Paper 40: Repairs and Retractions

The gap registry, the corrections, the refutations, and the discipline that produced them

Léon Fernando Vlegels

2026-06-11

Abstract

A corpus that can only confirm itself measures nothing. This paper is the canonical record of the corpus measuring itself. The canon carries 65 recorded gaps: 54 remarks typed *TBS* (to be shown) and 11 typed *Open*, written into the papers themselves as first-class, machine-findable objects. Addendum 277 extracted them into a registry (`verify/tbs_registry.json`), regenerated from the LaTeX sources on every verifier run by a single writer (`tbs_registry_gen.py`), under one hard rule: a status that cites a closing document fails the build if that document is absent from disk. A burn-down followed, on the record: A277 created the registry with 58 items unreviewed; A278 swept the prior addenda and ran computational kills (58 → 48); A279 settled the named load-bearing items (48 → 44); A280 forced the P07 decision and delivered the corpus's first formal refutation; A293 settled the severe items (44 → 38); A295 classified everything that remained (38 → 0 unreviewed). The final classification, quoted from the registry: 15 retired, 1 superseded, 21 scope, 4 refuted, 21 confirmed-load-bearing, 0 unreviewed; alongside these, 327 verifier-documented items record every expected-fail check in the paper-series verifiers, the registry in executable form. A later addendum (A344) added three flavor-sector refutations from Paper 01, bringing the refuted count to 7. The refuted items are told here as full case studies: P016_3 (the Mersenne shell statistics, irreproducible under all four natural readings of their own definition), P020_1 (the electroweak VEV conjecture, off by fourteen orders of magnitude as stated), P017_2 with P017_3_c (the Dirac spectra: the printed system is ill-posed, and the natural repair produces box spectra with no factor of seven and no 0.891), and P001_4/P001_5/P001_6 (the tri-bimaximal $\theta_{13} = 0$, the CP phase $\delta_{\mathrm{CP}} \approx 240^\circ$, and the topological strong-CP argument, none of which survive current data). The standing diagnoses, the distinction between retired, refuted, and scope, and the living-ledger property of this paper close the record. No claim in the corpus, including the seed identity itself, is protected from its own registry.

1 Introduction: Why a Corpus Should Audit Itself

A body of claims earns trust by what it is able to lose. A theory that confirms everything it touches has not been tested by anything it touches; its agreement with the world is a property of its bookkeeping, not of the world. This is an old point, and it is usually applied from outside: referees, replications, adversarial collaborations. The corpus applies it from inside. It carries its own list of weak points, written into the papers at the places where the weakness lives, in a fixed form that a program can find; it maintains a machine-readable ledger of those weak points and their fates; and it accepts, as a matter of standing policy, that some of its own headlines will die when the ledger is worked through. Several did. This paper records how, and what the verdicts were.

The need is structural, not decorative. The corpus makes three kinds of statement: derivations, which close or do not; identifications of corpus quantities with measured constants, which carry residuals; and conjectures, which are scaffolding until derived. All three kinds were written down, in places, with known soft spots: a boundary condition asserted and not derived, a constant matched and not explained, a table published and not recomputed. A corpus that kept such spots in margins or memory would slowly confuse the proven with the familiar. The discipline described here exists to make that confusion impossible: every soft spot is a tracked object, every tracked object has a status, and every status change cites the document that caused it or the build fails.

Two commitments shape this paper. First, it contains no new claims. Every number is quoted from the registry (`addenda/verify/tbs_registry.json`) or from the addendum that produced it; the paper describes the discipline and records the verdicts, exactly as the registry holds them. Second, the ledger is allowed its dates. The corpus's prose elsewhere avoids chronicle; here the dated entries and addendum numbers are the content, because a record of self-correction is precisely a record of when, what, and by which document.

2 The Discipline

2.1 The remark convention

Gaps are recorded in the canon papers themselves as theorem-class objects: a `remark` environment typed `[TBS]` (to be shown) or `[Open]`, carrying a label of the form `rmk:tbs:PNNN_k` or `rmk:open:PNNN_k`, where `PNNN` names the paper and `k` indexes the gap within it. A TBS remark marks a statement the paper uses but does not establish; an Open remark marks a question the paper raises and leaves. Both sit in the body of the paper, adjacent to the claim they qualify, in the same typographic register as the theorems. The convention makes the gap a first-class citizen of the document: it cannot be lost in a footnote, it survives every edit that preserves the label, and it is findable by a regular expression. Papers 00–34 carry 54 TBS remarks and 11 Open remarks: 65 recorded gaps. (Of the TBS remarks, the three latest — P001_4, P001_5, P001_6 — were added to Paper 01 by A344 to record three flavor-sector predictions as refuted.)

2.2 The registry and its single writer

Addendum 277 (2026-06-10) extracted the 62 remarks into `verify/tbs_registry.json`: one record per gap, holding the paper, the label, a 200-character excerpt, a status, the closing document where one exists, and a note. The registry is not edited by hand. It is regenerated from the LaTeX sources on every verifier run, so it cannot silently drift from the papers: if a remark is added, removed, or relabelled, the next run reflects it.

Ownership was deliberately narrowed as the registry matured. A277's verifier wrote the file first. The same day, ownership transferred to the Burn-down I verifier (`verify_P278.py`), with A277's verifier demoted to reading and asserting. A279 moved the status map and the generation logic into a shared module, `tbs_registry_gen.py`, which has been the single writer since: burn-down verifiers import it, regenerate the registry, and assert only their own contributions. No verifier owns the global counts, so two verifiers can never fight over them, and a later correction never breaks an earlier verifier's record. The expectations are evolution-tolerant by construction.

2.3 The closure rule

Proposition 2.1 (Registry discipline, adopted as canon (A277))**.** *(i) An addendum that closes a recorded gap updates the registry status in the same change. (ii) A status that cites a closing document fails the build if the cited document does not exist on disk. (iii) The unreviewed count is reported on every run. The canon's debt is therefore a number that can only change visibly: nothing is closed by forgetting, and nothing is closed by assertion without a citable document behind the assertion.*

The second clause is the load-bearing one. A ledger whose entries can point at nothing is a list of moods. Requiring the closing document to exist, and requiring the build to fail otherwise, converts every closure into a checkable cross-reference: the reader who doubts a retirement can open the addendum it cites, and the build system has already verified that there is an addendum to open.

3 The Burn-Down

The registry was created with 58 of its 62 items unreviewed. It was driven to zero in seven documented stages over two days, listed here by addendum number. The trajectory of the unreviewed count: $58 \to 48 \to 44 \to 44 \to 38 \to 0$.

A277 (2026-06-10), creation. The inventory: 51 TBS, 11 Open. Four closures already on record were registered retroactively: P028_2 (closed by A273, the unit bridge), P028_3_c (A263, the velocity identification), P014_1 (A266, the exact comma G_1 replacing $\kappa^{0.6}$), and Open P027_2 (A276, the section law replacing the invertibility objection). Unreviewed: 58.

A278 (2026-06-10), Burn-down I. Stage 1 swept addenda 35–255 for closures that had never been registered: P031_1_c had been explicitly closed by A237 on 2026-05-24; P015_3_c was retired by A266/A267 (the asserted 432 now derived, the $\pi/2$ correction identified as a 5.6%-wrong proxy for the exact comma gap $432 - T_{\text{breath}}$); P034_3_c was retired by A265/A273/A275 (the quoted "430.5 s" replaced by the dimensionless T_{breath} with the SI duration derived separately); Open P019_3 was superseded by P32's corner-residual route, with the equivalence of routes noted as unproved; and P021_2_c was reclassified *scope*, a permanent honesty marker rather than fixable debt. Stage 2 ran computational kills: P003_3 retired (the Neumann misstatement is real, the proof is Dirichlet, the spectrum is reproduced to five digits, downstream effect zero); P029_1 retired (the corrected resonance peak $\sqrt{\omega_0^2 - \gamma^2/2}$ confirmed to seven digits); P007_1/P007_2 confirmed load-bearing with quantified teeth (rebuilding the observation operator under the claimed versus de facto boundary conditions shifts $\lambda_{1,2,3}$ by -4.7%, -16.2%, -22.4%); and P016_3 confirmed failing as stated (the recomputed Mersenne counts are 1 and 9 against the claimed 14 and 19). Unreviewed: 48.

A279 (2026-06-10), Burn-down II. The three named load-bearing targets. P000_1_c: the deviation of $\Omega = 4\pi^3 + \pi^2 + \pi$ from the measured α^{-1} was computed precisely (2.2234 ppm), named the α-comma, and kept open with an objective. P014_2: retired by exact decomposition, $432/\pi^3 = (1 + G_1)\,\Omega/\pi^2$ identically; the claimed identification with E_{self} is structurally false (the additive excesses differ by a factor 2.16). P027_3_c: the mass-hierarchy formula exact-restated and sharpened, with the constant $3\pi/20$ isolated as the underived part. Bonus: P018_1 retired, its factor-$\sim$138 mystery resolved exactly as the $\mu_0 = \Omega$ normalization ($0.37385 \times \Omega = 51.23$, dressed 51.53). The status map moved to the shared generator in this addendum. Unreviewed: 44.

A280 (2026-06-10), Burn-down III. The two quantified debts from A278, finished. P007_1/P007_2 moved to *reworked*: the mass predictions and the stated boundary physics cannot both be true,

and the decision was filed as OI-280-1 rather than patched silently. P016_3 moved to *refuted*: the corpus's first formally refuted canon claim (Section 5.1). Unreviewed: 44 (reclassifications only).

A281–A283 (2026-06-10), repairs between burn-downs. A281 derived the observation operator's boundary condition from corpus primitives, resolving OI-280-1; P007_1 and P007_2 moved from reworked to retired, with the prose retraction recorded (observation vanishes at the universe boundary rather than being free there; the mass data select the corrected reading). A282 found the provenance of the mass-hierarchy formula in P03 and sharpened P027_3_c's note. A283 typed the α-comma's mechanism and sharpened P000_1_c's note.

A293 (2026-06-10), Burn-down IV. The severe items. P015_1 retired by repair: the claimed $\langle|\cos\theta|\rangle = 1/4$ is wrong (the spherical average is exactly $1/2$), but the paper's result $\bar{A}_{\text{shadow}} = S/4$ is Cauchy's projection formula, the missing factor of 2 being the front-back double count. P020_1 refuted (Section 5.2). P017_3_c confirmed with the diagnosis the remark lacked: the printed radial Dirac system has complex indicial exponents at the origin and is ill-posed as printed. P013_1, P013_2, P013_3_c confirmed: the quintic is a conjecture-schema (Section 6). Unreviewed: 38.

A295 (2026-06-11), pre-publication closure. The P17 repair was performed and the printed spectrum was not recovered: P017_3_c upgraded from diagnosed to refuted-as-stated, and P017_2 (the factor of seven) refuted with it (Section 5.3). P018_2 retired (its normalization conflict is the μ_0 factor documented in A279/A282). Every remaining item was classified: the application papers' honesty remarks to *scope*, the formalization gaps to *confirmed-load-bearing* with notes, two items settled with quantified diagnoses (P002_1: the flagged conflation is real and numerically benign at 1.2%; P033_2: the formula $\text{Vol}(S^{d-1}) = \pi^d$ is confirmed false at every d, with correct values supplied). Unreviewed: 0. Nothing in the corpus is both open and unexamined.

A344 (2026-06-15), flavor-sector refutations. After the burn-down closed, three predictions of Paper 01 were tested against current data and recorded as refuted (Section 5.4): the tri-bimaximal reactor angle $\theta_{13} = 0$ (P001_4), excluded by the measured $\theta_{13} \approx 8.5°$; the CP phase $\delta_{\text{CP}} \approx 240°$ (P001_5), now $\sim 3\sigma$ above the NuFIT-6.0 central value 177°; and the topological strong-CP argument (P001_6), which constrains the instanton winding number, not the Lagrangian θ-coupling. Three TBS remarks were added to Paper 01, raising the recorded-gap total to 65 and the refuted count to 7.

4 The Classification

The registry's final counts, quoted directly from `tbs_registry.json`:

status	count
retired	15
superseded	1
scope	21
refuted	7
confirmed-load-bearing	21
unreviewed	0
total recorded gaps	65
verifier-documented (separate class)	327

The statuses are not interchangeable shades of "closed." Each makes a different assertion, and the differences carry the paper's meaning.

4.1 Retired: the gap is closed, with a citable document

Retired means the defect was examined and resolved: either repaired (the statement corrected, the result preserved) or shown to have no downstream effect. Every retirement cites the document that closed it, and the build verifies the document exists. Retirement is the strongest closure the registry grants, and it subdivides in practice into two kinds.

The first kind is *repair*: the printed statement was wrong, the underlying result survives the correction. The repair table:

item	defect	repair
P003_3	Neumann claim, $\rho'(0) = 2\pi \neq 0$	proof used Dirichlet; spectrum intact (A278)
P007_1/2	discretization vs claimed BCs	Dirichlet derived twice over (A281)
P015_1	$\langle \lvert\cos\theta\rvert \rangle = 1/4$	true value $1/2$; $S/4$ is Cauchy's formula (A293)
P029_1	peak frequency formula	$\sqrt{\omega_0^2 - \gamma^2/2}$ confirmed (A278)
P018_1	factor-$\sim$138 discrepancy	the factor is $\mu_0 = \Omega$; normalization documented (A279)
P018_2	normalization conflict	the same μ_0 factor (A295, via A279/A282)
P014_1	$1 + \kappa^{0.6}$	exact comma $G_1 = 432/T_{\text{breath}} - 1$ (A266)
P014_2	$432/\pi^3 = E_{\text{self}}$	exact: $(4\pi + 1 + 1/\pi)(1 + G_1)$; identification false (A279)
P015_3_c	the $\pi/2$ correction	a 5.6‰-wrong proxy for the comma gap $432 - T_{\text{breath}}$ (A278)
P028_2	$\sqrt{1-\kappa}$ vs 2π	a unit mismatch; bridge derived (A273/P35)
P028_3_c	lapse identification	closed by the null-projection identification (A263/P39)
P027_2	2-to-1 invertibility objection	$C \circ P = I$ is a section law (A276/P38)
P031_1_c	monodromy base case	closed by A237 (2026-05-24, registered by A278)
P034_3_c	"430.5 s"	dimensionless T_{breath}; SI derived (A265/P35)

Each repair preserves the downstream result while correcting the statement; none required protecting a claim from data.

The second kind is *retraction of interpretation*: the mathematics stands, the prose around it is withdrawn. The recorded cases: "observation is free at the universe boundary" (P07), retracted because observation vanishes there, by substrate inheritance and by the nesting definition, with the mass data selecting the corrected reading at 9.1$\times$ (A280/A281); P27's mass-formula dressing $1/(1-\mu_1\alpha^2)$, a corrupted and uncited transcription of P03's derived $(1+\mu_1\alpha^2)$, with the original adopted as seven times closer to measurement (A282/P36); the Fatou/Julia dictionary, retired because the breath map is elliptic and its Julia set empty, the canonical/dark dichotomy being carried instead by Arnold tongues and kernel typing (A274/P37/P38); the tempered branch hypothesis, overturned within a day by the derived unit bridge (A273/P35), the overturning itself part of the record; and the one-scale halo's apparent rejection (the 7.6σ core-disk trend), a measurement artifact, quantified and collapsed by masking (A270/P39).

4.2 Superseded: replaced, not resolved

Superseded marks a gap whose question was overtaken by a different route rather than answered on its own terms. The single case: Open P019_3, superseded by P32's corner-residual route, which has been canonical since. The registry records, deliberately, that the equivalence of the two routes remains unproved. Supersession is weaker than retirement and the status says so.

4.3 Scope: a recorded limit, not fixable debt

Scope marks a defect that is real, acknowledged, and outside the corpus's repair obligations: a boundary of the claim rather than a hole in it. The 21 scope items divide into three groups. One is

P021_2_c, a permanent honesty marker: the paper itself states that it does not resolve the formal Clay problem, and no repair could change that without changing the paper's ambition. Sixteen are the recorded defects of the Millennium application papers P21–P26, classified by A295: those papers sit outside the load-bearing spine by package design, their defects stand on the record, and repair is not scheduled. Four are technical or scope notes (P008_3_c, P029_2_c, P029_3_c, P030_1): external anchors or theorem-hypothesis caveats whose corpus use is unaffected. A scope item is never silently promoted to "fine"; it stays in the ledger as a named limit.

4.4 Refuted: the claim is dead

Refuted means the claim was recomputed from its own definition and failed, under charitable readings, with the failure quantified and filed with the same verifier rigor as a theorem. The registry holds four refuted items, covering three cases (the two P17 items fall together). They are told in full in Section 5. Refutation is not retirement: nothing survives to be preserved. It is also not scope: the claim was not beyond the corpus's obligations, it was inside them and lost.

4.5 Confirmed-load-bearing: the gap is real and standing

Confirmed-load-bearing means the gap was examined and found to matter: it is a genuine dependency of something the corpus asserts, it is not yet closed, and the registry carries a quantified or structural note saying exactly what is missing. The 21 items include the corpus's two open hearts (the α-comma P000_1_c and the mass-hierarchy constant P027_3_c, both with named objectives), the three quintic items of P13, and a documented class of formalization gaps (Section 6). Load-bearing status is a debt acknowledgment, not a deferral: each note states what a closing document would have to contain.

5 The Refutations

Seven registry items are refuted: four from the burn-down, and three flavor-sector predictions of Paper 01 recorded by A344. Each case is recorded here in the same shape: the claim, the test, the verdict, and what survives.

5.1 P016_3: the Mersenne shell statistics

The claim. P16 asserted that Mersenne prime exponents prefer integer shells: that the cumulative quantities $S_n = \sum_{i \le n} p_i \kappa$, with $\kappa = \alpha^{5/4}$, land near integers far more often than chance allows. The published table reported 14 exponents with shell distance $d_n < 0.01$ and 19 with $d_n < 0.05$ out of 51, with mean distance 0.187.

The test. A278 recomputed the counts from the paper's own definition over all 51 known Mersenne prime exponents and obtained 1 and 9, mean 0.221, against expectations of approximately 1.0 and 5.1 under the uniform null. A280 then widened the test to four natural readings of the definition, on the principle that a claim should be given every charitable construction before a verdict: cumulative S_n as stated; per-prime $p\kappa$; the modular form $p \bmod \kappa^{-1}$; and the exponent variant $p\kappa^{4/5}$. The results: best counts 2 and 10 against the claimed 14 and 19, with means at or near the uniform null in every reading. The strongest excess anywhere (10 observed against 5.1 expected at the 0.05 threshold) carries $p \approx 0.03$ before any look-elsewhere correction across readings and thresholds, below the corpus bar.

The verdict. No reading reproduces the published counts; under every reading the headline is unsupported. P016_3 is the corpus's first formally refuted canon claim, filed by A280 with full verifier rigor. One forensic note is on the record: the per-prime reading reproduces the claimed *mean* (0.193 against 0.187, within 3%) while its counts still fail, which suggests the published table mixed a per-prime mean with counts from an unstated computation.

What survives. The downstream exposure was checked: nothing else in the corpus cites the shell counts quantitatively. P07's "three measurements of κ" analogy table inherits a caveat, not a collapse. The exponent 5/4 and the constant κ themselves are established elsewhere and are untouched.

5.2 P020_1: the electroweak VEV conjecture

The claim. P20 conjectured that the electroweak vacuum expectation value is $v = M_{\mathrm{Pl}}/\sqrt{\mu_0\mu_1}$, tying the electroweak scale to the corpus's moment constants.

The test. A293 evaluated it. With $\sqrt{\mu_0\mu_1} = 122.06$, the formula gives 1.0003×10^{17} GeV. The measured value is 246.22 GeV. The ratio is 4.06×10^{14}: the conjecture misses by more than fourteen orders of magnitude.

The verdict. Refuted as stated. There is no reading under which a factor of 10^{14} is a residual; the printed formula is dead.

What survives. Two flags were recorded in the same addendum, per the corpus's flag discipline (recorded, unpromoted, look-elsewhere acknowledged): $\ln(v/m_e) = 13.0854$ sits 0.70% from $E_{\mathrm{self}} = 13.177$, and the exponent fraction $\ln(M_{\mathrm{Pl}}/v)/\ln(M_{\mathrm{Pl}}/m_e) = 0.74605$ sits 0.53% from 3/4, the P10 correction. Either could seed a corrected conjecture; neither is one, and the registry says so in those words. A corrected conjecture would have to derive one of these and the residual. The refutation of P020_1 does not touch the mass-hierarchy formula of P27, which uses the moment constants differently and is tracked separately under P027_3_c.

5.3 P017_2 and P017_3_c: the Dirac spectra

The claim. P17 printed a radial Dirac system on the density background, quoted its spectrum as (0.891, 6.923, 7.763), and built on a factor-of-seven gap between levels (P017_2). P017_3_c recorded, originally as a TBS item, that the spectrum had not been independently reproduced.

The test, stage one (A293): the printed system. Near the origin the printed system reduces to $(f, g)' = r^{-1}M(f, g)$ with M having eigenvalues $-\frac{3}{2} \pm \frac{3}{2}i$ at $\kappa = 3/2$: a complex indicial pair. Both branches are singular and oscillatory; no regular solution exists; regularity at $r = 0$ cannot be imposed. The printed equations define no spectrum, so the quoted eigenvalues cannot be checked against them even in principle. The diagnosis identified the defect precisely: the printed system places κ/r off the diagonal, where the standard radial Dirac form places it on the diagonal. At this stage the item was confirmed load-bearing with the diagnosis filed, and P017_2 was marked as inheriting: unsupported pending the corrected system.

The test, stage two (A295): the natural repair. The standard-form system $f' = ((\kappa - \frac{3}{2})/r)f - Mg$, $g' = -((\kappa + \frac{3}{2})/r)g + Mf$ with $M(r) = m + \alpha\rho(r)$ restores a regular branch at the origin. Numerical shooting on $[0, 1]$ (RK4, $N = 4000$ to 20000, convergence checked), across both channels ($\kappa = 3/2,\ 5/2$) and both admissible boundary conditions, yields evenly spaced, box-like spectra with level spacing near π: (2.27, 5.90, 9.12), (4.00, 7.40, 10.62), (3.40, 7.20, 10.52), (5.10, 8.67). No reading produces the printed small ground state 0.891; no reading produces the close upper pair; the level ratios run 1.8 to 2.6, and no factor of seven appears anywhere.

The verdict. P017_3_c is refuted as stated on two independent grounds: the printed system is ill-posed, and the printed spectrum is not recovered by the natural repair. P017_2, the factor of seven, is refuted with it.

What survives. The corrected spectra are filed as the repaired operator's actual output, available as the starting point for any successor of P17's program. The diagnosis itself (where the sign convention went wrong, and what the corrected system is) is a permanent gain: the refutation leaves behind a working operator, just not the claimed numbers.

5.4 P001_4, P001_5, P001_6: the flavor-sector predictions

The claims. Paper 01 made three flavor-sector predictions from its A_4 tetrahedral structure: a tri-bimaximal reactor angle $\theta_{13} = 0$; a leptonic CP phase $\delta_{\rm CP} \approx 240°$, claimed within 2σ of an experimental $197° \pm 25°$; and a topological resolution of the strong-CP problem, $\theta_{\rm QCD} = 0$ from boundary conditions on B^4, with "no axion needed."

The test. Each was checked against current data. The reactor angle is measured at $\sin^2\theta_{13} = 0.02224$, i.e. $\theta_{13} \approx 8.5°$ (NuFIT-6.0, 2024, normal ordering), nonzero at more than 5σ since Daya Bay (2012); $\theta_{13} = 0$ sits more than 30σ from the measured value. The CP phase global fit has moved to $\delta_{\rm CP} \approx 177°\ ^{+19}_{-20}$ degrees, consistent with CP conservation within 1σ, placing the predicted 240° about 3.3σ above the central value; the paper's claimed agreement used the superseded $197° \pm 25°$ fit. The strong-CP argument fixes the instanton winding number $n = 0$, but the physical strong-CP parameter is the Lagrangian coupling $\bar\theta$, an independent input: the action enters as $\bar\theta\, n$, so a configuration with $n = 0$ contributes at every $\bar\theta$, and fixing the winding does not fix the coupling.

The verdict. All three are refuted. P001_4 ($\theta_{13} = 0$) is excluded by the data; P001_5 ($\delta_{\rm CP} \approx 240°$) fails on the current central value; P001_6 (topological strong CP) constrains the wrong object and does not solve strong CP, which remains an open problem of the Standard Model. The three are filed by A344, with TBS remarks added to Paper 01 adjacent to the claims they qualify.

What survives. The A_4 family structure and the lens-space account of three generations are untouched by these refutations; what falls is the specific quantitative flavor output — the zero reactor angle, the 240° phase, and the axion-free strong-CP claim. Paper 01's abstract was reframed to mark all three as not having survived current data and to flag them for this registry; the recording here makes that citation true.

6 The Standing Diagnoses

The 21 confirmed-load-bearing items are not undifferentiated debt; each carries a note stating exactly what is missing. The deepest is a structural diagnosis rather than a numerical one.

The P13 quintic is a conjecture-schema (P013_1, P013_2, P013_3_c; A293). P13's headline asserts that a certain quantity satisfies an irreducible quintic with computable coefficients. The diagnosis: Condition 1 invokes $\kappa_{\rm geom} = C_1\mu_2/\mu_0^2 + C_2\mu_1/\mu_0 + C_3$ with the constants C_1, C_2, C_3 never fixed; Condition 2 invokes $F(\lambda_1, \beta) = 0$ with the functional F never given; and the rationality premise (P013_3_c) is unproven and generically false, since Sturm-Liouville eigenvalues are generically transcendental in the parameter. The coefficients $a_5, \ldots, a_0$ are therefore not merely uncomputed but uncomputable as stated, and the degree bound (P013_2) is unverifiable until the construction is supplied. The headline stands reclassified as a schema: a shape a future derivation could fill, not a result. The corpus's 5/4 exponent and five-fold structure stand independently of P13.

The α-comma (P000_1_c; A279/A283/A284). The identity $\alpha^{-1} = \Omega$ is accurate to 2.2234 ppm against CODATA and is not exact at any edition, roughly 1.45×10^4 experimental standard

deviations away. The registry's note records the trajectory: typed as the time-average dressing of P03's oscillation, with deficit $1 - \sqrt{1 - A^2}$; the reading $A = \kappa$ predicts 2.2745 ppm, the right sign and 2.3% high; inverted, the residual becomes the first measurement of the oscillation amplitude, $A = 0.9887(2)\,\kappa$, stable across CODATA editions. A284 bracketed the correction ($1 - \pi/\Omega$ versus $1 - 1/44$, identification below bar) and filed the conditional edge prediction $\alpha^{-1} = 137.035999236$, 0.43 ppb from CODATA-2022, with a 0.1-ppb discriminator. The exact amplitude remains the objective (OI-283-1). This item is the corpus owning a residual in its most external-facing identity and converting it into a pre-registered prediction rather than absorbing it.

The mass-hierarchy constant (P027_3_c; A279/A282). The formula $\ln(M_{\mathrm{Pl}}/m_e)$ evaluates to 51.5299 against the measured 51.5278 (3.9×10^{-5} relative). Its provenance was found: derived in P03, uncited by P27, with P27's dressing a corrupted transcription of P03's form (the original is seven times more accurate and is adopted). What remains underived is the integer 10 in $3\pi/20 = (1/10)(3\pi/2)$, where $3\pi/2$ is the textbook one-loop QED coefficient, together with the correspondence postulate itself (OI-282-1). One underived integer buying a 3.9×10^{-5} hit on a 22-decade mass ratio is either deep or a fit, and the registry says the corpus does not yet know which.

The formalization-gap class (A295). The remaining load-bearing items are documented gaps of formalization or verification rather than suspected errors, each with a one-line note in the registry: the R-operator never formally specified (P000_3); the 5/4 global minimality unproven, local result only (P003_2); the moments-to-hierarchy correspondence postulate (P007_3); dimension not derived from transversality (P008_1); three recorded circularities (P008_2_c, P009_3_c, P010_3); the unexplained 4/3 step (P015_2); a drift-bound inconsistency that is testable against the GON implementation and not yet run (P030_2); two structural claims unverified (P030_3_c, P033_3_c); two open monodromy steps (P031_2_c, P031_3_c); an associativity proof gap (P034_2_c); the conflation in P002_1, confirmed real and numerically benign at 1.2%; and P033_2, where the printed formula $\mathrm{Vol}(S^{d-1}) = \pi^d$ is confirmed false at every d, the correct values are supplied, and the downstream argument is flagged for rework.

7 The Verifiers: the Registry in Executable Form

Every canon paper has a verifier, `verify_P000.py` through `verify_P034.py` and onward, which recomputes the paper's checkable claims at run time. The suite carries a convention that is the registry's executable counterpart: the *expected fail*. A check that a paper's claim does not survive recomputation is not deleted and not inverted; it is kept in the verifier, printed as a `[FAIL]` line with the claim and the reason, and documented as failing. The suite treats a documented failure as a recorded fact, not a build error: the verifier passes because it correctly reports the claim's status, which is the verifier's actual job.

These lines are harvested into the registry by the single writer's second function, `generate_vdc()`, which runs each paper-series verifier, parses its `[FAIL]` output, and files one *verifier-documented* item per check, with the verifier name, the claim, and the note. The registry currently carries 327 such items, in a section separate from the 62 remark-derived gaps so that existing consumers of the remark counts are undisturbed.

The point of the convention is that the corpus's defects are reproducible by execution, not only by reading. A reader who doubts the prose does not have to trust this paper or any addendum: running the suite reproduces the confirmations and the failures alike, and the registry's verifier-documented section is, at any moment, exactly what the suite reports. The closure rule binds in the same direction: a status citing a closing addendum is checked against the disk on every

regeneration. The ledger is not an account of the code; it is output of the code.

8 The Living Ledger

This paper is the prose face of a generated file, and the relationship runs in one direction: the registry is the source of truth, and the paper quotes it. Three consequences follow.

First, every count printed in an addendum is a dated snapshot. The burn-down addenda each printed the totals as they stood at filing, and the snapshots disagree with one another by construction, because the ledger moved between them. They also disagree, in places, with the final map: A295's printed totals predate the last status adjustments, under which P017_3_c moved from the load-bearing class to refuted (its repair having failed to recover the printed spectrum) and four technical notes settled as scope. Where a printed snapshot and the registry differ, the registry is current. The snapshots are not errors; they are the ledger's history, and the single-writer discipline exists precisely so that the history and the present state cannot be confused.

Second, this paper updates as the registry does. Its numbers are quoted from the registry at the date in its header; a future status change obligates a revision here, in the same change, under the same rule that binds every other closure. A repairs-and-retractions paper that froze while its registry moved would be an instance of the disease it documents.

Third, the zero matters but is not a terminus. Zero unreviewed means every recorded gap has been examined and classified; it does not mean the debt is paid. Twenty-one items stand confirmed-load-bearing with named objectives, and the registry will reopen, renumber, and reclassify as the corpus grows. The invariant is not that the ledger is ever finished; it is that the ledger is never silent.

9 Conclusion

A framework that cannot lose measures nothing. The statement is not rhetoric; it is the operating principle this paper documents. A measurement is a procedure that could have come out otherwise. A corpus whose every recorded check confirms its claims has either stopped checking or stopped recording, and in both cases its agreement with anything has ceased to carry information. The registry is the corpus's refusal of that state: 65 gaps recorded in the papers themselves, driven to zero unreviewed through six documented stages, with seven claims refuted (four from the burn-down and three flavor-sector predictions of Paper 01 recorded by A344), fifteen retired on citable repairs, twenty-one classified as standing debt with named objectives, and 327 expected-fail checks kept executable so that the defects reproduce on demand.

Remark 9.1 (The discipline, adopted as canon)**.** *(1) Every recorded gap lives in the registry; closing one updates the registry in the same change, and a closure must cite its closing document or the build fails. (2) Refutations and deflations receive the same verifier rigor as theorems: the registry has refuted canon headlines, deflated a hoped-for selection (A290), and overturned a working hypothesis within a day, and each is cited here as evidence the instrument works. (3) Same-day erratum and upgrade notes are appended to filed addenda rather than silently edited. (4) Flags are never promoted without a derivation (A245/A180); fragile fits are recorded as fragile (the* $\kappa/7$ *case); chains outrank fits because they falsify in more places. (5) No claim, including the seed identity itself, is protected from its own registry: the seed's* 2.2 *ppm deviation was named, typed, and turned into a pre-registered prediction (P36). The corpus corrects itself in public, on the record, by machine-checked rules; this paper exists so that the record is canon.*

Provenance and verification. Assembled without new claims from A277 (registry), A278–A280, A293, and A295 (burn-downs and closure), the flavor-sector refutations of A344, and the repair addenda A263, A266, A270, A273–A276, A279, A281–A284, A294. Verifiers `verify_P277/278/279/280/293/2` and the single-writer generator `tbs_registry_gen.py`, all green at filing; the registry JSON ships with the package as its quality certificate. All counts in this paper are quoted from `addenda/verify/tbs_registry.` 15 retired, 1 superseded, 21 scope, 4 refuted, 21 confirmed-load-bearing, 0 unreviewed, 327 verifier-documented.