The Eigenmode Decision

A Preregistered Protocol for the Completion-Density Ordering, the Single-Quantum Test, and the Bound Verdicts of the Strike Clause

Keith Taylor *VERSF Theoretical Physics Programme*

General Reader Summary

The appointment card

The Mass Hierarchy Theorem ended at a counter. The route's claim — that within every denomination the editions weigh in the watermark's order, blank lightest, both-struck heaviest, no ties — was pressed until everything stood on inherited structure or theorem except one small import: that each strike raises the bill, never lowers it, never rides free. And that import was carried to the one office whose scales can decide it: the assay office, whose instruments weigh not the pile of committed coin but the minting itself — how many completed stampings a pattern realizes per turn of the substrate's press. The office's computation was scheduled. The route's question — *is the strike on the bill?* — was left holding the last word.

This paper is the appointment card, filled out and notarized before the office opens. It does not report the weighing; it binds it. It fixes which scales are used, which patterns are placed on them, how the patterns are named before anyone reads the dial, what counts as the ladder winning, what counts as the ladder losing, what counts as the scales themselves failing, and the exact price of every outcome — signed, in public, before any number exists. A weighing whose victory conditions are written after the weighing decides nothing. This card exists so that, when the office reports, the report will mean what it says.

One sentence of translation, for the reader arriving from outside the route: the accounts are the particle types; the denominations are the columns of the world's table (up-type quarks, down-type quarks, charged leptons, neutrinos); the watermark's marks are the refinement positions; a strike is a held mark — one step of generation depth; the minting rate is completion density — completed cycles per substrate tick; the stamp is the fixed action quantum each completed cycle commits; and the appointment card is this protocol.

What the office's scales measure

The office's scales are blind in exactly the way the route requires. The computation solves for the substrate's admissible self-sustaining patterns — the eigenmodes — from closure structure alone: which denomination a pattern belongs to, how many marks it holds, and the dynamics of the press. No mass enters. Each pattern comes off the bench already labelled by its own structure, and beside each label, one number: its minting density. The ladder's claim is then one sentence read off the bench: **within every denomination, the deeper editions mint more densely.** If the computed densities land in the watermark's order in every denomination, the route's last import retires into machinery, and the fermion masses become what the route has claimed they are — the running cost of keeping structure closed, computed from the ground up. If any denomination mints out of order, the clause is struck, and the paper that carried it has already said, in print, what falls and what stands.

There is a second dial on the same scales. The office's foundation is one fixed stamp — every completed cycle commits the same whole quantum, in every pattern, always. The masses, divided by the computed densities, must therefore come out in one constant ratio across the entire table. That is the foundation's own test, named by its own author as the way it dies. The card binds it too.

The wager, sealed

Three denominations have already been weighed by the world; for those, the office's reading is a trial the route could fail but the world has already graded. The fourth — the neutrinos — is a sealed double envelope: the world's experiments have not yet decided the order, and the office has not yet computed it. Whatever the office finds there is a wager in numbers, lodged before the world's own count comes in. The card binds the wager's publication regardless of which way it points — including the uncomfortable case where the office and the route's earlier forecast disagree and must both wait for the world.

What the card cannot decide

The card is an instrument, not a result. It moves no marker, pays no claim, and strikes nothing. It adds not one line to the route's ledger of imports — every clause it tests is already on the books at a stated grade, and every price it quotes was published in the papers that carried the claims. What it adds is the one thing a decider must have and almost never does: a complete table of consequences, priced before the number exists. The office is open. The next word belongs to the number.

Abstract

The headline. The Mass Hierarchy Theorem pressed the route's residue to two clauses about one function — $\Delta(n) \geq 0$ and $\Delta(n) > 0$, the sign and strictness of the within-column cost grading $B(n)$ — and constructed their conditional discharge at the completion door: $m \propto C$ [ANCH-COMP, at its read grades], C strictly increasing in refinement depth [the density clause (α')], the clause stated in exactly the variables the anchoring source's eigenmode program returns. This paper is that program's appointment, bound before it reports: a **preregistered protocol** fixing the eigenproblem's requirements — **R1** construction from closure structure and the anchoring dynamics alone, the mass column and its K_c translations excluded; **R2** intrinsic labels, columns

by superselection and depth by the interface's own refinement labelling, with the run **void** for the ordering test if depth is not intrinsically labelled; **R3** outputs (p_v, K_c) per mode with a frozen uncertainty model, C = p_v/K_c; **R4** the input fence; **R5** the freeze on all discretionary numerics, one pre-unblinding implementation audit and nothing after; **R6** staged unblinding, the ordering test scored mass-blind before any mass is loaded; **R7** full publication of every cell — pass, fail, unresolved, void — with silent reruns voiding the instrument. Two tests with frozen margins: **T1, the ordering test**, deciding (α′) per column — $C(R_0) < C(R_1) < C(R_2)$, each adjacent log-gap resolved beyond κ·(relative δC) — and **T2, the uniformity test**, deciding the single quantum at two grades — per-column and global constancy of m_i/C_i within ε_col and ε_glob — with the **Lattice Lemma** [Proven — under the card's attested compatibility condition: ε_col plus the log-form margins strictly inside every adjacent charged-column mass splitting]: at a charged column, per-column proportionality entails the ordering — so no cell exists, under the attested condition, where the magnitudes fit and the order fails, the lemma's forbidden joint outcomes converting at runtime into implementation-audit triggers rather than verdicts. The **Verdict Table** runs at two grains — column verdicts {C-pass; C-prop, the proportionality strike; C-ord, the ordering strike; C-open, undecided and re-registrable at sharper numerics; C-void} and global verdicts {G-glob; G-quant, the quantum strike, its trigger frozen as a numerical criterion or declared the protocol's one discretionary seam} — with the distinguished configurations bound with their strike orders: **V1** (full payment — (α′) computation-paid; the completion door discharges REF-CLO(ii) and (iii) into the anchoring chain's conditionality; the completion-density corollary, the conditional Shared Strike theorem, and the third road's denominator closure all paid; Pℏ computation-consistent; the K_c numerics converted from translations to predictions; the R_0 densities delivered as computed base assays; the neutrino wager lodged in numbers); **V2a** (ordering paid, the cross-column ratio open at the source's named address — the within-column discharge complete, which is everything the Mass arc consumes); **V2b** (ordering confirmed as ordering, the door's proportionality failing at a column — REF-CLO stays imported, the residue relocated to the register's mass identification); **V3** (a resolved inversion at a charged column — one frozen-state implementation audit, then the strike lands jointly at (α′) and the depth-labelling as one seam clause; the door closes; the binding world's denominator road gains its exhibit in exactly the sparser-cycling shape the Mass paper named; the Mass Hierarchy Theorem stands at its import, the stock door untouched); **V4** (the quantum strike — T2 failing in the source's own named pattern, no single η under its stated bridge structure; Pℏ struck at its own surface, the door's substructure falling, the theorems beneath standing); **V5** (void — R2 unmet at the construction grain or a breach under the protocol's own refuters; undecided columns record as C-open, re-registering as numerics sharpenings under the same construction, while an R2 cure registers as a new construction under R1; nothing decided; publication required). The **neutrino cell** is the protocol's double-blind square: computed order against a world order not yet reported, its four outcomes priced in advance with the Mass paper's bound strike order (indexing clause first, identification second) integrated, the computed-inverted-while-the-world-waits cell expressly published as a standing disagreement rather than suppressed; the column registering for T1 and the square only — no absolute scale exists for its T2, which stands deferred as the prediction of record — and the fourth cell's bound re-scoring exempt from the re-identification breach, only the inherited comparison map moving, never a computed quantity. The **fences**: no cell decides the toll world, the spectrum's cost-reading [YUK-COST stays at conjecture in every cell — the door discharges through the anchoring register, and the firewall outlasts full payment], the stock door, CF, or the mixing arc; even at V1 the Mass arc consumes orderings only, the magnitudes the anchoring programme's own payment behind that arc's firewall. **W0 is unmoved in every cell** — V1 extends conditional structure and V3/V4 strike conditional structure, neither of which presses a verdict node. The protocol adds **no physical import**: its parameters are procedural and frozen at registration; its bound implications are [Proven — immediate from the inherited clauses, the binding the content, executed before the number exists]; its own refuters are the four breaches — fence, freeze, re-identification, selective publication — each voiding the instrument rather than the targets.

Epistemic markers: [Inherited] imported from prior VERSF papers; [Imported-External] imported from standard mathematics or physics outside the programme, carried at the external source's standing; [Proven] established here; [Conditional] holding under stated inputs; [Conjectural] motivated but unproven; [Open] undecided.

Table of Contents

---

1. Introduction — What This Paper Is, and What It Refuses to Be

The Mass Hierarchy Theorem closed its books with one import on the ledger and one sentence about it: *every door this paper leaves open opens onto a number someone can produce.* The primary door is the completion door — the Commitment Burden Theorem in density form, m ∝

C with C strictly increasing in refinement depth under the density clause (α′) — and the number behind it is the output of the anchoring source's eigenmode program: computed completion densities, per mode, derived without mass input. The Mass paper stated the clause in the computation's own variables precisely so that what the clause asserts and what the run reports would be one sentence. What it could not do, being the paper that carries the clause, is referee the run. A claimant who also writes the victory conditions after the contest is not refereeing; he is narrating.

This paper is the referee's instrument. It is a **preregistration**: a complete, public binding of the eigenmode computation's inputs, identifications, tests, margins, verdicts, payments, and strike orders, executed before the computation reports. It refuses to be three other things, and says so at the door. It is not a result — no density appears in it, and it pays no claim. It is not a derivation — it adds no physical content, proves nothing about the world, and consumes every clause it touches at the grade its source paper assigned. And it is not a hedge — the verdict table contains the route's defeat in the same ink as its victory, priced identically, with publication bound for every cell including the ones the route would rather not see. The route's standing discipline is that every import carries its refutation form in public; a computation that decides an import inherits the obligation in operational form. The strike conditions must exist, frozen, before the number does — otherwise the number, whichever way it points, is testimony coached after the fact.

One structural note fixes the paper's address. The computation is the anchoring programme's — its operator, its dynamics, its outputs. The stakes are the Mass arc's — (α′), the completion door, REF-CLO's retirement, the completion-density corollary. The protocol therefore lives at the seam between the two, which is exactly where a decider belongs: neither party owns the referee, and the verdict table binds both. The Seam Consistency Theorem made the Mass paper and the spectrum programme mutually falsifiable; this protocol does the same for the Mass paper and the anchoring programme, at the seam's third surface, with the addresses written down first.

2. The Stakes, Inherited — One Computation, Itemized

The Mass paper concentrated an unusual load onto one run, and the protocol begins by itemizing it, each item at its inherited grade, so that no payment and no strike can later be discovered to have been riding unannounced.

The density clause (α′) [the Mass paper's clause; Open]: *holding a refinement mark strictly increases flip-completion density — $C(R_0) < C(R_1) < C(R_2)$ within every populated column.* T1's target, verbatim.

The Commitment Burden Theorem at the completion door [Conditional — on (α′) and ANCH-COMP at its read grades]: m ∝ C, C strictly increasing, hence Δ(n) > 0 at every column. Discharges under T1 with per-column proportionality (T2-col); the discharge's residual conditionality is the anchoring chain's — P$\hbar$, the Layer B bridges [Imported-External at the source], the depth-labelling at R2.

REF-CLO, clauses (ii) and (iii) [Open — the sign and strictness of B(n)]: retire into the chain above under the door's discharge; stay imported in every other cell, the stock door standing independently.

The completion-density corollary [the third seam surface, logged at the Mass paper's §12]: C strictly increasing with refinement depth at every populated column — T1 read at the seam. Its violation was bound there to strike the seam; the protocol executes the binding.

The conditional Shared Strike theorem [filed at (α′)'s grade and scope]: sign universality at theorem form, confirmed column by column by the same run that decides (α′).

P*ħ*, the single quantum [Postulate at source — the source's own primary load-bearing assumption, expressly falsifiable there]: every completed cycle advances the ledger by the same whole +η*ħ*. T2's target, at the source's own falsification surface, hosted here.

The third road — the strike–base credit [open at REF-CLO(ii); the binding rival's cheapest road]: rides a deeper joint pattern of *lower* completion density — the denominator road. T1's per-column pass closes it at that column; a resolved inversion is its first exhibit, in exactly the shape the Mass paper named.

The Distinction Cost lean [argued, consumed by nothing]: a maintenance-silent mark is a flip pattern with no flips. Its limiting exhibit — a zero-marginal-density mark — is (α′)'s limiting failure; the run either displays it or does not.

The neutrino forecast [the Mass arc's standing prediction: normal ordering, under the indexing clause at O1, severable, strike order bound — indexing first, identification second]: the double-blind cell, §8.

The W0 promissory [recorded at the Mass paper's §13]: under the door's discharge, the chain extends from the maintenance accounting through the Generation census to the weighed table; *the re-pricing belongs to the paper that earns the link.* The protocol specifies the re-pricing; only the reporting instance executes it.

Nothing else rides. In particular — fenced now and again at §10 — the spectrum's cost-reading, the stock door, the toll world, CF, and the mixing arc ride nothing on this run in any cell.

3. Layer A, Restated — The Register the Run Reports In

The protocol consumes the anchoring register exactly as the Mass paper's executed reading graded it, and restates only what the run's discipline requires.

The register is **tick-prior and dimensionless** — the source's Layer A, where the tick index $\tau \in \mathbb{N}$ is the sole ordering primitive and "rate," "frequency," and "period" are Layer-B shorthand entering only through the emergent calibration Δt. A mode μ is a self-sustaining flip pattern: a closed cycle of micro-events over a finite set of flip channels, returning the pattern to itself. Its anchoring cycle has tick-length K_c(μ); its completions per cycle are p_v(μ) > 0 [ANCH-POS]; its completion density is

C(μ) = p_v(μ) / K_c(μ),

a dimensionless Layer-A quantity — completions per tick of substrate ordering. Each completed cycle advances the action ledger by the fixed positive quantum $\eta\hbar$ — whole, irreversible, mode-uniform [ANCH-QUANT, $P\hbar$ at postulate grade]; contributions add over independent channels [ANCH-ADD, with the source's caution that composite effective parameters are eigenmode outcomes, not additive bookkeeping]; and realized rest mass manifests as

$m(\mu) = \eta\hbar \cdot p\_v(\mu) / (c^2 \cdot \Delta t \cdot K\_c(\mu)) = (\eta\hbar / (c^2\Delta t)) \cdot C(\mu)$ [ANCH-COMP],

conditional at source on $P\hbar$, the Layer B bridges (action–energy; $E = mc^2$, both Imported-External there), and the tick calibration. Orderings are calibration-invariant: one Δt for the whole substrate, so the within-column mass order *is* the within-column completion order however, and whether, ticks map to time. The protocol consumes only the ordering induced by C for T1; the calibrated combination $\eta\hbar/(c^2\Delta t)$ appears only at T2, as the single constant the ratios m_i/C_i must realize.

The two-depths fence, inherited and load-bearing here. The source's *anchoring depth* K_c runs opposite to mass — heavier modes anchor shallower (electron $K\_c \sim 10^{23}$; Planck mass $\sim 2\pi$) — and must never be identified with refinement depth. The protocol constrains **C**, never K_c separately: a deeper refinement level may realize its greater density through p_v, through K_c, or both, and a reader who expects "deeper marks, deeper anchoring" has inverted the theorem before the run begins. The corollary under test is the correct composition of the two depths, and only that.

The translations, excluded. Under the source's calibration the inferred K_c values are constructed from the measured masses ($K\_c \propto 1/m$ — the source says so itself). They are translations, not predictions, and R4 excludes them from every stage upstream of T2's scoring. The run that consumed them would be weighing the table against itself.

4. The Eigenproblem — Requirements R1–R7 and the Scope Clause

The protocol does not re-derive the anchoring formalism; it binds what any run claiming to decide (α′) must satisfy. The operator's explicit construction is the source's, cited at the Registration Card; the requirements are this paper's, and a run that fails one is not a failed test but a non-test.

R1 — Construction. The interface operator $\mathfrak{A}$ is constructed from the anchoring dynamics over the catalogue's closure invariants — capacity k, winding tuple w, refinement state R_n — and from nothing else. Its admissible self-sustaining solutions are the eigenmodes. No quantity from the table's mass column, and no quantity derived from one (the K_c translations expressly included in the exclusion), enters the construction.

R2 — Intrinsic labels, or void. Each solution must emit its closure-invariant labels (k, w, n) from the formalism's own structure: columns by superselection — capacity and winding diagonal in $\mathfrak{A}$ — and depth by the interface's refinement labelling. The labels are computed, never assigned. If the formalism cannot label depth intrinsically, T1 has no mass-blind target and the run is **void for T1** — void, not failed: a formalism that cannot name its rungs cannot test their order, and the protocol forbids supplying the names from outside, because the only outside source of a depth ordering is the mass column, and that road is circular by construction.

R3 — Outputs. Per mode: (p_v, K_c), hence C = p_v/K_c, with a propagated numerical uncertainty δC under a stated error model — truncation, boundary data, solver tolerance — frozen with the rest under R5.

R4 — The input fence. The mass column and all its derivatives are excluded from R1–R3 and from T1's scoring entire. Masses enter the protocol at exactly one point: T2's right-hand side, as targets.

R5 — The freeze. Every discretionary numerical choice — basis truncation, boundary conditions, solver tolerances, the error model, the margin parameters of §6 — is fixed and logged before any unblinding. One implementation audit (for coding error, not model choice) is permitted before unblinding; after first unblinding, nothing changes. A post-freeze change is not a correction; it is a new registration.

R6 — Staged unblinding. Stage 1: T1 is scored with the mass column never loaded — the ordering verdict exists, logged, before any mass is in memory. Stage 2: the masses are loaded and T2 is scored. The stages and their timestamps are part of the published record. The point is operational, not ceremonial: a mass-blind ordering test must be blind in the process, not merely in the prose.

R7 — Full publication. Every cell publishes — pass, fail, unresolved, void — with the frozen parameters and the staged log. A rerun is a new registration with its own card. Silent reruns void the instrument: a protocol that reports only its passes is not a decider but an advertisement, and the route's standing rides on this clause more than on any verdict.

The scope clause. A registration names its columns. The populated columns are four — charged leptons (k = 1, w = −1), up-type (k = 3, w = +2), down-type (k = 3, w = −1), neutrinos — twelve modes at three per column. A run may register a subset; verdicts are per column; partial runs pay partially, and REF-CLO's full retirement requires every populated column across registrations. A registration that silently narrows scope after the freeze voids under R5.

5. The Identification Fence, the Freeze, and the Staged Unblinding

The requirements above contain the protocol's three disciplines; this section states why each is load-bearing, because a referee whose rules look arbitrary will be argued with after the result, and these must not be.

The identification fence (R2 + R4). The single deepest threat to the run's meaning is circular identification: deciding which eigenmode "is the muon" by its mass, then announcing that the masses come out in order. The fence closes the road at both ends. Modes are named by computed closure invariants only — the formalism's own superselection and refinement labels — and the mass column is excluded from every stage that could touch a name. T1 is then a genuinely internal question: *do the formalism's own depth labels grade its own computed densities?* The world's masses are not consulted; they could not be, having never been loaded. A T1 verdict obtained any other way is void, whichever way it points.

The freeze (R5). The eigenproblem has discretionary numerics, and discretionary numerics are degrees of freedom. Frozen before unblinding, they are method; adjusted after, they are steering. The single pre-unblinding audit exists because implementation error is real and catching it is honest; its placement before unblinding exists because catching it after is indistinguishable from disliking the answer. The same clause governs V3's audit allowance in §7: one pass, under the frozen state, for coding error only — and the allowance is written here, in advance, so that exercising it is compliance rather than rescue.

The staged unblinding (R6). T1's value to the route is precisely its mass-blindness — it is the one test of (α′) that cannot be accused of reading the table. The staging converts the blindness from a claim into a logged fact. It also produces, as a by-product, the cleanest possible exhibit for the circularity defense the Mass paper mounted at its §8: an ordering derived with the masses literally absent, agreeing or failing to agree with an ordering the world weighed — two instruments, no shared input, one answer owed.

6. The Tests — T1, T2, the Margins, and the Lattice Lemma

T1 — the ordering test (the decider of (α′)). Within each registered column, the computed densities must satisfy

$C(R_0) < C(R_1) < C(R_2)$,

with each adjacent inequality **resolved** in **log form** — the same language as the compatibility condition and the harness, so the Lattice Lemma's entailment needs no units conversion. An adjacent pair (R_n, R_{n+1}) scores *pass* only if

$\ln C(R_{n+1}) - \ln C(R_n) > \kappa \cdot (\delta C(R_n)/C(R_n) + \delta C(R_{n+1})/C(R_{n+1}))$,

i.e. the log-gap exceeds the propagated *relative* margin; *fail* only if the reversed inequality is resolved at the same margin, and *unresolved* otherwise — with κ a frozen protocol parameter (Registration Card; default κ = 3). A column passes T1 if both its adjacent pairs pass; it fails if any pair fails resolved; it is unresolved otherwise. Working in log-ratios throughout means T1, the compatibility condition (§6, also a log-splitting bound), and T2 (a constancy test on $\ln r_i$) all measure in the same unit, and the lemma's one-line entailment is an identity between log quantities rather than a glide across units. Unresolved is not failure: a run whose precision cannot separate two rungs has not weighed them, and the protocol forbids converting insufficient precision into a verdict in either direction. The margin discipline exists for one reason — to make post-hoc rescue impossible: a fail is a fail only when the freeze's own error model says the inversion is real, and a pass cannot be manufactured by declaring noise to be signal.

T2 — the uniformity test (the decider of the single quantum, and the magnitude payment). For each mode i with a measured mass, define the ratio $r_i \coloneqq m_i/C_i$. Under ANCH-COMP every r_i is the one constant $\eta\hbar/(c^2\Delta t)$. The test runs at two grades:

— **T2-col** (per column): within each registered column, max over mode pairs of $|\ln r_i - \ln r_j| \leq \varepsilon\_col$; — **T2-glob** (global): over all registered modes, max over mode pairs of $|\ln r_i - \ln r_j| \leq \varepsilon\_glob$;

with ε_col and ε_glob frozen at registration, set from the propagated uncertainties and the measured-mass scheme uncertainties, both stated on the card. T2-glob entails T2-col. The grades exist because their failures have different addresses: a per-column constant that varies *between* columns is a column factor, whose address at the source is a named open question — the bridge and channel structure that ANCH-ADD's own caution flags, or the quantum's universality — and the registration records, by citation to the source's construction, which attribution the source's text supports; a constant that fails *within* a column strikes the proportionality exactly where the completion door needs it, and the door closes at that column whatever the cross-column situation.

Lattice Lemma. *At a charged column — one whose measured masses stand in resolved depth order [Conjectural, at O1; scheme-robust] — and under the registration's attested compatibility condition, T2-col pass entails T1 pass, given R2's labelling; contrapositively, a resolved T1 failure at such a column entails T2-col failure there.* [Proven — under the compatibility condition: per-column proportionality within ε_col transfers the measured order to the computed densities with resolved gaps exactly when the tolerance and the margins fit inside the splittings.] **The compatibility condition, stated and attested in two parts:** at every adjacent pair of every registered charged column,

$$\varepsilon\_col + \kappa\cdot(\delta C(R_n)/C(R_n) + \delta C(R_{n+1})/C(R_{n+1})) < \ln(m_{n+1}/m_n),$$

with ε_col checked against the public adjacent splittings at the freeze, and the margin terms checked at Stage 2, before the lemma is invoked anywhere. The condition is not caution dressed as formality: without it, a loose ε_col can pass T2-col across an unresolved — in the pathological corner, inverted-within-tolerance — adjacent gap, and the entailment fails at exactly the grain the margins were built to police. In practice the charged columns' adjacent splittings are large (the smallest near ln 20), so any competent registration clears the condition with room — and "any competent registration" is exactly the kind of unstated assumption this protocol exists to eliminate, which is why the condition lives on the card rather than in the prose. Stated for three reasons. First, it disciplines the verdict table: under the attested condition, no cell exists at a charged column in which the magnitudes fit and the order fails, so an ordering strike (C-ord, §7) always arrives with its proportionality face attached. Second, it supplies the table's runtime consistency checks (§7): a joint outcome the lemma forbids is an implementation-audit trigger, not a verdict. Third, it locates the neutrino column as the one place the tests genuinely come apart: there the measured order is not yet on the books, the lemma's premise fails, and T1 stands alone as a forecast — which is §8's subject.

One secondary readout, logged in advance so it cannot be promoted later. Under a T1 pass, the computed densities have a shape beyond their order, and one shape is already spoken for: the infrastructure-reuse profile — compression without fold-back — which the spectrum programme's companion mechanism names and the Mass paper logged as consonance toward the source's cost-reading conjecture. The readout is defined on **ratios**, not additive increments, because ratio space is where the profile lives: with $\rho(n) \coloneqq C(R_{n+1})/C(R_n)$, the spoken-for shape is **ρ decreasing in n while staying above 1** — each step a smaller multiplicative factor, never a fold — the same manifestation-map translation the Mass paper's log-register consonance logs (additive structure in the cost register, multiplicative factors in mass). The additive form is named here only to be excluded: under a full payment the additive C-gaps track the additive mass gaps, which *grow* with depth at every charged column even as the ratios compress, so an additive reading would score the profile absent in the very cell where everything works. The protocol binds the readout's status now: whatever shape the run returns is **logged at firewall discipline and consumed by nothing** — it pays or withholds consonance toward YUK-COST at that conjecture's own grade, and it cannot be cited as support for (α′), which T1 has already decided by then on its own terms. One caveat the measured table forces, stated here so the readout is not read as a uniform expectation: the down column already *expands* rather than compresses — its log-gaps run 2.99 then 3.80, equivalently R_down < 0 (§15) — so the reuse profile is **expected to be mixed per column**, and a genuine V1 withholds the profile's consonance at the down sector while still paying (α′) and the door there. Compression without fold-back is one column's shape, not the table's.

7. The Verdict Table, Bound in Advance

The table is the paper — and it runs at two grains, because its conditions do. Per-column tests (T1; T2-col) yield **column verdicts**; cross-column tests (T2-glob; the quantum pattern) yield **global verdicts**; and the named configurations V1–V5, which the route's other papers cite, are distinguished configurations of the two layers rather than primitive cells. The implications are [Proven — immediate from the inherited clauses at their grades]; the content is the binding, executed before the number exists.

The column layer. Each registered column receives exactly one verdict of record:

— **C-pass.** T1 pass (both adjacent pairs resolved in order) ∧, at a measured column, T2-col pass. At an unmeasured column (the neutrinos, §8), T1 passes alone and the proportionality face is deferred — the harness names this sub-case **C-pass(T1-only)**: still a C-pass, its T2 standing as the prediction of record until an absolute scale exists. *Pays, per column:* (α′) at that column; the door's discharge there; the corollary's surface there; the denominator road closed there.

— **C-prop — the proportionality strike.** T1 pass ∧ T2-col fail. The door's one-line proof loses its premise at that column and the discharge does not complete there — REF-CLO stays imported. What the cell pays anyway, logged at honest weight: the formalism's own depth labels grade its own computed densities in exactly the order (α′) asserts, at a column where the masses were never loaded — strong structural consonance whose effect is to relocate the residue: the question stops being whether depth grades the density and becomes whether the register's mass identification holds at that column. The strike's address is ANCH-COMP's within-column face; the theorems beneath stand; the stock door stands.

— **C-ord — the ordering strike.** A resolved T1 failure at any adjacent pair. *Procedure, bound:* one implementation audit under the frozen state (R5's allowance — coding error only, no model change, no re-identification); a failure surviving the audit is the verdict. *Consequences:* the strike lands **jointly at (α′) and R2's depth-labelling, as one seam clause** — mirroring the Seam Consistency Theorem's address discipline: the violation strikes the identification of the formalism's depth with refinement depth before either programme's interior, because the run cannot distinguish a world in which deeper marks genuinely cycle sparser from one in which the formalism's rungs are not the census's rungs. Either way the completion door closes; the completion-density corollary is struck at the seam; the conditional Shared Strike theorem lapses with its condition. *The survivors, named:* the Mass Hierarchy Theorem stands at its import — it rides REF-CLO as import, and an import's failed discharge is not the import's failure; the stock door stands; the proven layer (Depth Content, Supervenience, Content Functional, Participation, Marginal) stands entire, exactly as the Mass paper's robustness paragraph priced. *The collector:* the binding rival's third road — the strike–base credit — receives its first exhibit, at the named column, in exactly the sparser-cycling shape the Mass paper said the denominator road must wear; and by the Lattice Lemma, under the attested compatibility, an exhibit at a charged column arrives with a within-column proportionality failure attached, so the rival's celebration and the register's embarrassment are the same cell — the honest price of letting one computation carry both.

— **C-open — the column undecided.** No resolved failure, but T1 not passed: at least one adjacent pair unresolved at the frozen margins. The column is undecided — not paid, not struck — publishes as such, and is eligible for re-registration at sharper numerics under the same construction. (At a charged column with attested compatibility, C-open cannot coexist with a T2-col pass — the lemma's runtime face, below.)

— **C-void.** A breach or numerical non-completion localized to the column. Publishes; decides nothing.

The global layer. The run additionally reports:

— **G-glob.** T2-glob pass or fail, scored over all measured registered modes.

— **G-quant — the quantum strike.** T2 failing in the source's own named pattern — the computed (p_v, K_c) admitting no single η under its stated bridge structure. The trigger freezes as a numerical criterion where the source's structure supports one — a test statistic over the (p_v, K_c) set with its threshold, stated on the card; where the text does not support a numerical form, the registration declares the trigger at citation grade and names it **the protocol's one discretionary seam**, binding that, should the seam ever be reached, both attributions' consequences — the column-factor reading and the quantum strike — publish side by side rather than the registrant's preferred one. *Consequences when triggered:* P$\hbar$ is struck at its own surface; ANCH-QUANT re-grades; the completion door's substructure falls — mode-uniformity was consumed by the proportionality itself, the variable-quantum world its named failure — and the door closes regardless of the column layer, since under a per-mode quantum the C-ordering decides nothing about the mass ordering. *The survivors:* everything beneath the door, on both sides of the seam. Column-layer T1 passes co-occurring with G-quant are logged as structural consonance for the depth-labelling and nothing more.

— **The square.** The neutrino cell, §8.

The distinguished configurations — the names the route's other papers cite:

— **V1 — full payment:** every registered measured column C-pass ∧ G-glob pass ∧ no G-quant ∧ (the neutrino column, if registered, scored at the square). The payment schedule is §9.

— **V2a — ordering paid, the cross-column ratio open:** every registered measured column C-pass ∧ G-glob fail ∧ no G-quant. The within-column discharge completes — everything the Mass arc consumes, its theorem within-column by clause (d)'s fence — and (α′), the corollary, the conditional Shared Strike theorem, and the denominator closure are paid as at V1. The cross-column factor is logged as an open question at the source's named address (bridge and channel structure against quantum universality, per the registration's citation), straining but not striking P$\hbar$: the protocol does not adjudicate an attribution the source's own structure must decide. The K_c numerics convert to predictions within columns only; the global magnitude payment waits.

— **V2b:** some column at C-prop — the door closed at that column, the configuration paying per column elsewhere.

— **V3:** some column at C-ord — the strike configuration, its consequences the column verdict's.

— **V4 — the quantum strike:** G-quant triggered, whatever the column layer.

— **V5 — void:** R2 unmet at the construction grain, or a breach under §11, voiding the registration entire — distinct from per-column C-open and C-void, which are records, not voids. Two re-registration roads, distinguished: a column left at C-open re-registers as a **numerics sharpening** — the same construction at tighter margins; an R2 void re-registers, if ever, as a **new construction under R1** — a formalism that develops intrinsic depth labelling is a different instrument, registered fresh, and the protocol forbids describing the second road as the first.

Exhaustiveness and consistency, restated at the right grain. Each registered column receives exactly one of {C-pass, C-prop, C-ord, C-open, C-void}; the run reports the two global verdicts and, if registered, the square; the configurations above are shorthand, and mixed configurations publish and pay per column — the scope clause's "partial runs pay partially," executed. Three consistency checks ride the Lattice Lemma's runtime face at charged columns with attested compatibility: (i) C-ord entails a T2-col failure there; (ii) T2-col pass entails C-pass; (iii) the joint outcome "T1 unresolved at Stage 1, T2-col pass at Stage 2" is an arithmetic impossibility in the run's own frozen numbers — its appearance is an implementation-audit trigger, not a verdict, and if it survives the audit the column voids, because a run inconsistent with its own arithmetic has not weighed anything.

8. The Neutrino Cell — The Double-Blind Square

The neutrino column is the protocol's one cell where neither instrument has reported: the world's oscillation programme has measured the splittings and not the ordering, and the run has computed nothing yet. The cell is therefore a genuine forecast surface — the only place T1 is prediction rather than retrodiction — and the protocol binds its outcomes as a square, with the Mass arc's already-bound strike order (the indexing clause first, the identification second; the clause's large-mixing exposure already on its books) integrated rather than re-litigated.

The computed verdict, under the transport-fixed index and R2's labels: **normal** (C in depth order), **inverted** (resolved reversal), or **unresolved** — with the resolution caveat preregistered: the neutrino column's measured splittings are the table's smallest, so partial resolution (the larger gap resolved, the smaller unresolved) is a foreseeable outcome, scored per pair, and an honest run may return the cell half-open.

One scoping note, executed at the configurations' conditions (§7): the neutrino column registers for **T1 and the square only**. T2 cannot run where no absolute mass is measured, so the column neither satisfies nor fails T2-col, and every configuration quantifying over T2 quantifies over measured columns. The proportionality face is deferred, standing: the day the world produces an absolute scale, the frozen densities' ratios — and, under a paid G-glob, their absolute values through the one constant — are already the prediction of record, lodged at this card.

The square, priced:

— *Computed normal × world reports normal:* double payment — the Mass arc's forecast and the protocol's wager paid by one announcement; under V1's global grade the run has, before the world reported, lodged the neutrino mass ratios in numbers (the computed C's, scaled by the one constant), and the wager's full face value is collected.

— *Computed normal × world reports inverted:* the Mass arc's bound order governs — the indexing clause is struck first; if the inversion survives a defensible indexing, the strike escalates to the identification, and the computed-normal verdict then stands against the world at that column, striking the register's neutrino-sector face (ANCH-COMP or the labelling there) in the same motion. The protocol prices this cell now because it is the one where the anchoring chain itself takes damage from a measurement, and a price quoted after the measurement is not a price.

— *Computed inverted × world undecided:* a standing public disagreement between the protocol's computation and the Mass arc's forecast — published as such, in numbers, with no reinterpretation and no suppression, resolution belonging to the world. The protocol states plainly what this cell is: two of the route's own instruments pointing opposite ways at a question the world will answer — which is not an embarrassment to be managed but the configuration in which a forecast is worth most, since whichever instrument the world refutes, the refutation lands on a preregistered surface.

— *Computed inverted × world reports inverted:* the indexing clause is struck per the bound order, and the computation is logged as having out-run it; the column's depth map then corrects under the surviving indexing, and the cell's T1 is **re-scored under the corrected map by the same frozen densities** — the re-scoring bound here, in advance, so that it is arithmetic rather than improvisation. The re-scoring is not §11's re-identification breach, and the reason is stated now so it cannot be litigated later: the struck object is the indexing clause — an inherited import of the Mass arc, the translation between the transport-fixed index and the world's named states, external to this registration; the correction is forced by the world's measurement, not chosen by the registrant; and nothing computed moves — no density, no intrinsic label, no frozen parameter — only the comparison map between the world's states and the catalogue's depth labels, which was never the registration's to freeze.

— *Unresolved pairs:* scored per pair as everywhere; the wager stands at whatever resolution the run achieves.

One fence inside the cell: the neutrino column's *base* — why its floor sits so far beneath every other column's — rides nothing here; it is the mass-origin question, fenced to its own arc by the Mass paper, and a computed neutrino density that surprises at the base is the base arc's datum, not this protocol's verdict.

9. The Payments Under V1, Line by Line

A full payment unstated in advance can be inflated after the fact; the schedule is therefore fixed now, each line with the conditionality that survives it, because the discharge retires an import into a chain, not into thin air.

1. **(α′)** — computation-paid: the density clause confirmed at every registered measured column — the neutrino column scored at the square (§8), its wager standing beside the payment — in the deciding computation's own variables.
2. **The completion door discharges**: $\Delta(n) > 0$ at every registered measured column. **REF-CLO(ii) and (iii) re-grade** from [Open] to [Discharged at the completion door — Conditional on the anchoring chain: P$\hbar$ now computation-consistent rather than bare postulate; ANCH-COMP at the Layer B bridges, Imported-External; the depth-labelling at R2]. The Mass Hierarchy Theorem's status line re-writes accordingly: its bridging import retired into a conditional chain whose every member is independently maintained — which is the only way the route ever shortens.
3. **The completion-density corollary** — the third seam surface, paid.
4. **The conditional Shared Strike theorem** — paid column by column: sign universality at theorem form, the sign the clause's own, confirmed per column by the run that decided the clause.
5. **P$\hbar$** — from postulate to computation-consistent: not proven — its falsifier survived at the stated tolerance, which is what a postulate can earn and all it can earn.

6. **The third road closed**: the strike–base credit's denominator shut at every column the run resolves; the binding rival's cheapest road ends at a computed number.
7. **The Distinction Cost lean's exhibit, absent**: no zero-marginal-density mark anywhere in the run — consonance for strictness, at the lean's own grade.
8. **The K_c numerics convert**: from translations (K_c ∝ 1/m, the table read backwards) to predictions (computed forward, matched at T2) — the anchoring programme's own payment, collected on its own books.
9. **The base assays delivered**: the computed R_0 densities are the Mass paper's "cleanest assays" in numbers — pure base costs per column, handed to the base-cost arc as data, consumed here by nothing.
10. **The neutrino wager lodged** — §8's cell, in numbers, before the world's count.
11. **REG-CONC consonance**: the completion register's ordering agreeing with the table's everywhere computed — logged toward the concordance at its conjectural grade, consumed by nothing.
12. **The W0 promissory comes due for re-pricing**: the chain extends from the maintenance accounting through the Generation census to the weighed table; the reporting instance executes the re-pricing the Mass paper assigned — and the marker still does not move, §12.

What V1 does **not** pay, said in the same breath: the spectrum's cost-reading (§10); the stock door's hypotheses (§10); the toll world (§10); the magnitudes *for the Mass arc*, which consumes orderings only in every cell, the magnitude payment landing on the anchoring programme's books behind that arc's firewall.

10. The Fences — What No Cell Decides

YUK-COST stays at conjecture in every cell — the firewall outlasts full payment. The completion door discharges through the anchoring register: m ∝ C, C ordered. At no point does the discharge consume the spectrum programme's reading of its eigenvalues as maintenance costs, and a V1 result therefore does not promote that reading — it pays consonance toward it at the conjecture's own grade (the two registers ordering alike is what one functional would look like), and nothing more. A route that let its decider quietly pay a sibling's conjecture would be running results through the firewall in the confusion of a celebration, which is when firewalls matter most.

The toll world is untouched. The run computes densities, not opposition. That the realized cost *opposes* transport-state change — rather than tolling it — is the opposition clause, at its own register with its own candidate, and no cell here reaches it. A V1 result is consistent with the toll world; the Mass paper said so and the protocol repeats it where the temptation will arise.

The stock door is untouched in every cell. ENT's tests are the neutrino sector's fold arithmetic and the THz line, at that register's own addresses. V3 and V4 close the completion door and leave the stock door standing; V1 pays the completion door and leaves the stock door unconsumed. The two doors' complementary characters — strong register with the load in a density clause; strong inference on a register at hypothesis grade — are unchanged by any cell.

CF and the class-grain reading are untouched. The run neither tests nor consumes the Content Functional Lemma's condition; a token-mass anomaly is the table's to exhibit, not this computation's.

The mixing arc is untouched. The depth-distance candidate and its neutrino tension ride the spectrum programme's overlap structure, not the anchoring register's densities.

The magnitudes, for the Mass arc, are fenced even at V1. The Mass Hierarchy Theorem consumes orderings; it consumed orderings before this protocol and consumes orderings after V1. The computed magnitudes are the anchoring programme's payment, and the seam between the programmes stays a seam.

11. Against the Protocol Itself

A decider has its own refuters, and they differ in kind from its targets': a breach voids the instrument, deciding nothing about the world.

The fence breach. Any mass-column quantity — a mass, a translation, a ratio — found upstream of T2's scoring voids the registration entire, T1's verdict first. The breach is structural, not moral: a contaminated ordering test has no meaning to salvage.

The freeze breach. Any post-unblinding change to a frozen parameter, the error model, or the margins voids the run from the change forward. The one audit allowance is pre-unblinding, once, logged.

The re-identification breach. Any post-freeze adjustment of the mode-to-class map — re-labelling a mode's column or depth after a density is known — voids the registration. R2's labels are computed once, under the freeze, and a formalism whose labels can be argued with after the fact has failed R2, which is V5, which publishes. The square's fourth cell's bound re-scoring is not this breach: §8 states the exemption and its reason — the moving object is an inherited comparison map, never a computed label.

The publication breach. A withheld cell — a fail or a void that does not publish — voids more than the run: it voids the instrument's standing, and with it the meaning of every future pass. The protocol therefore binds the breach's own consequence: a registration is presumed breached if its card exists and no report follows within the card's stated horizon, and the presumption publishes.

The four breaches share one address — the registration's integrity, never the physics — and one consequence: V5, with the strike landing on the run, the targets standing exactly as imported.

12. W0, Unmoved in Every Cell

No cell moves the marker. V1 extends a conditional chain — P$\hbar$ computation-consistent is still a postulate that survived, the bridges are still Imported-External, the Generation chain is still at its conditionality — and conditional structure cannot press on a verdict node, exactly as the route has held at every instalment. V3 and V4 strike conditional structure, which presses nothing either. V5 decides nothing. The promissory recorded at the Mass paper's §13 comes due for re-pricing under V1, and the re-pricing — executed by the reporting instance, not here — re-states a longer conditional chain; it does not convert one. The marker stays [Open], at the Born-arc slot, in every cell of the table. [The pressure: unchanged by construction; the construction: the point.]

13. Position in the Programme

What this paper adds to the ledger: no physical line. Its parameters (κ, ε_col, ε_glob, the error model, the scope) are procedural, frozen at the Registration Card. Its bound implications are [Proven — immediate from the inherited clauses at their stated grades; the content is the binding, executed before the number exists]. Its inherited exposure is the stakes of §2, each at its source's grade, consumed as stakes rather than premises: the anchoring chain as read (ANCH-COMP, $P\hbar$, ANCH-ADD, ANCH-POS, ANCH-IRR, the two-depths fence), the Mass arc's clauses ((α′), REF-CLO, the door, the corollary, the conditional Shared Strike form), O1 at the neutrino indexing, and R2's labelling requirement standing where MODE-identification questions live — as a void condition, not an import.

The seam, extended by one instrument:

```
Mass Hierarchy Theorem                  Anchoring programme
  (α′) [Open] — the density clause        eigenmode program —
  REF-CLO at two clauses                  (p_v, K_c) per mode,
  the completion door [Conditional]       mass-blind [R1–R4]
          |                                      |
          └──────────────────────────────────────┘
                        |
            THE EIGENMODE DECISION
            [this paper — the referee's instrument]
            R1–R7; the identification fence;
            the freeze; staged unblinding;
            T1 (ordering, mass-blind) — decides (α′);
            T2 (uniformity, two grades) — hosts Pℏ's
            own falsifier; Lattice Lemma;
            VERDICT TABLE at two grains — column verdicts
            {C-pass, C-prop, C-ord, C-open, C-void},
            global verdicts {G-glob, G-quant}, the
            distinguished configurations V1–V5, strike
            orders bound;
            the neutrino DOUBLE-BLIND SQUARE;
            payments and fences priced in advance;
            W0 unmoved in every cell
                        |
            the run reports → one column verdict per
            column + the global verdicts + the square;
            every cell publishes [R7]
```

The division of labor, restated at the instrument's grain. The Mass paper supplies the clause and its price list; the anchoring programme supplies the operator and the run; this paper supplies the only thing neither party may supply for itself — the conditions under which the number means something. Three documents, one seam, and the seam now has a referee whose rules were public before the contest.

14. The Protocol Exercised — Public Data and a Referee Harness

A protocol that decides a number ought to be exercised before the number arrives — not on the physics, which awaits the explicit operator, but on the **referee itself**: does the verdict logic return the right cell for a passing run, a failing run, an underpowered run, and a non-test? This section reports three exercises. None tests VERSF; none can, the operator 𝔄 not yet being explicit enough to produce (p_v, K_c) from closure structure. What they test is the instrument — and one of them refuses a result the careless reading expects, which is the exercise earning its place.

The discipline observed in the exercise itself. A toy hand-tuned to ascend would pass T1 by construction — densities chosen to climb are R1/R4 violated, the answer inserted by hand — and a harness that only ever sees a rigged pass is itself unaudited. The exercise therefore runs *blind* mode-generators (each emitting (label, p_v, K_c, δC) without knowledge of the verdict logic) across scenarios built to **pass, to strike, to void**, so that each verdict is earned by a generator's structure rather than chosen.

14.1 The public-data sanity check

The protocol retrodicts the visible mass ladder, and the check confirms it at O1 — orderings only, scheme-robust, as R4 requires. PDG 2024 central values, MeV:

Column	R_0	R_1	R_2	Ordering	$\ln(R_1/R_0)$	$\ln(R_2/R_1)$
Charged leptons	e = 0.511	μ = 105.66	τ = 1776.86	**pass**	5.33	2.82
Up-type quarks	u ≈ 2.16	c ≈ 1270	t ≈ 172570	**pass**	6.38	4.91
Down-type quarks	d ≈ 4.70	s ≈ 93.5	b ≈ 4183	**pass**	2.99	3.80

The ladder holds in every charged column. The smallest adjacent splitting anywhere — τ/μ at ln ≈ 2.82, a factor 16.8 — is the binding case for the **Lattice Lemma's compatibility condition** (§6): any margin budget below it clears the condition, and the public splittings show the budget clears with room. This is the condition checked as a *fact* rather than assumed, exactly where §6 said the card must carry it. The check's scope is its own confession: it confirms the empirical claim and tests no mechanism.

14.2 The referee harness

The harness implements T1 (resolved-gap test at margin $\kappa\cdot(\delta C_n + \delta C_{n+1})$), T2-col (per-column proportionality of m/C within ε_col), the column-verdict logic of §7, and staged unblinding — T1 scored with the mass column never loaded, T2 scored only after. Frozen for the exercise: κ = 3, ε_col = 0.20, ε_glob = 0.40. Five blind generators, five outcomes:

Scenario	Generator rule	Verdict	What it demonstrates
S1 ascending toy	p_v↑, K_c↓ with depth (one rule, all columns)	**C-prop**	ordering alone is not full payment — see below
S2 inverted	p_v↓, K_c↑ with depth	**C-ord**	the ordering strike is caught, both pairs resolved-failed
S3 mis-scaled	C ∝ a wrong mass set	**C-prop**	the V2b configuration — order confirmed, proportionality struck
S4 underpowered	gaps ≈ 0.01, error ≈ 10%	**C-open**	thin precision yields no verdict, in either direction
S5 unlabelled	ascending, R2 unmet	**C-void(R2)**	a non-test, not a failed test

The instructive outcome is S1. A toy whose densities ascend — the very construction a careless reading expects to "pass" — lands at **C-prop, not C-pass**. Its densities climb, so T1 passes; but they are not proportional to the measured masses, so T2-col fails (global ln(m/C) spread ≈ 10.9, far outside ε_glob). The referee is reporting, correctly, that ordering is not the full payment: a run reaches V1 only by reproducing the measured *ratios*, which a toy cannot assert into being. The exercise that an ascending toy "passes" would have mis-certified the protocol; the harness refuses it, which is the harness working.

14.3 A mass-blind toy construction — steps 1–3 executed, and what they earn

The harness above scores generators; this exercise runs the three construction steps the run itself must perform — (1) construct modes without mass input, (2) obtain intrinsic depth labels, (3) compute C = p_v/K_c — under a deliberately transparent rule, to show what a mass-blind construction *can* and *cannot* deliver before the real operator exists. The rule, chosen for legibility and emphatically not fitted to any mass:

p_v = (k + |w|)·(n + 1), K_c = 10/(n + 1), C = p_v/K_c,

with the column labels (k, w) intrinsic and depth n ∈ {0, 1, 2} the refinement label. The densities, and the referee's verdict:

Column	(k, w)	$C(R_0)$	$C(R_1)$	$C(R_2)$	T1	T2-col	Verdict	ln(m/C) spread
Charged leptons	(1, −1)	0.20	0.80	1.80	pass	fail	**C-prop**	5.96
Up-type quarks	(3, +2)	0.50	2.00	4.50	pass	fail	**C-prop**	9.09
Down-type quarks	(3, −1)	0.40	1.60	3.60	pass	fail	**C-prop**	4.59

The toy passes T1 in every column — ordering confirmed, mass-blind, on intrinsic labels — which is the genuine demonstration: steps 1–3 execute, and an internal rule with no mass input reproduces the ladder's *direction* at the census grade. But the referee returns **C-prop, not the full pass**, and the closed form says why. $C = (k + |w|)(n + 1)^2/10$, so within every column $C \propto (n + 1)^2$, giving identical depth ratios everywhere — $C(R_1)/C(R_0) = 4$, $C(R_2)/C(R_1) = 2.25$ in all three columns — while the measured ratios diverge wildly (μ/e ≈ 207 against τ/μ ≈ 17; c/u ≈ 588 against $t/c \approx 136$). The toy gets the *order* right and the *magnitudes* wrong, so T2-col fails and the discharge does not complete: the V2b configuration, reached by a toy.

This is the more instructive outcome, and it is favorable rather than embarrassing for the same reason the §8 forecast is. A two-line rule that returned full V1 would be the alarming result — it would mean a hand-chosen formula had reproduced the actual mass ratios, which is the fitting R4 forbids and a skeptic would immediately allege. Instead the toy demonstrates the arc's own division of labor in miniature: **the ladder is cheap and the rungs are dear.** Many rules give the ordering; the skeleton paper claims exactly that and no more. Reproducing the spacings is the content that cannot be assigned by a chosen rule and must be *derived* from the operator — which is precisely the work C-prop points at, and precisely what the real eigenmode run has to earn that this toy cannot. Steps 1–3 are therefore confirmed as *executable under intrinsic labels with no mass input*, and confirmed as *insufficient for the magnitude payment* until p_v and K_c come from $\mathfrak{A}$ rather than a formula — which is the honest boundary, drawn by the referee rather than asserted.

14.4 The fence, made executable

A genuine V1 needs densities whose ratios reproduce the measured masses — $C \propto m$. The harness accepts such a ladder (global spread → 0, configuration V1). But it then asks R4's question, and the answer flips the verdict on identical numbers:

How the proportional densities were built	Verdict
from closure structure, mass column absent (R1)	**legitimate V1**
by reading the mass column (e.g. $K_c \propto 1/m$)	**void — fence breach (R4)**

The same density ladder is the route's full payment or a voided run depending **only on provenance** — and the numbers cannot reveal which. Only the staged, mass-blind Stage-1 log (R6) can certify that the ladder was produced before any mass was loaded. This is the operational content of R6 stated as a worked case: blindness is a logged fact precisely because a proportional ladder is indistinguishable, after the fact, from a fitted one. The harness makes the indistinguishability concrete, and the staging is what resolves it.

14.5 An attempt at the operator, and a counting bound on what the charged sector can decide

§14.3 located the bottleneck — the ordering comes free, the spacing does not — and the natural next move is to ask whether the spacing can be *derived*: whether a map $(k, w, n) \to (p_v, K_c)$ with constants fixed by closure structure reproduces the columns' rung-spacing without fitting it. The spacing is well captured, per column, by a three-parameter form

$$C_{c,n} = B_c \cdot \exp(A_c \cdot n - \tfrac{1}{2} R_c \cdot n(n-1)),$$

with B_c the column's base floor (the base-cost arc's, fenced from the ordering test), A_c the first refinement jump, and R_c a compression/reuse correction — the form's shape encoding the

§14.3 insight that depth gives the ladder and column geometry gives the spacing. The honest question is not whether this *fits* — three parameters fit three masses per column exactly, by construction, with zero content — but whether (A_c, R_c) *follow from* (k, w) under a law lean enough to be falsifiable.

The attempt fails, and the failure is a result. A genuinely over-determined ansatz — $A_c = \alpha(k + s_A \cdot w)$, $R_c = \rho(k + s_R \cdot w)$, four constants for six column-values — is refused by the data: its best fit returns A = [1.6, 5.5, 5.1] against the targets [5.3, 6.4, 3.0], missing the lepton column badly and mis-ordering the rest, at normalized residuals near 0.8 (a form capturing nothing sits near 1.0). A scan over simple structural features (k, |w|, k+|w|, 1/k, and others) finds no two-constant law reaching even normalized 0.7 on both A and R. There is no simple law of (k, w) hiding in the charged spacings.

And a counting argument shows there cannot be one *the charged sector could validate*. Three columns supply six target numbers (three A, three R) with two labels each; a label-map with six or more constants reproduces any six numbers — including random ones, as a scrambled-target check confirms — so a fit there is uninformative by construction; a map with fewer constants is testable but, as the attempt shows, the simple ones fail. Three measured columns therefore cannot, even in principle, distinguish a derived spacing law from a fitted one: the sector has exactly enough freedom to absorb any law, and the fourth column is unmeasured. This is the same wall the protocol meets from the verdict side, reached here from the construction side, and it sharpens the program's real requirement rather than weakening it: **the operator's constants must be fixed by the substrate's cycle dynamics, not by the columns** — derived from what a self-sustaining cycle *is*, with (B, A, R) read off and checked afterward. Provenance again decides everything, exactly as §14.4's fence shows for the density ladder: constants fitted to columns are absorbed and prove nothing; constants forced by the operator turn the charged columns into a genuine test and the neutrino column into a real forecast. The bottleneck is not the ansatz; it is the operator, and no parametrization substitutes for it.

14.6 What the exercise establishes, and what it does not

The harness returns the correct cell for pass, ordering-strike, proportionality-strike, underpowered, and non-test inputs; it refuses a rigged pass; it voids a fence-breaching construction; run on a mass-blind construction of steps 1–3 it confirms the ladder while striking the magnitudes at C-prop; and an honest attempt to derive the rung-spacing from (k, w) both fails for the simple laws and is shown, by a counting bound, to be undecidable on the charged sector at all — the division of labor and its evidential limit made executable. The protocol mechanics — blind labelling, the resolution margins, the five-verdict logic, the compatibility condition, and the provenance fence — behave as §§4–7 specify, debugged before the deciding number exists, which is the order the protocol insists on. What remains untested is everything the protocol was built to decide: the real operator's output, and therefore (α′), the door, and the route's claim. The exercises locate the bottleneck precisely and bound it: the ordering comes free, the magnitudes do not, and the charged columns alone cannot adjudicate a spacing law — so the operator $\mathfrak{A}$ must be made explicit enough to *derive* p_v and K_c with constants fixed by the substrate rather than fitted to columns. The harness is the referee waiting at the bench; Appendix B is its reference implementation, to be re-frozen with the true κ, ε, and error model on a Registration Card before any run that reports.

15. The Operator's Requirements — A Specification, Not a Fit

§14 established, from both the verdict side and the construction side, that the charged sector cannot adjudicate a spacing law: any (k, w) → (A_c, R_c) map with as many constants as columns fits anything, and the tight maps fail. The constructive response is not another ansatz but a specification of what the operator 𝔄 must *be*, written so that p_v and K_c are derived counts rather than tuned parameters — the only route that turns the charged columns from an absorbing fit into a genuine test. This section states the requirements at [Conjectural]/[Open] grade; it derives nothing about the world, and it deliberately fits nothing.

The degree-of-freedom target, stated as the bright line. The charged sector offers six numbers — three A_c, three R_c, the base floors B_c fenced to the base-cost arc. A (k, w) → (A, R) map with six or more constants reproduces any six numbers (§14.5's scrambled-target check), so it is empty by construction. A non-empty operator must therefore deliver all six with **strictly fewer than six free constants**, the surplus being predictive content — the cleanest target being **zero constants tuned to columns**, every constant a substrate quantity fixed elsewhere in the corpus, so that {lep, up, down} are all tests and the neutrino column is a forecast. An honest intermediate target is **at most two shared substrate constants**, reproducing six numbers from two and yielding four genuine predictions.

The obligations, each gradeable, none fitted:

- **D1 — the cycle.** A self-sustaining cycle is a closed orbit of the anchoring map on the closure state of (k, w, n): a finite sequence of micro-events returning the configuration to itself. The eigenmodes are these orbits; this is what "self-sustaining flip pattern" (§3, ANCH-IRR) means made operational.
- **D2 — K_c as a count.** The tick-length K_c is the number of substrate ticks the closed orbit takes to return — a count on the orbit, not a parameter. The protocol's R3 output is thereby a property of the solution, satisfying R1's no-free-input demand at the orbit level.
- **D3 — p_v as a count.** The completion number p_v is the number of irreversible commitment events the orbit logs per return — also a count on the orbit, each event one $+\eta\hbar$ on the ledger (ANCH-QUANT, P$\hbar$).
- **D4 — C without freedom.** C = p_v/K_c follows once the orbit is fixed; no knob remains. This is where the toy of §14.3 and the parametrizations of §14.5 differ from a derivation in kind: there C was assigned, here it is read off.
- **D5 — admissibility selects the depths.** Which orbits are admissible eigenmodes is fixed by closure balance — the corpus's BCB/ledger condition — and this selection is what must *yield* the three depths n = 0, 1, 2 rather than impose them, tying the operator back to the Generation census it is meant to price.
- **D6 — labels through structure only.** Capacity k enters as the number of channels the orbit threads; winding w as how they thread; neither enters as a free coefficient. This is

R2's intrinsic-labelling requirement satisfied at the construction grade — the labels are computed, never assigned.

The falsifiable consequence, pre-registered so it cannot be retrofitted. If K_c and p_v are both counts on one orbit whose length grows with held marks, then the orbit geometry that sets the first log-gap also sets the second — so the two are linked by a one-parameter relation rather than free. Stated correctly, the constraint is on the **two genuinely independent observables**, the adjacent log-gaps $G_1 = \ln C_1 - \ln C_0$ and $G_2 = \ln C_2 - \ln C_1$: an orbit-count operator must predict G_2 from G_1 by a one-parameter law, which three columns constrain and the fourth tests. The constraint must *not* be stated as "A_c and R_c are correlated," and the reason is a trap worth naming because the section's authority is that it refuses to fall into it. Under the form $C = B\cdot\exp(A\cdot n - \frac{1}{2}R\cdot n(n-1))$, the gaps are $G_1 = A$ and $G_2 = A - R$, so $R \equiv A - G_2$ *by construction*: A and R share the A term definitionally, and corr(A, R) carries a floor near 0.6 under any hypothesis whatever, random gaps included. The charged columns' corr(A, R) ≈ 0.81 is that manufactured floor plus the gaps' own mild coupling — it reads a number the parametrization invents, not orbit geometry. The honest number is corr(G_1, G_2), and across the three charged columns it is **≈ 0.34 — no signal**. That is the right result: the genuinely independent observables show the operator's $G_2(G_1)$ relation is unconstrained-to-mild by the charged data, leaving the fourth column a clean forecast surface rather than a pattern already half-read. The discipline records the gap correlation as the forecast surface it is, and flags the A–R correlation explicitly as the artifact a careless reading would mistake for evidence.

A known expansion, owned. One charged column already fails the reuse-profile shape §6 speaks of, and §15 prints it plainly: the down column *expands* — $G_2 = 3.80$ exceeds $G_1 = 2.99$, equivalently $R\_down = -0.81 < 0$, negative compression — so its ratio sequence rises rather than compresses. This is no contradiction (the §6 readout is consumed by nothing, and a non-compressing column simply withholds reuse-profile consonance), but it must be owned rather than glossed: the reuse profile is **expected to be mixed per column**, the down sector is a known expansion in the measured table, and a genuine V1 would withhold the profile's consonance there while still paying (α′) and the door. The spoken-for "compression without fold-back" is one column's shape, not the table's; the operator must reproduce the expansion as readily as the compressions, and the $G_2(G_1)$ relation is where that obligation lives.

What this section is, and is not. It is a target with a counting bound, a set of construction obligations, and one pre-registered constraint (A–R dependence) that the eventual operator can satisfy or violate — the constructive complement to §14's demonstration that fitting cannot decide the matter. It is not the operator: $\mathfrak{A}$'s explicit form, the proof that closure balance selects exactly three depths, and the derivation of the orbit counts remain [Open], the genuine next paper's burden. The specification's value is that it makes that paper's success checkable in advance — a derived $G_2(G_1)$ relation reproducing the charged columns (including the down sector's expansion) from substrate constants, with G_2 for the neutrino column then forced rather than fitted, is earning the spectrum; a relation tuned to columns, or one carrying as many constants as gaps, is fitting it, and the difference is now stated before the attempt rather than after.

16. Conclusion

The route has been asked, at every instalment, what would change its mind, and it has answered with named refuters, priced fallbacks, and bound strike orders. This paper is that answer carried to its operational limit: an entire decision, notarized before the deciding number exists. The assay office's scales weigh the minting — completed stampings per turn of the substrate's press — and the ladder's claim is one sentence on the card: *deeper editions mint more densely, in every denomination.* If the timed mintings land in the watermark's order, the route's last import retires into machinery, and the table's strangest fact — that the repetitions weigh in order, everywhere, always, with no ties — becomes arithmetic performed by the substrate and audited by the office. If a denomination mints out of order, the clause is struck at a seam whose address was published first, the rival that has waited three papers for an exhibit collects one in exactly the shape it was promised, and the theorems built to survive the strike survive it in the configuration their own robustness paragraph specified. If the stamp itself proves not to be one stamp, the office's foundation falls by its author's own stated test, and everything beneath it stands. And in one denomination the card is a sealed wager: the office will compute an order the world has not yet weighed, and the number will wait for the world the way the route has taught its claims to wait — in public, at a price, with the strike order already bound.

The books are kept; the markers hold; the import stays on the ledger until the number says otherwise; and the protocol adds nothing to the world except the one thing the route still owed it — a referee. The question that built the last paper kept the last word it had earned: *is the strike on the bill?* This paper does not answer it. It does something a route can only do once per question, and must do before answering: it makes the answer expensive to fake, in both directions. The card is filled in. The office is open. The next word belongs to the number.

Appendix A — The Registration Card

A registration is the frozen instance of this protocol. One card per run; the card publishes with the report; an unreported card publishes as a presumed breach at its horizon (§11).

Registration identity — Registration ID: [ ] — Date of freeze: [ ] — Reporting horizon: [ ] (default: one year from the freeze) — Code identity (repository and hash): [ ]

The construction (R1) — Operator 𝔄: constructed per the anchoring source at [section/equation citation], from closure invariants (k, w, n) and the anchoring dynamics; mass-column exclusions attested. — The source's stated bridge structure governing T2's attribution (V2a vs V4): [citation]. — P$\hbar$'s falsification criterion, as stated at the source: [citation].

Labels and scope (R2, scope clause) — Depth-labelling mechanism, intrinsic: [citation / description]. If absent: the run registers as void for T1 in advance. — Columns registered this run: [ℓ / u / d / ν].

Outputs and error model (R3) — Uncertainty model for δC (truncation, boundary, solver): [stated].

Frozen parameters (R5) — Resolution margin κ = [ ] (default 3). — Per-column tolerance ε_col = [ ] (set from the propagated δC and the measured-mass scheme uncertainty; method stated). — Global tolerance ε_glob = [ ]. — Compatibility condition (the Lattice Lemma's premise, §6): $\varepsilon\_col + \kappa\cdot(\delta C(R_n)/C(R_n) + \delta C(R_{n+1})/C(R_{n+1})) < \ln(m_{n+1}/m_n)$ at every adjacent pair of every registered charged column — ε_col against the public splittings, attested at freeze [ ]; the margin terms, attested at Stage 2 before the lemma is invoked [ ]. — G-quant trigger (§7): frozen numerical criterion — test statistic over the computed (p_v, K_c) under the cited bridge structure, with threshold: [ ]. If the source supports citation grade only: the discretionary-seam declaration, with both attributions bound to publish side by side: [ ]. — Measured-mass scheme and values used at T2, with uncertainties: [stated; loaded at Stage 2 only].

Attestations — Input fence (R4): no mass-column quantity upstream of T2's scoring. [ ] — Freeze (R5): all discretionary numerics fixed above; one pre-unblinding implementation audit, logged if exercised. [ ] — Staged unblinding (R6): Stage 1 (T1, mass-blind) timestamp [ ]; Stage 2 (T2) timestamp [ ]. — Full publication (R7): every cell — pass, fail, unresolved, void — publishes with this card. [ ]

The verdicts (completed at report) — T1, per registered column and adjacent pair: [ ] — T2-col, per registered measured column: [ ] — T2-glob: [ ] — Column verdict of record, per registered column (C-pass / C-prop / C-ord / C-open / C-void): [ ] — Global verdicts (G-glob; G-quant or its non-trigger): [ ] — The neutrino square's cell: [ ] — Secondary readout ($\rho(n) = C(R_{n+1})/C(R_n)$ against the reuse profile — decreasing, above 1; logged, consumed by nothing): [ ] — Distinguished configuration, if any (V1 / V2a / V2b / V3 / V4 / V5): [ ]

Appendix B — The Referee Harness (reference implementation)

The exercise of §14, in full — Parts 1–6: public data, the five-scenario referee, the honest-pass/fence demonstration, the §14.3 toy, the §14.5 fit/scan/scramble, and the §15 gap constraint. The harness implements the referee, not a passing answer: mode-generators are blind to the verdict logic, T1 is scored before any mass is loaded, and the scenarios include runs built to strike and to void. It tests the protocol mechanics only — the real operator 𝔄 is not yet explicit. Re-freeze κ, the tolerances, and the error model on a Registration Card (Appendix A) before any run that reports.

```
import math, statistics
import numpy as np

# ==============================================================
# THE EIGENMODE DECISION — REFERENCE HARNESS
# Reference implementation for section 14. Tests the PROTOCOL
# (labelling, margins, verdict logic, fences), never VERSF: the
# real operator 𝔄 is not yet explicit. Mode-generators are blind
# to the verdict logic; T1 is scored before any mass is loaded.
# ==============================================================
```

```
# ---- frozen protocol parameters (procedural; real run re-freezes) ----
KAPPA   = 3.0     # resolution margin (log form, see t1_pair)
EPS_COL = 0.20    # per-column ln-ratio tolerance
EPS_GLB = 0.40    # global ln-ratio tolerance

# ---- Stage-2 masses (loaded ONLY for T2) ----
MEAS = {
    "charged leptons": [0.51099895, 105.6583755, 1776.86],
    "up-type quarks":  [2.16, 1270.0, 172570.0],
    "down-type quarks":[4.70, 93.5, 4183.0],
    # neutrinos: no absolute scale -> T2 deferred
}

# ============================================================
# T1 in LOG FORM throughout (one language with the compatibility
# condition and the harness). An adjacent pair is resolved when the
# log-gap exceeds the propagated relative margin.
# ============================================================
def t1_pair(c_lo, c_hi, d_lo, d_hi):
    if c_lo <= 0 or c_hi <= 0:
        return "void"
    gap = math.log(c_hi) - math.log(c_lo)             # ln ratio
    margin = KAPPA*(d_lo/c_lo + d_hi/c_hi)            # propagated rel.
margin
    if gap >  margin: return "pass"
    if -gap > margin: return "fail"
    return "unresolved"

def t1_column(modes):
    Cs  = [pv/Kc for (_,pv,Kc,_) in modes]
    dCs = [dC for (*_, dC) in modes]
    p01 = t1_pair(Cs[0],Cs[1],dCs[0],dCs[1])
    p12 = t1_pair(Cs[1],Cs[2],dCs[1],dCs[2])
    if "fail" in (p01,p12):          return "C-ord", Cs, (p01,p12)
    if p01=="pass" and p12=="pass":  return "T1pass", Cs, (p01,p12)
    return "C-open", Cs, (p01,p12)

def t2_col(col, Cs):
    if col not in MEAS: return "deferred", None
    r = [math.log(m) - math.log(C) for m,C in zip(MEAS[col], Cs)]
    return ("pass" if (max(r)-min(r)) <= EPS_COL else "fail"), r

def column_verdict(col, modes, labels_intrinsic=True):
    if not labels_intrinsic:
        return "C-void(R2)", None, None
    t1, Cs, pairs = t1_column(modes)
    if t1 == "C-ord":   return "C-ord", Cs, pairs
    t2, _ = t2_col(col, Cs)
    if t1 == "C-open":  return "C-open", Cs, pairs
    if t2 == "deferred":return "C-pass(T1-only)", Cs, pairs   # unmeasured
sub-case of C-pass
    if t2 == "pass":    return "C-pass", Cs, pairs
    return "C-prop", Cs, pairs

def lattice_compatible(col, Cs, modes):
```

```
    """Log-form compatibility: eps_col + relative margins < adjacent ln-
splitting."""
    if col not in MEAS: return None
    ms = MEAS[col]; ok = True
    for n in (0,1):
        rel = (modes[n][3]/Cs[n]) + (modes[n+1][3]/Cs[n+1])
        budget = EPS_COL + KAPPA*rel
        if not (budget < abs(math.log(ms[n+1]/ms[n]))): ok = False
    return ok

# ================================================================
# PART 1 — PUBLIC-DATA SANITY TEST (retrodictive, O1)
# ================================================================
print("="*64); print("PART 1 — public-data mass ordering (O1)");
print("="*64)
allp=True; smallest=1e9
for col,ms in MEAS.items():
    ok = ms[0]<ms[1]<ms[2]; allp &= ok
    ls=[math.log(ms[1]/ms[0]), math.log(ms[2]/ms[1])]
    smallest=min(smallest,*ls)
    print(f"  {col:18s} ordering {'PASS' if ok else 'FAIL'}   "
          f"ln-splits {ls[0]:.2f},{ls[1]:.2f}")
print(f"  ALL: {'PASS' if allp else 'FAIL'}; smallest splitting
{smallest:.2f} "
      f"(={math.exp(smallest):.1f}x) -> compatibility clears.")

# ================================================================
# PART 2 — REFEREE HARNESS across blind scenarios
# ================================================================
print(); print("="*64)
print(f"PART 2 — referee (kappa={KAPPA}, eps_col={EPS_COL},
eps_glob={EPS_GLB})")
print("="*64)
def gen_ascending(col):
    base={"charged leptons":1.0,"up-type quarks":1.2,"down-type
quarks":0.9,"neutrinos":0.15}[col]
    return [(f"R{n}", base*(1+n), 10.0/(1+0.4*n),
0.02*(base*(1+n)/(10.0/(1+0.4*n)))) for n in range(3)]
def gen_inverted(col):
    return [(f"R{n}", 3-n, 5.0+2.0*n, 0.02*((3-n)/(5.0+2.0*n))) for n in
range(3)]
def gen_misscaled(col):
    fake=[1.0,2.0,3.0]; return [(f"R{n}", 0.1*fake[n]*10, 10.0,
0.01*0.1*fake[n]) for n in range(3)]
def gen_unresolved(col):
    return [(f"R{n}", (0.5+0.01*n)*10, 10.0, 0.10*(0.5+0.01*n)) for n in
range(3)]
scen=[("S1 ascending toy",gen_ascending,True),("S2
inverted",gen_inverted,True),
      ("S3 mis-scaled",gen_misscaled,True),("S4
error>>gap",gen_unresolved,True),
      ("S5 no labels",gen_ascending,False)]
for name,gen,intr in scen:
    print(f"\n--- {name} ---")
    gr=[]
    for col in ["charged leptons","up-type quarks","down-type
quarks","neutrinos"]:
```

```
        modes=gen(col);
v,Cs,pairs=column_verdict(col,modes,labels_intrinsic=intr)
        extra=""
        if Cs and col in MEAS and v in ("C-pass","C-prop"):
            r=[math.log(m)-math.log(C) for m,C in zip(MEAS[col],Cs)]; gr+=r
        if Cs: extra=f"  T1={pairs}"
        print(f"  {col:18s} -> {v}{extra}")
    if gr: print(f"  G-glob ln(m/C) spread={max(gr)-min(gr):.2f} -> "
                 f"{'pass' if (max(gr)-min(gr))<=EPS_GLB else 'fail'}")

# ============================================================
# PART 3 — honest C-pass requires proportionality; the fence
# ============================================================
print(); print("="*64); print("PART 3 — honest pass needs C∝m; provenance
fence"); print("="*64)
def gen_prop(col):
    ms = MEAS[col] if col in MEAS else [1.0,4.5,20.0]
    return [(f"R{n}", 0.001*m*10, 10.0, 0.01*0.001*m) for n,m in
enumerate(ms)]
gr=[]
for col in MEAS:
    modes=gen_prop(col); v,Cs,_=column_verdict(col,modes)
    gr+=[math.log(m)-math.log(C) for m,C in zip(MEAS[col],Cs)]
    print(f"  {col:18s} -> {v}")
print(f"  G-glob spread={max(gr)-min(gr):.3f} -> configuration V1 (IF built
mass-blind)")
print("  fence (R4): built from closure structure -> legitimate V1;")
print("              built by reading the mass column -> VOID. Same
numbers,")
print("              opposite verdict; only the staged log (R6) certifies.")

# ============================================================
# PART 4 — mass-blind toy construction of steps 1-3
# ============================================================
print(); print("="*64); print("PART 4 — toy construction (steps 1-3, mass-
blind)"); print("="*64)
toy_kw={"charged leptons":(1,-1),"up-type quarks":(3,2),"down-type
quarks":(3,-1)}
def gen_toy(col):
    k,w=toy_kw[col]
    return [(f"R{n}", (k+abs(w))*(n+1), 10.0/(n+1),
0.02*((k+abs(w))*(n+1)/(10.0/(n+1)))) for n in range(3)]
print("  rule: p_v=(k+|w|)(n+1), K_c=10/(n+1), C=p_v/K_c  (closed form
C∝(n+1)^2)")
for col in toy_kw:
    modes=gen_toy(col); v,Cs,pairs=column_verdict(col,modes)
    r=[math.log(m)-math.log(C) for m,C in zip(MEAS[col],Cs)]
    print(f"  {col:18s} C={[round(c,2) for c in Cs]} -> {v}  ln(m/C)
spread={max(r)-min(r):.2f}")
print("  identical depth ratios every column (4.00, 2.25); real ratios differ
->")
print("  T1 PASS, T2-col FAIL -> C-prop. Ladder free, magnitudes not.")

# ============================================================
# PART 5 — operator attempt + counting bound (rung-spacing)
# ============================================================
```

```
print(); print("="*64); print("PART 5 — deriving the rung-spacing (k,w)-
>(A,R)"); print("="*64)
spec={"lep":(1,-1,5.332,2.509),"up":(3,2,6.377,1.465),"down":(3,-1,2.990,-
0.810)}
k=np.array([c[0] for c in spec.values()],float); w=np.array([c[1] for c in
spec.values()],float)
A=np.array([c[2] for c in spec.values()]); R=np.array([c[3] for c in
spec.values()])
def best_scaled(t):
    bb=None
    for s in np.linspace(-3,3,6001):
        x=k+s*w; g=np.dot(x,t)/np.dot(x,x); rms=np.sqrt(np.mean((g*x-t)**2))
        if bb is None or rms<bb[0]: bb=(rms,g,s)
    return bb
rA,gA,sA=best_scaled(A); rR,gR,sR=best_scaled(R)
print(f"  ansatz A=alpha(k+sA w): pred {(gA*(k+sA*w)).round(2)} vs {A}
rms={rA:.2f}")
print(f"  ansatz R=rho(k+sR w):   pred {(gR*(k+sR*w)).round(2)} vs {R}
rms={rR:.2f}")
print("  -> over-determined form REFUSED (rms ~ 0.8 of a no-skill
baseline).")
# feasibility scan (the "no two-constant law reaches 0.7" claim)
absw=np.abs(w)
feats={"k":k,"|w|":absw,"w":w,"k+|w|":k+absw,"k*|w|":k*absw,"k-|w|":k-
absw,"1/k":1/k,"k^2":k**2}
best=[]
for fa in feats:
    for fb in feats:
        xa=feats[fa]; ga=np.dot(xa,A)/np.dot(xa,xa);
na=np.sqrt(np.mean((ga*xa-A)**2))/A.std()
        xb=feats[fb]; gb=np.dot(xb,R)/np.dot(xb,xb);
nb=np.sqrt(np.mean((gb*xb-R)**2))/(R.std() or 1)
        best.append((na+nb,fa,fb,na,nb))
best.sort()
tot,fa,fb,na,nb=best[0]
print(f"  best 2-constant law: A~{fa}(nrms {na:.2f}), R~{fb}(nrms {nb:.2f}) -
> none reach 0.7")
# counting bound: 6-constant square law fits anything, incl random
K=np.column_stack([np.ones(3),k,w]); rng=np.random.default_rng(1); ok=True
for _ in range(200):
    t=rng.uniform(-5,8,3); ok &= abs(K@np.linalg.solve(K,t)-t).max()<1e-9
print(f"  counting bound: 6-const (k,w)-law fits 200/200 RANDOM targets "
      f"({'confirmed' if ok else 'FAILED'}) -> charged sector cannot validate
a law.")

# ============================================================
# PART 6 — operator specification: DOF target + the GAP constraint
# (restated in independent observables G1, G2 -- not A, R)
# ============================================================
print(); print("="*64); print("PART 6 — operator spec: DOF target + gap
constraint"); print("="*64)
G1=A.copy(); G2=A-R   # the two genuinely independent log-gaps
print("  charged sector: 6 numbers (3 A,3 R); B fenced. Non-empty operator")
print("  needs < 6 constants; ideal 0 tuned -> 3 columns tests, nu
forecast.")
print(f"  independent observables: G1(=lnC1-lnC0)={G1.round(2)}  G2(=lnC2-
lnC1)={G2.round(2)}")
```

```
print(f"  corr(A,R)   = {np.corrcoef(A,R)[0,1]:.2f}  <- DEFINITIONAL: R=A-G2
shares A,")
print(f"                 floor ~0.6 even for independent gaps -> NOT a
signal.")
print(f"  corr(G1,G2) = {np.corrcoef(G1,G2)[0,1]:.2f}  <- the honest number:
NO signal,")
print(f"                 a clean forecast surface for the 4th column.")
print("  pre-registered constraint: orbit-count operator must predict G2
from")
print("  G1 by a one-parameter relation; 3 columns constrain it, nu tests
it.")
print("  NOTE: down column EXPANDS (G2_down=3.80 > G1_down=2.99, R_down<0) -
>")
print("  the reuse-profile readout is expected MIXED; V1 withholds consonance
there.")
```