The Stable-Spectrum Gap Functional

A Spectral Reduction of the Generation Census to a Computable Spectral-Sign Question — Provided the Frozen Audit Determines the Operator

Keith Taylor

VERSF Theoretical Physics Programme

General Reader Summary

Ordinary matter comes in three near-copies — the generations: electron, muon, tau, and the same threefold pattern twice more among the quarks. Earlier papers in this programme showed that a two-mark refinement architecture *generates* three levels, matching that count. They were careful not to claim more. The question they left open, called D5, was sharper: granting that the marks generate three, why could the underlying substrate not admit a *fourth* level the marks do not produce — a hidden generation that the world has simply not shown?

Two previous attempts to close that question failed in the same way, and the failure is instructive rather than embarrassing. Both attempts tried to argue, from the assumptions the programme already holds, that no fourth level can exist. But the programme already contains two internally consistent pictures: one in which the count stops at three, and one in which a fourth admissible level exists alongside the three. When two consistent pictures disagree about a statement, that statement cannot be proved from the assumptions they share — this is a standard fact of logic. The bare count of three is, in this precise sense, *independent* of the current assumptions. No rewording reaches it. New content is needed.

This paper does not supply the missing answer to the question as originally posed. It does something more careful, and more useful: it argues that the question worth answering is a slightly different and stronger one, and it builds the instrument that will answer that. The shift is this. A "level" in the loose sense — anything the substrate admits and that recurs — is not yet a generation. A generation is a *distinct species*: a mode that anchors its own persistent record, separate from the others, and stays separate as the description is refined. The paper defines a property a mode can have — being *spectrally isolated* — that captures exactly that staying-separate, and is a feature of the substrate's own dynamics making no reference whatever to the marks or to how many of them there are. The census that matters for physics is then not "how many recurrent levels" but "how many isolated ones" — how many genuine species.

The paper proves two things about this property. First, it is well-defined: whether a mode is isolated can be settled for any candidate without circular dependence on the answer. Second, it reduces the species-census to a single, sharp, spectral-sign question: *is the disputed fourth mode spectrally isolated, or does its separation from the other three close as the description is refined?* If it closes, the fourth mode is not an independent species but a shadow of the three, and the species-count is three. If it stays open, there are four genuine species, and the world is owed an explanation for why one is unseen. Whether that question is genuinely *computable* — answerable from the programme's frozen content rather than from a tilt secretly built into the substrate model — depends on one further thing being true: that the frozen audit which fixes the substrate model determines this sign rather than leaving it free. The paper does not establish that; it names it as an explicit open condition and states the headline accordingly: the census reduces to a computable sign-question *provided the frozen audit determines the operator*. (One scope caveat in plain terms: this settles only how many *levels* there are — the "depth" of the family tree. A second, separate question, how many *seats* each level admits — its "capacity" — is left to a companion argument and is not touched here.)

The decisive feature is that the instrument can return either answer. A test that can only return three has assumed three. This one cannot, because the consistent picture with a fourth recurrent level does not, by itself, fix whether that level is *isolated* — that is left to be computed from the substrate. What remains is to build the substrate model and read off the sign — and, for that reading to count as an answer the programme earns rather than one it chose, to confirm the model is fixed by the frozen audit rather than tuned. This paper hands that work a single sharp obligation in place of a vague one: not "rule out a fourth generation," but "build the audited model, compute the sign of the gap on the disputed set, and report it honestly whichever way it falls."

Two load-bearing premises are named openly at the outset, because the headline rests on both. The first: that the *isolated* modes are exactly the generations — so that the isolated-count, not the recurrent-count, is the census physics cares about. This is not proved here; it is carried as a stated reading, expected to be discharged from already-established content (the programme's persistence ontology) in a companion result. The second: that the frozen audit fixing the substrate model determines the disputed sign rather than leaving it free — the condition under which "reduces to a computable question" is genuine rather than a relocation of the original difficulty into the choice of model. This one has no discharge route in hand and is left explicitly open, to be settled (or not) when the substrate model is built. Both are flagged in the status table and re-flagged wherever they do work. The paper proves everything that follows once they are granted; it does not smuggle either in as though it were free, and it does not claim the headline reduction except conditionally on the second.

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. There are four grades, ordered **Proven ≻ Conditional ≻ Conjectural ≻ Open** (strongest to weakest); "[Inherited]" marks an imported result carrying its source grade. Where a grade is **Open**, a gate tag names the undischarged hypothesis it waits on — **Gate: D5(depth)** (stable-level exhaustiveness), **Gate:**

D5(capacity) (the capacity catalogue), or **Gate: O1** (the world-to-framework map). A **definitional choice** (marked **[Def]**) is not a truth-apt claim and sits off the grade ordering, like a definition (see §2). Tags and [Def] are presentation, not grades.

Claim	Grade	Gating dependency
Recurrence-level exhaustiveness 4e_rec (no beyond-mark *recurrent* level) is independent of H1–H7 + ontology, and is refuted by M_census	**[Inherited, Proven]**	M_0 / M_census independence (§1)
The sign of the gap functional $\mathfrak{G}$ (isolated vs gap-closed) is well-defined on every recurrent commitment-supported mode	**[Proven]**	resolvent/annulus formulation (§5)
$\mathfrak{G}$ does not read the mark sectoring (a beyond-mark verdict is not a restatement of the mark count)	**[Proven]**	mark-independence certification (§6)
The three mark-levels R_0, R_1, R_2 are spectrally isolated stable modes	**[Inherited]**	Generation Theorem; Orbit Count Theorem
The isolated modes are exactly the generations — so 4e_stab, not 4e_rec, is the operative census (the **generation-identification reading**)	**[Conditional · carried reading; pending companion characterization]**	persistence ontology (R2); staged characterization (§7)
The frozen audit S1–S5 with I1–I5 determines the operator's beyond-mark spectral sign (the **operator-determinacy premise**) — without it, "computable" collapses back toward "independent"; this is a **precondition of Gate D5(depth)**, since that gate's "operator construction" is not a *framework* discharge unless the audit fixes the operator	**[Open · Gate: audit-determinacy]**	whether S1–S5 + I1–I5 fix the operator up to its beyond-mark sign (§1, §8, §11)
Reduction: 4e_stab (N_stab = 3) $\Leftrightarrow$ $\mathfrak{G}$ vanishes on the beyond-mark recurrent set	**[Proven]** (predicate-level; physical force via gen-id row)	on inherited ontology + carrier reduction
The beyond-mark recurrent set is finite, so the reduction is to finitely many signs	**[Inherited · on Orbit Count Theorem]**	I2 (return count K_c bounds closed orbits)
The reduction is non-vacuous: M_census fixes m_4's recurrence but not the sign of $\mathfrak{G}(m_4)$	**[Proven]**	undeterminacy lemma (§8)
N_stab = 3 is not predetermined by I1–I5 (the assumptions do not force it)	**[Proven]**	Lemma 8.1
$\mathfrak{G}(m) > 0 \Leftrightarrow$ eventual-uniform-gap	**[Proven]**	definition of liminf (§4.3)

Claim	Grade	Gating dependency
Def 4.2 takes the liminf, not the limsup	**[Def]**	persistence ontology (§4.3)
Stable-level exhaustiveness 4e_stab — N_stab = 3 (the species-census)	**[Open · Gate: D5(depth); Cond: gen-id reading]**	sign of 𝔊 on the beyond-mark recurrent set; operator construction (Gate D5(depth) **folds the audit-determinacy precondition** — see that row); generation-identification reading (§7)
Capacity termination — the residue-family catalogue	**[Open · Gate: D5(capacity)]**	not addressed by this functional
Observed generations occupy the three levels	**[Open · Gate: O1]**	occupancy map O1
Gauge group, hypercharge, chirality, CKM, PMNS, exact spectrum	**[Open]**	not addressed here

The single most consequential non-definitional premise in the paper is the **generation-identification reading**, now given its own row. Without it the central theorem is a near-tautology about the definition of N_stab; with it, the theorem reduces a physics question to a computation. It is graded **[Conditional]** — contingent on a stated reading the programme expects to discharge from the persistence ontology already in hand — rather than [Proven]. The downstream census claim 4e_stab inherits the meet of its derivation path and so cannot rise above [Conditional] until the reading is discharged; its row therefore carries **two** dependency tags, mirroring the way an Open grade names its gate: **Gate: D5(depth)** for the still-missing operator computation, and **Cond: gen-id reading** for the carried reading. The grade shown is the meet — **[Open]**, the weakest on the path — but the path is now self-documenting: discharging the computation alone lifts the row to **[Conditional]** (not to [Proven]), and discharging the generation-identification reading then lifts it from [Conditional] to [Proven]. The two debts are distinct — one is a missing calculation, the other a missing characterization — and the tag records both so that neither can be silently retired by settling the other. A second, structurally parallel premise — the **operator-determinacy premise** (that the frozen audit fixes the operator's beyond-mark sign), without which "computable" relocates rather than resolves — has its own row immediately below, graded [Open · Gate: audit-determinacy] and developed in §1; it is the twin the "computable, not independent" contrast rests on, named here so the two load-bearing premises are surfaced together rather than one silently. A bookkeeping point this raises, by the same "distinct debts get distinct tags" principle: operator-determinacy also bears on the 4e_stab discharge path — a vanishing-signs result is a *framework* fact only if the audit fixed the operator — so it is arguably a third debt on that row. It is handled as a **precondition of Gate D5(depth)** rather than a third coordinate tag: the gate's "operator construction" is not a framework discharge unless the audit determines the operator, so audit-determinacy is folded *into* the meaning of that gate rather than listed beside it. The census row's gating-dependency cell says so explicitly, so the debt is transparent, not dropped — Gate D5(depth) is read as "construct the operator *and* discharge audit-determinacy," and settling the computation without the latter does not retire the gate.

1. The Independence Diagnosis — and the Two Census Predicates It Forces Apart

The generation census question has resisted two distinct closure strategies. The resistance is not accidental, and naming its source determines both what a genuine discharge must contain and *which* census a discharge can reach.

The Matter Skeleton Theorem exhibited two closure models over the standing hypothesis set H1–H7 and the operational ontology:

- **M_0**, the intended model, in which the admissible levels are exactly the three mark-levels; and
- **M_census**, in which all of H1–H7 hold, the two marks still generate exactly three mark-levels — so the architectural count 4c is intact — and a further admissible *recurrent* level exists that no mark-combination produces, characterized there as admissible and recurrent (i.e. as a member of the recurrent set of §3), so level exhaustiveness fails.

M_census was certified consistent in that paper by independence-model construction.

The two predicates. What M_census refutes must be stated precisely, because the loose phrase "no fourth level" names two different propositions, and the entire soundness of this paper turns on keeping them apart.

4e_rec (recurrence-level exhaustiveness). No admissible *recurrent* level beyond the three mark-levels: $\mathrm{Rec}\_\ell(\mathscr{R}) = \{R_0, R_1, R_2\}$. **4e_stab (stable-level exhaustiveness).** No spectrally *isolated* mode beyond the three: N_stab = 3, i.e. $\mathrm{Stab}\_\ell(\mathscr{R}) = \{R_0, R_1, R_2\}$.

Since the isolated modes are a subset of the recurrent ones ($\mathrm{Stab}\_\ell(\mathscr{R}) \subseteq \mathrm{Rec}\_\ell(\mathscr{R})$, §4), and all three mark-levels lie in both (I2), the predicates are ordered: **4e_rec ⟹ 4e_stab, but not conversely.** A recurrent level beyond the three that is *gap-closed* falsifies 4e_rec while leaving 4e_stab intact. As a census claim, therefore, 4e_stab is the *weaker* statement — it closes the smaller set — even though the predicate "isolated" is *stronger* than "recurrent." This inversion is easy to misread and is stated once, plainly, so that no later claim inflates it.

What is independent, and what is reachable. M_census's defining content is a single non-mark *recurrent* level. So what M_census refutes is **4e_rec**, and it is **4e_rec** that is independent of H1–H7 together with the operational ontology: $M_0 \vDash$ 4e_rec, M_census $\vDash \neg$4e_rec, same theory, hence by the standard model-theoretic fact (a sentence true in one model and false in another is not a theorem of the shared theory) 4e_rec is not derivable from that theory. [Inherited, Proven] No rewording reaches it; the two prior closure attempts each asserted, at their final step, a proposition equivalent to 4e_rec, and the assertion was unavoidable because the content is not in the shared theory.

This paper does **not** claim to reach 4e_rec. It cannot, and nothing below tries. What it does is argue that 4e_rec is not the physically operative census — a recurrent level that is not a distinct species is not a generation — and that the operative census is **4e_stab**, which is a strictly stronger *predicate* applied to the count, is genuinely Open rather than independent, and is computable. The contribution is therefore a *replacement*, not a conversion: an unreachable, independent census (4e_rec) is set aside in favour of a stronger-predicate census (4e_stab) whose value is open and decidable by computation. One consequence must be absorbed honestly and is stated here so no reader infers more than is delivered: a clean discharge of 4e_stab (World A below) closes the species-census at three while **leaving 4e_rec refuted and independent** — M_census still stands, with its recurrent-but-gap-closed fourth level. The species count closes; the recurrence count does not.

The premise the replacement rests on. The replacement is legitimate only if the isolated modes are in fact the generations — only if a gap-closed recurrent level is genuinely not a fourth species but a resonance of the three. Call this the **generation-identification reading**. It is the hinge of the entire paper, and it is *not* free: it is a truth-apt proposition that could fail (some gap-closed recurrent mode might, against the reading, be a physical generation). It is therefore carried explicitly as a **[Conditional]** premise — contingent on a stated reading the programme expects to discharge from the persistence ontology of (R2) already in hand — and not asserted as proved. The companion characterization staged in §7 is what would lift it to [Proven]; until then every census conclusion downstream of it inherits at best the [Conditional] grade. The reading is stated here, graded in the table, and re-flagged at each point it does work, precisely so that the paper's central move cannot be mistaken for a free lunch.

A second premise the contrast rests on. The headline contrast — 4e_rec *independent*, 4e_stab *computable* — carries a second load-bearing premise that is structurally parallel to the first and must be named and graded alongside it, on pain of a double standard. To say 4e_stab is "computable rather than independent" is to claim more than that some concrete operator, once built, has a definite spectral sign on the beyond-mark set; *any* concrete operator trivially does. The substantive claim is that the **frozen audit S1–S5 together with the inherited inputs I1–I5 determines that sign** — that the construction is pinned by frozen content, not chosen with the answer already inside it. Call this the **operator-determinacy premise**. Its necessity is sharpened, not softened, by §8: Lemma 8.1 proves the inherited theory does *not* fix the sign, which is exactly why the sign-fixing content must come entirely from the audit construction — making that construction's determinacy from frozen content the thing the whole "computable" claim leans on. If S1–S5 + I1–I5 admit operators of both signs, then the sign is independent of the frozen content too, and the relocation of §1 is merely from "4e_rec given H1–H7" to "which operator given the audit" — incomplete, not achieved. The §11 firewall ("operator derived from the audit; sign read off; no backflow") is the *mechanism* that would respect determinacy; it is not the *premise* that determinacy holds. Because this paper does not construct the operator, it cannot here establish determinacy, and so the premise is carried as **[Open · Gate: audit-determinacy]** in the status table. Until it is discharged — as [Inherited] if the corpus's Selection Audit already fixes the operator up to its beyond-mark sign, or as a proved lemma in the operator-construction paper — the "two kinds of open" contrast is itself conditional on it, and the paper says so wherever the contrast is drawn.

Why the two premises, billed as parallel, carry different grades. The generation-identification reading and operator-determinacy are structurally parallel — each is a single load-bearing proposition on which a headline claim rests — yet the first is graded [Conditional] and the second [Open], and the reason is principled, not a double standard. Under the §2 calculus a bare undischarged proposition is [Open]; it is only legitimately *carried* as [Conditional] when there is an **in-hand expected discharge route** — a specific, already-held piece of content the programme expects to complete the discharge. The generation-identification reading has one: the persistence ontology of (R2) is already inherited, and the staged characterization is expected to derive "isolated = generation" from it; so the reading is carried, [Conditional]. Operator-determinacy has no such in-hand route — whether S1–S5 + I1–I5 fix the operator's beyond-mark sign awaits either the construction paper or an inherited audit property the corpus has not (here) been shown to supply; with no route in hand, it is a bare undischarged proposition and is graded [Open], not carried. The grade gap is therefore exactly the availability of a discharge route, and the two are parallel in *role* while differing in *epistemic maturity* — which is what the grades are for.

---

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. A conclusion inherits the **meet** (the weakest) of the grades on its derivation path. Open grades carry a gate tag naming the undischarged hypothesis they wait on.

Definitional choices are not truth-apt and sit off the grade axis. A stipulation — a choice of which object to define, such as taking a liminf rather than a limsup — is not the kind of thing that is true or false, so it carries no grade and does not enter the meet. Such choices are marked **[Def]** and justified by motivation, not proof; what *follows* from a definition, once made, is graded normally. The generation-identification reading of §1 is, by contrast, *not* a definitional choice: it asserts that a particular mathematical predicate (spectral isolation) coincides with a physical kind (generation), which is a substantive claim that can be true or false. It therefore carries a grade — [Conditional] — and does enter the meet on every path that uses it.

The following are imported and used, not re-proven.

- **I1 — Refinement architecture.** [Inherited] The active refinement architecture contains two binary exchangeable marks, generating activation counts 0, 1, 2 and the three mark-levels R_0, R_1, R_2 (Generation Theorem). I1 fixes what the marks generate; it is silent on whether the recurrent admissible set contains a level no mark-combination produces — that silence is the census question.
- **I2 — Recurrent structure.** [Inherited] The closure evolution admits a recurrent structure whose closed orbits are counted by the Orbit Count Theorem (return count K_c, commitment count p_v over one return). The three mark-levels possess closed recurrent orbits and are spectrally isolated within the commitment-supported spectrum. The orbit

count being finite, the recurrent admissible set Rec_ℓ(ℛ) is finite; this finiteness is used in §7 to ensure the reduction is to finitely many signs.

- **R2 — Persistence (the inherited criterion, stated verbatim).** [Inherited] *A generation-relevant mode survives admissible refinement — it neither decays into nor merges with another mode's orbit as depth n increases; an eventual condition.* This is orbit-persistence and **nothing more**: the criterion carries no positive-gap term. The distinction between this inherited orbit-persistence and the strictly stronger gap-persistence ($\mathfrak{G} > 0$) introduced in §4 is load-bearing and is drawn out in the note to (R2) in §3; the paper inherits the former and adds the latter, it does not weaken the former.
- **I3 — Operational ontology.** [Inherited] A distinction counts as substrate structure only if it is operationally witnessable at or above the distinguishability floor; physical individuation is internal, not externally labelled (Internality).
- **I4 — Carrier reduction.** [Inherited, Proven on I3] Every admissible generation distinction is a refinement-sector carrier: the scalar, external-label, and transport/cohomological sectors are excluded as generation carriers, so any candidate generation beyond the three mark-levels lies in the refinement sector and hence in the recurrent admissible set of §3. (This is the reduction the carrier analysis supports; it is *not* a discharge, and it leaves the within-refinement question open. It is used here only to guarantee that the recurrent set is the right place to look.)
- **I5 — Commitment functional.** [Inherited] The commitment functional ℓ on the admissible closure state space 𝒜 supplies an intrinsic positive form ⟨·,·⟩_ℓ, induced by the commitment dynamics and independent of any basis decomposition of 𝒜.

The witness model **M_census** of §1 is carried throughout as the object against which non-vacuity and falsifiability are tested.

---

3. The Recurrent Admissible Set

Let (𝒜, ℓ, ℛ) be the closure operator system: 𝒜 the admissible closure state space, ℓ the commitment functional, ℛ the refinement-and-closure evolution operator. Write $\mathscr{R}_n$ for the action of ℛ at refinement depth n, and ⟨·,·⟩_ℓ for the intrinsic commitment form of I5.

Definition 3.1 (Recurrent admissible commitment-supported set). The set Rec_ℓ(ℛ) consists of the modes m ∈ 𝒜 satisfying all three:

(R1) m is recurrent under ℛ — its orbit closes in finitely many ticks; (R2) m persists under admissible refinement — its closed orbit does not decay into, or become literally identical with, another mode's orbit as n increases (inherited orbit-persistence; see the note below); (R3) m couples nontrivially to commitment — ℓ(m) ≠ 0.

Note on (R2) — orbit-persistence, not gap-persistence (inherited criterion verified verbatim). (R2) is the inherited persistence clause. Its formal statement in the companion corpus is: *a generation-relevant mode survives admissible refinement — it neither decays into nor merges with another mode's orbit as depth n increases; an eventual condition.* The operative word "merge" denotes **orbit coincidence**: two modes' closed orbits becoming the literally same orbit, so that one mode ceases to exist as a separate recurrent object. Crucially, the inherited criterion contains **no positive-gap term** — no clause requiring that the spectral distance between two distinct orbits stay bounded below by some $\Delta > 0$. It is orbit-persistence ("the winner remains the winner"), not gap-persistence ("the winner remains *robustly* the winner"); the second implies the first, the first does not imply the second, and the inherited definition asserts only the first. This paper therefore does **not** weaken the inherited criterion: it inherits orbit-persistence exactly as stated and introduces gap-persistence ($\mathfrak{G} > 0$) as a *strictly stronger, new* property layered on top — which is precisely the content of $\mathrm{Stab}_\ell(\mathscr{R}) \subsetneq \mathrm{Rec}_\ell(\mathscr{R})$. A mode can satisfy (R2) at every depth — remaining a genuinely distinct closed orbit, never coinciding with another — while the distance from its orbit to the nearest neighbour liminf-vanishes; that mode is *recurrent and orbit-persistent* (in Rec) yet *gap-closed* ($\mathfrak{G} = 0$). This is the middle case the whole paper turns on; its coherence is established in Remark 3.3, and the terminology is kept disjoint throughout: **"merge / coincide"** is reserved for the (R2)-violating orbit-identity event ($\notin$ Rec), and **"spectrally isolated" / "uniform gap-persistence"** for $\mathfrak{G} > 0$ — the word "survive" is never used for the gap condition.

These are the *candidate* generation-relevant modes. They are not yet counted as independent generations: (R1)–(R3) are necessary but, as the next sections establish, not sufficient.

Remark 3.2 (the candidates include the disputed mode, and the set is finite). By I1 and I2 the three mark-levels lie in $\mathrm{Rec}_\ell(\mathscr{R})$. By I4 any candidate fourth generation also lies in $\mathrm{Rec}_\ell(\mathscr{R})$, since it must be a refinement-sector carrier and conditions (R1)–(R3) are the refinement-sector admissibility conditions. In particular the M_{census} witness mode m_4 — non-mark, admissible, recurrent, commitment-coupled — satisfies (R1)–(R3) and lies in $\mathrm{Rec}_\ell(\mathscr{R})$. **The recurrent set is therefore the correct and complete arena: every place a fourth generation could hide is inside it.** This is precisely why (R1)–(R3) cannot themselves close the count — they admit m_4 by construction, which is exactly the content of $M_{census} \vDash \neg 4e_{rec}$. By the Orbit Count Theorem (I2), the closed recurrent orbits are finite in number, so $\mathrm{Rec}_\ell(\mathscr{R})$ is a finite set; the beyond-mark recurrent set $\mathrm{Rec}_\ell(\mathscr{R}) \setminus \{R_0, R_1, R_2\}$ is therefore also finite. This finiteness is what makes the reduction of §7 a reduction to *finitely many* sign evaluations rather than to a quantifier over an unbounded family.

The work of the species-census cannot be done by recurrence. It must be done by a property that some elements of $\mathrm{Rec}_\ell(\mathscr{R})$ can *fail*.

Remark 3.3 (the middle case is coherent, and the paper is empty without it). The entire advance depends on the existence of a non-empty middle band: modes that satisfy (R1)–(R3) — hence lie in Rec — yet have $\mathfrak{G} = 0$. If no such mode is possible, then $\mathrm{Stab}_\ell(\mathscr{R}) = \mathrm{Rec}_\ell(\mathscr{R})$, $4e_{stab} \Leftrightarrow 4e_{rec}$, and the independence of §1 is re-inherited wholesale, with the gap functional adding nothing. So the coherence of "recurrent and persistent but gap-closed" is not a curiosity; it is the load-bearing possibility, and it must be exhibited as mathematically consistent, not merely asserted. It is consistent. Consider a spectral point whose orbit $\tau_m^{(n)}$ is, at every finite depth n, strictly distinct from every other orbit in the commitment-supported spectrum — so (R2) holds, no coincidence ever occurs — while the nearest-neighbour distance $d_n(m) = \mathrm{dist}(\tau_m^{(n)}, \text{rest of spectrum})$ satisfies $d_n(m) > 0$ for all n but $\liminf_{n\to\infty} d_n(m) = 0$. Such a sequence is routine: take $d_n(m) > 0$ with a subsequence $d_{n_k}(m) \to 0$ and the complementary subsequence bounded below. The link from the controlled quantity d_n to the functional Δ_n is one clause: an isolating annulus centred on $\tau_m^{(n)}$ cannot extend past the nearest other point of the spectrum, so its radius is bounded by the nearest-neighbour distance — $\Delta_n(m) \le d_n(m)$ at every depth. Hence $d_{n_k}(m) \to 0$ forces $\Delta_{n_k}(m) \to 0$, and so $\liminf_{n\to\infty} \Delta_n(m) = 0$ (the moat closes along the subsequence n_k), while at every depth the mode is a separate orbit (R2 satisfied, never merged). By the eventual-uniform-gap reading (§4.3) $\mathfrak{G}(m) = \liminf_n \Delta_n(m) = 0$. The mode is thus recurrent, persistent, and gap-closed simultaneously — distinct at every depth, yet not uniformly separated in the limit. This is exactly the World-A resonance: a fourth recurrent level that is no fourth species. The middle band is non-empty as a matter of spectral possibility, which is all the architecture requires; whether the *actual* operator places m_4 in it is the World-A/World-B question §8 proves the inherited theory leaves open.

---

4. The Gap Functional

The discriminating property is spectral isolation. It is introduced in three steps — the finite-depth separation (§4.1), the depth limit (§4.2), and the consequence of the limit (§4.3) — so that each clause that carries weight is visible. The reading under which spectral isolation is the *generation* condition — the generation-identification reading of §1, carried as [Conditional] — is applied in §7; this section defines the functional and its sign, and depends on no such reading.

4.1 Finite-depth separation, defined without pre-individuation

A naive separation — "the distance from m's spectral value to the rest of the spectrum" — is ill-posed exactly where it must be sharp. When a candidate mode is nearly merged with its neighbours, the assignment of spectral weight to that mode versus to the tails of the others is ambiguous, and any answer to the isolation question would inherit that ambiguity. The functional is therefore defined so that it never assigns m a spectral component.

Let $\zeta \mapsto (\zeta - \mathscr{R}_n)^{-1}$ be the resolvent of $\mathscr{R}_n$ on the commitment-supported subspace (the closure of the span of $\mathrm{Rec}_\ell(\mathscr{R})$ under $\langle\cdot,\cdot\rangle_\ell$). For a recurrent mode m with orbit value region $\tau_m^{(n)}$ (the support of m's closed orbit in the spectral parameter), define an **isolating annulus** for m at depth n to be an open annulus A in the spectral parameter, centred on $\tau_m^{(n)}$, on which the resolvent is analytic — equivalently, which contains no other point of the commitment-supported spectrum — and which separates $\tau_m^{(n)}$ from the remainder of that spectrum.

Definition 4.1 (finite-depth separation). The separation of m at depth n is

$$\boldsymbol{\Delta_n(m) = \sup \{ r(A) : A \text{ is an isolating annulus for m at depth n} \},}$$

the supremum of isolating-annulus radii, with $\boldsymbol{\Delta_n(m) = 0}$ when no isolating annulus exists (m is an embedded or accumulation point of the commitment-supported spectrum at depth n).

The radius r(A) is measured in some fixed metric on the spectral parameter. **Only the sign of the resulting functional is load-bearing** (§4.2, §7), and the sign — whether an isolating moat of positive radius exists at all — is invariant under any bi-Lipschitz-equivalent choice of that metric. The numerical value of $\Delta_n(m)$, and hence of $\mathfrak{G}(m)$, is therefore well-defined only up to such a choice; no result below uses the value, only its vanishing or non-vanishing. (Naming a canonical spectral metric induced by ℓ via the functional calculus would fix the value as well, but that structure is not constructed here and is not needed.)

Definition 4.1 measures isolation as the existence and size of the largest resolvent-analytic moat around m's orbit. It never decomposes the spectrum into "m's share" and "the rest"; an accumulation point simply admits no moat and scores zero. The well-definedness of the sign this buys is proved in §5.

4.2 The depth limit

Definition 4.2 (gap functional). The stable-spectrum gap functional is

$\mathfrak{G}(m) = \liminf_{n\to\infty} \Delta_n(m)$. [Def: liminf, not limsup — see §4.3]

A recurrent commitment-supported mode m is **spectrally isolated** iff $\mathfrak{G}(m) > 0$, and **gap-closed** iff $\mathfrak{G}(m) = 0$. The admissible stable spectrum is

$\mathrm{Stab}\_\ell(\mathscr{R}) = \{ m \in \mathrm{Rec}\_\ell(\mathscr{R}) : \mathfrak{G}(m) > 0 \}$, $N\_stab = | \mathrm{Stab}\_\ell(\mathscr{R}) |$.

Since isolation is the extra condition layered on recurrence, **$\mathrm{Stab}\_\ell(\mathscr{R}) \subseteq \mathrm{Rec}\_\ell(\mathscr{R})$**, and by I2 the three mark-levels lie in $\mathrm{Stab}\_\ell(\mathscr{R})$, so $N\_stab \geq 3$. Since $\mathrm{Rec}\_\ell(\mathscr{R})$ is finite (Remark 3.2), N_stab is finite.

4.3 What the liminf means, and why liminf rather than limsup

[Proven] The condition $\mathfrak{G}(m) > 0$ is *equivalent* to the eventual-uniform-gap condition — there exist N and $\delta > 0$ with $\Delta_n(m) \geq \delta$ for all $n \geq N$. This is not a reading imposed on the functional; it is the definition of liminf, since $\inf_{n\geq N} \Delta_n(m)$ is nondecreasing in N with limit equal to $\liminf_{n\to\infty} \Delta_n(m)$. There is nothing to stipulate here.

[Def] The genuine choice is one step earlier: Definition 4.2 takes the **liminf**, not the limsup. These are different functionals, not two readings of one. The limsup condition (isolation infinitely often) would admit a mode whose moat closes on a cofinal subsequence of depths; the liminf condition excludes it. The choice of liminf is justified by the persistence ontology of (R2): a generation is a species-level structure that must be *eventually and uniformly* separated, and a separation that lapses infinitely often has not survived **as an isolated species, though the mode persists as a recurrent orbit** (it remains in Rec — its orbit never coincides with another; what lapses is the uniform gap, not the orbit's distinctness). The distinction is exactly the one drawn in the note to (R2) and in Remark 3.3: orbit-persistence is an (R2) fact about the mode staying a separate orbit; uniform gap-persistence is the stronger $\mathfrak{G} > 0$ fact, and it is the latter, not the former, that liminf-vanishing destroys. The limsup, average-gap, and persistence-density alternatives are recorded and set aside on the same ground: species distinctness is an eventually-categorical fact, not a statistical or intermittent one — a mode isolated only on a subsequence has the *recurrence* of a species but not the *separation* of one. (That this liminf reading is *forced* by the existing eventual form of (R2), rather than merely motivated by it, is the subject of a characterization staged separately; pending that, the choice is carried as [Def] with the persistence motivation above.)

---

5. Well-Definedness of the Sign

Proposition 5.1 (the sign of $\mathfrak{G}$ is well-defined on Rec_ℓ($\mathscr{R}$)). For every m ∈ Rec_ℓ($\mathscr{R}$), the dichotomy $\mathfrak{G}(m) = 0$ versus $\mathfrak{G}(m) > 0$ is determined and is independent both of any decomposition of the commitment-supported spectrum into per-mode components and of the bi-Lipschitz choice of spectral metric. **[Proven]**

Proof. At each depth n, the set of isolating annuli for m is determined by the analyticity domain of the resolvent $(\zeta - \mathscr{R}_n)^{-1}$ on the commitment-supported subspace — a property of the operator $\mathscr{R}_n$ and the form ⟨·,·⟩_ℓ alone, by I5. Whether that set is non-empty ($\Delta_n(m) > 0$) or empty ($\Delta_n(m) = 0$) is a metric-independent fact; only the *size* of a non-empty moat depends on the metric, and bi-Lipschitz equivalence preserves positivity. The sequence $n \mapsto \Delta_n(m)$ is thus well-defined in [0, ∞] up to bi-Lipschitz rescaling, and its liminf is positive or zero independently of that rescaling. No step assigns m a spectral weight or excises a component: when m is an accumulation point the annulus set is empty and $\Delta_n(m) = 0$ directly, with no labelling required. Hence the sign of $\mathfrak{G}(m)$ is determined.

Corollary 5.2 (the disputed case is well-posed). In particular the sign of $\mathfrak{G}(m_4)$ is well-defined: the question "$\mathfrak{G}(m_4) = 0$ or $\mathfrak{G}(m_4) > 0$?" is a well-posed binary property of the operator ($\mathcal{A}$, ℓ, $\mathscr{R}$), not an artefact of how m_4's spectral share is apportioned or of a metric choice. **[Proven]**

This is the clause the naive separation lacked. The functional is sharp, in its sign, exactly on the mode it exists to judge.

---

6. Mark-Independence — The No-Circularity Certification

The discharge route fails the moment the mark count re-enters the definition of the discriminating property, for then "N_stab = 3" would be "three mark-levels" rephrased. This section certifies that it does not, by the same independence-model method the corpus uses for tightness.

Proposition 6.1 ($\mathfrak{G}$ does not read the mark sectoring). The value $\mathfrak{G}(m)$ is computable from ($\mathcal{A}$, ℓ, $\mathscr{R}$) alone and is invariant under any relabelling or re-decomposition of the refinement-mark structure on $\mathcal{A}$ at fixed $\mathscr{R}$. In particular, the mark sectoring is not an argument of $\mathfrak{G}$: the sign of $\mathfrak{G}$ on a beyond-mark mode is not computed from how $\mathcal{A}$ is partitioned into mark-activation sectors. **[Proven]**

Proof. Every ingredient of $\mathfrak{G}$ — the resolvent $(\zeta - \mathscr{R}_n)^{-1}$, the commitment-supported subspace (the ℓ-closure of Rec_ℓ($\mathscr{R}$)), the spectral parameter, the analyticity domain defining isolating annuli — is a function of $\mathscr{R}$ and ℓ. The two-mark architecture I1 is a *separate* structure on $\mathcal{A}$ (a distinguished decomposition into mark-activation sectors); it appears nowhere in the construction of §4. Toggling the mark decomposition — relabelling the marks, or replacing the two-mark sectoring with any other admissible sectoring while holding $\mathscr{R}$ and ℓ fixed — leaves the resolvent, the form, and hence the sign of $\Delta_n(m)$ and $\mathfrak{G}(m)$ unchanged. Therefore the sign $\mathfrak{G}$ assigns to a beyond-mark mode is not read off the mark sectoring, and a beyond-mark verdict is not a restatement of the mark count.

Scope of Proposition 6.1 — what it does and does not say. It does *not* say $\mathfrak{G}$ is logically independent of the number of mark-levels. The number of mark-levels is itself a spectral feature of $\mathscr{R}$ — I1 says the refinement architecture *generates* R_0, R_1, R_2, and I2 says they are spectrally isolated in $\mathscr{R}$'s spectrum — so $\mathfrak{G}$, being a functional of $\mathscr{R}$, of course depends on what spectrum $\mathscr{R}$ carries, three isolated mark-levels included; indeed N_stab ≥ 3 is exactly that dependence (Remark 6.2). The claim that is both true and sufficient is the narrower one: the mark *sectoring* — the chart that labels which modes are mark-activations — is not an input, so the sign $\mathfrak{G}$ returns on a *beyond-mark* mode cannot be a relabelling of "there are two marks." That narrow claim is all §7 requires, and it is what is proved.

Remark 6.2 (what mark-independence licenses). I2 already supplies, as inherited content, that the three mark-levels are spectrally isolated — so N_stab ≥ 3 is handed to the paper by the resolvent spectrum, not independently rediscovered. The operative discharge is therefore not a coincidence of two independently computed counts both landing on three; it is the *vanishing of the sign* of $\mathfrak{G}$ on the beyond-mark recurrent set — the absence of any beyond-mark isolated mode — exactly as Corollary 7.3 and §11 state. What Proposition 6.1 buys is narrower and sufficient: because the mark count is provably not an input to $\mathfrak{G}$, a computed verdict that $\mathfrak{G}$ vanishes there is a genuine fact about the substrate's spectrum and not a restatement of "there are two marks." Mark-independence licenses reading the sign as a discharge rather than a tautology; it does not require, and the paper does not claim, a two-count coincidence.

---

7. The Reduction Theorem

Theorem 7.1 (Spectral reduction of the species-census). Assume I1–I5. Then stable-level exhaustiveness is equivalent to the vanishing of $\mathfrak{G}$ on the beyond-mark recurrent set:

4e_stab ⇔ N_stab = 3 ⇔ $\mathfrak{G}(m) = 0$ for every $m \in$ Rec_ℓ($\mathscr{R}$) \ $\{R_0, R_1, R_2\}$.

[Proven, on inherited ontology + carrier reduction]

Proof. By I2 the three mark-levels lie in Stab_$\ell(\mathscr{R})$: each has a closed recurrent orbit isolated within the commitment-supported spectrum, so $\mathfrak{G}(R_i) > 0$ for i = 0, 1, 2, giving $\{R_0, R_1, R_2\} \subseteq$ Stab_$\ell(\mathscr{R})$ and N_stab ≥ 3.

For the equivalence, N_stab = 3 means Stab_$\ell(\mathscr{R}) = \{R_0, R_1, R_2\}$, since those three always lie in Stab_$\ell(\mathscr{R})$. By Definition 4.2, Stab_$\ell(\mathscr{R}) = \{R_0, R_1, R_2\}$ holds iff no m ∈ Rec_$\ell(\mathscr{R}) \setminus \{R_0, R_1, R_2\}$ has $\mathfrak{G}(m) > 0$ — that is, iff $\mathfrak{G}$ vanishes on the beyond-mark recurrent set. This is the stated three-way equivalence, and 4e_stab is by definition N_stab = 3.

Remark 7.2 (where the physical content sits, and what grade it carries). The equivalence above is, given the definitions, near-immediate — its work is done by Definition 4.2, not by a long argument. The substantive claim is the **generation-identification reading** carried from §1 and applied here: that the *isolated* modes are the generations, so that 4e_stab — and not the unreachable 4e_rec — is the census physics cares about. Under that reading, a beyond-mark recurrent mode m with $\mathfrak{G}(m) = 0$ is not a fourth species but a resonance or accumulation point of the three: it recurs and stays a distinct orbit, but its isolating gap from them closes in the refinement limit, so it individuates no distinct species. This reading is what gives Theorem 7.1 its bite, and it is graded **[Conditional]** (status table; §1): it is a truth-apt identification of a mathematical predicate with a physical kind, expected to be dischargeable from the persistence ontology of (R2) but not discharged here. Consequently the physically loaded form of Theorem 7.1 — "the species-census equals the isolated-count" — inherits the meet and is itself **[Conditional]** pending the companion characterization; only the predicate-level equivalence "N_stab = 3 ⟺ $\mathfrak{G}$ vanishes on the beyond-mark set" is [Proven] outright, because that part is true by the definition of N_stab regardless of the reading. The present paper proves the predicate-level equivalence and carries the identification.

Corollary 7.3 (the question reduced to finitely many signs, with m_4 the disputed representative). Under I1–I5, the discharge of 4e_stab reduces to the sign of $\mathfrak{G}$ on the beyond-mark recurrent set Rec_$\ell(\mathscr{R}) \setminus \{R_0, R_1, R_2\}$, which is *finite* by Remark 3.2 (Orbit Count Theorem, I2). Explicitly:

$\mathfrak{G} \equiv 0$ on Rec_$\ell(\mathscr{R}) \setminus \{R_0, R_1, R_2\}$ ⟹ no beyond-mark isolated mode; the species-count closes at three (4e_stab discharged, World A). $\mathfrak{G}(m) > 0$ for some such m ⟹ m is a further stable mode; the species-count is at least four (4e_stab refuted, World B).

The M_census witness m_4 is the *canonical disputed representative* of that set — the mode whose status the prior independence result puts in question — and the worlds are named through it for concreteness. But the reduction is to the sign of $\mathfrak{G}$ on the whole finite beyond-mark set, not to m_4 alone; should the operator's recurrent set contain further beyond-mark modes, each contributes a sign, and 4e_stab holds iff all of them vanish. The finiteness from I2 is what guarantees this is a finite computation rather than an unbounded quantifier. In either world, 4e_rec remains refuted and independent (§1): M_census stands, with m_4 recurrent. The grade of "N_stab = 3" is thus **[Open · Gate: D5(depth); Cond: gen-id reading]** — the gate naming the missing computation, the Cond tag naming the carried generation-identification reading on which its physical force depends — now in *computational* form, the gate being the evaluation of finitely many signs against the operator construction, not an unreachable derivation. **[Proven that the reduction holds; the signs are Open; the physical identification is Conditional.]**

---

8. Non-Vacuity and Falsifiability — Why This Is Not Three by Assumption

A reduction is worthless if the operator's structure already forces the answer, and a discharge route is worthless if it cannot fail. Both are addressed by examining what M_census does and does not fix about m_4.

Lemma 8.1 (M_census leaves the sign of $\mathfrak{G}(m_4)$ undetermined). The witness model M_census specifies that m_4 is admissible, recurrent (R1), refinement-persistent (R2 — its orbit stays distinct, never coinciding with another), and commitment-coupled (R3) — i.e. that $m_4 \in \mathrm{Rec}\_\ell(\mathscr{R})$, which is exactly its content M_census ⊨ ¬4e_rec. It does *not* specify the sign of $\mathfrak{G}(m_4)$: the construction that makes m_4 a consistent fourth element of $\mathrm{Rec}\_\ell(\mathscr{R})$ constrains its membership in the recurrent set, not the eventual behaviour of its isolating annulus. **[Proven]**

Proof. M_census was built to break recurrence-level exhaustiveness 4e_rec while preserving H1–H7 and the mark count 4c; its defining content is the existence of a non-mark admissible *recurrent* level. Spectral isolation in the refinement limit is a strictly finer datum than membership in $\mathrm{Rec}\_\ell(\mathscr{R})$, and the gap formalized in the note to (R2) is exactly why: (R1)–(R3) constrain m_4's *orbit* — that it closes (R1), that it never coincides with another orbit (R2), that it couples to commitment (R3) — while the sign of $\mathfrak{G}(m_4)$ is whether the orbit's isolating moat stays uniformly open in the limit (§4.3), a property of the *separation* between distinct orbits that (R2) does not touch. By Remark 3.3 a mode can satisfy all of (R1)–(R3) with either sign of $\mathfrak{G}$. Nothing in the specification of M_census fixes that sign — its non-mark recurrent level is compatible both with a persistent moat (isolated, $\mathfrak{G}(m_4) > 0$) and with a moat that closes in the limit while the orbit stays distinct (a resonance of the three, $\mathfrak{G}(m_4) = 0$). Hence both signs are consistent with M_census, and 4e_stab is not entailed either way by the content M_census carries.

Corollary 8.2 (the reduction is non-vacuous and the instrument is not predetermined). Because M_census fixes m_4's recurrence but not the sign of $\mathfrak{G}(m_4)$:

1. **Non-vacuity.** $\mathfrak{G}(m_4) = 0$ is not entailed by the inherited ontology — it is genuine new content, available only from the operator. The reduction does not secretly already hold. **[Proven]**
2. **Not predetermined to three.** The assumptions I1–I5 do not force $\mathfrak{G}(m_4) = 0$; the value $N\_stab = 3$ is not fixed in advance by the shared theory. A return of $N\_stab \geq 4$ is not precluded by anything assumed. **[Proven]**

The inference in (2) uses one step worth stating explicitly: **M_census ⊨ I1–I5.** This holds because I1–I5 are inherited theorems over the base theory H1–H7 + ontology, and M_census is a model of that base theory (§1), so it satisfies every theorem of it, I1–I5 included. That is what licenses a fact about M_census — that it leaves the sign of $\mathfrak{G}(m_4)$ free (Lemma 8.1) — to bound what I1–I5 can entail: if I1–I5 forced $\mathfrak{G}(m_4) = 0$, that would hold in every model of I1–I5, hence in M_census, contradicting Lemma 8.1. So I1–I5 do not force it.

Note the precise content of (2): what is proven is *non-predetermination* — the assumptions do not force three — not the exhibition of an operator that realizes four. No operator is constructed (§10), so the paper does not claim to display a four-mode case; it claims that nothing assumed rules one out. This is the same discipline the corpus applies throughout: non-entailment, not exhibited existence.

A scope note, since §7 reduced 4e_stab to the sign of $\mathfrak{G}$ on the whole finite beyond-mark set while the present lemma speaks only of m_4. Non-predetermination is established *via m_4 as the representative beyond-mark mode* — and m_4 alone suffices, because a single undetermined sign in the set is enough to show the verdict N_stab = 3 is not forced by I1–I5. M_census supplies exactly that one mode and constrains exactly its recurrence; it says nothing about any further beyond-mark modes, because it does not posit them. Should the operator's recurrent set turn out to contain beyond-mark modes other than m_4, their signs are additional evaluations the construction will report — further entries in the finite computation of Corollary 7.3 — not further quantities M_census leaves undetermined. The asymmetry is therefore intended, not an oversight: §7 widens the net to the whole finite set because 4e_stab requires every beyond-mark sign to vanish, whereas §8 needs only one undetermined sign to defeat predetermination, and m_4 is the one M_census hands us.

This non-predetermination is the decisive property and the difference from the two prior attempts. Those attempts asserted, from the shared theory, a proposition (equivalent to 4e_rec) that could only come out one way; the assertion was the assumption. Here the discriminating sign is provably undetermined by the shared theory (Lemma 8.1) and provably not read off the mark sectoring (Proposition 6.1) — so a beyond-mark verdict cannot be a relabelling of "there are two marks," and its eventual value is information, not stipulation. **A test whose three-or-four verdict is not fixed in advance by the assumptions, and whose beyond-mark verdict is not read off the sectoring it would otherwise merely echo, is an instrument; this one is.**

Remark 8.3 (the two worlds, named). Two computational outcomes are open and both are admissible results of the programme. *World A:* $\mathfrak{G}$ vanishes on the beyond-mark recurrent set (in particular $\mathfrak{G}(m_4) = 0$) — the non-mark recurrent levels are gap-closed resonances, the species-count is three, 4e_stab is discharged, and 4e_rec remains refuted (m_4 still recurs). *World B:* $\mathfrak{G}(m) > 0$ for some beyond-mark m (e.g. m_4) — that level is spectrally isolated, the species-count is at least four, and the programme owes either a fourth generation or an account of why it is unobserved (an O1 question for the fourth mode). The present paper does not predict which. It guarantees only that the operator construction will decide, and that whichever it decides is reportable without rewriting the instrument.

---

9. Scope — Depth, Not Capacity

The gap functional takes the refinement-*depth* limit $n \to \infty$ and adjudicates whether the admissible *stable levels* exhaust at three (4e_stab). It is silent on the *capacity* face of the census — whether the admissible capacities {1, 2, 3} terminate, which gates the residue-family catalogue. [Open · Gate: D5(capacity)]

A clean result that $\mathfrak{G}$ vanishes on the beyond-mark recurrent set therefore discharges the depth face — stable-level exhaustiveness — and no more. The catalogue and the capacity census remain behind their own gate, and would require an analogous capacity-side argument, a distinct object not constructed here. The paper is scoped to claim exactly the half it reaches.

---

10. What This Paper Does Not Do

Stated plainly, against the temptation to read the reduction as the discharge.

It does not close the census. It constructs no operator ($\mathscr{A}$, ℓ, $\mathscr{R}$) concretely, computes no resolvent, and evaluates no sign of $\mathfrak{G}$ on the beyond-mark set. The discharge is the signs; the signs are owed to the operator construction.

It does not prove the generation-identification reading. That the isolated modes are exactly the generations is carried as a [Conditional] premise (§1, §7), expected to be discharged from the persistence ontology in a companion characterization; until that is done, the physical reading of the central theorem rests on it, and the paper says so wherever the reading does work.

It does not establish the operator-determinacy premise. That the frozen audit S1–S5 with I1–I5 fixes the operator's beyond-mark spectral sign is carried as [Open · Gate: audit-determinacy] (§1, §11); the paper builds no operator and so cannot here show the audit pins one. Until that is discharged, the "computable, not independent" contrast is conditional on it: the relocation of the census from independence to computation is achieved only to the extent the audit determines the sign rather than leaving a both-signs choice of operator.

It does not reach 4e_rec. Recurrence-level exhaustiveness is independent of H1–H7 and refuted by M_census, and stays so under either computational outcome. The paper replaces it as the operative census, it does not derive it.

It does not engage the frozen Selection Audit S1–S5. The audited closure the programme committed to is the evaluation of the operator's recurrent spectrum against S1–S5; this paper supplies the functional that computation will report, not the computation.

It does not address the capacity catalogue [Gate: D5(capacity)], the world-identification [Gate: O1], or any of the gauge, chirality, mixing, or absolute-mass structure. It reduces one face of one gate.

It does not refute M_census. M_census remains a consistent model of H1–H7; what the paper supplies is a *strictly stronger admissibility predicate* — spectral isolation — under which the species-census becomes computable. M_census is the witness that keeps the predicate honest, not a target to be eliminated.

---

11. Position in the Programme

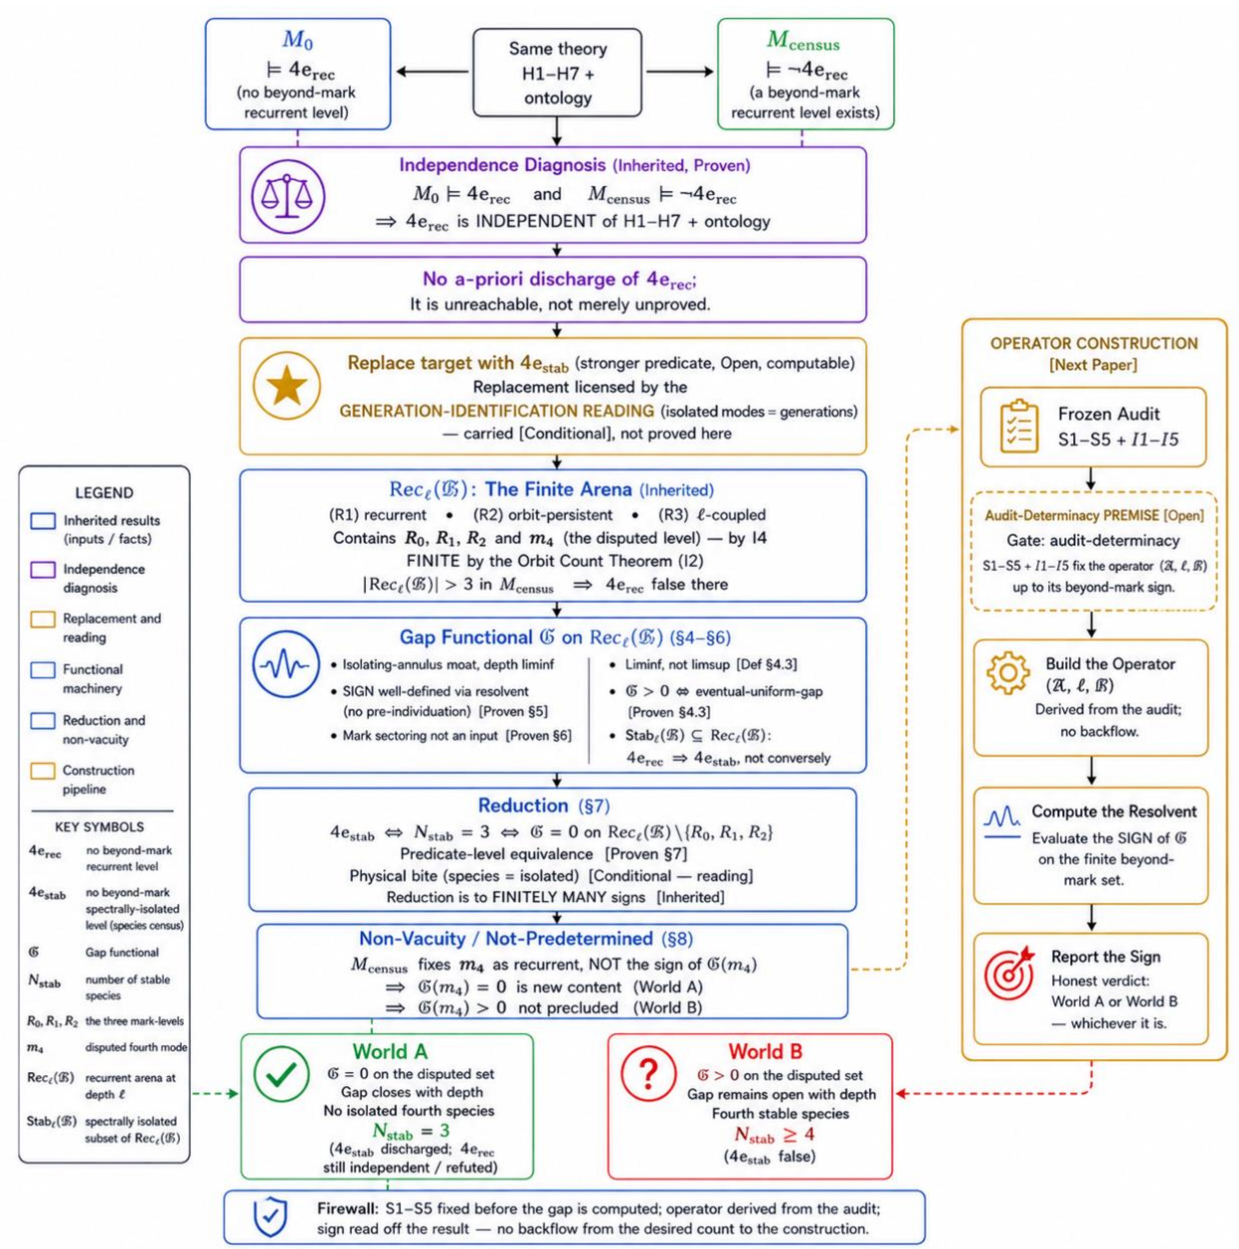

The principal achievement, without inflation. One functional — spectral isolation, defined on the substrate operator alone — with its *sign* proven well-defined and proven not read off the mark sectoring (call this **sectoring-blind**: the beyond-mark verdict is not computed from how $\mathcal{A}$

is partitioned into mark-activation sectors — *not* the stronger and false claim that $\mathfrak{G}$ is independent of how many isolated mark-levels $\mathscr{R}$ carries, since $N_{stab} \geq 3$ is exactly that dependence, §6), and one reduction theorem identifying the stable-level census 4e_stab with the vanishing of that sign on the (finite) beyond-mark recurrent set, where 4e_stab is a strictly-stronger-predicate replacement for the independent, unreachable recurrence-census 4e_rec. The predicate-level equivalence is [Proven]; the physical reading that makes it a *generation* census rests on the generation-identification reading, carried [Conditional]; and the claim that the resulting question is *computable rather than relocated* rests on the operator-determinacy premise, carried [Open · Gate: audit-determinacy]. The discriminating sign is provably not predetermined by the inherited assumptions and can come out either way. No operator is built and no sign is produced; the species-count of three is not closed. What is delivered is the instrument, the proof that its sign is well-posed and sectoring-blind, the finiteness that makes the reduction a finite computation, and the relocation of the census from an independence obstruction (4e_rec, unreachable) to a computability obligation (the signs of $\mathfrak{G}$ on the beyond-mark set) — a relocation that is genuine *to the extent the two named premises hold*, and conditional on them where they do not yet. Those are different kinds of open, and the paper claims only the second, with both premises flagged at every step they carry weight.

---

12. Conclusion

The two earlier closures failed at the same joint because they could not have succeeded: recurrence-level exhaustiveness 4e_rec is independent of the hypotheses they argued from, and an independent proposition is not reachable by rephrasing its premises. The honest response is not a third reframing of 4e_rec but a change of target — to identify the census physics actually cares about, the count of distinct *species*, and to make that one computable. That change of target is legitimate only under the generation-identification reading, which the paper carries openly as a [Conditional] premise rather than burying.

The stable-spectrum gap functional is that change of target made precise. It names a property of the substrate's own dynamics — that a mode's separation from the rest of the commitment-supported spectrum stays uniformly open under refinement — defined through the resolvent so that its sign is well-posed even on the disputed mode, and built from the operator and the commitment form alone so that the mark sectoring is provably not among its inputs. Under this functional the species-census 4e_stab becomes a single finite question with a computable answer: does $\mathfrak{G}$ vanish on the beyond-mark recurrent set, or is some beyond-mark level spectrally isolated? The reduction theorem makes the predicate-level equivalence exact; the non-vacuity lemma proves the sign is not already fixed by what the programme assumes; the mark-independence proposition proves a beyond-mark verdict would be a fact about the spectrum rather than a restatement of the mark count; the finiteness of the recurrent set makes the obligation a finite computation rather than an unbounded quantifier.

Two kinds of open should not be conflated, and the paper's value is in separating them. The recurrence-census 4e_rec is *independent* — unreachable from the axioms, and it stays unreached here. The species-census 4e_stab is *computable* — open only because an operator has not yet been built and finitely many signs not yet read — *conditional* on the generation-identification reading for its physical force, and computable-rather-than-relocated only insofar as the operator-determinacy premise holds (that the frozen audit fixes the operator's beyond-mark sign rather than leaving a both-signs choice). Both premises are named, graded, and on the table; neither is buried. This paper does not close 4e_stab; it relocates the question from the first kind of open to the second, names the two premises the relocation rests on, and proves the instrument that will close it is well-posed and sectoring-blind. What remains is the operator construction and the evaluation of the signs of $\mathfrak{G}$ against the frozen audit, behind the firewall that keeps the construction from being tuned to the answer — and, with it, the discharge of audit-determinacy that certifies the firewall has frozen content to read from. That computation can close the species-census at three or open it to four, and the value of this paper is precisely that it cannot do so in advance. The instrument is built and certified. The signs are the thing left to compute — and they are now, given the two named premises, things that can be computed.