional, or the mark count. In particular it is evaluable in M_census directly, on the disputed mode m_4, without first computing $\mathfrak{G}(m_4)$. This independence from the spectral apparatus is precisely what the discharge of §1 requires, and it is what makes the coincidence proved in §6 a theorem rather than a definition.

---

4. The Two Premises the Bridge Requires

The equivalence of Gen(m) with spectral isolation does not hold for free. It hold under two premises, stated here in full and graded [Conditional]; the honesty of the paper is in naming them as the entire residue on which the species reading now rests, rather than burying them.

Premise B1 (projector realizability). Stated as two halves of distinct weight, because they are graded differently and §8(i) withdraws only one of them.

B1-forward. If m has an isolating annulus at depth n ($\Delta_n(m) > 0$), the Riesz projection $P\_m(n) = (1/2\pi i) \oint\_\Gamma (\zeta - \mathscr{R}_n)^{-1} d\zeta$ (Γ a contour in the annulus) lies in $\mathcal{O}(\mathscr{R}, \ell)$ and resolves m's record above $\delta\star$. **B1-reverse.** Any admissible observable that witnesses a positive margin for m furnishes a spectral separation of m from the remainder — equivalently, no observable in $\mathcal{O}(\mathscr{R}, \ell)$ separates spectrally inseparable modes.

Status — the two halves are not equal. B1-forward is close to inherited: P_m(n) is built by the bounded functional calculus of $\mathscr{R}_n$, so it lies in $\mathcal{O}(\mathscr{R}, \ell)$ by I6 (the functional calculus is contained in $\mathcal{O}$), and §5 shows it cleanly separates m. **[Conditional, near-inherited]**

B1-reverse is the genuine analytic load of the paper, and the grade reflects that without softening. Because I6 makes $\mathcal{O}(\mathscr{R}, \ell)$ *strictly larger* than the functional calculus — closed under composition and ℓ-pairing — B1-reverse is a real closure constraint: it asserts that the extra observables those operations build still cannot out-separate the resolvent. This is not nearly definitional. The candidate route would derive it from the record metric being ℓ-induced: if sep_ℓ is controlled by ⟨·,·⟩_ℓ and every admissible observable acts within the commitment-supported subspace, then record-separation is bounded by spectral separation and no observable outruns the resolvent. But that route must rule out composition/ℓ-pairing manufacturing separation that no single resolvent function achieves, and it does not obviously do so; it is aspirational, not closed. B1-reverse is therefore carried as a substantive premise, not a near-free one. **[Conditional · the substantive premise · candidate route from the ℓ-induced metric, not closed]**

The whole of B1 is the conjunction; its operative grade is that of B1-reverse, [Conditional]. The reader should hold B1-reverse, specifically, as the analytic debt — and note (§8) that it is exactly B1-reverse, not B1-forward, that separability withdraws.

Premise B2 (gap quantization). [Conditional] A persistent positive moat is bounded below uniformly in depth: if $\Delta_n(m) > 0$ for all $n \geq N$, then $\liminf_{n\to\infty} \Delta_n(m) > 0$. Equivalently, the commitment-supported spectrum admits no eventually-positive-but-vanishing moat — a separation that stays open yet shrinks to zero.

Status. The per-depth content is inherited: by H1 the distinguishable states at bounded refinement are finite, so at each fixed depth the commitment-supported spectrum is finite and any existing moat has positive radius — $\Delta_n(m)$ takes no infinitesimal positive values at fixed n. The carried content is the *depth-uniform* lower bound: finiteness at each depth does not by itself prevent the moat from narrowing toward zero as $n \to \infty$. Section 9 discharges most of B2 — it proves the competing spectrum is entirely persistent and (under E-stab) that persistent positions are eventually stationary, which reduces B2 to a single static residue: the persistent positions do not accumulate at an isolated value (¬ACC). What remains carried is ¬ACC alone, narrower and more plausible than B2 as first stated. **[Reduced in §9 to ¬ACC · candidate route from a no-accumulation (uniform-spacing) bound]**

The two premises do different work. B1 connects the operational switch (margin positive or zero) to the spectral switch (moat present or absent) at each depth. B2 connects eventual-positivity of the moat to the gap functional's liminf-positivity. Neither is the species reading itself; together they are exactly what the reading reduces to.

---

5. The Riesz Extractor and Sign-Agreement

This section establishes the depth-by-depth half of the bridge: a contamination-free witness for m exists iff m has an isolating moat. The forward direction is essentially functional analysis; the reverse is where B1 is spent.

Proposition 5.1 (moat ⟹ clean witness). If $\Delta_n(m) > 0$, then $\mu_n(m) \geq \delta\star$. **[Conditional on B1, forward half]**

Proof. Let A be an isolating annulus for m at depth n with a contour Γ enclosing m's orbit value region τ_m^(n) and no other point of the commitment-supported spectrum. The resolvent $(\zeta - \mathscr{R}_n)^{-1}$ is analytic on A, so the Riesz projection P_m(n) = (1⁄2πi) ∮_Γ $(\zeta - \mathscr{R}_n)^{-1}$ dζ is well-defined, idempotent, and commutes with $\mathscr{R}_n$; it projects onto m's spectral subspace and annihilates every other recurrent mode, whose spectral values lie outside Γ. Its operator norm is controlled by the resolvent on the contour,

‖P_m(n)‖ ≤ (|Γ|⁄2π) · sup_{ζ∈Γ} $\|(\zeta - \mathscr{R}_n)^{-1}\|$ ≤ (|Γ|⁄2π) · (1⁄dist(Γ, spec)),

and the isolating moat bounds dist(Γ, spec) below by a positive multiple of $\Delta_n(m)$. By B1 (forward) P_m(n) ∈ 𝒪(ℛ, ℓ) is an admissible witness; it reads m's record and suppresses the rest, so it achieves a positive separation in the ℓ-record metric. By Lemma 3.2 any positive margin is ≥ δ⋆. Hence $\mu_n(m) \geq \delta\star$.

Proposition 5.2 (no moat ⟹ no clean witness). If $\Delta_n(m) = 0$, then $\mu_n(m) = 0$. **[Conditional on B1, reverse half]**

Proof. $\Delta_n(m) = 0$ means m admits no isolating annulus: m's orbit value is an embedded or accumulation point of the commitment-supported spectrum at depth n, so every neighbourhood of τ_m^(n) contains other recurrent modes' spectral content. A contamination-free witness for m would, by B1 (reverse), furnish a spectral separation of m from the remainder — an isolating moat — contradicting $\Delta_n(m) = 0$. So no admissible observable separates m's record from the rest above the floor, and $\mu_n(m) = 0$ by Lemma 3.2.

Corollary 5.3 (sign-agreement). Under B1 (forward and reverse), for every m ∈ Rec_ℓ(ℛ) and every depth n,

$\mu_n(m) > 0 \Leftrightarrow \Delta_n(m) > 0.$

[Conditional on B1]

Sign-agreement is the precise sense in which the operational and spectral notions track each other: at each depth, a clean witness exists exactly when an isolating moat exists. It is a statement about presence-or-absence, not about magnitudes — which is all the floor-quantized margin needs.

---

6. The Characterization Theorem

Theorem 6.1 (Individuation–Isolation Characterization). Under B1 and B2, a recurrent commitment-supported mode is a generation iff it is spectrally isolated:

$\mathrm{Gen}(m) \Leftrightarrow \mathfrak{G}(m) > 0$, for every $m \in \mathrm{Rec}\_\ell(\mathscr{R})$.

[Conditional on B1, B2]

Proof. Gen(m) is, by Definition 3.3, the eventual positivity of the witness margin: $\exists$ N, $\mu_n(m) > 0$ for all $n \geq N$. By sign-agreement (Corollary 5.3, under B1) this holds iff $\Delta_n(m) > 0$ for all $n \geq N$ — the moat is eventually present. By B2 an eventually-present moat is eventually bounded below: $\liminf\_\{n\to\infty\}\ \Delta_n(m) > 0$, i.e. $\mathfrak{G}(m) > 0$. Conversely $\mathfrak{G}(m) > 0$ gives $\liminf \Delta_n(m) > 0$, so $\Delta_n(m) > 0$ for all sufficiently large n, whence by sign-agreement $\mu_n(m) > 0$ eventually, i.e. Gen(m). The two implications give the equivalence.

Corollary 6.2 (the isolated modes are exactly the species). Under B1 and B2, $\mathrm{Stab}\_\ell(\mathscr{R}) = \{ m \in \mathrm{Rec}\_\ell(\mathscr{R}) : \mathfrak{G}(m) > 0 \}$ is exactly the set of generations, and N_stab is the number of distinct species. The companion paper's stable-level census 4e_stab — N_stab = 3 — is therefore the operative species-census, and its carried reading is discharged **to** the premises B1-reverse and ¬ACC: relocated from an unstated assumption to a dependence on two named, independently meaningful propositions, and narrowed — not removed. **[Conditional on B1-reverse, ¬ACC]**

This is the result the companion paper deferred. It does not make the species reading free; it converts it from an unstated assumption into a theorem resting on two named, independently meaningful premises. The reader who accepts those premises must accept that the spectrally isolated modes are the generations; the reader who doubts the reading now has exactly two propositions to contest, both stated in the open — and the companion paper's reading row, post-this-paper, reads "[Conditional on B1-reverse, ¬ACC]", not "[Proven]".

Remark 6.3 (the operative premise set, after §9). Section 9 discharges most of B2: the competing spectrum is shown to be entirely persistent and persistent positions eventually stationary (under E-stab), so B2 holds except through one static mechanism, ¬ACC (the persistent positions do not accumulate at an isolated value). Substituting that reduction, the operative premise set of Theorem 6.1 is **B1-reverse and ¬ACC** — one premise about the realizability of the extractor (no observable out-separates the resolvent), one about the persistent spectrum having no accumulation point. The body above is stated with B1, B2 for continuity with §4; the reader should carry forward B1-reverse + ¬ACC as the true residue, and §9.4 on the possibility that the two share a discreteness root.

---

7. The liminf Discharged — Persistence Forces the Reading

The companion paper marked one choice [Def]: its gap functional takes a liminf, not a limsup, and it justified this by the persistence motivation of R2 while conceding the choice was a stipulation pending a separate argument. The operational definition supplies that argument and upgrades the status.

Proposition 7.1 (the liminf is forced). Given the operational definition of a generation as durable distinguishability (Definition 3.3), the functional whose positivity characterizes generations must read its limit as a liminf (eventual positivity), not a limsup (positivity infinitely often). **[Proven, on Definition 3.3 + Lemma 3.2]**

Proof. A generation is, by R2, a mode that *survives* refinement — durable from some depth onward; this is the eventual reading $\mu_n(m) > 0$ for all $n \geq N$. Suppose instead one read durability as the limsup condition: $\mu_n(m) > 0$ infinitely often, with $\mu_n(m) = 0$ also infinitely often (the margin is floor-quantized by Lemma 3.2, so "lapsing" means dropping to exactly 0, not merely shrinking). Such a mode is contamination-free witnessable at infinitely many depths and operationally indistinct at infinitely many depths: its separate record appears and vanishes endlessly as the description is refined. By R2 this is precisely a transient, not a survivor — a distinction that does not persist. Reading it as a species would count an intermittently-present record as a permanent one, contradicting the persistence that defines a species. Hence the limsup reading misclassifies transients as generations, and the liminf is the unique reading consistent with Definition 3.3. The choice the companion paper carried as [Def] is thereby forced by what durability means; it is not a free stipulation.

The floor-quantization of Lemma 3.2 is what makes this argument clean rather than merely suggestive: because the margin cannot take intermediate values, "lapsing infinitely often" is the sharp event of the witness dropping to zero infinitely often, and a record that genuinely vanishes infinitely often has, unambiguously, not survived.

---

8. Non-Circularity — The Two Notions Are Separable

The discharge is worthless if Gen(m) and $\mathfrak{G}(m) > 0$ are secretly the same definition wearing two notations. This section certifies that they are distinct objects whose coincidence (Theorem 6.1) is content delivered by the carried premises — exactly as the corpus certifies tightness, by exhibiting that the equivalence can fail when a premise is withdrawn. Crucially, each withdrawal removes exactly the premise graded substantive, leaving the near-inherited parts (B1-forward, Propositions 9.1–9.2) untouched — so the thing shown contingent is precisely the thing carried [Conditional], and no cheaply-established claim is put at risk.

Proposition 8.1 (separability). The predicates Gen(m) and "$\mathfrak{G}(m) > 0$" are a priori distinct: there are consistent closure specifications in which one holds and the other fails. Specifically — (i) withdrawing **B1-reverse** (not B1-forward) admits a specification with an isolating moat ($\Delta_n(m) > 0$ eventually) but no admissible witness realizing it ($\mu_n(m) = 0$), so $\mathfrak{G}(m) > 0$ while ¬Gen(m); and (ii) withdrawing **¬ACC** (the residue of B2 after §9) admits a specification with a clean witness at every depth ($\mu_n(m) \geq \delta\star$, hence Gen(m) by the eventual reading) whose underlying moat shrinks to zero ($\mathfrak{G}(m) = \liminf \Delta_n(m) = 0$), so Gen(m) while $\mathfrak{G}(m) = 0$. **[Conditional · realizability]**

Proof. (i) Withdraw B1-reverse alone — leave B1-forward intact, so the functional calculus still lies in $\mathcal{O}$, but allow the record metric to fail to register a spectral separation, so a contamination-free witness need not exist even where a moat does. Specify a system with an eventual isolating moat for m in which no admissible observable witnesses a separation above the floor. There $\Delta_n(m) > 0$ eventually so $\mathfrak{G}(m) > 0$, while $\mu_n(m) = 0$ (Lemma 3.2: nothing witnessed) and ¬Gen(m). This is exactly a withdrawal of the substantive half; B1-forward is untouched, and so is its near-inherited status. (ii) Withdraw ¬ACC. Specify a system in which infinitely many persistent modes accumulate at τ_m so that m's moat is positive at every depth but $\Delta_n(m) \to 0$; by B1 a clean witness exists at every depth (the moat is positive at each n), so $\mu_n(m) \geq \delta\star$ for all n and Gen(m) holds by Definition 3.3, while $\mathfrak{G}(m) = \liminf \Delta_n(m) = 0$. Each specification is logically consistent — it contains no internal contradiction — which is all the separability claim requires.

Remark 8.2 (what "· realizability" qualifies). As with the corpus tightness witnesses, Proposition 8.1 exhibits consistent *specifications*, not realized closure machines: it shows the predicates are not logically identical, by displaying premise-withdrawn worlds where they diverge. What it does not supply is a concrete operator ($\mathcal{A}$, ℓ, $\mathcal{R}$) instantiating either divergence — that realization is owed to the operator construction, as everywhere in the programme. The grade [Conditional · realizability] tags exactly this: sound as a logical-independence argument, pending an instantiating operator — the same modulo every carried premise in the corpus carries, written uniformly rather than as a one-off. This is sufficient for the present purpose: it certifies that Theorem 6.1 is not a tautology, because its two sides come apart precisely when B1-reverse or ¬ACC is removed. The coincidence is the content; B1-reverse and ¬ACC are what buy it.

Corollary 8.3 (the failure mode is avoided). Because Gen(m) is defined without reference to the spectral apparatus (Remark 3.4) and is provably separable from spectral isolation (Proposition 8.1), the characterization is not the relabelling the discharge had to avoid (§1). "Generation" was not reverse-engineered to mean "isolated"; it was defined as durable distinguishability and *proven* to coincide with isolation under stated premises. **[Conditional · realizability]**

---

9. Discharging Most of B2 — Persistence Confines the Competing Spectrum

B2 was carried in §4 as a global assertion: every persistent positive moat is bounded below uniformly in depth. This section proves the bulk of it from inherited content and isolates the single residue that remains genuinely assumed. The result is that B2 is not an independent premise of the weight B1-reverse carries; it is almost a theorem, lacking only one static exclusion.

The argument turns on a fact about *which* spectrum the moat competes against. The moat $\Delta_n(m)$ is the distance from m's orbit value to the nearest other point of the **commitment-supported spectrum**, and that spectrum lives on the commitment-supported subspace — by the companion paper's construction, the ℓ-closure of the span of Rec_ℓ(ℛ). Two consequences follow, both proven, and together they leave only one way for B2 to fail.

Proposition 9.1 (the competing spectrum is entirely persistent). Every point of the commitment-supported spectrum that can close m's moat is the orbit value of a recurrent (R2-persistent) mode. Transient modes — those that decay or merge under refinement, failing R2 — contribute no competing spectrum. **[Proven, on the subspace definition + R2]**

Proof. The commitment-supported subspace is the ℓ-closure of span Rec_ℓ(ℛ), and Rec_ℓ(ℛ) consists by definition only of modes satisfying recurrence, refinement-survival (R2), and nontrivial ℓ-coupling. A transient mode fails R2 and so is not in Rec_ℓ(ℛ); its content is not in the span whose closure defines the subspace, and the resolvent of $\mathscr{R}_n$ restricted to that subspace sees no spectral point from it. Hence the only spectral content within any neighbourhood of τ_m^(n) — the only content that can shrink the moat — comes from persistent modes.

Proposition 9.2 (persistent positions are eventually stationary). Assume the spectral embedding is depth-stable on fixed integer data: the map from a mode's orbit data (K_c, p_v) to its orbit value τ in the spectral parameter does not itself drift with n once that data is fixed (E-stab). Then for each persistent mode m′, the orbit value τ_{m′}^(n) is eventually constant in n: there is N_{m′} with τ_{m′}^(n) = τ_{m′} for all n ≥ N_{m′}. **[Proven on I2 + R2 + E-stab]**

Proof. By I2 a recurrent mode's orbit value is the image, under the spectral embedding, of its integer orbit data — the return count K_c and commitment count p_v over one return (Orbit Count Theorem). R2 is precisely the condition that the mode neither decays nor merges as refinement deepens; a mode whose return or commitment count changed with depth would be reorganizing its orbit — merging with or splitting from neighbouring structure — which is the failure of R2. So for a mode satisfying R2 the integer pair (K_c, p_v) is eventually constant. By E-stab the embedding does not move that fixed data to a drifting spectral location; hence the orbit value τ it determines is eventually constant, not merely the data behind it.

Remark 9.2a (E-stab is a real clause, not free). The stabilization of the *integer data* follows from R2 alone, but the stabilization of the *spectral value* does not — the same non-merging orbit can in principle sit at a drifting spectral location if the refinement reparameterizes the spectral embedding with depth. Theorem 9.3(⟸) uses stationarity of the values, not of the labels, so this is load-bearing rather than pedantic. E-stab is the clause that the embedding is depth-independent once the integer data is fixed; it is carried here, and is the natural property of a refinement that sharpens resolution without relocating already-settled orbits. It should be confirmed against the companion paper's construction of $\mathscr{R}_n$, where it is either a property of the embedding or a further premise; this paper names it rather than assuming it silently. **[Carried · to be confirmed inherited from the $\mathscr{R}_n$ construction]**

These two facts almost close B2, and — this is the referee-sharpened point — Proposition 9.2 has already done more than it appears: it converts the *dynamical* picture ("modes activating at increasing depth") into a *static* one. Once every persistent mode has an eventually-fixed spectral position, the competing spectrum near τ_m is, in the limit, a fixed point set — the set of stationary positions of the persistent modes. Suppose m has an eventually-positive moat ($\Delta_n(m) > 0$ for all $n \geq N$) and ask how liminf $\Delta_n(m)$ could nonetheless be zero. By Proposition 9.1 the competitors are all persistent; by Proposition 9.2 their positions converge to fixed values. The only way the moat can shrink to zero while staying positive is for those fixed positions to *accumulate* at τ_m — infinitely many distinct persistent modes whose stationary values cluster arbitrarily close to τ_m without coinciding with it. Stated as a property of the point set rather than of the dynamics:

¬ACC (no accumulation at an isolated value). The set of stationary positions of the persistent modes has no finite accumulation point at the value τ_m of an otherwise-isolated persistent mode m.

This static phrasing is exact, not merely suggestive: by 9.2 "activation at increasing depth converging to τ_m" *is* an accumulation point of the persistent position set at τ_m, and conversely. The two formulations coincide because the eventually-stationary lemma has already frozen the dynamics into a point set.

Theorem 9.3 (B2 reduces to ¬ACC). Given Propositions 9.1–9.2, B2 holds if and only if ¬ACC holds. Equivalently: for a mode with an eventually-positive moat, liminf $\Delta_n(m) > 0$ iff the persistent position set does not accumulate at τ_m. **[Proven that B2 ⟺ ¬ACC]**

Proof. (⟸) Assume ¬ACC and $\Delta_n(m) > 0$ for $n \geq N$. The competitors are persistent (9.1) with positions converging to fixed values (9.2). No accumulation at τ_m means a fixed punctured neighbourhood of τ_m contains only finitely many stationary positions, all at fixed positive distances; beyond some depth the nearest competitor is one of that finite set, so the moat is eventually bounded below and liminf $\Delta_n(m) > 0$. (⟹) If the position set accumulates at τ_m, infinitely many distinct stationary positions lie arbitrarily close to τ_m; the distance from τ_m to the nearest competitor tends to 0, so liminf $\Delta_n(m) = 0$, while each finite-depth moat may remain positive — exactly an eventually-positive-yet-vanishing moat, the failure of B2.

The reduction is the useful content. B2 began as an unbounded claim about all moats at all depths; it is now equivalent, given inherited structure, to a single static exclusion at a single value — the persistent spectrum does not accumulate at an isolated mode. The lattice route is now *visibly* sufficient rather than asserted: if the stationary positions sit on a set with a uniform minimum spacing — a lattice condition on where persistent species can sit — then no accumulation point can exist, and ¬ACC follows immediately. Whether the spacing is uniform is the residual discreteness question, now localized to a point-set condition rather than a global moat assertion. **[¬ACC: Open · candidate route from a uniform-spacing (no-accumulation) bound on persistent positions]**

Remark 9.4 (what this buys — and the shared root the framing must not hide). With B2 reduced to ¬ACC, the operative premise set of the characterization (Theorem 6.1, Remark 6.3) is B1-reverse and ¬ACC. Of the two, B1-reverse carries the analytic weight — that no admissible observable out-separates the resolvent — while ¬ACC is a single combinatorial exclusion: the persistent positions do not accumulate. The species reading the companion paper carried unproved now rests on one substantive analytic premise and one local combinatorial one, against the broad assumption it began as.

But the framing "one analytic, one combinatorial" must not be allowed to imply the two are *independent* debts. They are not obviously independent: B1-reverse's candidate route appeals to the commitment-induced metric resolving separation, and ¬ACC's candidate route appeals to a uniform spacing on persistent positions — and both, pressed to their root, appeal to the substrate's discreteness (finitely many distinguishable states, H1). If a single discreteness principle discharges both, that is an economy worth surfacing in the operator-construction paper. But it also means a reader who doubts that root doubts both premises at once: they are two debts with, plausibly, one underwriting collateral. This shared-root claim is itself a **[Conjectural · structural observation]** — a possibility to be settled in the operator paper, not a position this paper takes; the paper presents B1-reverse and ¬ACC as two named residues for clarity, while flagging — not asserting — that they may share a single point of failure.

---

10. What This Settles for the Census

The characterization changes the standing of the companion paper's reading and nothing downstream of the still-uncomputed sign.

The operative census is identified. Under the operative premises (B1-reverse and, after §9, ¬ACC), the species-census is 4e_stab — the count of spectrally isolated modes — and not the recurrence-census 4e_rec. This is the discharge the companion paper deferred: 4e_stab is the count physics cares about *because* the isolated modes are exactly the durable species, by Theorem 6.1. The recurrence-census 4e_rec remains independent of H1–H7 and refuted by M_census, exactly as before; this paper does not touch it. A gap-closed recurrent mode — m_4 in World A — is now not merely *declared* a non-species but *characterized* as one: it lacks a durable contamination-free record, so the substrate cannot keep it told apart, so it is not a generation though it recurs.

Falsifiability is preserved intact. The characterization says *which* modes are generations — the isolated ones — but not *whether* m_4 is among them. That remains the sign of $\mathfrak{G}(m_4)$, uncomputed and uncomputable without the operator. World A ($\mathfrak{G}(m_4) = 0$, three species) and World B ($\mathfrak{G}(m_4)$

> 0, four species) are both open after this paper exactly as before it. The characterization sharpens what each world *means* — World B is now a genuine fourth durable species, not merely a fourth recurrent level — without making either more or less likely. A test that could only return three would have assumed three; this characterization does not constrain the sign, only its interpretation.

The conditionality is located precisely. After this paper, the species reading rests on B1-reverse and ¬ACC and on nothing else — one substantive analytic premise (no observable out-separates the resolvent) and one combinatorial one (the persistent positions do not accumulate at an isolated value), the latter being all that survives of B2 after §9. A sceptic of the reading has these two named propositions to attack; a defender has two to establish. The candidate routes — B1-reverse from the record metric being ℓ-induced (§4), ¬ACC from a no-accumulation (uniform-spacing) bound on persistent positions (§9) — are the natural next targets, and either would lift the characterization toward [Proven]. As §9.4 flags, the two may not be independent: pressed to their root both appeal to substrate discreteness, so a reader who doubts that root doubts both at once.

---

11. What This Paper Does Not Do

Stated plainly, against over-reading.

It does not fully discharge the premises. B1-forward is near-inherited; B1-reverse — that no admissible observable out-separates the resolvent — is carried [Conditional] with a candidate route only, and is the substantive analytic weight. Most of B2 is discharged in §9 (Propositions 9.1–9.2, Theorem 9.3), leaving the single static residue ¬ACC (persistent positions do not accumulate at an isolated value), also carried with a candidate route only. The characterization (Theorem 6.1) and the census identification (Corollary 6.2) inherit [Conditional on B1-reverse, ¬ACC]. The paper reduces the species reading to these; it does not close them, and notes (§9.4) they may share a discreteness root.

It does not construct the operator or compute the sign. No $(\mathcal{A}, \ell, \mathcal{R})$ is exhibited, no resolvent evaluated, no value of $\mathfrak{G}(m_4)$ produced. Whether m_4 is a generation is untouched and remains [Open · Gate: D5(depth)].

It does not reach 4e_rec. Recurrence-level exhaustiveness stays independent and refuted; the paper changes which census is operative, not the standing of the one it sets aside.

It does not establish a canonical magnitude for the gap or the margin. Only the sign is load-bearing, and only the sign is characterized; the numerical value of $\mathfrak{G}$ and of μ remains metric-dependent and is used nowhere.

It does not address the capacity catalogue [Gate: D5(capacity)], the world-occupancy [Gate: O1], or any gauge, chirality, mixing, or absolute-mass structure.

12. Conclusion

The principal achievement, without inflation. One operational definition of a generation — durable contamination-free distinguishability — built from inherited internality and persistence with no spectral vocabulary, proven well-defined and floor-quantized (the floor applied before the supremum, §3); one characterization theorem identifying that notion with spectral isolation; one by-product discharging the companion's liminf stipulation as forced rather than chosen; one separability certificate proving the equivalence is content and not a relabelling, by withdrawing exactly the substantive premises; and one premise reduction (§9) discharging most of B2 from inherited persistence and integer orbit data, leaving the single static residue ¬ACC. The species reading the companion paper carried unproved is reduced to one substantive analytic premise (B1-reverse) and one combinatorial one (¬ACC) — possibly sharing a discreteness root. No operator is built, no sign computed, no census closed; the conditionality is named, located, narrowed, and handed forward.

The companion paper reduced the generation census to a single computable sign and was candid that the reduction *meant* something only under a reading it could not prove — that the spectrally isolated modes are the generations. A reading carried but unestablished is a debt, and the temptation in paying it is to define a generation as an isolated mode and call the debt settled. That settles nothing; it relabels.

This paper pays the debt the only way that carries content. It defines a generation as physics already understands the word — a distinct species, a mode the substrate can durably tell apart, built from the inherited principles that a distinction is real only if witnessable and that a species must persist — and it mentions no spectra in doing so. It then proves that this independent notion coincides with spectral isolation, exhibits the two notions diverging when either supporting premise is withdrawn so that the coincidence is shown to be a theorem rather than a synonym, and as a by-product shows that the companion's one stipulated choice, the liminf, is forced by what durability means.

The coincidence is conditional, and the conditions are named without hedging. It rested initially on two premises: B1, that the clean spectral extractor is a realizable measurement (its forward half, near-inherited) and that no admissible measurement out-separates the spectrum (its reverse half, the substantive weight); and B2, that the substrate's discreteness keeps a persistent separation from vanishing. Most of B2 is then discharged from inherited content — the competing spectrum is shown to be entirely persistent and (under E-stab) persistent positions eventually stationary — leaving B2 equivalent to a single static exclusion, ¬ACC: the persistent positions do not accumulate at an isolated value. So the residue on which the species reading now rests is one substantive analytic premise (B1-reverse) and one combinatorial one (¬ACC), where before there was an unbounded assumption. The honest caveat is that these two may not be independent: pressed to their candidate routes — B1-reverse from the commitment-induced record metric, ¬ACC from a no-accumulation bound on persistent positions — both appeal to substrate discreteness, and may share a single point of failure. Discharging them, and settling whether they share a root, is the work this paper hands forward, alongside the operator construction that will, behind its firewall, finally return the sign.

What the census means is therefore settled, conditional on those two stated premises; what the census *is* — three durable species or four — remains the one sign left to compute. The instrument was built by the companion paper; this paper certifies that, when it reports, it reports about generations and not about an artefact of its own definition — and narrows to two named premises, possibly one root, the discreteness assumption on which that certification depends.