Electroweak Coherence Selection in VERSF

TPB Projection, Interface Matching, and the Residual Hierarchy Problem

Keith Taylor *VERSF Theoretical Physics Programme*

Plain-Language Summary

A version of this paper for readers who want the picture without the mathematics.

This paper is one of two within the VERSF programme (the Void Energy–Regulated Space Framework) addressing the hierarchy problem — the longstanding puzzle of why the Higgs particle (the carrier of mass for the other particles, discovered at the Large Hadron Collider in 2012) weighs about 125 GeV rather than something close to 10^{19} GeV, the energy at which gravity itself becomes quantum. The standard reasoning says the Higgs ought to be "infected" by quantum effects from that high energy, with the universe fine-tuned to about one part in 10^{34} to keep it light. Half a century of theoretical physics has been spent trying to explain this away.

The companion paper does the reframing. A separate paper in the programme, *The Hierarchy Problem as a Category Error* (Taylor, VERSF), argues that the puzzle is misdiagnosed at the level of what is being compared. The Planck scale, it argues, is not a higher energy at which more particles still propagate — like a deeper part of the same ocean — but a *boundary* on what reality can distinguish at all, more like the surface of the ocean than a region below it. The seventeen-order gap between the Higgs and the Planck scale is then not fine-tuning. It is the natural separation between two scales belonging to different *layers* of reality: the substrate (the **closure layer**, which governs what configurations are admissible at all) and the field-theoretic level where ordinary physics happens (the **record layer**, where particles propagate). The companion paper identifies a natural energy scale around **5–6 TeV** — at or just above the upper edge of current direct-search reach across most LHC channels — as the maximum scale at which stable, organized structure can be sustained by the substrate. It calls this the **coherence band**.

This paper does the second-order analysis. With the category-error reframing taken as given, the obvious follow-up question is: *if the seventeen-order puzzle dissolves, what is the structure that actually sets the electroweak scale within that coherence band?* The Higgs vacuum (the scale that sets the masses of W and Z particles, about 246 GeV) lives at roughly **4%** of the 5–6 TeV band. The companion paper explains why the band sits where it sits. It does *not* explain why the electroweak vacuum stabilizes at this particular fraction rather than at, say, 1% or 40%. That is the question the present paper answers.

The proposal is that the 4% — denoted χ_TPB — factorizes into a product of three structurally distinct ingredients:

a **channel multiplicity** of 6, from how the substrate's closure constraints organize themselves (six local constraints with equal weights),

times **α ≈ 1/137**, the fine-structure constant of electromagnetism — here playing a more fundamental role as the efficiency with which individual closure channels transfer distinguishability across surfaces,

times a small **efficiency factor** close to 1 (about 0.97), capturing the loss to second-order closure competition between channels.

The product $6 \times \alpha \times 0.97 \approx 0.042$ matches the observed 4%. This is the paper's central equation, the **master occupancy relation**. It is offered as a *factorization ansatz*: a proposal about the structural form of the answer, to be tested by an independent derivation of the small efficiency factor from substrate microphysics.

How the two papers relate. The companion paper performs the *reframing* — converting a seventeen-order fine-tuning crisis into a moderate coherence-selection problem at the few-TeV scale, with the dissolution resting on the Ontological Separation Principle, the Closure Non-Propagation Principle, the admissibility kernel, and the geometric-mean scale $E_{geo} \approx 5$–6 TeV. The present paper performs the *second-order analysis* — accepting that reframing and asking what specific structural form the residual occupancy fraction takes within the coherence band. The companion paper produces the picture; this paper attempts to organize the percent-level structure inside it. They are complementary, and either can be read first, though readers unfamiliar with the category-error framing may prefer the companion paper as orientation.

The paper is careful about what it claims, derives, and leaves open. The multiplicative factorization is proposed, not derived. The channel multiplicity of 6 is derived elsewhere in VERSF, conditional on a deeper principle called *closure democracy* (no closure constraint is privileged over any other). The structural role of α is argued from a separate analysis of geometry and finite distinguishability. The small efficiency factor near 1 is currently a *target value* for future derivation. The κ-field scale that anchors the geometric-mean coherence band sits *near* the cosmological-constant scale of about 2.3 meV — but the framework's own self-consistency requires it to sit slightly above by a structurally specific factor (about 1.20 in the working numerics), not coincide with it; this is itself a target the framework predicts and a future derivation must reproduce. None of this is presented as settled.

What the paper does claim is that, with the category-error reframing of the companion paper in hand, the residual question "why 4% specifically?" has a candidate structural answer: a product of multiplicity, rate, and competition, each independently meaningful within the VERSF programme.

An appendix records exploratory work on the Higgs mass itself (about 125 GeV, distinct from the vacuum scale of 246 GeV). One tempting reading is rejected on structural grounds. A second — **$m_H / v = 32 / 63 \approx 0.508$** — is now backed by three layered structural arguments using the same closure-normalised dimension $K + 1 = 8$. A **Closure-Deformation Selection Principle**, named explicitly and an instance of a recurring architectural pattern in the VERSF admissibility programme (the same admissibility-equivalence-then-quotient pattern that organises entropy partition uniqueness, record-current transport, and gauge-sector structure, applied here to the deformation-algebra sector), states that corrections to dimensionless ratios count admissible traceless internal deformations rather than raw configurations — this selects the central denominator $63 = 8^2 - 1$. The Principle's load-bearing move is distinguishing two ontologically distinct kinds of dimensionless content of the same architecture: bare couplings (which inherit the *space* dimension $K + 1 = 8$) and corrections (which inherit the *traceless deformation-algebra* dimension $(K+1)^2 - 1 = 63$), with trace removal applying only to the latter. A mean-field closure-fluctuation lemma fixes the leading correction at 1/63; a closure-normalised quartic-democracy lemma fixes the baseline 1/2 via $\lambda_{baseline} = 1/(K + 1) = 1/8$. Each argument is structurally motivated but not yet derived from the framework's master equations, so the appendix records this as a *structurally derived candidate* conditional on (i) rigorous master-action derivation of the two supporting lemmas and (ii) a universality check that is currently *methodologically pre-registered but operationally awaiting candidates* — m_H/v is the only known SM-internal observable of the relevant kind, so the test sample is currently one. The framework neither promotes the result nor disowns it; it sits in the appendix because that is where partially-derived candidates honestly belong.

Readers who want only the picture can stop here. The technical abstract and body of the paper lay out the framework, the equations, and the open problems in detail.

Table of Contents

Abstract

We present an account of electroweak coherence selection within the Void Energy–Regulated Space Framework (VERSF), in which the conventional electroweak hierarchy problem is reinterpreted as a residual coherence-selection problem at the geometric-mean scale $E_{geo} = \sqrt{M_P \cdot m_\kappa}$. The account integrates recent results from the closure, TPB (Ticks-Per-Bit), finite-distinguishability, interface-matching, and fine-structure programmes into a single formal structure.

The central claim is that the conventional hierarchy problem arises from extending record-layer renormalization through a constitutive closure boundary across which the assumptions defining effective field theory cease to possess physical meaning. Within VERSF the Planck scale is interpreted not as a propagating ultraviolet field regime but as a closure threshold associated with finite distinguishability and admissibility structure. The Higgs sector is correspondingly reinterpreted as a stabilized coherence eigenmode of a deeper record-field organization rather than a fundamental ultraviolet scalar.

The electroweak vacuum then corresponds not to a bare scalar parameter but to a stabilized coherent occupancy fraction within a closure–CCC coherence band of natural scale

$$E_{geo} = \sqrt{M_P \cdot m_\kappa},$$

where M_P is the closure threshold and m_κ is the κ-field coherence scale associated with Causal–Coherence Compatibility (CCC). The observed electroweak scale satisfies

$$v = \chi_{TPB} \cdot \sqrt{M_P \cdot m_\kappa},$$

where χ_{TPB} is a TPB coherence projection factor describing the efficiency with which substrate-level distinguishability stabilizes into persistent electroweak coherence. Numerically $\chi_{TPB} \approx 0.0424$, which lies close to 6α, where α is the fine-structure constant.

Using recent VERSF interface-matching results we argue that α is not merely an electromagnetic coupling but a finite-distinguishability normalization coefficient governing interface-level coherence transfer. The TPB projection factor is correspondingly interpreted as a coherence-occupancy efficiency built from binary closure rarity, democratic channel redistribution, finite distinguishability, and second-order closure competition. The structural content of the framework is captured by a single **master occupancy relation**, proposed here as a factorization ansatz:

$\chi_{TPB} = N_{part} \cdot \alpha \cdot \eta_{closure}$,

where $N_{part} = 6$ is the effective channel multiplicity of the closure participation structure (independently derived as 1/IPR), α is the per-channel interface distinguishability-transfer efficiency (structural, low-energy value), and $\eta_{closure} \approx 0.97$ is the second-order closure occupancy efficiency, currently defined as the residual and proposed as a target for independent derivation. The geometric-mean scale $\sqrt{(M_P \cdot m_\kappa)} \approx 5.8$ TeV connects v to a κ-field scale $m_\kappa \approx 1.20 \cdot \Lambda^{(1/4)}$ — *specifically above the cosmological-constant scale by a factor required by the framework's own $\eta_{closure} \leq 1$ self-consistency*, not coincident with it. The wider VERSF programme anchors m_κ structurally; the precise multiplier ≈ 1.20 is a structurally specific target that a future master-action derivation must reproduce (§4.3). Setting $m_\kappa = \Lambda^{(1/4)}$ exactly would force $\eta_{closure} > 1$ and falsify the master relation.

The hierarchy problem is consequently transformed: not into a problem of seventeen-order ultraviolet cancellation, but into a moderate coherence-selection problem at the E_{geo} scale governed by multiplicity, rate, and competition, with a residual $E_{geo}^2 / m_H^2 \approx 2150$ ratio that the present paper does not drive to unity but reinterprets as qualitatively distinct from the conventional hierarchy.

An appendix records two candidate extensions to m_H itself. The first — a half-channel sub-band relation — is rejected on Goldstone-counting grounds. The second proposes $m_H / v = (1/2)(1 + 1/D_{corr}) = 32/63$ in the Mexican-hat curvature avenue, with $D_{corr} = \dim \mathfrak{su}(K + 1) = 63$ selected by a named **Closure-Deformation Selection Principle** (an instance of the wider VERSF programme's admissibility-equivalence-then-quotient architectural family). Two supporting lemmas — *Uniform Closure-Fluctuation Distribution* for the coefficient and *closure-normalised quartic democracy* for the baseline — remain structurally argued rather than master-action-derived; a universality check is methodologically pre-registered but operationally awaits curvature-vs-minimum candidates beyond m_H/v itself (§A.7). Neither extension is promoted to a result of the framework; the appendix is exploratory.

---

1. Introduction

1.1 The conventional hierarchy problem

Scalar masses in quantum field theory receive quadratic ultraviolet corrections

$$\delta m_H^2 \propto \Lambda_{UV}^2.$$

If the Standard Model is taken to remain valid up to the Planck scale $M_P \approx 10^{19}$ GeV, the observed Higgs mass $m_H \approx 125$ GeV appears unnaturally small by roughly seventeen orders of magnitude in mass-squared. Conventional resolutions invoke supersymmetry, compositeness, extra dimensions, asymptotic safety, or anthropic selection. All four classes share an unstated premise: that the same propagating field ontology extends continuously from the electroweak scale to the Planck scale.

1.2 The VERSF reinterpretation

VERSF rejects that premise. Within the framework:

- physical structure emerges from irreversible commitment;
- distinguishability is fundamental rather than derived;
- spacetime geometry is emergent rather than primitive.

VERSF distinguishes two ontologically distinct layers:

Layer	Function	Supports
Closure layer	Governs admissibility, finite distinguishability, coherence constraints	Selection rules, not propagating excitations
Record layer	Supports renormalized field structure and gauge transport	Effective field theory, observable coherent physics

The hierarchy problem becomes the question: *can record-layer renormalization legitimately be extended through a constitutive closure boundary?* The present paper argues that it cannot, and that doing so is the source of the apparent fine-tuning.

1.3 What this paper adds

This paper does four things. First, it consolidates the projection-operator picture of electroweak vacuum formation into a single formal statement, $v = \chi\_TPB \cdot \sqrt{(M\_P \cdot m\_\kappa)}$. Second, it identifies the geometric-mean scale $\sqrt{(M\_P \cdot m\_\kappa)} \approx 5\text{–}6$ TeV as a candidate structurally selected coherence band, conditional on the wider VERSF programme's anchoring of m_κ at a value $\approx 1.20 \cdot \Lambda^{(1/4)}$ — a deviation from the cosmological-constant scale specifically required by the framework's $\eta\_closure \leq 1$ self-consistency rather than a numerical convenience (Section 4.3). Third, it proposes a factorization of the residual projection factor $\chi\_TPB \approx 0.0424$ as the **master occupancy relation** $\chi\_TPB = N\_part \cdot \alpha \cdot \eta\_closure$, with $N\_part = 6$ supplied by the closure participation structure, α by interface distinguishability transfer, and $\eta\_closure \approx 0.97$ as the second-order closure occupancy efficiency — the first two carrying independent structural definitions, the third currently defined as the residual and proposed as a target for derivation. Fourth, it shows that within this factorization the conventional ten-orders-of-magnitude hierarchy problem reduces to a coherence-selection problem at E_geo with a residual ratio $E\_geo^2 / m\_H^2 \approx 2150$ of qualitatively different character (Section 7).

---

2. Closure, Distinguishability, and Interface Matching

2.1 Gauge-invariant structure lives on interfaces

Gauge-invariant electromagnetic observables are holonomies, Wilson loops, and flux integrals. Wilson loops are 1D objects in the loop-variable sense, but via Stokes' theorem they admit equivalent surface-supported representations as flux integrals over 2D surfaces bounded by the loop. The minimal coherence-carrying object in the surface-flux formulation — the formulation VERSF uses to define interface-level distinguishability — is the 2D interface. This is the structure singled out once surface-flux gauge invariance is imposed at the operational level, and it fixes the dimensionality on which the interface-matching machinery of §2.2 acts.

Status of this subsection. The Stokes-side argument here is bookkeeping rather than a derivation of interface minimality; the substantive surface-extensivity argument singling out the 2D interface as the minimal gauge-invariant coherence carrier is developed in the interface-matching paper of the wider VERSF programme. §2.1 records the conclusion in the form used downstream by §2.2 and §2.3.

2.2 Finite distinguishability

The interface-matching framework introduces finite throughput, finite phase resolution, and bit-commitment normalization. The gauge kinetic normalization is then fixed by finite distinguishability structure rather than left as a continuous free parameter:

$$\beta(\Lambda) = \ln 2 / (1 - \cos(2\pi/N_\varphi)),$$

where N_φ is the finite interface phase resolution. This converts a continuous gauge normalization into a discrete interface-selection quantity. The continuum limit $N_\varphi \to \infty$ is recovered smoothly, but the framework asserts that physical interfaces operate at finite N_φ.

Status of this subsection. N_φ and the explicit $\beta(\Lambda)$ form do load-bearing work in the interface-matching paper of the wider programme; they are recorded here for completeness and to establish the structural origin of α used in §2.3, but they are not computed with directly in the body of this paper. The exact relationship between N_φ and the closure structure (K, N_{loop}) is flagged as an open problem in §9.3.

2.3 Physical meaning of α

Within this picture the fine-structure constant ceases to be an empirical coupling and acquires a structural definition. The impedance representation

$$\alpha = Z_0 / (2 R_K),$$

with Z_0 the vacuum impedance and R_K the von Klitzing quantum resistance, is rewritten as a distinguishability-transfer efficiency between discrete quantum transport and the continuum electromagnetic coherence channel. In TPB language, α measures the fraction of bit-level commitment that propagates as coherent electromagnetic structure across an interface.

This reinterpretation is the conceptual bridge to the rest of the paper: if α governs interface-level transfer, then any electroweak-scale quantity built from interface coherence should inherit an α-weighted structure.

3. TPB and Coherence Transfer

3.1 Ticks and bits

The TPB framework separates reversible substrate exploration ("ticks") from irreversible committed distinguishability ("bits"). The substrate-level entropy is

$$S = k_B \cdot \ln(1 / \mathrm{TPB}),$$

so that low TPB corresponds to efficient distinguishability stabilization and high TPB corresponds to inefficient stabilization. Electroweak vacuum formation, as a paradigmatic case of persistent coherent occupancy, must therefore depend on TPB efficiency rather than on any single bare parameter.

3.2 The coherence projection operator

We define a TPB projection operator

$$\Pi_{coh} : \mathscr{H}_{sub} \to \mathscr{H}_{record}$$

mapping substrate distinguishability structure into stabilized record-layer coherence. The coherence-transfer efficiency for a substrate state Ψ is

$$\eta_{TPB}(\Psi) = \langle \Pi_{coh} \Psi \mid \Pi_{coh} \Psi \rangle / \langle \Psi \mid \Psi \rangle.$$

This is the fraction of substrate distinguishability surviving as persistent coherent structure. The electroweak occupancy fraction χ_{TPB} introduced below is the value of η_{TPB} evaluated on the stabilized electroweak coherence mode. The operator Π_{coh} is not a perturbative object: it acts between layers rather than within a single Hilbert space, and its admissibility structure is fixed by closure rather than by Hamiltonian dynamics.

Status of this subsection. The operator-theoretic formulation here is **conceptual rather than computational** in this paper: Π_{coh} is defined to fix the meaning of χ_{TPB} as a coherence-transfer efficiency, but no matrix elements of Π_{coh} are computed in the body, and the master relation of §5.3 does not require explicit construction of the operator. A full matrix-element treatment of Π_{coh} — needed for an independent derivation of $\eta_{closure}$ from substrate microphysics — is part of the open work flagged in §9.3.

4. The Closure–CCC Coherence Band

4.1 The Planck scale as a closure threshold

Within VERSF the Planck scale is the threshold at which admissible distinguishability density saturates. It is not the scale at which a propagating field ontology breaks down — there is no propagating ontology to break down at that scale, because the closure layer does not host propagating excitations. M_P is therefore a structural ceiling, not a UV cutoff in the EFT sense, and corrections of the form $\delta m_H^2 \propto M_P^2$ do not arise: there is no record-layer process that integrates a propagating mode up to M_P.

4.2 The κ-field and the geometric-mean scale

The κ-field governs admissible record transport, coherence persistence, and Causal–Coherence Compatibility. Its scale m_κ is the lower edge of admissible coherence. The natural coherence-band scale formed from the two endpoints is the geometric mean

$$E_{geo} = \sqrt{(M_P \cdot m_\kappa)}.$$

This is dimensionally fixed by the two endpoints, requires no tuning, and is the unique scale-invariant combination of M_P and m_κ under multiplicative scaling — any other combination introduces an additional dimensionful parameter that the framework does not contain.

4.3 The status of m_κ

The value of E_{geo} depends on m_κ, and we are explicit about the current epistemic status of that input. Within this paper m_κ is **not** independently derived; it is anchored to the κ-field scale developed elsewhere in the VERSF programme (most directly, the Two-Planck framework's analysis of the cosmological constant). At present this anchoring is structural rather than fully numerical: a derivation of m_κ from closure microphysics decoupled from electroweak data is open.

The κ↔Λ identification is approximate, not exact. Using the cosmological-constant energy scale $\Lambda^{(1/4)} \approx 2.3$ meV (Planck Collaboration 2020) as a direct anchor — i.e. setting $m_\kappa = \Lambda^{(1/4)}$ exactly — gives

$$E_{geo}|_{m_\kappa = \Lambda^{(1/4)}} = \sqrt{(M_P \cdot \Lambda^{(1/4)})} \approx \sqrt{(1.22 \times 10^{19} \times 2.3 \times 10^{-12})}\ \text{GeV} \approx \mathbf{5.30\ TeV},$$

and then $\chi_{TPB} = v/E_{geo} \approx 246/5300 \approx$ **0.0464**, against $6\alpha \approx 0.0438$, so

$$\eta_{closure}|_{m_\kappa = \Lambda^{(1/4)}} \approx 0.0464 / 0.0438 \approx \mathbf{1.06}.$$

That is, $\eta_{closure}$ *over unity*. The §5.3 master relation interprets $\eta_{closure}$ as a second-order closure occupancy *efficiency* — a fractional loss to closure competition — and the framework therefore requires $\eta_{closure} \leq 1$ on structural grounds: a residual cannot enhance the channel-summed transfer beyond $N_{part} \cdot \alpha$ (§5.4 argues this from the normalized-occupancy-fraction interpretation; the dynamical-microphysics derivation remains open, §9.3). Identifying m_κ with $\Lambda^{(1/4)}$ exactly *violates* that constraint by about 6%.

Two distinct commitments must be separated. Pulled apart cleanly:

1. **Hard structural bound** (from the $\eta\_closure \leq 1$ inequality alone): $E\_geo \geq v / (6\alpha) \approx 5.62$ TeV $\Longrightarrow m\_\kappa \geq (5.62 \text{ TeV})^2 / M\_P \approx 2.59$ meV $\approx$ **1.13 · Λ^(1/4)**. This is the *bound* the framework commits to as a structural inequality. It is not a prediction of a specific multiplier, only of "at least this much above Λ^(1/4)."
2. **Working empirical anchor** (from reading $\eta\_closure \approx 0.97$ as a tight target inside that bound): $m\_\kappa \approx 2.76$ meV $\approx$ **1.20 · Λ^(1/4)**, yielding $E\_geo \approx 5.80$ TeV. The ≈ 1.20 multiplier is what falls out when the §5.3 soft target $\eta\_closure \approx 0.97$ is taken at face value as the working value. It is not itself a structural prediction.

These two commitments are different in kind: (1) is hard and follows from the master relation's structure (granted $\eta\_closure \leq 1$, which is itself an open derivation problem — §9.3); (2) is soft and follows from reading the target value as tight. The framework's hard prediction is $m\_\kappa \geq 1.13 \cdot \Lambda$^(1/4); the working anchor $m\_\kappa \approx 1.20 \cdot \Lambda$^(1/4) is calibrated by $\eta\_closure \approx 0.97$ and would shift if the empirical target changed.

This is a load-bearing honesty point. The framework cannot treat 3% as substantive in η_closure (where it serves as the falsifiable target for derivation; §5.3) while treating 20% as "slightly" in the κ↔Λ relation. We commit explicitly: **m_κ is not Λ^(1/4)** — the $\eta\_closure \leq 1$ bound forces m_κ above Λ^(1/4) by *at least* ~13%. The specific working multiplier ≈ 1.20 is not itself the structural commitment; it is a calibration to the empirical η_closure target.

Working values used in this paper. Where numerical values are needed:

Quantity	Value	Source / Status
E_geo	5.80 TeV	Working anchor, consistent with $\eta\_closure \approx 0.97$
m_κ	2.76 meV	Derived from E_geo via $m\_\kappa = E\_geo^2 / M\_P$
Λ^(1/4)	2.30 meV	Planck 2018 (independent cosmological measurement)
m_κ / Λ^(1/4) ≥ 1.13	hard	From $\eta\_closure \leq 1$ structural inequality
m_κ / Λ^(1/4) ≈ 1.20	soft	From $\eta\_closure \approx 0.97$ empirical anchor

Numerical conclusions remain sensitive to E_geo (and therefore to m_κ) at the few-percent level. The *bound* is what the framework commits to structurally; the specific *anchor* is what calibrates the working numerics.

The retraction tightens the bound, not the anchor. Acknowledging $m\_\kappa \neq \Lambda$^(1/4) might at first reading look like a weakening — exchanging a clean identification for a deviation. The opposite is true *for the bound*. The framework now commits to $m\_\kappa \geq 1.13 \cdot \Lambda$^(1/4) as a hard inequality (granted $\eta\_closure \leq 1$), which several distinct outcomes for a future closure-microphysics derivation would falsify:

- A derivation yielding $m\_\kappa = \Lambda$^(1/4) exactly forces $\eta\_closure > 1$, falsifying the bound and the master-relation framing of §5.3.
- A derivation yielding $m\_\kappa \approx 1.5 \cdot \Lambda$^(1/4) is *inside* the bound but shifts the empirical anchor: it forces $\eta\_closure \approx 0.87$, requiring §5.3's "~3% deviation from unity" framing

to be revised to ~13% as a structurally distinct claim about second-order closure competition.
- A derivation yielding m_κ much larger than $\sim 2 \cdot \Lambda^{1/4}$ is also inside the bound but pushes $\eta_{closure}$ into a regime where the second-order closure interpretation loses physical content.

Only the first scenario falsifies the framework's *hard* commitment; the second and third would require revision of §5.3's empirical anchor but not of the bound itself. Distinguishing these is the point of separating the two commitments. Honesty about the deviation has converted a soft anchoring into one hard inequality plus one soft target, both flagged with their respective epistemic status.

---

5. Electroweak Occupancy as TPB Projection

5.1 The projection relation

We propose that the electroweak vacuum scale is the TPB-projected fraction of the coherence band:

$$v = \chi_{TPB} \cdot \sqrt{M_P \cdot m_\kappa}.$$

Solving numerically with $v \approx 246$ GeV and $E_{geo} \approx 5.8$ TeV gives

$$\chi_{TPB} \approx 0.0424.$$

Approximately 4.2% of the available coherence-band capacity stabilizes into persistent electroweak vacuum structure. This is the structural content of the hierarchy problem within VERSF: not a seventeen-order cancellation, but a moderate occupancy-selection problem with a residual projection factor of order a few percent.

5.2 The relation $\chi_{TPB} \approx 6\alpha$

Numerically 0.0424 lies close to

$$6\alpha \approx 6 \times 0.00730 \approx 0.0438.$$

The agreement is at the ~3% level. This is **not** presented as a derivation of v from α; it is presented as a structurally suggestive coincidence whose components have independent VERSF interpretations:

- α is the interface-level distinguishability-transfer efficiency (Section 2.3);
- the factor 6 is the inverse participation ratio of the local closure participation structure (Section 6).

The proposed reading is therefore that χ_{TPB} is a closure-weighted interface-transfer efficiency, with α supplying the per-channel transfer rate and the closure participation structure supplying the multiplicity factor. The factorized form of this reading is the master relation introduced in Section 5.3.

Choice of α. Throughout this paper α denotes the low-energy fine-structure constant $\alpha \approx 1/137.036$, not a running value at the electroweak scale. This is consistent with the interpretation in §2.3 of α as a *structural* finite-distinguishability normalization coefficient — fixed by the impedance ratio $\alpha = Z_0 / 2R_K$ and the interface phase resolution N_φ — rather than as an RG-running coupling evaluated at a renormalization scale. Using $\alpha(M_Z) \approx 1/127.9$ would give $6\alpha(M_Z) \approx 0.047$ and $\eta_{closure} \approx 0.90$, a numerically different relation; the structural reading commits to the low-energy value as the one that enters the master relation.

The distinction is ontological as well as numerical. In the present framework, α enters *prior to* record-layer renormalization, as an interface-level distinguishability-normalization coefficient. The quantity appearing in the master occupancy relation is therefore not the effective running coupling of electroweak perturbation theory evaluated *inside* the record layer, but the underlying closure/interface normalization inherited by all lower-energy realizations.

Renormalization-group flow describes how propagating record-layer excitations reorganize effective interaction strengths as a function of probe scale — it is a statement about how scattering amplitudes change when measured at different energies *within* the record layer. The master occupancy relation, by contrast, concerns the substrate-to-record transfer efficiency itself: the normalization of distinguishability transport across the closure/interface boundary. These are two different physical quantities. They coincide numerically at low energy not by accident but by construction: the closure/interface normalization is the *IR boundary value* that the running coupling approaches as the probe energy goes to zero. The running coupling realizes, in the IR limit of its scale-dependence, the structural value fixed by the interface normalization; the two quantities diverge at higher energies because the running coupling reorganizes within the record layer while the structural quantity, defined at the interface boundary, does not run.

The structural quantity entering the master relation is therefore the low-energy normalization coefficient associated with the impedance form

$$\alpha = Z_0 / (2 R_K),$$

not the scale-dependent electroweak running value $\alpha(M_Z)$. Using $\alpha(M_Z)$ in the master relation would be a category error: it would treat the relation as if it were an intra-record-layer perturbative statement, when the framework's whole claim is precisely that it is not. The master relation is a constraint on the closure/interface layer, not on record-layer running content; the appropriate coupling is the one defined at that layer, which is the low-energy impedance-fixed value.

5.3 The master occupancy relation

The verbal structure "multiplicity × rate × competition" admits a compact statement. We propose — **as a structural ansatz to be verified by independent derivation of η_closure** — that χ_TPB factorizes as the **master occupancy relation**

$$\boldsymbol{\chi\_TPB = N\_part \cdot \alpha \cdot \eta\_closure},$$

where the three factors have the following epistemic status:

Factor	Value	Status
N_part = 6	Effective channel multiplicity of the local closure participation structure	**Independently derived** as 1/IPR from the closure participation weights (Section 6.1)
$\alpha \approx$ 1/137.036	Interface-level distinguishability-transfer efficiency per channel	**Independent structural definition** via $\alpha = Z_0 / 2R\_K$ (Section 2.3)
η_closure ≈ 0.97	Second-order closure occupancy efficiency	**Currently defined as the residual** χ_TPB / (N_part · α); derivation open (Section 9.3)

The honest current state. Until η_closure has an independent derivation from closure microphysics, the master relation reduces to $\chi\_TPB = (6\alpha) \cdot (\chi\_TPB / 6\alpha)$, which is trivially true. The nontrivial content of §5.3 is therefore *not* the equation itself but two distinct claims:

1. **A factorization claim.** That χ_TPB has the multiplicative structure (multiplicity × rate × competition), with the first two factors carrying independent structural definitions.
2. **A target.** That $\eta\_closure \approx 0.97$, derived here as a residual, is the value an independent derivation from second-order closure microphysics must reproduce.

The factorization claim is what the framework asserts; $\eta\_closure \approx 0.97$ is its current empirical target, not a derived quantity. We are explicit about this because the relation is too easy to misread as established otherwise.

Numerics. With $\chi\_TPB \approx 0.0424$ and $6\alpha \approx 0.04378$,

$$\eta\_closure \approx 0.0424 / 0.04378 \approx 0.969.$$

Approximately three percent of the channel-summed coherence capacity is lost to second-order closure competition. This is the falsifiable prediction: a future derivation of η_closure from closure microphysics must reproduce this ~3% deviation from unity, or the factorization claim fails.

Combined relation. Substituting into $v = \chi\_TPB \cdot \sqrt{(M\_P \cdot m\_\kappa)}$ gives the electroweak vacuum scale in fully factored form:

$$\boldsymbol{v = N\_part \cdot \alpha \cdot \eta\_closure \cdot \sqrt{(M\_P \cdot m\_\kappa)}}.$$

Read right to left: a geometric-mean coherence band (conditional on m_κ, Section 4.3), weighted by interface transfer efficiency, multiplied by closure participation multiplicity, modulated by second-order closure competition. That is what the framework asserts the electroweak vacuum is.

5.4 Why $\eta_{closure} \leq 1$

The quantity $\eta_{closure}$ is not introduced as an arbitrary correction coefficient but as an occupancy-efficiency factor multiplying the channel-summed interface-transfer structure:

$$\chi_{TPB} = N_{part} \cdot \alpha \cdot \eta_{closure}.$$

Its interpretation is therefore constrained by the same admissibility logic governing all normalized occupancy measures in the wider VERSF programme.

The quantity $N_{part} \cdot \alpha$ represents the *maximal* channel-summed interface-transfer capacity permitted by the closure participation structure and distinguishability-transfer normalization. The factor $\eta_{closure}$ measures the fraction of that capacity surviving second-order closure competition. If $\eta_{closure} > 1$, the residual sector would contribute *positive* occupancy beyond the channel-normalized transfer bound itself — that is, closure competition would *generate* net distinguishability-transfer capacity rather than redistributing or suppressing it.

Within the VERSF admissibility programme, normalized occupancy factors are structurally bounded: redistribution among admissible channels may preserve or reduce effective occupancy, but it cannot create occupancy beyond the normalized support on which the quantity is defined. The admissible range is therefore

$$0 \leq \eta_{closure} \leq 1.$$

What this changes about the upper bound. In earlier sections the inequality $\eta_{closure} \leq 1$ was flagged as "asserted, requiring derivation" (§4.3, §9.3). §5.4 sharpens the status: the upper bound is not yet derived from a microscopic closure-action calculation, but it is not an independent assumption either — it follows from the *interpretation* of $\eta_{closure}$ as a normalized occupancy fraction rather than as an independent enhancement channel. A reader unwilling to grant this interpretation must reject the master-relation factorization entirely; one who grants it must accept $\eta_{closure} \leq 1$.

The open derivation gap therefore shifts: it is not "derive $\eta_{closure} \leq 1$ from closure microphysics" but rather "derive from closure microphysics that the residual factor in χ_{TPB} / ($N_{part} \cdot \alpha$) deserves the normalized-occupancy-fraction interpretation in the first place." A future closure-microphysics derivation must reproduce this inequality dynamically — by showing that the residual sector cannot contribute positive occupancy — or the master-relation interpretation requires revision.

This sharpens the §4.3 falsification chain: the hard bound $m_\kappa \geq 1.13 \cdot \Lambda^{(1/4)}$ is now conditional on a structural interpretation (occupancy-fraction reading of $\eta_{closure}$) rather than on a bare assertion. Conditional structural commitments are stronger than bare assertions but weaker than first-principles derivations; the upper bound now sits at exactly that intermediate epistemic level.

---

6. Second-Order Closure Competition

6.1 Local participation structure

Recent second-order closure work in the programme derived six local closure constraints, one global closure constraint, and equal participation weights

$$w_i = 1/6, \quad i = 1, \ldots, 6,$$

yielding the inverse participation ratio

$$\mathrm{IPR} = \Sigma_i w_i^2 = 6 \cdot (1/6)^2 = 1/6,$$

so that 1/IPR = 6.

What is and is not being claimed here. The IPR identity 1/IPR = N for equal weights $w_i = 1/N$ is trivial bookkeeping: it just recovers the number of constraints one started with. The IPR machinery only earns nontrivial work when the weights are unequal and the effective participation count is strictly less than the raw count. The substantive structural claim is therefore **not** the IPR calculation itself but the antecedent assertion that the six local closure constraints carry *equal* weights — closure democracy. The IPR identity then confirms that under democracy the effective channel multiplicity is the raw constraint count, giving $N_{part} = 6$.

Closure democracy is asserted on structural grounds in the wider VERSF programme (no closure constraint is privileged over any other once the participation structure has been fixed); a derivation of closure democracy itself from deeper admissibility microphysics is open. The identification $N_{part} = 6$ in the master relation $\chi_{TPB} = N_{part} \cdot \alpha \cdot \eta_{closure}$ of §5.3 inherits its status from closure democracy: it is as well-established as democracy is, and no better.

6.2 Physical interpretation

The electroweak vacuum is then read as a globally stabilized coherent occupancy mode surviving three filters in sequence: binary closure filtering (admitting only structurally rare commitment configurations), democratic redistribution across six local channels, and second-order closure competition between channels for the global mode. The factor of six does not arise as a normalization choice — it is fixed by the closure participation structure derived independently elsewhere in the programme.

In this reading the master relation $\chi_{TPB} = N_{part} \cdot \alpha \cdot \eta_{closure}$ of Section 5.3 acquires its physical content: $N_{part} = 6$ is supplied by the participation structure derived here; α is the per-channel interface transfer efficiency; and $\eta_{closure}$ is the residual efficiency of second-order closure competition. The electroweak hierarchy reduces to the question of why these three structures combine to produce a percent-level occupancy fraction rather than an order-unity one. The framework's answer is structural: N_{part} is bounded above by the closure participation count, α is bounded above by interface distinguishability transfer, and $\eta_{closure}$ is bounded above by unity — and the product cannot saturate the coherence band without violating closure admissibility.

---

7. The Higgs as a Coherence Eigenmode

Within this framework the Higgs is not a fundamental ultraviolet scalar. It is a coherence eigenmode of the stabilized record field, where the record field is the field-theoretic manifestation of the deeper scalar-entropic commitment structure developed throughout the wider VERSF programme.

The various descriptive languages used elsewhere in the programme — entropy-field language, commitment-density language, record-field language, coherence-mode language — describe the same underlying ontology at different levels of description. The electroweak vacuum is therefore a stabilized coherent occupancy state, not a bare ultraviolet mass insertion.

A consequence is that the standard quadratic-divergence diagrams contributing to δm_H^2 are reinterpreted as record-layer processes whose UV behaviour is bounded by the coherence-band edge E_{geo}, not by M_P. The diagrams retain their record-layer meaning; what they do not retain is the right to integrate through the closure boundary.

The residual tuning. This move reduces but does not eliminate the tuning. The conventional ratio $M_P^2 / m_H^2 \approx 10^{34}$ is replaced by

$$E_{geo}^2 / m_H^2 \approx (5800 / 125)^2 \approx 2150.$$

A factor of ~30 dB rather than ~340 dB — qualitatively different, but not order-unity. We confront this directly: the framework as developed here does *not* drive the residual to one. There are three available structural responses, and we flag them rather than choose between them:

(i) the residual is itself a coherence-selection effect with its own factorization, plausibly involving the same N_{part} and α structure applied at a sub-band level (an exploratory observation in this direction is recorded in Appendix A, together with the structural objection — Goldstone counting in the SM Higgs sector — that rules out the simplest "half-channel" version of this reading, and the identification of a more physical curvature-based avenue);

(ii) the residual is absorbed by a record-layer feature not yet identified in the programme — for example, the relationship between v and m_H mediated by the quartic coupling λ (Section 9.3);

(iii) the framework accepts a modest residual tuning at E_{geo} as the irreducible content of the hierarchy problem, with the claim being that this residual is qualitatively unlike the conventional one, not absent.

The current paper neither chooses among these nor resolves the residual; it claims only that the catastrophic ten-orders-of-magnitude version of the problem is a category error, and that what remains is a coherence-selection question at the E_{geo} scale of a fundamentally different character.

---

8. Why Naturalness Appeared Compelling

Effective field theory works extremely well within a fixed dynamical layer, and the conventional naturalness intuition was a correct extrapolation *given the assumption that no constitutive boundary intervenes between the electroweak and Planck scales*. VERSF rejects only that assumption, not EFT itself: within the record layer, RG flow, vacuum polarization, and loop corrections retain their usual content. What changes is the right to integrate through the closure boundary. The seventeen-orders-of-magnitude version of the hierarchy problem was therefore not a calculation error — it was the correct answer to the wrong question.

---

9. What Is and Is Not Claimed

9.1 Structural claims

The following are structural claims of the framework, asserted but not derived in this paper:

- closure–CCC coherence bands exist as ontologically distinct from propagating record-layer structure;
- electroweak structure corresponds to coherent occupancy on such a band;
- α functions as a distinguishability-transfer efficiency at the interface level;
- TPB projection governs stabilized occupancy fractions;
- the projection factor factorizes multiplicatively in the form $\chi_{TPB} = N_{part} \cdot \alpha \cdot \eta_{closure}$, with N_{part} and α carrying independent structural definitions and $\eta_{closure}$ currently defined as the residual.

9.2 Derived content (from prior VERSF papers)

The following are derived elsewhere in the programme and used here:

- finite-distinguishability interface matching, including $\beta(\Lambda) = \ln 2 / (1 - \cos(2\pi/N_\varphi))$;
- binary closure rarity;
- $K = 7$ closure structure and $N_{loop} = 14$ channel structure;

- second-order closure competition with IPR = 1/6 (and therefore N_part = 1/IPR = 6);
- interface normalization formulas.

9.3 Open problems

The following are explicitly open:

- **Independent derivation of m_κ.** Currently anchored to the Two-Planck framework's κ-scale. The $\eta_{closure} \leq 1$ constraint of §5.3 requires $m_\kappa \approx 1.20 \cdot \Lambda^{1/4}$ — not the cosmological-constant scale $\Lambda^{1/4} \approx 2.3$ meV itself, but a structurally specific multiplier above it. This multiplier is itself a target that a future closure-microphysics derivation must reproduce; producing $m_\kappa = \Lambda^{1/4}$ exactly would falsify the master relation by forcing $\eta_{closure} > 1$ (§4.3).
- **Derivation of η_closure ≈ 0.97** from second-order closure microphysics, including whether the ~3% deviation from unity matches the correction predicted by the participation structure. Until this is supplied, the master relation $\chi_{TPB} = N_{part} \cdot \alpha \cdot \eta_{closure}$ is a factorization ansatz with the third factor defined as the residual.
- **Derivation of η_closure ≤ 1 as a structural inequality.** The §4.3 falsification chain relies on $\eta_{closure} \leq 1$. §5.4 argues for this bound from the *normalized-occupancy-fraction interpretation* of η_closure: a residual fraction of capacity cannot exceed unity by definition, since "residual" presupposes the full capacity is the upper limit. This shifts the open derivation gap from "derive $\eta_{closure} \leq 1$ from microphysics" to "derive from closure microphysics that the residual factor $\chi_{TPB} / (N_{part} \cdot \alpha)$ deserves the normalized-occupancy-fraction interpretation." The latter is a sharper, more localized derivation target. A reading in which η_closure represents net coherence with possible constructive-interference enhancement would not impose ≤ 1, and would also reject §5.4's interpretive premise.
- **Derivation of closure democracy.** §6.1 makes N_part = 6 conditional on the equal-weight assignment $w_i = 1/6$ over the six local closure constraints. Closure democracy is asserted on structural grounds in the wider programme; a derivation from admissibility microphysics is open. This is the keystone of the master relation: the entire factor N_part = 6 inherits its status from closure democracy and is no better-established than democracy is.
- **The physical Higgs mass m_H = 125 GeV.** The framework as developed here predicts the vacuum scale v but not the quartic coupling λ that sets $m_H^2 = 2\lambda v^2$; either an independent prediction of λ or a direct coherence-mode interpretation of m_H within the record-field organization is required. An exploratory observation on $m_H/v \approx 1/2$, together with structural objections to a "half-channel" reading and identification of the Mexican-hat curvature route as the more physical avenue, is recorded in Appendix A.
- **The residual $E_{geo}^2 / m_H^2 \approx 2150$ tuning** (Section 7): whether absorbed by a sub-band coherence-selection factorization, by the v↔m_H structure, or accepted as the irreducible residual.
- **Exact relationship between N_φ and (K, N_loop)** — i.e., the interface phase resolution and the closure/channel structure.
- **The exact mechanism selecting the electroweak sector** specifically over other possible coherence modes.

- **The structural scale of α.** This paper uses the low-energy value $\alpha \approx 1/137.036$ on the interpretation that α is a structural normalization rather than a running coupling; a derivation showing why the structural value coincides with the low-energy limit of the running coupling is required.
- **Experimental tests** distinguishing $E_{geo} \approx$ 5–6 TeV as a coherence-band edge from a conventional new-physics threshold.

9.4 A note on the framework's epistemic architecture

This paper labels its commitments at three distinct epistemic levels, and reads coherently only if the levels are kept distinct.

1. *Derived from structure* — first-principles or near-first-principles results from the wider VERSF programme: the K = 7 closure number, the (K + 1)-dimensional closure-normalised space, $N_{part} = 6$ from the IPR (conditional on closure democracy), the impedance form $\alpha = Z_0 / 2R_K$.
2. *Conditional structural* — forced by interpretive or structural commitments that the framework has made explicit, but not derived from first principles. This level includes the master-relation factorization $\chi_{TPB} = N_{part} \cdot \alpha \cdot \eta_{closure}$ (§5.3); the $\eta_{closure} \leq 1$ bound (forced by the occupancy-fraction interpretation, §5.4); the $m_\kappa \geq 1.13 \cdot \Lambda^{1/4}$ hard bound (forced jointly by $\eta_{closure} \leq 1$ and the master relation, §4.3); the Closure-Deformation Selection Principle and Lemmas A and B in Appendix A; the pre-registered curvature-vs-minimum universality form (§A.7).
3. *Working empirical anchors* — target values calibrated to data and flagged as such: $\eta_{closure} \approx 0.97$; $m_\kappa \approx 1.20 \cdot \Lambda^{1/4}$; the §A.5 candidate $m_H/v = 32/63$.

The bulk of the framework's nontrivial content sits at the conditional-structural middle level: stronger than bare assertions because forced by named interpretive commitments, weaker than first-principles derivations because those commitments are themselves open derivation targets (§9.3). Skeptical readers should evaluate the framework by examining whether the interpretive commitments are themselves defensible — whether $\eta_{closure}$ deserves the normalized-occupancy-fraction reading, whether the deformation algebra of a (K + 1)-dimensional closure-normalised space really is $\mathfrak{su}(K + 1)$, whether the IR boundary value of α really is the structural quantity entering the master relation — not by treating conditional structural claims as if they were bare assertions or fully derived theorems. The framework's argument is that those interpretive commitments are themselves principled rather than arbitrary; the open work is converting that principled status into rigorous master-action derivation.

This is, incidentally, the epistemic level at which most working theoretical physics actually operates. First-principles derivations are rarer than commonly claimed; commitments forced by named interpretive choices, with the choices themselves defended on structural grounds, are the modal case. The present paper does not apologize for occupying that level — it labels it explicitly so the reader can locate each claim correctly.

---

10. Conclusion

The central relation of this paper is the master occupancy relation

$$\chi_{TPB} = N_{part} \cdot \alpha \cdot \eta_{closure},$$

proposed as a structural factorization ansatz, with $N_{part} = 6$ independently derived from the closure participation structure, α the per-channel interface distinguishability-transfer efficiency (low-energy value $\alpha \approx 1/137.036$), and $\eta_{closure} \approx 0.97$ currently defined as the residual second-order closure occupancy efficiency. Combined with the geometric-mean coherence band $\sqrt{(M_P \cdot m_\kappa)}$, this gives the electroweak vacuum scale in fully factored form:

$$v = N_{part} \cdot \alpha \cdot \eta_{closure} \cdot \sqrt{(M_P \cdot m_\kappa)}.$$

Two factors carry independent structural definitions; the third is a target for derivation. The value of m_κ is currently anchored to the Two-Planck framework's κ-scale; granted the structural assumption $\eta_{closure} \leq 1$ (itself an open derivation, §9.3), the master relation forces the hard bound $m_\kappa \geq 1.13 \cdot \Lambda^{(1/4)}$, and reading the §5.3 soft target $\eta_{closure} \approx 0.97$ as tight gives the working empirical anchor $m_\kappa \approx 1.20 \cdot \Lambda^{(1/4)}$ (§4.3 separates the two). The framework therefore makes three commitments — a factorization, an empirical target $\eta_{closure} \approx 0.97$, and a *conditional* hard inequality $m_\kappa \geq 1.13 \cdot \Lambda^{(1/4)}$ (conditional on $\eta_{closure} \leq 1$ being derivable rather than asserted) — each of which is independently falsifiable.

On this reading the electroweak vacuum is not a finely tuned ultraviolet scalar but a stabilized coherent occupancy fraction surviving closure filtering and TPB projection. The conventional hierarchy problem $M_P^2 / m_H^2 \approx 10^{34}$ is correspondingly transformed: not into a problem of seventeen-order ultraviolet cancellation, but into a moderate coherence-selection problem at the E_{geo} scale governed by multiplicity, rate, and competition. A residual ratio $E_{geo}^2 / m_H^2 \approx 2150$ remains and is not driven to unity by the present analysis (§7); the framework's claim is qualitative reorganization, not the elimination of all tuning.

The framework predicts that $E_{geo} \approx$ 5–6 TeV is structurally selected rather than accidental (conditional on the κ-field anchoring of the wider VERSF programme, which the $\eta_{closure} \leq 1$ constraint requires to deviate from $\Lambda^{(1/4)}$ by a structurally specific factor ≈ 1.20 in the working numerics — see §4.3); that no propagating new-physics threshold sits between v and E_{geo} to absorb the quadratic-divergence problem in the conventional sense; and that $\eta_{closure} \approx 0.97$ should match a value computable from second-order closure microphysics, with the ~3% deviation from unity as the empirical target.

The hierarchy problem, on this reading, was never a tuning problem of the conventional kind. It was a category error about which layer the Planck scale belongs to.

Appendix A. Exploratory observation: $m_H/v \approx 1/2$ and the question of sub-band coherence selection

Status: exploratory. The content of this appendix records a numerical observation and the structural objections to its premature promotion. It is not to be cited as a result of the framework.

A.1 The single empirical observation

The main paper predicts the electroweak vacuum scale $v \approx 246$ GeV but not the physical Higgs mass $m_H \approx 125$ GeV (§9.3). The empirical fact connecting them is

$$m_H / v \approx 0.508.$$

This is a measured Standard Model feature; the framework currently has nothing to say about why it takes this value.

A previous version of this exploratory material presented three separate numerical agreements — $m_H/E_{geo} \approx 3\alpha$ (1.6%), $m_H/v \approx 1/2$ (0.6%), and $\lambda_{predicted} \approx \lambda_{observed}$ (1.5%) — as independent corroborating evidence for a "half-channel" sub-band factorization. **Conditional on the main relation $v = 6\alpha \cdot \eta_{closure} \cdot E_{geo}$ holding, they are not independent.** The identity

$$m_H / E_{geo} = (m_H / v) \cdot (v / E_{geo}) = (m_H / v) \cdot 6\alpha \cdot \eta_{closure}$$

is exact under that assumption, and $\lambda = (m_H / v)^2 / 2$ is SM algebra. The three agreements are therefore algebraically locked, given the main relation, to a single empirical fact:

$$\mathbf{m_H / v \approx 0.508.}$$

The conditionality matters: the reduction uses the main relation to rule out independent informational content in the m_H/E_{geo} and λ agreements. If the main relation is itself off at the few-percent level — because m_κ is mis-anchored (§4.3) or because $\eta_{closure}$ departs from 0.969 — then m_H/E_{geo} and m_H/v would recover some independent measurement weight. Within the precision the framework currently supports, the conditional collapse holds, and the empirical content of this appendix is one coincidence, not three.

A.2 The candidate sub-band relation

Under the main relation, $m_H / v \approx 1/2$ implies

$$m_H \approx 3\alpha \cdot \eta_{closure} \cdot E_{geo},$$

formally identical to the main relation $v = 6\alpha \cdot \eta_{closure} \cdot E_{geo}$ with the channel multiplicity halved: 3 instead of 6. The candidate sub-band relation is therefore

$m\_H = N\_part\_sub \cdot \alpha \cdot \eta\_closure \cdot E\_geo, \quad N\_part\_sub = 3,$

using the *same* closure residual η_closure ≈ 0.97 as the main relation rather than a distinct η'_closure. (A previous version of this material introduced η'_closure ≈ 0.98 alongside η_closure ≈ 0.97; within the m_κ uncertainty of §4.3 these values are indistinguishable, and treating them as two physically distinct residuals attached to two distinct partitions reads as overinterpretation. The data is consistent with a single closure residual doing both jobs, and this appendix proceeds on that assumption.)

If this sub-band relation held, v and m_H would both be coherence-occupancy fractions of the same E_geo, with the same α and the same closure residual, differing only in their effective channel multiplicity.

A.3 Why simple channel counting fails

The half-channel reading does not survive contact with SM Higgs structure. The SM Higgs is a complex doublet — four real degrees of freedom — that after electroweak symmetry breaking decomposes as **three Goldstone modes** (eaten by W±, Z longitudinals) plus **one physical Higgs h**. That is 3 + 1 inside 4, not 3 + 3 inside 6. The "6" of the main relation refers to closure participation indices, not to Higgs field components, so a partition of the six participation indices into two 3-subsets — one associated with m_H, one with the rest — does not naturally correspond to the actual structural split inside the Higgs sector.

Promoting the candidate sub-band relation therefore requires not only deriving "3" from closure microphysics but deriving that the partitioned 3-subset is the **right 3** — specifically, that it tracks the physical Higgs mode and not the eaten Goldstones. A closure-derived 3-subset that happens to coincide with the three Goldstone modes would be a wrong prediction (eaten modes have no physical mass), not a successful derivation. The partition has to be the right partition.

A.4 The Mexican-hat avenue

A more physical route to m_H exists but it is not a half-channel story. In the SM,

$m\_H = \sqrt{(2\lambda)} \cdot v,$

so m_H expresses the curvature of the Higgs potential at its minimum, while v is the radial location of the minimum. The factor of m_H / v = √(2λ) ≈ 0.51 (for the observed λ ≈ 0.13) reflects the **geometry of the potential** — curvature versus minimum-location — rather than a multiplicity ratio of any kind.

A VERSF-native derivation of m_H along this avenue would have the form: *the Higgs mass is the curvature of a closure-stabilized effective potential at its closure-stabilized minimum, with both the location of the minimum and the curvature at that point fixed by the same closure microphysics.* This would predict m_H / v as a specific functional relation between curvature and minimum-location under closure constraint, and would have nothing to do with halving a channel count. It would also distinguish the radial mode (physical Higgs) from the azimuthal modes (Goldstones) naturally, addressing the §A.3 objection.

This is a distinct and more demanding programme than the sub-band master relation. §A.5 records a specific candidate within it.

A.5 A candidate λ-baseline proposal within the Mexican-hat avenue

A specific candidate form for λ within the §A.4 Mexican-hat avenue has been proposed. The candidate is

$m_H / v = (1/2) \cdot (1 + 1/(K^2 + 2K))$,

with K = 7 the VERSF closure number (Fano plane PG(2,2) / minimal-fact architecture, locked independently by the K = 7 no-go theorem). Substituting K = 7 gives $K^2 + 2K = 49 + 14 = 63$, so

$m_H / v = (1/2) \cdot (64/63) =$ **$32 / 63 \approx 0.50794$**.

The observed value $m_H / v \approx 0.50849$ (using PDG 2024 $m_H = 125.20 \pm 0.11$ GeV and $v \approx 246.22$ GeV) agrees with the prediction at the ~0.11% level — within roughly 1.2σ of the Higgs-mass measurement uncertainty. This is an order of magnitude tighter than the half-channel reading of §A.2 (~0.6%).

Translation to λ. In SM language $m_H = \sqrt{2\lambda} \cdot v$, so the prediction is equivalent to

$\sqrt{2\lambda} = (1/2) \cdot (1 + 1/63) \Longrightarrow \lambda \approx 0.1290$,

against the observed $\lambda \approx 0.129$.

Why this is not the same as the §A.2 half-channel reading. This candidate is a Mexican-hat-avenue proposal: it factorises m_H/v as curvature-at-minimum rather than as a halved channel multiplicity. It does not invoke any 3+3 partition of the six closure participation indices and does not therefore face the Goldstone counting objection of §A.3. It is a genuinely distinct candidate, not a relabelling of the rejected one.

A second structural advantage: the §A.5 relation is **pure SM-internal**. It involves only m_H, v, and the closure number K = 7 — and *does not depend* on E_{geo}, on m_κ, on the κ↔Λ identification of §4.3, or on $\eta_{closure}$. Where §A.2 inherits all the conditional anchoring of the main relation $v = 6\alpha \cdot \eta_{closure} \cdot E_{geo}$, §A.5 is decoupled from that anchoring entirely. The §A.1 conditional-collapse argument (that the three apparent agreements between m_H/E_{geo}, m_H/v, and λ are algebraically locked given the main relation) does not apply to §A.5, because §A.5 makes no use of the main relation. This is a feature, not a bug: §A.5 either stands or falls on its own structural derivation, independent of the m_κ-anchoring status of §4.3.

Two structural ingredients require independent derivation.

(i) The 1/2 baseline as closure-normalised quartic. The SM relation $m_H = \sqrt{(2\lambda)} \cdot v$ fixes the baseline $m_H/v = 1/2$ to a specific bare quartic coupling

$$\lambda_{baseline} = 1/8 = 1 / (K + 1) \quad (\text{for } K = 7).$$

The structural candidate identification — *not yet a derivation* — is that this is the natural value of the Higgs quartic under democratic dilution across the K + 1 closure-normalised electroweak dimensions: a single overall coupling budget distributed evenly over the (K + 1)-dimensional closure-normalised space. Lemma B of §A.5(iii) supplies the proposed closure-normalised quartic argument; until that argument is master-action-derived, the baseline $\lambda_{baseline} = 1/(K + 1)$ is to be read as a candidate structural identification rather than a derivation, and the 1/2 is empirical at the SM level. The previous-round phrasing "binary fold symmetry at the half-closure point" is dropped here: it suggested numerological motivation when the cleaner structural claim is closure-normalised quartic democracy over K + 1 = 8 dimensions.

(ii) The combination $K^2 + 2K = \dim \mathfrak{su}(K + 1)$. The denominator 63 admits multiple algebraically equivalent readings ($K(K+2)$, $(K+1)^2 - 1$, dim SU(8), $|PG(5, 2)| = 2^6 - 1$), but the structural argument that selects it among them is the following.

If the electroweak closure potential is defined on a (K + 1)-dimensional **closure-normalised space**, then infinitesimal internal deformations of that space are generated by the unitary algebra $\mathfrak{u}(K + 1) = \mathfrak{u}(8)$, with

$$\dim \mathfrak{u}(K + 1) = (K + 1)^2 = 64.$$

On the (K + 1)-dimensional closure-normalised space. The K = 7 closure number is the total constraint count from the wider VERSF programme's K = 7 minimal-fact architecture, parsed in §6.1 as 6 local closure constraints + 1 global closure constraint. The "+1" extension to K + 1 = 8 in §A.5(ii) is *not* a second global constraint — it is the overall closure-norm scale dimension, the dimension along which the $\mathfrak{u}(1)$ trace generator acts. The closure-normalised space is therefore K closure constraints supplemented by one overall scale, totalling K + 1 dimensions. This identification of the dimensionality is itself a load-bearing premise inherited from the wider closure-architecture programme: §A.5(ii)'s argument runs on top of the K = 7 closure structure (which inherits the closure-democracy gap of §6.1 — see §9.3) and the (K + 1)-dimensional closure-normalised space identification (which inherits a closure-dimensionality gap that a future derivation of the master closure potential must close).

The overall trace generator of $\mathfrak{u}(K + 1)$ — the $\mathfrak{u}(1)$ direction — corresponds to a uniform rescaling of the closure norm. This shifts the vacuum location v but *does not change the dimensionless ratio* m_H/v, which is normalisation-invariant by construction. The trace mode is therefore **inadmissible** for the leading correction to m_H/v. The admissible correction algebra is the traceless subalgebra

$$\mathfrak{su}(K + 1) = \mathfrak{u}(K + 1) / \mathfrak{u}(1),$$

of dimension

$D\_corr = \dim \mathfrak{su}(K + 1) = (K + 1)^2 - 1 = K^2 + 2K = 63$ (for K = 7).

This selects 63 on structural grounds against the numerical alternatives of the table below:

- **$K(K + 1) = 56$ is excluded by closure failure.** It counts the off-diagonal entries of a $(K + 1) \times (K + 1)$ matrix and so does not close under commutators — it is not the dimension of any classical simple Lie algebra ($\mathfrak{su}(N)$ gives $N^2 - 1$, $\mathfrak{so}(N)$ gives $N(N-1)/2$, $\mathfrak{sp}(2n)$ gives $n(2n+1)$; none yields 56 for integer N). $K(K + 1)$ cannot serve as the algebra of admissible deformations, regardless of its numerical proximity to the data.
- **$(K + 1)^2 = 64$ is excluded by trace inadmissibility.** It is the dimension of the closed algebra $\mathfrak{u}(K + 1)$, but includes the $\mathfrak{u}(1)$ trace generator. Deformations along that direction rescale the closure norm uniformly and so contribute to corrections of normalisation-dependent quantities (v itself, m_H itself), not to the dimensionless ratio m_H/v.

The denominator 63 is therefore selected as the dimension of the **unique algebra of admissible internal deformations preserving closure normalisation**. $K(K + 1)$ fails algebraic closure; $(K + 1)^2$ fails normalisation invariance; only $\mathfrak{su}(K + 1)$ survives both filters.

The Closure-Deformation Selection Principle

The argument just given deserves to be named. Across the wider VERSF programme, structural arguments in the same architectural family recur in different sectors, and pulling the common content out into a single principle clarifies what is being applied here.

Closure-Deformation Selection Principle. *On the $(K + 1)$-dimensional closure-normalised space, two ontologically distinct kinds of dimensionless content must be distinguished:*

*(a) **Bare couplings** — scalar parameters defined on the space before deformations are turned on (e.g. the bare Higgs quartic λ_baseline). These admit a natural dilution by the space dimension K + 1.*

*(b) **Corrections to dimensionless ratios** — quantities generated by infinitesimal internal deformations of the closure-normalised space. These are carried by elements of the unitary deformation algebra $\mathfrak{u}(K + 1)$. The $\mathfrak{u}(1)$ trace mode acts as a uniform rescaling of the closure norm; it shifts only the absolute scale of the vacuum, not the dimensionless ratio. The trace mode is therefore physically redundant for dimensionless ratios and is quotiented out of the admissible correction sector. The correction algebra is the traceless subalgebra:*

$$\mathfrak{d}\_corr = \mathfrak{u}(K + 1) / \mathfrak{u}(1) = \mathfrak{su}(K + 1), \quad \dim \mathfrak{d}\_corr = (K + 1)^2 - 1.$$

*Corrections inherit the algebra dimension dim $\mathfrak{d}$_corr; bare couplings inherit the space dimension K + 1. The Principle applies to corrections; it does not collapse the two cases into one count, and the trace-removal logic does **not** carry over to bare couplings — these live on the space before any deformation is invoked.*

Cross-programme positioning: same architectural family, specific Principle. The Principle as stated is specific: it applies to corrections-to-dimensionless-ratios generated by deformations of a closure-normalised space, where the relevant counting is the traceless deformation-algebra dimension. The wider VERSF admissibility programme contains independent applications of a *related architectural pattern* — admissibility-equivalence followed by normalisation-mode quotient — in different specific settings:

- **Entropy partition uniqueness** (Taylor, *Admissibility Closure and the Uniqueness of Physical Entropy*): physical entropy reduces to a uniquely admissible closure-equivalence class, with redundancy in absolute normalisation quotiented out via the $\eta = 1$, $\Phi\_c = 1$ No-Alternative Partition Theorem.
- **Record-current closure transport** (Theorem 12.2 of the record-current uniqueness chain): admissible record-current deformations preserve total commitment-density transport; uniform rescalings of the current normalisation are quotiented out as physically equivalent.
- **Closure-preserving gauge structure** ($K = 7$ hexagonal closure / Standard Model gauge group derivation): the gauge sector picks out traceless internal structure as the carrier of dynamical content, with the abelian trace mode treated as overall normalisation.
- **Operational law closure**: admissible operational equivalence classes systematically remove redundant normalisation modes before identifying physical content.
- **Generation-space transport and CKM/PMNS mixing** (Taylor, *Deriving Flavour Mixing from Closure Geometry*): the $K = 7$ closure architecture is reduced via gauge-mode quotient, hub-constraint removal, admissibility filtering on the channel sector, and rank-1 matching for non-adjacent transport, with composite/normalisation modes quotiented at each reduction step.

These applications share the *family-level* form "physically meaningful structure on closure-normalised spaces is admissibility-equivalence-quotient" but are not all instances of the specific deformation-algebra Principle of §A.5(ii). The Principle is the specific manifestation of the family in the curvature-correction sector; the family-level recurrence supplies independent grounds for thinking the trace-removal architecture is recurrent rather than *ad hoc*, without claiming each cited application is itself an instance of the same Principle. Promoting the §A.5(ii) Principle to fully cover all four applications would require widening its statement (e.g. to "physically meaningful structure on closure-normalised spaces lives in the traceless quotient of the unitary deformation algebra"); we do not do that here, preferring to keep the specific Principle precise and the cross-programme claim honest about scope.

This addresses the correction-algebra gap of the previous version. With the Closure-Deformation Selection Principle in hand, the $K^2 + 2K$ denominator is no longer a numerically convenient choice but is forced by an architecture-wide selection rule — instantiated specifically in the curvature-correction sector — applied to the $(K + 1)$-dimensional closure-normalised space.

What remains open at this level. Two distinct gaps remain even after §A.5(ii) selects the algebra:

(a) **The coefficient form.** Selecting $D\_corr = 63$ as the relevant algebraic dimension does not by itself fix the form $1/D\_corr$ for the leading correction. Alternative perturbative weightings — $1/D^2\_corr$, $D\_corr/(D\_corr + 1)$, $1/\sqrt{D}\_corr$, and others — are dimensionally and structurally consistent with $\mathfrak{su}(K + 1)$ being the deformation algebra. The candidate $(1/2)(1 + 1/D\_corr)$ form requires a closure-perturbation-theory derivation showing that $1/D\_corr$ is the natural leading-order coefficient. The numerical match at 0.11% is consistent with this form but does not force it among the alternatives.

(b) **The 1/2 baseline (§A.5(i)) is unchanged.** The $\mathfrak{su}(K + 1)$ argument selects the *correction* algebra; it does not derive the baseline. The $\lambda\_baseline = 1/(K + 1)$ candidate of §A.5(i) remains as flagged.

The denominator 63 is now structurally selected by the Closure-Deformation Selection Principle. The coefficient form and the baseline still need work — and §A.5(iii) records two candidate lemmas that close those gaps to the same structurally-argued (not master-action-derived) standard as the Principle itself.

(iii) Two candidate lemmas closing the residual gaps. The two open gaps in (i) and at the end of (ii) admit candidate closures of the same epistemic status as the Closure-Deformation Selection Principle: structural arguments motivated by physical regime considerations, ruling out the natural alternatives, but not derived from the VERSF master action.

Lemma A (Uniform Closure-Fluctuation Distribution). *If the leading finite-resolution correction to a dimensionless curvature ratio is carried by a single unresolved admissible closure fluctuation distributed uniformly across the physically admissible traceless deformation algebra $\mathfrak{d}\_corr = \mathfrak{su}(K + 1)$, then the first-order correction per mode is*

$$\varepsilon\_lead = 1 / D\_corr,$$

i.e. the uniform mean-field share of one closure-fluctuation quantum across the admissible deformation algebra.

The dismissals are by regime:

- $1/D^2\_corr$ would correspond to a *two-quantum* pair correction at leading order — i.e. a connected pair probability. The lemma assumes the leading correction is one-quantum, which is what "leading" means in a perturbative expansion in $1/D\_corr$.
- $D\_corr / (D\_corr + 1)$ is a saturation factor ($\to 1$ as $D\_corr \to \infty$) appropriate to *filling* available capacity, not to perturbing around a stable mode. The lemma assumes the perturbative regime $1/D\_corr \ll 1$ (here $1/63 \approx 0.016$), where saturation is not the relevant mechanism.

Lemma A is a **first-order mean-field closure hypothesis**, not yet a theorem. It is structurally argued, not master-action-derived. Its load-bearing assumptions are (i) that the leading correction is one-quantum rather than two-quantum, and (ii) that closure-symmetry acts transitively on the D_corr deformation modes so that uniform distribution is forced.

Positioning relative to large-N. The form (1/2)(1 + 1/N) is structurally familiar from large-N expansions in gauge theory and matrix models, where leading corrections to tree-level dimensionless ratios scale as 1/N with N tied to the gauge or matter content. Lemma A's prediction $\varepsilon_{lead} = 1/D_{corr}$ converges with that familiar form, and an experienced reader will pattern-match accordingly. This convergence is *structurally suggestive* but should not be over-read in either direction:

- *Not a special case of large-N.* Standard large-N is justified by diagrammatic counting (planar-diagram dominance in the $N \to \infty$ limit, with non-planar corrections suppressed by $1/N^2$). Lemma A invokes no such counting: it asserts uniform mean-field distribution of one closure-fluctuation quantum across the admissible deformation algebra, a distinct mechanism. The physical content of N also differs — large-N's N is a gauge/matter rank, while D_{corr} is the dimension of the closure-deformation algebra of a (K + 1)-dimensional closure-normalised space.
- *Not unrelated to large-N either.* The shared parametric form suggests that closure-fluctuation distribution and planar-diagram dominance may be limit cases of a more general first-order democratic mechanism on admissible algebraic structure. A future master-action derivation that produced Lemma A either as a large-N limit of a closure-microphysics action, or as a genuinely distinct first-order mechanism, would settle the relationship.

The position recorded here is the honest one: Lemma A's form is recognisable from a well-established family of 1/N corrections, but is not currently derived as a member of that family. Pattern-matching to large-N is neither grounds for accepting Lemma A nor grounds for dismissing it; the lemma should be evaluated on its own structural argument (one-quantum mean-field plus closure-symmetry transitivity), with the large-N convergence treated as suggestive rather than supporting.

Lemma B (closure-normalised quartic). *If the bare quartic coupling of the Higgs potential is diluted democratically over the K + 1 closure-normalised electroweak coherence dimensions with overall closure-norm unity, then*

$$\lambda_0 = 1 / (K + 1),$$

giving the baseline

$$m_H / v |_{baseline} = \sqrt{2\lambda_0} = \sqrt{2 / (K + 1)} = 1/2 \quad (\text{for } K = 7).$$

This is a 't Hooft-style democracy argument: the bare coupling per closure dimension is the total normalised coupling divided by the dimension count. In gauge theory, the 't Hooft limit g^2N held fixed as $N \to \infty$ is justified by diagrammatic counting (planar dominance). The analogue here is asserted rather than derived from closure-microphysics diagrammatic counting; the structural picture is that K + 1 closure-normalised dimensions share a fixed total quartic budget democratically.

Internal consistency and the space-vs-algebra asymmetry. Lemma A and Lemma B both use "democratic dilution" rhetoric but divide by different counts — Lemma B by $K + 1 = 8$, Lemma A by $(K + 1)^2 - 1 = 63$. The asymmetry is not arbitrary: it is exactly what the refined Closure-Deformation Selection Principle (§A.5(ii)) predicts. A bare coupling parameter is a scalar defined on the $(K + 1)$-dimensional vector space before any deformation is invoked, and inherits the space dimension as its natural dilution count (Lemma B). A correction is generated by infinitesimal deformations *of* that space, lives in the traceless deformation algebra $\mathfrak{su}(K + 1)$, and inherits the algebra dimension $(K + 1)^2 - 1$ as its natural count (Lemma A). The combination (1/8 baseline × 1/63 correction) is therefore the joint application of the Principle's two ontologically distinct counts to the same closure-normalised architecture — not a tuned pair of independent denominators chosen because they happen to give 32/63.

A hostile reader noticing that "democratic distribution" appears twice with two different counts is observing exactly what the Principle predicts. The sharper challenge — *why not apply trace removal to the baseline too, getting $\lambda_{baseline} = 1/63$ and $m_H/v|_{baseline} \approx 0.178$?* — fails because trace removal is the apparatus that filters which *deformations* contribute to corrections, and a coupling defined before deformations are turned on is not a deformation to be filtered. The Principle's ontological distinction between space-defined and algebra-generated content is what prevents this miscategorisation.

Combined result. With Lemmas A and B accepted as candidate closures, §A.5's central relation reads

$$m_H / v = (1/2) \cdot (1 + 1 / D_{corr}) = (1/2) \cdot (1 + 1 / (K^2 + 2K)) = \mathbf{32 / 63 \approx 0.50794},$$

equivalent in SM language to an effective quartic coupling

$$\lambda_{eff} = (1 / (K + 1)) \cdot (1/2) \cdot (1 + 1 / (K^2 + 2K))^2 = 0.1290,$$

against the observed $\lambda \approx 0.129$. The numerical match is at the ~0.11% level (within current m_H uncertainty), and is now backed by:

- §A.5(ii): the Closure-Deformation Selection Principle selects the algebra-dimension count $D_{corr} = \dim \mathfrak{su}(K + 1) = 63$ for corrections to dimensionless ratios.
- Lemma A (Uniform Closure-Fluctuation Distribution): selects the coefficient form $1/D_{corr}$ as the per-mode share of one closure-fluctuation quantum across the admissible deformation algebra.
- Lemma B (closure-normalised quartic democracy): selects the baseline $\lambda_{baseline} = 1/(K + 1) = 1/8$ as the bare coupling diluted democratically across the space dimensions (the asymmetry with Lemma A's algebra-dimension count is the Principle's prediction, not a tuned pair; see the consistency check above).

What this does and does not establish. Accepting Lemmas A and B closes the §A.5 gaps at the *lemma level*. §A.5 becomes a structurally derived candidate, conditional on the two lemmas surviving rigorous derivation. What is *not* yet established:

- **Master-action derivation of Lemmas A and B.** Both are structural arguments motivated by physical regime considerations (one-quantum, perturbative; democratic distribution; 't Hooft-style dilution). Neither is derived from the VERSF master action's closure-potential functional. A future master-action derivation must reproduce both forms; if it produces $\varepsilon_{lead} = 1/D^2_{corr}$ at leading order, or $\lambda_0 \neq 1/(K + 1)$, the lemmas fail and §A.5 reverts to "structurally selected denominator with empirical baseline and coefficient."
- **Universality (§A.7).** The lemmas might be internally self-consistent for m_H/v while failing to predict any other electroweak ratio. Universality is independent of the lemmas being correct, and remains the highest-leverage diagnostic.

§A.5 is therefore now structurally derived *up to* the lemma-level closure and the universality check. That is meaningfully stronger than the previous round's "structurally selected denominator with empirical baseline and coefficient," but it is still appropriately short of "completed derivation from first principles."

Numerical alternatives within the K = 7 architecture. To make the previous point concrete: with K = 7 fixed, a short list of "natural" combinatorial denominators in the form (1/2)(1 + 1/D) gives

D	(1/2)(1 + 1/D)	deviation from 0.50849
K(K+1) = 56	0.50893	**+0.087%**
$K^2 = 49$	0.51020	+0.336%
$(K+1)^2 - 1 = K^2 + 2K = 63$	0.50794	−0.108%
$(K+1)^2 = 64$	0.50781	−0.133%
$K^2 + 3K = 70$	0.50714	−0.266%
$2K^2 = 98$	0.50510	−0.668%

The data alone do not select $K^2 + 2K$. Three of these readings — K(K+1), $K^2 + 2K$, and $(K+1)^2$ — sit within roughly 0.1% of the observed value, with K(K+1) = 56 *closer to the data than $K^2 + 2K = 63$*. The current m_H measurement uncertainty (~0.09% at 1σ on $m_H = 125.20 \pm 0.11$ GeV) is comparable to the spread among these three. The numerical fit therefore does not pick $K^2 + 2K$ out from the alternatives; if anything it weakly favours K(K+1). **The structural argument of §A.5(ii) supplies the selection rule that the data cannot:** K(K+1) is excluded by algebraic closure failure (it is not the dimension of any Lie algebra), and $(K+1)^2$ is excluded by trace inadmissibility (it includes a normalisation-rescaling mode that cannot contribute to corrections of a dimensionless ratio). The framework therefore commits to the −0.108% deviation as a *prediction* — not a fitted preference — and a future closure-microphysics derivation that produced K(K+1) numerically would falsify the dim $\mathfrak{su}(K + 1)$ commitment, regardless of how close it sat to the data.

The structural argument is necessary, not merely sufficient. This point deserves explicit emphasis. (1/2)(1 + 1/63) = 0.50794 and (1/2)(1 + 1/64) = 0.50781 agree to four decimal places — they differ by ~0.025%, well within the current m_H uncertainty of ~0.09%. No amount of improved experimental precision on m_H, short of an order-of-magnitude reduction, will distinguish them. The trace-inadmissibility argument of §A.5(ii) is therefore not *supplementing* a marginal data preference for 63 over 64; it is the *only available mechanism* for adjudicating between observationally degenerate forms. This sharpens the framework's commitment: if a future derivation produces $(K+1)^2 = 64$ instead of $\dim \mathfrak{su}(K + 1) = 63$ — for reasons that override the trace-inadmissibility argument — the framework is falsified at a level the data cannot detect, by an internal structural constraint rather than by experiment. §A.5(ii)'s argument is doing work that no experiment can do, and that work is load-bearing.

Epistemic status. The candidate is recorded here as a structurally derived proposal within the Mexican-hat avenue (§A.4), conditional on three layered structural arguments. The numerical agreement at ~0.11% is at the edge of what the current m_H measurement uncertainty can distinguish, and the K = 7 input is independently locked. The **denominator** $D_{corr} = \dim \mathfrak{su}(K + 1) = 63$ is selected by the Closure-Deformation Selection Principle of §A.5(ii) — an architecture-wide rule with independent applications elsewhere in the VERSF admissibility programme. The **coefficient** $1/D_{corr}$ is selected by Lemma A (§A.5(iii)) — Uniform Closure-Fluctuation Distribution across the admissible deformation algebra, a first-order mean-field closure hypothesis. The **baseline** 1/2 is selected by Lemma B (§A.5(iii)) — closure-normalised quartic democracy across the K + 1 dimensions, giving $\lambda_{baseline} = 1/(K + 1)$. All three arguments share the same (K + 1)-dimensional closure-normalised structure, which is an internal consistency check rather than independent fits. §A.5 has graduated from "tight number" through "structurally selected denominator with empirical baseline and coefficient" to "**structurally derived candidate conditional on Lemmas A and B and on the universality check, supported by a named architectural Principle**." It is offered at exactly that status: more than a coincidence, less than a fully-derived theorem, with the master-action work that would close the remaining gap clearly identified.

A.6 What promoting either candidate would commit the framework to

Promoting $m_H = 3\alpha \cdot \eta_{closure} \cdot E_{geo}$ (§A.2) to the status of a structural relation would commit the framework to predicting $m_H / v = 1/2$ *exactly* (modulo the same $\eta_{closure}$ that already appears in the main relation), when no closure-microphysics argument currently supports that value. The honest case against promotion is not that two residuals would double — they would not, since within resolution they are the same residual (§A.2) — but that the framework would acquire a prediction it cannot yet defend.

Promoting the §A.5 binary-baseline-plus-correction proposal would commit the framework to a specific functional form, $m_H/v = (1/2)(1 + 1/(K^2+2K))$, conditional on the remaining derivations: the rigorous master-action derivation of Lemmas A and B (§A.5(iii)); and a universality check against at least one other electroweak ratio (§A.7). The **algebraic dimension** $D_{corr} = \dim \mathfrak{su}(K + 1)$ is structurally selected by §A.5(ii) (it selects *which algebra*, not the coefficient form against that algebra); the **coefficient** $1/D_{corr}$ is structurally argued by Lemma A (§A.5(iii)); the **baseline** 1/2 is structurally argued by Lemma B (§A.5(iii)). Both lemmas operate at the lemma level rather than from the master action. With the three structural layers in place, §A.5 is a *structurally derived candidate* — not a fitted relation, and not a fully-derived theorem either, but a relation where every numerical factor is identified with a closure-architecture quantity. Until the lemmas are master-action-derived and the universality check is performed, §A.5 remains in the appendix; with both, it would graduate to the body of a follow-up paper.

A.7 Diagnostic for future work

For the §A.2 half-channel reading, two questions must be answered together before promotion:

1. **Derivation.** Can $N_{part\_sub} = 3$ be derived from closure microphysics *without using* m_H *as input?*
2. **Identification.** Does the derived 3-subset coincide with the physical Higgs mode, or is it orthogonal to it — e.g. tracking the eaten Goldstones?

A "yes" to (1) with a "no" to (2) would be worse than no derivation: the framework would be predicting a mass for the wrong modes. The Goldstone counting objection of §A.3 already shows that the simplest partition fails identification.

For the §A.5 binary-baseline-plus-correction proposal, the diagnostic — updated to reflect that §A.5(ii) closes the algebra gap and §A.5(iii)'s Lemmas A and B close the coefficient and baseline gaps at the lemma level — is:

1. **Universality check.** Does the same $K = 7$ closure architecture, with the same structurally-derived form $(1/2)(1 + 1/D_{corr})$, predict *other* dimensionless ratios with comparable precision? In principle this is the diagnostic that matters most — independent of whether Lemmas A and B survive master-action derivation, and the test that would convert "structurally derived candidate" into "framework working." In practice, however, the pre-registered form is narrower than it first appears, and the sample of candidate observables is currently very limited. Both points need explicit acknowledgement.

Pre-registration of the universality form. To prevent ex post selection, the form a successful prediction must take is specified in advance:

> *Any dimensionless electroweak ratio with curvature-vs-minimum structure analogous to* m_H/v *should take the form > >* $R = (a / b) \cdot (1 + 1 / D_{corr})$, *> > for some small-integer a, b tied to closure architecture via a named lemma analogous to Lemma B of §A.5(iii) — supplying a/b as a bare-coupling democratic dilution over the K + 1 closure-normalised dimensions — and with the* ***same*** $D_{corr} = (K + 1)^2 - 1 = 63$ *supplied by the Closure-Deformation Selection Principle.*

The form is parametric — its content is the *combination* of three structural constraints: (i) the baseline a/b is supplied by a pre-stated closure-architecture lemma in the Lemma B family, not reverse-engineered from the data; (ii) the correction is $1/D_{corr}$ in coefficient form (the Lemma A family); (iii) D_{corr} is the *same* 63 from §A.5(ii). A successful universality prediction must match all three.

Operationalising "tied to closure architecture." This qualifier is doing real work and must be operationalised, not left rhetorical. (a, b) qualifies as closure-architectural if and only if there exists a named lemma — pre-stated, with the same internal structure as Lemma B — that supplies a/b as a closure-architectural quantity (specifically, a bare-coupling democratic dilution over the K + 1 closure-normalised dimensions of the same architecture used in §A.5). Any (a, b) introduced *after* observing the target ratio, however plausibly justified, fails the test. Future ratios get their own pre-stated lemmas or they don't qualify.

The "curvature-vs-minimum" qualifier excludes most obvious candidates. The §A.5 architecture is tailored to ratios where the numerator and denominator arise from curvature vs. minimum-location of the same closure potential (the m_H/v structure). Most well-measured electroweak ratios do *not* have this structure:

- *m_W / m_Z* ≈ 0.8814 is set by gauge mixing ($\cos \theta_W$ at tree level) — a different closure mechanism altogether. The pre-registered form does not apply; numerical exploration confirms there is no clean small-integer (a, b) producing 0.8814 via $(a/b)(1 \pm 1/63)$. The framework's prediction for m_W/m_Z, if any, would come from a *different* closure-architectural mechanism with its own pre-stated lemma, not from the §A.5 form.
- *Lepton and quark mass ratios* (e.g. $m_e/m_\mu \approx 0.00484$) are set by Yukawa couplings — yet another mechanism, with no curvature-vs-minimum structure of the §A.5 kind.
- *CKM mixing angles* are set by the diagonalisation structure of the Yukawa matrices — again a distinct mechanism.

Each of these is "electroweak" in the loose sense but none is a curvature-vs-minimum observable in the sense the §A.5 pre-registration specifies.

The current test sample. Within the SM, m_H/v is *the only known dimensionless electroweak ratio with curvature-vs-minimum structure of the §A.5 kind*. The pre-registered universality test therefore has a sample of exactly one observable — the one that motivated the derivation. This is not a small problem: it means the universality diagnostic, while methodologically sound as a *form*, currently has no immediate external falsification candidate. The framework can claim only that the predicted form fits m_H/v, which is the same observable used to motivate the form, and cannot yet test the form's universality in practice.

Where new candidates would come from. The universality test becomes operational once additional curvature-vs-minimum observables are identified. Plausible sources:

- *Identifying additional curvature-vs-minimum structure in the SM*, in observables currently set by other mechanisms. For example, if part of the Higgs self-coupling λ structure can be parsed as a separate curvature-vs-minimum sub-observable, that would qualify.
- *BSM scalar sectors*. If the closure architecture predicts additional scalar degrees of freedom with their own potential minima, the corresponding curvature-vs-minimum ratios would be testable.
- *Closure-architectural predictions for other mechanisms*. A separate closure-architectural derivation in a different sector — using the same K = 7 closure architecture but a different reduction mechanism appropriate to that sector — would extend the framework's reach. This is methodologically the *broad-reading* universality strategy: the

architecture predicts something for each sector, but the form differs by sector. The methodological worry with the broad reading is that it lets the framework choose form per sector, weakening falsification weight. The discipline that prevents this is pre-registration: each sector's form must be derived from a named structural argument *before* the data check, not reverse-engineered.

A working example of disciplined cross-sector closure-architectural derivation is the VERSF flavour-mixing programme [Taylor, *Deriving Flavour Mixing from Closure Geometry*], which derives the CKM mixing hierarchy and PMNS structure from generation-space transport on the same K = 7 hexagonal closure architecture. Its structural inputs — dim R = 5, $\eta = 3/5$, $\chi = 2/5$, $D = \mathrm{diag}(1, 2, 4)$, $\beta = \sqrt{2}/6$, and a curvature-response coupling $\alpha\_curv = 1/10$ (not to be confused with the fine-structure constant α used in the main body of the present paper) — are each supplied by a named structural argument with conditional or postulate status explicitly tabulated in that paper's epistemic table, rather than fitted to flavour-mixing data. *That paper is not a universality check on the §A.5 form* — the mechanisms and the resulting parametric forms differ (transport-amplitude reduction versus curvature-correction trace removal) — but it demonstrates that the same closure architecture can support disciplined pre-registered derivations in multiple sectors.

The broad reading is more permissive than the narrow §A.5 form. This must be flagged openly. The CKM/PMNS form is small-integer ratios from rank-counting in a different reduction; it is not of the §A.5 form $R = (a/b)(1 + 1/D\_corr)$ at all. The framework's broad-reading universality is therefore more permissive than its narrow §A.5 form: the discipline that prevents per-sector form-choice from collapsing into post-hoc accommodation is *pre-registration* (each sector's form derived before any data check), not *formal constraint* (a single functional form across all sectors). The CKM/PMNS programme respects pre-registration discipline within its own sector — that is the substantive datum. Whether the broad-reading strategy across the programme as a whole accumulates enough sectors with disciplined pre-registration to constitute a non-trivial constraint, rather than effectively unlimited per-sector form freedom, is itself an open methodological question about the framework's predictive reach. The two interpretive questions are therefore distinct: (i) does this count as universality of the §A.5 form? — no, the forms differ; (ii) does this count as evidence that the K = 7 closure architecture supports disciplined cross-sector predictive content? — the CKM/PMNS programme is a positive datum on this, subject to the open question about per-sector form freedom.

Honest verdict on universality. As currently developed, the §A.5 architecture is *narrowly predictive* and its universality is *methodologically pre-registered but operationally awaiting candidates*. The framework's current empirical reach is limited to the m_H/v ratio; the form is structurally sharp but the test sample is sample-of-one. This is itself a feature, not just a limitation: a framework whose universality test has an empty external sample is one whose predictive content cannot currently be falsified externally, and the framework should not claim universality as an established property until the sample is non-empty. The §A.5 candidate's standing therefore remains "structurally derived from a named principle and two lemmas, with universality formally pre-registered but currently untestable in practice." This is more honest than claiming a universality check is "the diagnostic that matters most" while having no observable to apply it to.

2. **Master-action derivation of Lemma A.** Can the one-quantum mean-field form $\varepsilon_{lead} = 1/D_{corr}$ be derived from a closure-perturbation-theory expansion around the unperturbed potential, with the deformation algebra $\mathfrak{su}(K + 1)$ summed over? A derivation showing $1/D^2_{corr}$ or $D_{corr}/(D_{corr} + 1)$ instead would falsify the lemma. With the universality test currently sample-of-one, this becomes the highest-leverage diagnostic in practice.
3. **Master-action derivation of Lemma B.** Can the closure-normalised quartic dilution $\lambda_0 = 1/(K + 1)$ be derived from the closure-potential normalisation, with a closure-diagrammatic argument analogous to (or replacing) the 't Hooft-style democracy assumption? A derivation showing $\lambda_0 \neq 1/(K + 1)$ would falsify the lemma.

The diagnostic is sharper for §A.5 than §A.2 because §A.5's structural argument is layered: the algebra (§A.5(ii)), the coefficient (§A.5(iii) Lemma A), and the baseline (§A.5(iii) Lemma B) must each survive independent derivation, *and* the resulting form would need to extend to other ratios with the relevant curvature-vs-minimum structure — for which the test sample is currently empty (see the universality discussion above). With the algebra-selection gap of §A.5(ii) closed, the lemma-derivation gaps of §A.5(iii) identified, and the universality check methodologically pre-registered but operationally awaiting candidates, the load-bearing next steps are the master-action derivations of Lemmas A and B. These are higher-leverage than universality *in practice*, despite universality being higher-leverage *in principle*, because the universality test has no immediate external observable to apply it to while the lemma derivations are work that can be done now.

A.8 Summary

The single empirical fact $m_H / v \approx 0.508$ is not currently predicted by the framework and is acknowledged as an open problem (§9.3). Two candidate forms are recorded in this appendix:

- The §A.2 **half-channel reading** $m_H = 3\alpha \cdot \eta_{closure} \cdot E_{geo}$. *Structurally unsupported*: agreement at ~0.6%, distinct objection (Goldstone counting, §A.3) that rules out the simplest version, and a partition-identification gap (§A.7).
- The §A.5 **binary-baseline-plus-correction reading** $m_H / v = (1/2)(1 + 1/(K^2 + 2K)) = 32/63$. *Structurally derived at the lemma level*: a named architectural principle and two supporting lemmas select every factor. The **Closure-Deformation Selection Principle** (§A.5(ii)) selects the *algebraic dimension* D_{corr} as dim $\mathfrak{su}(K + 1) = 63$ by quotienting the redundant $\mathfrak{u}(1)$ trace mode from the unitary deformation algebra; this principle is an instance of a recurring architectural family across the wider VERSF admissibility programme (entropy partition uniqueness, record-current transport, gauge-sector trace removal). §A.5(iii) Lemma A (Uniform Closure-Fluctuation Distribution) selects the *coefficient form* $1/D_{corr}$ as the per-mode share of one closure-fluctuation quantum. §A.5(iii) Lemma B (closure-normalised quartic democracy) selects the *baseline* 1/2 via $\lambda_{baseline} = 1/(K + 1)$. The structural argument is also necessary, not merely sufficient: the alternative $D = 64$ sits at 0.50781, indistinguishable from $D = 63$ at 0.50794 within m_H precision, so the trace-inadmissibility argument is the only mechanism that can adjudicate. Tight agreement at ~0.11%, $K = 7$ independently locked, in the Mexican-hat avenue (§A.4) so not subject to the Goldstone objection, and pure SM-internal (no

dependence on m_κ anchoring or $\eta_{closure}$). *Remaining gaps:* master-action derivation of Lemmas A and B (currently structural arguments rather than first-principles derivations); inherited dependencies on closure democracy (§9.3) and on the (K + 1)-dimensional closure-normalised space identification; and a universality check that is methodologically pre-registered but currently sample-of-one — m_H/v is the only known SM-internal curvature-vs-minimum observable, so the framework's predictive form has no immediate external falsification candidate (§A.7). With the Principle, the two lemmas, *and the honest acknowledgement that universality is operationally untestable until new candidate observables are identified*, §A.5 is a **structurally derived candidate conditional on lemma derivation and closure-architecture inheritance, with universality formally pre-registered but currently sample-of-one** — not a fully-derived theorem and not yet a framework with external predictive reach beyond the observable that motivated the derivation.

The Mexican-hat curvature derivation (§A.4) is identified as the avenue both candidates ultimately address. §A.5 is the leading exploratory proposal within that avenue. Neither candidate is promoted: this appendix records the observations, the objections, and the diagnostics. The full derivation, when it comes, will be a separate paper.

See versf-eos.com for the wider VERSF programme, including the closure structure papers (K = 7, N_loop = 14), the interface-matching and finite-distinguishability papers, the TPB and bit-conservation papers, the fine-structure constant series, and the record-current uniqueness theorems referenced above.

References

Internal: VERSF programme

The full catalogue of papers in the Void Energy–Regulated Space Framework programme, including titles, version histories, dates, and download links, is maintained at **versf-eos.com**. The following programme components are referenced in this paper.

- **Companion paper.** Taylor, K., *The Hierarchy Problem as a Category Error: A VERSF Interpretation of Scale Separation, Emergent Mass, and Closure Dynamics*. Establishes the two-layer ontology (closure vs. record), the Ontological Separation Principle (OSP) and Closure Non-Propagation Principle (CNPP), the admissibility kernel $F_{adm}(k)$ regulating record-layer correlation functions, the derivation of the κ-field mass $m^2_\kappa = (4/3)\cdot\xi^{-2}$ from the K=7 minimal-fact architecture and the CCC threshold, and the geometric-mean scale $E_{geo} = \sqrt{(M_P\cdot m_\kappa)} \approx 5$–6 TeV. The present paper takes the category-error reframing of this companion as its starting point and develops the second-order structure of the occupancy fraction within the coherence band.

- **Closure structure.** Specifically: Taylor, K., *A No-Go Theorem for Non-Simplicial Relational Substrates – K=7* (establishes K = 7 minimal-fact architecture via the Fano plane PG(2,2)); and the companion papers on the N_loop = 14 channel structure and the hexagonal closure geometry derivation of the Standard Model gauge group. The K = 7 no-go theorem is the load-bearing reference for §6.1 and §A.5.
- **Flavour mixing from closure geometry.** Taylor, K., *Deriving Flavour Mixing from Closure Geometry: CKM and PMNS Structure in the VERSF Framework*. Derives the CKM mixing hierarchy and PMNS structure from generation-space transport on the K = 7 hexagonal closure architecture, with structural fractions $\eta = 3/5$ (channel-survival) and $\chi = 2/5$ (non-adjacent continuation) obtained from rank-counting in the five-dimensional closure residual space, and the stiffness hierarchy D = diag(1, 2, 4) derived from closure-depth mode doubling. Cited in §A.5(ii) as a parallel application of the architectural family and in §A.7 as a working example of disciplined cross-sector derivation; see §A.7 for the precise scope of what this citation establishes and what it does not.
- **Interface matching and finite distinguishability.** The interface-matching framework, including the gauge kinetic normalization $\beta(\Lambda) = \ln 2 / (1 - \cos(2\pi/N\_\varphi))$ and the impedance representation $\alpha = Z_0 / 2R\_K$ as a structural distinguishability-transfer coefficient.
- **Second-order closure competition.** The participation-structure analysis yielding IPR = 1/6 over six local closure constraints with one global closure constraint, and the closure-democracy assumption underwriting N_part = 6. Load-bearing reference for §6.1; specific paper title to be supplied from the versf-eos.com catalogue.
- **Ticks-Per-Bit (TPB) and Bit Conservation and Balance (BCB).** The framework distinguishing reversible substrate exploration ("ticks") from irreversible committed distinguishability ("bits"); the entropy relation $S = k\_B \cdot \ln(1/TPB)$; the Maxwell admissibility paper.
- **Fine-structure constant series.** The six-paper sequence deriving $\alpha^{-1} \approx 137.034$ from finite-distinguishability interface matching.
- **Two-Planck framework.** Specifically: Taylor, K., *The Two-Planck Framework*, together with Taylor, K., *A Hidden "Middle Scale" in the Universe* (which establishes the mesoscopic CCC coherence scale $\xi \approx 8.2 \times 10^{-5}$ m used as input to the κ-field mass). Load-bearing reference for §4.3.
- **κ-field mass.** Taylor, K., *The κ-Field Mass as the Physical Hessian of the Commitment Constraint Surface*, proving $m^2\_\kappa = \lambda\_eff \cdot \xi^{-2} = (4/3) \cdot \xi^{-2}$ as a theorem under the K = 7 architecture and CCC threshold.
- **Record-current uniqueness.** Theorem 12.2 (record-current functional uniqueness) and the surrounding theorem chain on record-layer rank selection, current uniqueness, and metric emergence.
- **Born rule derivation series.** The Double Square Rule, Admissibility Closure, and Overdetermined Born Rule routes to the quantum probability rule.
- **PAR/CC/IAC chain.** Pre-Factual Algebraic Reversibility, Compositional Completeness, and Internal Admissible Closure as the substrate-level admissibility framework.
- **Facts as Structural Necessities.** The metaphysical foundation paper establishing record-layer facticity.
- **VERSF Guide and Results Catalogue.** Front-door overview and structured classification of programme results.

External: standard literature

The hierarchy problem and naturalness.

- 't Hooft, G. (1980). "Naturalness, chiral symmetry, and spontaneous chiral symmetry breaking." In *Recent Developments in Gauge Theories* (Cargèse Lectures), NATO Sci. Ser. B, vol. 59, pp. 135–157.
- Susskind, L. (1979). "Dynamics of spontaneous symmetry breaking in the Weinberg-Salam theory." *Physical Review D* **20**, 2619–2625.
- Giudice, G. F. (2008). "Naturally Speaking: The Naturalness Criterion and Physics at the LHC." In *Perspectives on LHC Physics*, G. Kane and A. Pierce, eds., World Scientific.

Renormalization group and effective field theory.

- Wilson, K. G. (1971). "Renormalization group and critical phenomena. I. Renormalization group and the Kadanoff scaling picture." *Physical Review B* **4**, 3174–3183.
- Wilson, K. G., & Kogut, J. (1974). "The renormalization group and the ε-expansion." *Physics Reports* **12**, 75–199.

Higgs discovery and Standard Model parameters.

- ATLAS Collaboration (2012). "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC." *Physics Letters B* **716**, 1–29.
- CMS Collaboration (2012). "Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC." *Physics Letters B* **716**, 30–61.
- Particle Data Group / Navas, S. *et al.* (2024). "Review of Particle Physics." *Physical Review D* **110**, 030001.

Higgs mechanism (foundational).

- Englert, F., & Brout, R. (1964). "Broken symmetry and the mass of gauge vector mesons." *Physical Review Letters* **13**, 321–323.
- Higgs, P. W. (1964). "Broken symmetries and the masses of gauge bosons." *Physical Review Letters* **13**, 508–509.
- Guralnik, G. S., Hagen, C. R., & Kibble, T. W. B. (1964). "Global conservation laws and massless particles." *Physical Review Letters* **13**, 585–587.

Goldstone modes (referenced in Appendix A.3).

- Goldstone, J. (1961). "Field theories with 'superconductor' solutions." *Il Nuovo Cimento* **19**, 154–164.
- Goldstone, J., Salam, A., & Weinberg, S. (1962). "Broken Symmetries." *Physical Review* **127**, 965–970.

Gauge invariance and Wilson loops (referenced in §2.1).

- Wilson, K. G. (1974). "Confinement of quarks." *Physical Review D* **10**, 2445–2459.
- Polyakov, A. M. (1987). *Gauge Fields and Strings*. Harwood Academic Publishers.

Fine-structure constant.

- Tiesinga, E., Mohr, P. J., Newell, D. B., & Taylor, B. N. (2021). "CODATA recommended values of the fundamental physical constants: 2018." *Reviews of Modern Physics* **93**, 025010.

Quantum Hall effect and impedance (referenced in §2.3).

- von Klitzing, K., Dorda, G., & Pepper, M. (1980). "New method for high-accuracy determination of the fine-structure constant based on quantized Hall resistance." *Physical Review Letters* **45**, 494–497.

Cosmological constant (referenced in §4.3).

- Planck Collaboration / Aghanim, N. *et al.* (2020). "Planck 2018 results. VI. Cosmological parameters." *Astronomy & Astrophysics* **641**, A6.