

The Fine-Structure Constant as a Geometric State Count

(Information-Geometric Physics System XI)

Pruk Ninsook

Ratchaburi, Thailand

prukns26@gmail.com

June 2026

Abstract

We show that the fine-structure constant follows from the information-geometric phase structure (IGPS) of the single Oloid seam under one explicitly stated physical postulate, with no parameters fitted to α . The postulate—which we name the Distinguishability Coupling Principle and prove is not forced by any standard principle (unitarity, gauge invariance, positivity, or the quantum speed limit)—is that the strength of the electromagnetic interaction *is* the operational distinguishability of the coupled charge sectors. The specific information-geometric machinery it is sometimes phrased in—the Fisher–Rao form, the Bures metric, and the square-root (length) measure—is not posited but *derived* beneath the principle, by the classification theorems of Chentsov and Petz (given charge-preserving coarse-graining), an optimal-measurement selection, and the counting no-go of Appendix C. Granting it, the tree-level value is fixed topologically by the seam phase area $\Phi_{\text{total}} = 8\pi$ and the geometric stiffness $\beta = 1/(\sqrt{3}\pi)$, giving $\alpha_{\text{tree}}^{-1} = \Phi_{\text{total}}/\beta = 8\sqrt{3}\pi^2 \approx 136.757$. A self-consistent one-loop (Dyson) correction, whose coefficient $c = 12 + \frac{1}{2\pi}$ combines the Wess–Zumino–Witten vertex $12 = (k_{\text{SU}(2)} + h_{\text{SU}(2)}^\vee)/h_{1/2}^{(1)}$, fixed by the $\text{SU}(2)_1$ conformal weight $h_{1/2}^{(1)} = \frac{1}{4}$, with the Atiyah–Patodi–Singer index term $1/2\pi$, yields

$$\alpha^{-1} = 8\sqrt{3}\pi^2 + \left(12 + \frac{1}{2\pi}\right)\pi\alpha \approx 137.036,$$

matching the measured value to a relative error of 3×10^{-6} . Each ingredient—the geometric tree term, the WZW vertex, and the APS index—is derived independently in earlier work, so that, *given the postulate*, the result is a zero-parameter consequence rather than a fit. The choice the postulate makes is sharp and falsifiable: the length reading gives $\alpha^{-1} \approx 137$, whereas the integrated susceptibility of standard field theory would give ≈ 237 ; experiment selects the former. We are careful to separate two claims of different strength. The high-sensitivity content is the tree term: it reproduces α^{-1} to 0.2% with no inputs, and because it supplies 99.8% of the total, the geometry-fixed stiffness β sits on a knife edge—it must be correct to $\sim 0.02\%$ for the prediction to hold at all—so the agreement is genuine evidence for the derived β . The correction coefficient, by contrast, is only loosely constrained by the data (the small factor $\pi\alpha$ leaves a wide admissible band), so the 3×10^{-6} figure is a low-sensitivity test of c , not a high-precision one. It is, however, a *well-defined* prediction: we show the relation is exact at the level of the photon two-point function—the seam being a free boson, the $O(\alpha^2)$ correction to the coupling vanishes identically by the Schwinger-model exactness of the vacuum polarisation—so no higher-order term has been dropped. The linear (length, Hooke-type) accumulation that the prediction requires is precisely the content of the Distinguishability Coupling Principle stated above: we prove that, given a count of distinguishable states, the measure must be the Bures *length* $\oint \sqrt{F}$ rather than the integrated information $\oint F$, while the factor $\sqrt{3}$ is proved from three-dimensional isotropy.

We further show that the tree-level seam theory is a two-dimensional free-boson conformal field theory ($c = 1$), which explains structurally why the framework fixes dimensionless

quantities—the quartic coupling $\lambda_H = 31/240$ (hence m_H/v to 0.08%) and the strong-coupling fixed-point value $\alpha_{s,\text{kink}}^{-1} = k + h^\vee = 4$ —while leaving the absolute mass scales undetermined: a conformal theory carries no intrinsic scale, so one dimensionful anchor remains, exactly as Λ_{QCD} and v are inputs in the Standard Model. The framework also yields a tree-level weak mixing angle $\sin^2 \theta_W = \frac{1}{4}$ ($\theta_W = 30^\circ$), which agrees with the measured value only at the 8% level and which we flag as an open relation rather than a precision result; notably the α^{-1} prediction is independent of it, resting on the $\text{SU}(2)_1$ conformal weight rather than on the mixing angle. As a by-product the colour and generation multiplicities are unified, $N_c = N_{\text{gen}} = 3$, through a single seam $\mathbb{Z}_3$ (the trefoil knot group projected onto the $\text{SU}(3)$ centre, with Alexander polynomial $\Delta(t) = t^2 - t + 1$). We delineate precisely what is proved, what is conditional, and what remains open—chief among the latter the stabilization of the conformal size modulus, which would generate the absolute scale but whose required quantum (Casimir) term has an unsecured sign tied to the unresolved Standard-Model $\rightarrow$ seam map.

Contents

1	Introduction	3
1.1	Context	3
1.2	Why the Oloid	3
1.3	The question	4
1.4	Result and claims	4
1.5	What is and is not new	5
2	The Tree-Level Result $\alpha^{-1} = 8\sqrt{3}\pi^2$	6
2.1	The phase area $\Phi_{\text{total}} = 8\pi$	6
2.2	Factorization of the partition function	6
2.3	The topological energy $E_{\text{topo}} = \pi$ from Gauss–Bonnet	6
2.4	The strain factor $\sqrt{3}$ from three-dimensional isotropy	7
2.5	The stiffness $\beta^{-1} = \sqrt{3}\pi$	7
2.6	Why the accumulation is linear, not quadratic	8
2.7	Assembly	8
3	The One-Loop (Dyson) Correction	9
3.1	The self-consistent form	9
3.2	The quadratic piece $c = 12$: the WZW vertex	9
3.3	The linear piece $\frac{1}{2\pi}$: the APS index term	10
3.4	Why the correction is self-consistent	10
3.5	Two distinct spectral objects	10
3.6	Result	10
4	Zero Parameters: the Epistemic Claim	11
4.1	Accounting	11
4.2	Why the result is not reverse-engineerable	12
4.3	Rejected near-coincidences (a single, uniform standard)	12
5	The Strong Coupling α_s	13
5.1	The kink value	13
5.2	Interpretation as a conformal fixed point	13
5.3	What is open: the absolute scale	14
5.4	Status of the weak mixing angle	14

6	Conformal Structure and the Absolute-Scale Question	15
6.1	The tree theory is a two-dimensional conformal field theory	15
6.2	Why dimensionless data is derivable and absolute scales are not	15
6.3	Where the one anchor might come from: the size modulus (open)	16
7	Structural By-Products	17
7.1	Colour-generation unification: one $\mathbb{Z}_3$	17
7.2	Unified spin-dependent symmetry: D_2 representation theory	18
7.3	One seam operator, APS data by chirality	19
8	Open Problems	19
9	Conclusion	21
A	Numerical Summary	22
B	Gauss–Bonnet Gives $E_{\text{topo}} = \pi$	23
C	Why the Accumulation is Linear, Not Quadratic	24
C.1	The node count is a dimensionless functional	24
C.2	The two scale-free count functionals are linear	25
C.3	The quadratic reading is excluded by scale-invariance	25
C.4	The Distinguishability Coupling Principle	27
C.5	The node operator from the seam current algebra	29
C.6	A correction to earlier framing	32
C.7	Status and the interpretive hinge	32
D	The Dyson Self-Consistent Equation	33
E	APS Anomaly Inflow and the Dyson Coefficient	34

1 Introduction

1.1 Context

The information-geometric phase-structure (IGPS) programme seeks to derive Standard-Model observables from the geometry and topology of a single Oloid manifold and its seam. The Oloid—the convex hull of two unit circles in perpendicular planes, each passing through the other’s centre [33]—is a developable surface ($K_G = 0$ almost everywhere) whose two interlocking circles meet along a closed *seam*. Earlier papers in the series established the framework’s vocabulary and its leptonic, baryonic, and gauge sectors: the informational phase current $Q_\mu = \partial_\mu \Phi$, where Φ is a $U(1)$ phase and Q_μ is its gradient (a phase current, not an internal multiplet); the geometric stiffness β that weights the seam action $S = \int \frac{1}{2} \beta |Q_\mu|^2 d\Sigma$, with seam distances measured by a Connes-type spectral metric [17]; and the total phase area $\Phi_{\text{total}} = 8\pi$ carried by a spin- $\frac{1}{2}$ field over the surface. The consistent direction of the programme has been *seam geometry* $\rightarrow$ *Standard-Model observables*.

1.2 Why the Oloid

Because this paper can be read independently of the rest of the series, we state at the outset why the carrier manifold is the Oloid and not some other surface; the full argument is given in Paper I [1], and we summarize the part that the present results rely on. The selection rests on two defining premises and one exact geometric fact.

The two premises define the admissible class. First, the manifold must be *developable* ($K_G = 0$ almost everywhere), so that the phase field $Q_\mu = \partial_\mu \Phi$ maps onto it without intrinsic curvature distortion; this excludes the sphere ($K_G > 0$) and is the property that makes the tree-level theory conformal (Section 6). Second, the manifold must support a spin- $\frac{1}{2}$ closure—an $SU(2)$ double cover consistent with 4π rotations—which excludes the cylinder. These are premises of the framework, adopted because the object of study is a spin- $\frac{1}{2}$ matter node, not theorems derived from a deeper principle; we flag them as such.

Within the resulting class—compact developable surfaces built from the convex hull of two linked unit circles (a Hopf-link frame)—the *regular* Oloid, in which the circles are perpendicular and each passes through the other’s centre (centre separation $d = r$), is distinguished by an exact and consequential property: its surface area is

$$A_{\text{Oloid}} = 4\pi r^2, \quad (1)$$

exactly that of a sphere of the same radius—a non-trivial closed-form theorem [33]. This is not a property of the two-circle roller at other separations (for instance the constant-height roller at $d = r\sqrt{2}$ has a different area). Equation (1) is precisely what the framework requires: measured in units of r^2 , the surface area is 4π , and a spin- $\frac{1}{2}$ field (which needs a 4π rotation to close, $\pi_1(SO(3)) = \mathbb{Z}_2$) doubles this to a total phase area $\Phi_{\text{total}} = 2 \times 4\pi = 8\pi$ (Section 2.1), the single geometric input to the tree-level α^{-1} . The Oloid is therefore selected not by aesthetic preference but because it is the developable, spinorial carrier whose area furnishes exactly the phase budget 8π .

We add one observation, marked as suggestive rather than load-bearing. Exactly one third of each generating circle’s perimeter lies interior to the hull, so the Oloid is equivalently the convex hull of two arcs each spanning $4\pi/3$ [33]; the resulting three-fold angular unit ($2\pi/3$) resonates with the $\mathbb{Z}_3$ structure that governs colour and generation number in Section 7.1—though that $\mathbb{Z}_3$ is established independently, from the trefoil topology of the seam, and does not depend on this arc-length coincidence. We do not claim the Oloid frame is itself a trefoil; it is a Hopf link of two circles, and the trefoil enters separately as the knot type of the seam.

All of the present work concerns *dimensionless* structure. As we make precise in Section 6, this is not incidental: the tree-level seam theory is conformal, so the quantities it can fix are exactly the dimensionless ones, and the absolute mass scales require a separate input. We state this limitation at the outset so that the claims below are read in the right light.

1.3 The question

The fine-structure constant α is a pure number. Its measured inverse, $\alpha^{-1} = 137.035999177(21)$ [35], is among the most precisely known constants in physics, yet the Standard Model offers no account of its value: α is an input. The question this paper addresses is whether the seam geometry already developed in Papers I–X fixes α [1, 2, 3, 4, 5, 7] *with no parameter adjusted to reproduce it*—and, if so, what the residual disagreement and the remaining assumptions actually are.

1.4 Result and claims

The tree-level value follows from two quantities each established elsewhere in the series,

$$\alpha_{\text{tree}}^{-1} = \frac{\Phi_{\text{total}}}{\beta} = \frac{8\pi}{1/(\sqrt{3}\pi)} = 8\sqrt{3}\pi^2 \approx 136.757 \quad (-0.20\%), \quad (2)$$

with $\Phi_{\text{total}} = 8\pi$ topologically forced ($\pi_1(SO(3)) = \mathbb{Z}_2$, so a spin- $\frac{1}{2}$ field needs a 4π rotation, doubling the $4\pi R^2$ surface area) and $\beta = 1/(\sqrt{3}\pi)$ the Oloid stiffness. A self-consistent one-loop (Dyson) correction with coefficient $c = 12 + \frac{1}{2\pi}$ —the Wess–Zumino–Witten vertex

$12 = (k_{\text{SU}(2)} + h_{\text{SU}(2)}^\vee)/h_{1/2}^{(1)}$, fixed by the $\text{SU}(2)_1$ conformal weight $h_{1/2}^{(1)} = \frac{1}{4}$ [8, 9] together with the Atiyah–Patodi–Singer index term $\frac{1}{2\pi}$ [12, 13]—closes the gap:

$$\alpha^{-1} = 8\sqrt{3}\pi^2 + \left(12 + \frac{1}{2\pi}\right)\pi\alpha \approx 137.036 \quad (\text{relative error } +3 \times 10^{-6}). \quad (3)$$

Table 1 shows the approach to the measured value as each piece is included.

Formula	α^{-1}	Relative error
$8\sqrt{3}\pi^2$ (tree)	136.757	-2.0×10^{-3}
$\sqrt{3}(8\pi^2 + \frac{1}{2\pi})$ (closed form)	137.033	-2.3×10^{-5}
$8\sqrt{3}\pi^2 + 12\pi\alpha$ (Dyson, $c = 12$)	137.032	-2.7×10^{-5}
$8\sqrt{3}\pi^2 + (12 + \frac{1}{2\pi})\pi\alpha$	137.036	$+3 \times 10^{-6}$

Table 1: Successive approximations to $\alpha^{-1} = 137.035\,999\,177(21)$. Each ingredient in the final row is derived independently (Section 2, 3); none is fitted to α .

We stress the epistemic status of Eq. (3). Every ingredient—the geometric tree term $8\sqrt{3}\pi^2$, the vertex 12, and the index $\frac{1}{2\pi}$ —is fixed by earlier results, so that, *given the one physical postulate stated below*, the result is a *zero-parameter consequence*, not a fit. We do not claim to *derive* α^{-1} from the seam dynamics with no further input; we claim the sharper and more defensible thing—that a single, named, falsifiable postulate about the meaning of the coupling, together with the independently-fixed Oloid geometry, yields the value with nothing tuned to it. Two claims of different strength must then be separated. The tree term alone reproduces α^{-1} to 0.2% and supplies 99.8% of the total; the formula is correspondingly sensitive to the stiffness (β must be correct to $\sim 0.02\%$ even allowing the correction coefficient to vary over its admissible band, and to $\sim 3 \times 10^{-4}\%$ to pin the headline), so a geometry-fixed β landing there is genuine evidence the derived β is correct. The correction coefficient c , contributing only the residual 0.2% through the small factor $\pi\alpha$, is by contrast only loosely constrained by the data; the 3×10^{-6} is therefore a *low-sensitivity* test of c , not a high-precision one. It is nonetheless a well-defined, truncation-free prediction: the relation is exact at the photon two-point level, the $O(\alpha^2)$ correction to the coupling vanishing identically because the seam is a free boson (Appendix E). The single foundational assumption that remains is a physical one about the meaning of the coupling, which we name the Distinguishability Coupling Principle (Appendix C.4): that the strength of the electromagnetic interaction *is* the operational distinguishability of the coupled charge sectors. The Fisher–Rao form and the square-root (length) measure are then derived, not posited—by the Chentsov–Petz classification of monotone metrics and by the counting no-go—so that the stiffness accumulates the strain density $\sqrt{3}$ *linearly* ($\beta^{-1} = \int_0^\pi \sqrt{3} d\phi = \sqrt{3}\pi$) rather than quadratically. We argue this postulate is irreducible—no unitarity, gauge, positivity, or quantum-speed-limit principle forces it, as each fixes only that distinguishability is a length, not that the coupling is the distinguishability—and we therefore state it as the framework’s one physical axiom rather than leaving it implicit. It is also falsifiable: it selects the length (137) over the integrated susceptibility (237), and experiment favours the former. Given it, the factor $\sqrt{3}$ is proved from three-dimensional isotropy (Section 2). We are careful throughout to separate what is proved from this one conditional step.

1.5 What is and is not new

The inputs $\Phi_{\text{total}} = 8\pi$, the stiffness β , the WZW vertex, and the APS index are inherited from Papers I–X [1, 3, 7]. New to this paper are: (i) the closed self-consistent (Dyson) form of Eq. (3) and the identification of its coefficient as vertex-plus-index; (ii) the observation that the tree-level seam theory is a two-dimensional free-boson conformal field theory ($c = 1$), which explains *why* the framework fixes dimensionless data while leaving absolute scales open (Section 6); and

(iii) the unification of colour and generation multiplicities, $N_c = N_{\text{gen}} = 3$, through a single seam $\mathbb{Z}_3$ —the trefoil knot group projected onto the $\text{SU}(3)$ centre, with Alexander polynomial $\Delta(t) = t^2 - t + 1$ (Section 7.1). We also state plainly what is not achieved: the absolute mass scales are not derived, because a conformal theory carries no intrinsic scale and one dimensionful anchor must remain—exactly as Λ_{QCD} and the electroweak vacuum expectation value v are inputs in the Standard Model (Section 6).

The remainder of the paper is organised as follows. Section 2 derives the tree-level result (2) and isolates the one conditional step. Section 3 establishes the Dyson correction and its coefficient. Section 4 makes the zero-parameter claim precise and records the coincidences set aside by the same standard. Section 5 treats the strong coupling. Section 6 proves the conformal character of the tree theory and discusses the absolute-scale question, including the size-modulus problem. Section 7 collects the dimensionless structural by-products, and Section 8 states the open problems.

2 The Tree-Level Result $\alpha^{-1} = 8\sqrt{3}\pi^2$

The tree-level value (2) is the ratio of two geometric quantities of the seam: the total phase area Φ_{total} and the stiffness β . We derive each in turn and then state precisely the single step that is conditional.

2.1 The phase area $\Phi_{\text{total}} = 8\pi$

The IGPS field carries spin $\frac{1}{2}$ over the Oloid surface, whose area is exactly $4\pi R^2$ —equal to the area of a sphere of the same radius, a classical result for the Oloid established by Dirnböck and Stachel and confirmed by subsequent work [33, 34]. We rely only on this surface area, which is a verified geometric invariant of the surface; in particular the construction does *not* use the Oloid’s centre-of-mass rolling behaviour (which in fact meanders rather than staying at constant height), so no claim here depends on that motion. Because $\pi_1(\text{SO}(3)) = \mathbb{Z}_2$, a spin- $\frac{1}{2}$ configuration returns to itself only under a 4π rotation, doubling the effective phase area swept by the field:

$$\Phi_{\text{total}} = 2 \times 4\pi = 8\pi. \quad (4)$$

This is topologically forced and carries no free parameter.

2.2 Factorization of the partition function

Write the seam phase as a Hodge decomposition $\phi = \phi_{\text{harm}} + \phi_{\text{exact}}$ into harmonic and exact parts, which are L^2 -orthogonal, $\langle d\phi_{\text{harm}}, d\phi_{\text{exact}} \rangle_{L^2} = 0$. The action is then additive, $S[\phi] = S_{\text{harm}} + S_{\text{exact}}$, and the Gaussian measure factorizes over the orthogonal subspaces,

$$Z = \underbrace{\int D\phi_{\text{harm}} e^{-S_{\text{harm}}}}_{Z_{\text{topo}}} \times \underbrace{\int D\phi_{\text{exact}} e^{-S_{\text{exact}}}}_{Z_{\text{space}}}. \quad (5)$$

This factorization is exact. The topological factor is the relevant one for α ; the spatial factor decouples and does not enter the fine-structure constant.

2.3 The topological energy $E_{\text{topo}} = \pi$ from Gauss–Bonnet

On the $\mathbb{Z}_2$ fundamental domain Σ_+ (the half-Oloid, a disk with $\chi = 1$), the Gauss–Bonnet theorem reads

$$\int_{\Sigma_+} \underbrace{K_G}_{=0} dA + \oint_{\partial\Sigma_+} \kappa_g ds + \sum_i \theta_i = 2\pi\chi(\Sigma_+) = 2\pi, \quad (6)$$

where $K_G = 0$ because the Oloid is developable. The half-seam is a semicircle of geodesic curvature $\kappa_g = 1/R$ and arc length πR , so $\oint \kappa_g ds = \pi$, and the corner terms contribute the remaining π . The minimal energy of one $\mathbb{Z}_2$ winding on Σ_+ is therefore

$$E_{\text{topo}} = \left| \oint_{\partial\Sigma_+} \kappa_g ds \right| = \pi. \quad (7)$$

2.4 The strain factor $\sqrt{3}$ from three-dimensional isotropy

The field is $Q_\mu = \partial_\mu \Phi$, the gradient of a $U(1)$ phase—a phase current, not an internal multiplet. The factor $\sqrt{3}$ is the strain density of an isotropic three-dimensional configuration. For an isotropically oriented direction $\langle \cos^2 \theta \rangle = \frac{1}{3}$, so the ratio of the root-mean-square magnitude to a single component is

$$\sqrt{\frac{1}{\langle \cos^2 \theta \rangle}} = \sqrt{3}. \quad (8)$$

This derivation is elementary and self-contained.

The strongest cross-check is that this same $\sqrt{3}$ is a measured modulus of the seam surface itself. The Oloid’s developable generator—the line segment along which it contacts a plane as it rolls—has the exact constant length $\sqrt{3}r$ for circles of radius r , a classical result of Oloid geometry [33, 34] that we have verified directly from the convex-hull construction. The strain factor $\sqrt{3}$ is therefore not only the isotropy L^2 norm of Eq. (8) but coincides with the contact-line modulus of the very surface on which the seam lives; the two routes give the same number, and the stiffness $\beta^{-1} = \sqrt{3}\pi$ inherits a geometrically grounded $\sqrt{3}$.

As a further, weaker cross-check, the same value is the circumradius-to-midradius ratio of the regular tetrahedron, the minimal three-dimensional simplex,

$$\frac{R_v}{R_e} = \frac{L\sqrt{6}/4}{L\sqrt{2}/4} = \sqrt{3}, \quad (9)$$

which is exact. We caution, however, that the interpretation of this ratio as a “stress-tensor eigenvalue” is heuristic: the tetrahedron admits the ratios $\{1/3, 1/\sqrt{3}, \sqrt{3}, 3\}$, and both the selection of R_v/R_e and the identification with a stress eigenvalue are asserted rather than derived. We therefore take the isotropy argument (8) and the verified Oloid contact-line modulus as the primary basis for $\sqrt{3}$, and Eq. (9) as a consistent but secondary cross-check.

2.5 The stiffness $\beta^{-1} = \sqrt{3}\pi$

The stiffness is the accumulation of the strain density along the field’s circulation path. The Q_μ flow traverses a semicircular phase interval $\Delta\phi = \pi$ —the transition between the Oloid’s two base circles—and the constant strain density $\sqrt{3}$ is integrated over it:

$$\beta^{-1} = \int_0^\pi \sqrt{3} d\phi = \sqrt{3}\pi \equiv E_{\text{node}}. \quad (10)$$

Here $\sqrt{3}$ is a constant coefficient pulled outside the integral (it does not vary along the seam) and $\int_0^\pi d\phi = \pi$ is the path length; no fibred internal space or cross-term arises, since $\sqrt{3}$ is a coefficient, not a separate metric direction.

The accumulation is linear. Equation (10) accumulates the strain density to the *first* power. The following subsection shows that this is not a modeling choice but is forced by the conformal (scale-free) character of the tree-level seam theory; the alternative quadratic accumulation,

$$\int_0^\pi (\sqrt{3})^2 d\phi = 3\pi \implies \beta = \frac{1}{3\pi}, \quad \alpha^{-1} = 24\pi^2 \approx 237, \quad (11)$$

is excluded. We note in passing that this is *not* a Bogomolny/BPS self-duality property: $Q_\mu = \partial_\mu \Phi$ is a phase gradient, and the Bogomolny bound saturates at an amplitude-independent winding number, a different statement from the linear accumulation at issue here.

2.6 Why the accumulation is linear, not quadratic

The single nontrivial feature of Eq. (10) is the power: why is the strain density $\sqrt{3}$ accumulated linearly, giving $\alpha^{-1} = 8\sqrt{3}\pi^2 \approx 137$, rather than quadratically as an energy density, which would give $\alpha^{-1} = 24\pi^2 \approx 237$ [Eq. (11)]? Naively the action density $\frac{1}{2}\beta|Q|^2$ is quadratic in the field, so one might expect the node cost to scale as $|Q|^2$. We show that the opposite is forced, by an argument that uses only two facts, both established independently: (i) the tree-level seam theory is conformal (Section 6), hence scale-free; and (ii) α^{-1} is a dimensionless count of informational nodes (the defining interpretation of the framework, and in any case α^{-1} is a pure number).

The argument is the following. A node count is a *dimensionless* functional of the seam configuration. In a scale-free theory the only configuration data available are dimensionless: the phase winding and the geometry. Two scale-free functionals can serve as the node count, and both are linear in the field amplitude:

- the *topological winding* $\oint d\Phi$, which counts homotopy sectors and is linear in $|Q| = |\partial\Phi|$ by construction; and
- the *information-geometric length* $\int ds$ with $ds = \sqrt{\beta}|Q|$, the geodesic distance in the metric whose squared line element is the action density, $ds^2 = \beta|Q|^2$.

The information-geometric reading is the natural one for this framework: the action density $\frac{1}{2}\beta|Q|^2 = \frac{1}{2}\beta g_{ab} d\Phi^a d\Phi^b$ is the squared line element of the seam's information metric (with β the metric coefficient), and by the Fisher–Rao/Fubini–Study distinguishability theorem [26, 27] the number of distinguishable states along a path equals its metric *length* $\int ds$ —the square root of the quadratic form, hence one power of $|Q|$, not two. Both scale-free readings therefore give the linear law of Eq. (10), with the three-dimensional isotropy of Section 2.4 entering the length as the L^2 norm of the three equal components, $\sqrt{1^2 + 1^2 + 1^2} = \sqrt{3}$.

The quadratic alternative is excluded precisely because the tree theory is scale-free. To use the energy $|Q|^2$ as a node count one must convert an energy into a pure number, which requires dividing by an energy gap or cutoff; the conformal seam theory ($c = 1$ free boson, Section 6) is gapless and supplies no such scale, so the energy reading does not define a dimensionless count at tree level at all. The 237 value is thus not a competing branch but an ill-defined one: it presupposes a scale the scale-invariant theory does not possess. A full statement, including the cross-check that the spectral (Weyl) mode count $N \sim |Q|^2$ fails for the same reason, is given in Appendix C.

2.7 Assembly

Combining $\Phi_{\text{total}} = 8\pi$ from (4) with $\beta = 1/(\sqrt{3}\pi)$ from (10),

$$\alpha_{\text{tree}}^{-1} = \frac{\Phi_{\text{total}}}{\beta} = \frac{8\pi}{1/(\sqrt{3}\pi)} = 8\sqrt{3}\pi^2 \approx 136.757, \quad (12)$$

a relative deviation of -2.0×10^{-3} from the measured $\alpha^{-1} = 137.035\,999\,177(21)$ [35]. Every quantity in Eq. (12) is geometric and none is adjusted to reproduce α . The residual two-parts-in-a-thousand is closed by the one-loop correction of Section 3; we will see that it is not a free counterterm but is itself fixed by independently derived data.

3 The One-Loop (Dyson) Correction

The tree-level value (12), $8\sqrt{3}\pi^2 \approx 136.757$, falls short of the measured $\alpha^{-1} \approx 137.036$ by $\Delta \approx 0.279$, a relative -2.0×10^{-3} . We now show that this gap is closed by a one-loop correction whose coefficient is not a free counterterm but is fixed by data derived independently in Papers III and X.

3.1 The self-consistent form

The correction takes a self-consistent (Dyson) form, with α appearing on both sides as in any self-energy resummation $\alpha_{\text{full}} = \alpha_{\text{tree}} + \text{loops}(\alpha_{\text{full}})$:

$$\frac{1}{\alpha} = 8\sqrt{3}\pi^2 + \left(\frac{1}{\alpha_{\text{kink}}} + \frac{1}{2\pi} \right) \pi \alpha = 8\sqrt{3}\pi^2 + c \pi \alpha, \quad c = 12 + \frac{1}{2\pi}. \quad (13)$$

Equivalently, $c \pi \alpha^2 + 8\sqrt{3}\pi^2 \alpha - 1 = 0$. The coefficient c has two pieces, from two distinct spectral quantities of the same seam Dirac operator: a quadratic piece $12 = \alpha_{\text{kink}}^{-1}$ (the Wess–Zumino–Witten vertex) and a linear piece $\frac{1}{2\pi}$ (the Atiyah–Patodi–Singer index boundary term). We derive each, then explain why the form is self-consistent rather than a constant shift, and finally distinguish the two spectral objects involved.

3.2 The quadratic piece $c = 12$: the WZW vertex

The photon couples to the electromagnetic current of the GKO coset. In the Paper X normalization, the EM two-point current coefficient is

$$K_{\text{EM}} = \sin^2 \theta_W k_{\text{SU}(2)} + \cos^2 \theta_W k_Y = \frac{1}{4}(1) + \frac{3}{4}(\frac{1}{2}) = \frac{5}{8}. \quad (14)$$

We stress that $K_{\text{EM}} = \frac{5}{8}$ is the current *normalization*, not the Dyson coefficient. The coefficient is the *vertex* factor, the inverse Schwinger–Dyson coupling at the seam fixed point, proved in Paper X [7, 22, 8, 9, 10]. It is fixed by the conformal weight of the doublet primary, not by the weak mixing angle:

$$\frac{1}{\alpha_{\text{kink}}} = \frac{k_{\text{SU}(2)} + h_{\text{SU}(2)}^{\vee}}{h_{1/2}^{(1)}} = \frac{1+2}{1/4} = 12, \quad h_{1/2}^{(1)} = \frac{j(j+1)}{k+2} \Big|_{j=1/2, k=1} = \frac{1}{4}, \quad (15)$$

where $h_{1/2}^{(1)} = \frac{1}{4}$ is the Knizhnik–Zamolodchikov conformal weight of the spin- $\frac{1}{2}$ primary of $\text{SU}(2)_1$, a robust rational-CFT datum [11]. Numerically this weight coincides with the tree-level weak mixing angle $\sin^2 \theta_W = \frac{1}{4}$ of Paper X, but the two are *independent* quantities with different origins— $h_{1/2}^{(1)}$ is a conformal weight at $k = 1$, while $\sin^2 \theta_W$ is a generation-summed Callan–Harvey level ratio (Section 5.4). The vertex, and hence the α^{-1} prediction, rests on the conformal weight alone; it does *not* depend on the value or the derivation of $\sin^2 \theta_W$. In the self-energy $\Pi = (\text{vertex}) \times (\text{propagator}) \times (\text{loop})$, the vertex carries the ultraviolet coupling $\alpha_{\text{kink}}^{-1} = 12$ while the loop propagator carries the infrared coupling α , so that

$$\Pi_{\text{quad}} = \frac{1}{\alpha_{\text{kink}}} \cdot \pi \cdot \alpha = 12\pi\alpha, \quad (16)$$

where π is the half-Oloid Gauss–Bonnet phase of Section 2.3. The same result follows from the two-dimensional Schwinger enhancement $(e^2/\pi) \times (k + h^{\vee}) \times \pi = 3e^2 = 12\pi\alpha$, using $k + h^{\vee} = 3$ and $e^2 = 4\pi\alpha$.

3.3 The linear piece $\frac{1}{2\pi}$: the APS index term

Paper III [3] supplies the APS index of the seam Dirac operator together with the Callan–Harvey anomaly-inflow mechanism [12, 13],

$$\text{Ind}(\mathcal{D}) \approx \frac{1}{2\pi} \oint_{\Sigma} A ds + \frac{1}{2} \eta(0). \quad (17)$$

Evaluated on the photon holonomy $a = \frac{1}{2\pi} \oint A$, the boundary term contributes

$$\Pi_{\text{lin}} = \frac{1}{2\pi} \times \pi \times \alpha = \frac{\alpha}{2}, \quad (18)$$

i.e. the $\frac{1}{2\pi}$ addition to c , again with the half-Oloid phase π .

3.4 Why the correction is self-consistent

The seam carries chiral zero modes (Callan–Harvey). Their exact one-dimensional Dirac determinant in a holonomy background $a = \frac{1}{2\pi} \oint A ds$ is

$$W[a] = -\log[2 \sin(\pi a)] = -\log(2\pi a) + \frac{\pi^2}{6} a^2 + \mathcal{O}(a^4). \quad (19)$$

The coefficient $\pi^2/6$ is the universal one-dimensional CFT Casimir constant—fixed, not adjustable. Because the holonomy a is sourced by the α -coupled photon, $a \propto \sqrt{\alpha} A$, the determinant’s contribution is *linear* in α ; this is what makes the correction a Dyson term $\Pi = c\pi\alpha$ rather than a constant shift.

3.5 Two distinct spectral objects

It is essential not to conflate the objects involved. The coefficient $c = 12 + \frac{1}{2\pi}$ does *not* come from the determinant’s Casimir $\pi^2/6$. It is built from two different spectral quantities of the same seam operator: the *vertex* $12 = \alpha_{\text{kink}}^{-1}$ (a coupling, from the WZW data of Paper X) and the *index* boundary term $\frac{1}{2\pi}$ (a zero-mode/spectral-asymmetry quantity, from Paper III). The determinant (the product of all eigenvalues, whose expansion yields the Casimir) and the index (a count of zero modes, or spectral asymmetry) are genuinely different objects; the determinant’s role here is solely to establish that the correction is linear in α , not to fix the coefficient. Identifying the Casimir $\pi^2/6$ with the index $\frac{1}{2\pi}$ would be an error.

For the same reason we set aside a constant alternative. The shift $\Delta(\alpha^{-1}) = \pi^2/36 \approx 0.274$ also matches numerically (to -3.4×10^{-5}) and factors as $\frac{\pi^2}{6} \cdot \frac{1}{4} \cdot \frac{2}{3}$, but the factor $\frac{2}{3}$ does not survive scrutiny: the seam is a one-dimensional curve, whose isotropic projection factor is $\frac{1}{3}$, not $\frac{2}{3}$, which would give $\pi^2/72$ and undershoot. The value $\pi^2/36$ is therefore a numerical snapshot of $12\pi\alpha$ at the physical α , not an independent mechanism; the Dyson form (13) is primary.

3.6 Result

Solving the quadratic (13):

$$c = 12 : \quad \alpha^{-1} = 137.032 \, (-2.7 \times 10^{-5}), \quad c = 12 + \frac{1}{2\pi} : \quad \alpha^{-1} = 137.036 \, (+3 \times 10^{-6}). \quad (20)$$

The index term sharpens the agreement by nearly two orders of magnitude. An equivalent closed form, treating the correction as a constant rather than self-consistently, factors the whole result through a single $\sqrt{3}$,

$$\frac{1}{\alpha} = \sqrt{3} \left(8\pi^2 + \frac{1}{2\pi} \right) \approx 137.033 \quad (-2.3 \times 10^{-5}), \quad (21)$$

with $8\pi^2$ (phase area) and $\frac{1}{2\pi}$ (anomaly boundary term) inside. Both pieces of c are thus grounded in proved IGPS structure (Papers III and X); the only remaining technical step is the half-Oloid loop integral that supplies the Gauss–Bonnet phase π , a geometric computation rather than new physics. This closes the question (XI-G6) of whether the APS term enters the photon self-energy.

A caveat on the “phase = 8” shortcut. An alternative route to the loop factor requires a phase ≈ 7.984 , close to the gluon/phase-area number 8. However, 7.984 differs from 8 by the same 0.20% as the tree-level gap, so setting it to 8 re-encodes that gap rather than deriving it independently. We note the coincidence as structurally suggestive but do not rely on it.

4 Zero Parameters: the Epistemic Claim

The agreement of Eq. (13) with experiment is meaningful only to the extent that nothing in it was adjusted to reproduce α . We make that claim precise, explain why it cannot be evaded by hidden fitting, and—to hold ourselves to a single standard—record the near-coincidences that were considered and rejected by the same rule.

4.1 Accounting

Every quantity entering the final formula is fixed before α is computed:

- $\Phi_{\text{total}} = 8\pi$ is topological (Section 2.1);
- $\sqrt{3}$ is the isotropy strain factor (Section 2.4);
- $\beta = 1/(\sqrt{3}\pi)$ follows from $\sqrt{3}$ and the linear accumulation of the stiffness, which is itself a consequence of the Distinguishability Coupling Principle—the count being a Bures length, not an integrated information (Theorem C.1)—rather than a separate assumption (Section 2.5);
- $12 = \alpha_{\text{kink}}^{-1} = (k_{\text{SU}(2)} + h_{\text{SU}(2)}^{\vee})/h_{1/2}^{(1)}$ is the WZW vertex of Paper X, fixed by the $\text{SU}(2)_1$ conformal weight $h_{1/2}^{(1)} = \frac{1}{4}$ (not by $\sin^2 \theta_W$; Section 3.2);
- $\frac{1}{2\pi}$ is the APS index boundary term of Paper III (Section 3.3).

None is a free parameter, and none is tuned against the measured α . Given the Distinguishability Coupling Principle (Appendix C.4)—the one physical axiom, from which the linear (length) form follows as a theorem—the full result (13) is therefore a zero-parameter consequence, with nothing fitted to α . Two distinct claims must, however, be kept separate, because they have very different strengths.

What is tested: the stiffness, at a knife edge. The tree term $8\sqrt{3}\pi^2$ supplies 99.8% of α^{-1} , so the prediction is acutely sensitive to the geometry-derived stiffness β . Even allowing the correction coefficient c to range over the entire band that reproduces α^{-1} to 10^{-4} , the (un-fitted) β must still be correct to $\sim 0.02\%$ for any in-band c to match; pinning the headline 3×10^{-6} would require β correct to $\sim 3 \times 10^{-4}\%$. A geometry-fixed β landing on that knife edge—and the tree alone already reproducing α^{-1} to 0.2% with no inputs—is the genuine, high-sensitivity evidence the result offers. This evidence is for β (equivalently, the tree).

What is not a precision test: the correction coefficient. The correction $c\pi\alpha$ contributes only the residual 0.2%, and $\pi\alpha \approx 0.023$ is small, so the data constrain c only loosely: the window reproducing α^{-1} to 10^{-4} is $c \in [11.6, 12.8]$, within which several simple values (e.g. 12, $\frac{37}{3}$, 4π) all “work.” The headline 3×10^{-6} is reached only because $c = 12 + \frac{1}{2\pi}$ happens to hit the second decimal, and the data cannot, on their own, distinguish that forced value from a nearby one. Each ingredient of c is nonetheless *derived*, not fit: 12 is the WZW vertex (Paper X) and $\frac{1}{2\pi}$ the APS index boundary term (Paper III), from independent sources. And the relation is now known to be *exact* at the level of the photon two-point function: the seam being the $c = 1$ free boson of Section 6.1, the $O(\alpha^2)$ correction to the coupling vanishes identically (Appendix E, by the Schwinger-model exactness of the vacuum polarisation). The 3×10^{-6} is therefore a *well-defined* prediction with no dropped higher order—but it is a *low-sensitivity* test of the correction coefficient, not a high-precision one. We state it as such and do not claim it as a 3×10^{-6} test of the framework.

4.2 Why the result is not reverse-engineerable

The decisive feature is the absence of a dimensionful knob. A single dimensionful parameter can be tuned to hit any one target: a mechanism of the form $m = \sqrt{\sigma/\kappa}$ or $m = M_{\text{Pl}} e^{-N}$ reproduces any single mass by choice of σ or N , and dressing that knob as a “pure number” is empty, because the space of simple expressions built from π , $\sqrt{3}$, α^{-1} , and small integers is dense at the percent level—one can match almost any target to $\sim 1\%$ and announce a “derivation.” This is precisely why absolute mass values, which always retain one such anchor (Section 6), cannot serve as evidence for or against a framework.

The constant α is different. Equation (13) contains no free coefficient into which a discrepancy could be absorbed; every symbol is fixed externally. The force of this is sharpest where the data are sensitive: the tree term reproduces α^{-1} to 0.2% with no adjustable quantity, and the geometry-fixed β sits on a knife edge (Section 4.1)—a mistuned β would miss by a thousandfold the residual it is supposed to leave. That cannot be reproduced by an incorrect theory through tuning, because there is nothing to tune. This is the strongest epistemic point of the paper, and it is the reason we place the dimensionless α^{-1} result—not any absolute mass—at the centre of the claim. We are careful, per Section 4.1, to rest this on the tree-level/ β sensitivity rather than on the 3×10^{-6} figure, which (while a well-defined, truncation-free prediction) tests the correction coefficient only loosely.

4.3 Rejected near-coincidences (a single, uniform standard)

We apply one rule throughout: a numerical match counts only if it follows from a forced derivation, never because it is close. Several attractive near-coincidences fail this test and are recorded here, set aside, precisely so that the standard is visibly uniform. Numerical closeness is *not* listed as a point in their favour.

For completeness we also discarded three structural attempts: the formula $87\pi/2 \approx 136.66$ (it requires $I_{\text{SM}} = 29/4$, obtainable only by weighting leptons and quarks inconsistently); a probability-amplitude argument $P_{\text{space}} \times P_{\text{topo}}$ (circular—it *defines* β to equal the product rather than deriving it); and integer “snapping” to $\lfloor 8\sqrt{3}\pi^2 \rfloor = 137$ (no dynamical equation forces the rounding, so it is numerology unless a WZW level or Chern number is shown to equal 137). Listing these alongside the kept results is the point: the same criterion that admits $8\sqrt{3}\pi^2$, 12, and $\frac{1}{2\pi}$ is what excludes the candidates above, and it is applied without regard to numerical proximity.

Candidate	Value / close- ness	Why it is rejected
$\sqrt{m_c m_b}$ (a mass scale)	2.30 GeV (−24% vs tar- get 3.02)	Mixes generation and type; no IGPS mecha- nism selects the geometric mean. Far off in any case.
$24\sqrt{3} = 8\sqrt{3} N_c$	+0.3%	Inside the $\pm 4.6\%$ α_s band, so not even a sharp target; and the photon-sector factor $8\sqrt{3}$ is imported into another sector with no forced reason—post-hoc.
$2\sqrt{3}$ (warp ratio)	matches kL to 0.1% <i>only with the wrong pref- actor</i>	Arises from a Goldberger–Wise stabilization with a $1/\epsilon^2$ prefactor; the correct prefactor is $4/\epsilon^2$ [19], which gives $kL = 147$, not the required ≈ 74 —wrong by a factor of two.
$29/4 \approx 4\pi/\sqrt{3}$	+0.07%	The numerator 29 has no IGPS derivation; 29/4 is merely the best denominator-4 ratio- nal approximant to $4\pi/\sqrt{3}$ (continued frac- tion $[7; 3, 1, 11, \dots]$) and cannot be made ex- act ($48\pi/(29\sqrt{3}) \approx 3.0022 \neq 3$).

Table 2: Near-coincidences considered and rejected by the uniform standard (no forced derivation), independent of how close they are. None plays any role in the α result.

5 The Strong Coupling α_s

The same WZW level structure that fixes the electroweak vertex $\alpha_{\text{kink}}^{-1} = 12$ (Section 3.2) gives a clean analogue for the strong sector. The dimensionless coupling at the seam fixed point is fully determined; the absolute scale, as for Λ_{QCD} in the Standard Model, is not. We separate the two carefully.

5.1 The kink value

For $\text{SU}(3)$ at level $k = 1$, the dual Coxeter number is $h^\vee = 3$, so the strong-coupling vertex at the kink scale is

$$\boxed{\frac{1}{\alpha_{s,\text{kink}}} = k + h^\vee(\text{SU}(3)) = 1 + 3 = 4} \quad (22)$$

exactly parallel to the electroweak case $\alpha_{\text{kink}}^{-1} = (k_{\text{SU}(2)} + h_{\text{SU}(2)}^\vee)/h_{1/2}^{(1)} = 3/(1/4) = 12$, where $h_{1/2}^{(1)} = \frac{1}{4}$ is the conformal weight of the $\text{SU}(2)_1$ doublet primary (Section 3.2). A clean corollary is the kink-scale ratio

$$\left. \frac{\alpha_s}{\alpha} \right|_{\text{kink}} = \frac{12}{4} = 3 = N_c, \quad (23)$$

the number of colours. Both Eq. (22) and the ratio $N_c = 3$ follow from the WZW levels with no free parameters.

5.2 Interpretation as a conformal fixed point

The value $\alpha_{s,\text{kink}}^{-1} = 4$ is the inverse Sugawara coupling of the $\text{SU}(3)_1$ current algebra at the seam fixed point. With the standard WZW normalization $g_*^2 = 4\pi/(k + h^\vee)$ [11] and $\alpha_s = g^2/4\pi$,

$$\frac{1}{\alpha_{s,*}} = \frac{4\pi}{g_*^2} = k + h^\vee = 4, \quad (24)$$

the same identification (coupling = inverse Sugawara level) that gives the electroweak vertex $\alpha_{\text{kink}}^{-1} = 12$ (Section 3.2), applied to the strong sector. This is the strong-sector analogue of how $\alpha^{-1} = 8\sqrt{3}\pi^2$ is the scale-free topological value for electromagnetism: being a conformal theory (Section 6), the seam fixes the coupling *value* but is scale-invariant and carries no intrinsic GeV scale.

5.3 What is open: the absolute scale

Connecting Eq. (22) to the observable $\alpha_s(m_Z) = 0.1179$ requires the SU(3) kink *scale* in GeV, which current IGPS does not determine. The electroweak sector is anchored at a physical kink—the Higgs at $m_H = 125$ GeV—but QCD has no Higgs: its scale Λ_{QCD} is generated by dimensional transmutation [20] and requires one dimensionful input. This mirrors the ordinary Standard Model, in which $\alpha_s(m_Z)$ (equivalently Λ_{QCD}) is measured, not derived from group theory. The seam CFT fixes the coupling value (4); being scale-invariant, it cannot fix the absolute scale.

Given $\alpha_s(m_Z)$, the crossover $\mu_{\text{color}} \approx 3$ GeV—where four-dimensional running meets $\alpha_s^{-1} = 4$, i.e. $\alpha_{s,4\text{D}}(\mu_{\text{color}}) = \alpha_{s,*} = \frac{1}{4}$ —is located. We emphasise that this is a genuine, non-circular matching point forced by the CFT value 4: it marks where the geometric seam description hands off to running QCD, and it is *located* only once $\alpha_s(m_Z)$ is supplied, not predicted from it. Numerical near-coincidences for μ_{color} (for instance $m_H/(24\sqrt{3})$, which lands inside the $\pm 4.6\%$ α_s band) are set aside by the uniform standard of Section 4.3; none follows from a forced derivation.

In summary, the dimensionless content— $\alpha_{s,\text{kink}}^{-1} = 4$ and the ratio $N_c = 3$ —is derived, while the absolute scale is open (XI-G4), a dimensional-transmutation input exactly as in the Standard Model.

5.4 Status of the weak mixing angle

Paper X derives the tree-level weak mixing angle as a Callan–Harvey level ratio,

$$\sin^2 \theta_W = \frac{k_Y}{k_Y + k_{\text{SU}(2)}} = \frac{1/2}{1/2 + 3/2} = \frac{1}{4}, \quad (25)$$

equivalently $\theta_W = 30^\circ$. We record its status carefully, because the present work depends on it only in a limited way, and the relation carries a known gap.

It is a tree-level value, 8% from experiment. The measured effective angle is $\sin^2 \theta_W \simeq 0.231$ at m_Z , so Eq. (25) is high by 8.1%. Within the Standard Model the running of $\sin^2 \theta_W$ between the electroweak scale and a few TeV is only $\sim 1\%$, so the offset is not removed by the renormalisation flow available near the kink; the value $\frac{1}{4}$ would coincide with the running angle only around ~ 4 TeV, far above the Higgs kink scale m_H at which Paper X’s other tree-level relations (notably $\lambda_H = 31/240$) match to $\lesssim 0.2\%$. The 8% discrepancy is therefore a genuine open item, not an artifact resolved by an RG-running argument; we flag it as such rather than presenting Eq. (25) as a precision success.

The coupling axiom does not improve it, and this is informative. It is natural to ask whether the Distinguishability Coupling Principle that fixes α^{-1} (Appendix C.4) also fixes the weak angle. Applied to the electroweak sector it would read $\sin^2 \theta_W = L_Y/(L_Y + L_2)$, a ratio of the Bures (distinguishability) lengths of the hypercharge and weak-isospin currents, with $L_a \propto \sqrt{\langle Q_a^2 \rangle}$. Evaluating this requires the Standard-Model charge assignments Y, T_3 of the fermion multiplet as input, and the result is normalisation dependent: the generation-summed ratio gives $\sin^2 \theta_W \simeq 0.56$ (Standard-Model hypercharge normalisation) or 0.50 (grand-unified

normalisation), matching neither the measured 0.231 nor the level-ratio $\frac{1}{4}$. The axiom therefore does *not* improve the tree value. We regard this as structurally informative rather than as a failure: α^{-1} is the length of a single abelian U(1) charge loop, for which the seam geometry supplies the value with no Standard-Model input, whereas the weak angle is a ratio across a non-abelian multiplet whose charge content is itself Standard-Model input the seam does not generate. The success at α^{-1} reflects this distinction—a single geometric distinguishability length—rather than a general ability to fix electroweak mixing.

The α^{-1} prediction does not depend on it. Crucially, the electromagnetic vertex $\alpha_{\text{kink}}^{-1} = 12$ of Section 3.2 is fixed by the conformal weight $h_{1/2}^{(1)} = \frac{1}{4}$ of the $\text{SU}(2)_1$ doublet primary—a Knizhnik–Zamolodchikov datum at level $k = 1$ —and *not* by the weak mixing angle. The two equal $\frac{1}{4}$ numerically but arise from different constructions (a $k = 1$ conformal weight versus a generation-summed level ratio). Consequently the 8% gap in $\sin^2 \theta_W$ does not propagate into the α^{-1} result: the 3×10^{-6} relation of Section 3 would stand unchanged even if the weak-mixing derivation were revised. We therefore treat $\sin^2 \theta_W = \frac{1}{4}$ as a separate, weaker result of the framework (an approximate relation, $\theta_W = 30^\circ$, good to 8%), and do not place it alongside the α^{-1} prediction as a comparable success.

6 Conformal Structure and the Absolute-Scale Question

We now give the structural reason that the framework fixes dimensionless quantities but not absolute scales, and we examine—honestly—how far a scale-generation mechanism can presently be carried. The first two subsections are results; the third is an open problem stated as such.

6.1 The tree theory is a two-dimensional conformal field theory

The IGPS field is $Q_\mu = \partial_\mu \Phi$, the gradient of a U(1) phase, with effective action

$$S_{\text{eff}} = \int_{\Sigma} \frac{1}{2} \beta |Q_\mu|^2 d\Sigma. \quad (26)$$

Dimensionally, $[|Q_\mu|^2] = 1/\text{length}^2$ and $[d\Sigma] = \text{length}^2$, so the length dimensions cancel and S_{eff} is scale-invariant. Three independent signals confirm this: a phase gradient carries no intrinsic scale; on a developable surface ($K_G = 0$) the equation of motion reduces to the flat two-dimensional Laplacian $\Delta \Phi = 0$, which is conformally invariant; and $\frac{1}{2} \beta |\partial \Phi|^2$ on a two-dimensional surface is precisely the two-dimensional free-boson CFT ($c = 1$, exactly marginal) [11]. The tree-level seam theory is therefore a two-dimensional conformal field theory, which by definition has no intrinsic mass scale.

6.2 Why dimensionless data is derivable and absolute scales are not

This single fact organizes the programme, once two kinds of scale-free datum are distinguished. A conformal theory fixes *intrinsic* dimensionless data—central charges, anomalous dimensions, conformal weights, fixed-point couplings, topological phases—and these supply π (the current normalization), the conformal weight $h_{1/2}^{(1)} = \frac{1}{4}$ behind the vertex 12, the quartic $\lambda_H = 31/240$, and $\alpha_{s,\text{kink}}^{-1} = 4$. The remaining ingredients $\sqrt{3}$ and 8π are not intrinsic $c = 1$ outputs but *geometric moduli* of the Oloid sigma-model on which the seam CFT lives—the contact-line modulus and the surface area (Appendix C.5, primitive (C))—and they too are scale-free, which is why a conformal framework can carry them. The value α^{-1} then follows from these moduli *under the Distinguishability Coupling Principle* (Appendix C.4), as the Bures length of the charge loop, rather than from the conformal structure alone; the tree-level $\sin^2 \theta_W = \frac{1}{4}$ the CFT fixes in form though it lands 8% from the measured value (Section 5.4). What conformal

invariance secures is the broad fact that all of this content is *dimensionless*: the framework succeeds precisely on the quantities a conformal seam theory, together with its sigma-model geometry and the one coupling postulate, *can* fix.

Conversely, mass enters only through Paper I’s [1] $m \sim (\hbar/c) R_{\text{node}}^{-1} \Xi(\delta, \beta, \alpha)$, in which the correlation length R_{node} is the lone dimensionful quantity. A conformal theory cannot fix R_{node} : no scale-invariant dynamics produces a scale. The absolute-scale gaps (Sections 5.3 and 8) are thus not flaws to be patched but the inevitable signature of a conformal framework that requires exactly one dimensionful input—the same structure as the Standard Model, in which Λ_{QCD} (by dimensional transmutation) and the electroweak vacuum expectation value v are measured inputs, not group-theoretic outputs.

6.3 Where the one anchor might come from: the size modulus (open)

Conformal invariance leaves the overall Oloid size R a flat direction—a modulus, or breathing mode, the analogue of the radion of a warped extra dimension [18, 19]. An unstabilized modulus is an instability: the seam would either collapse ($R \rightarrow 0$) or run away ($R \rightarrow \infty$). The framework must therefore stabilize R , and the value R^* at which it does would be the sought absolute scale—obtained not by hand but from the requirement that the modulus be stable. The intuition that “the system cannot be stable without generating a scale” is sound and identifies a genuine requirement. We report how far it can currently be carried, and where it stops.

The classical terms do not stabilize. The natural candidate is a balance $V(R) = \sigma R + A/R$ between a growing and a falling term, with a stable minimum at $R^* = \sqrt{A/\sigma}$. The framework supplies the growing term—the seam worldsheet tension σ from $S_{\text{seam}} = \int d^2\xi \sqrt{-g} [\sigma + \dots]$ (Papers III/IV [3, 4, 21]), giving $E_\sigma \sim \sigma R$. But a careful computation of the winding-phase profile shows there is no A/R term. The phase winds n times around the seam—a vortex with $|\partial\Phi| \sim n/\rho$ in the transverse coordinate ρ —so the energy per unit seam length is the standard logarithmic vortex energy $dE/d\ell = \pi\beta n^2 \ln(R/a)$ (a = seam core thickness), and multiplying by the seam length $\sim 2\pi R$,

$$E_{\text{phase}} \sim 2\pi^2 \beta n^2 R \ln(R/a), \quad (27)$$

which *grows* as $R \ln R$. (A reading of Paper I’s mass formula as $\sim A/R$ assumed the transverse scale was R ; the vortex core fixes it at a , so the seam energy grows with the seam length rather than falling—the $1/R$ was an artifact of the wrong transverse scaling.) With both the tension and the winding energy growing,

$$V(R) \sim [\sigma + 2\pi^2 \beta n^2 \ln(R/a)] R + \text{const}, \quad \frac{dV}{dR} > 0, \quad (28)$$

is monotonically increasing: no interior minimum, and the modulus is driven to the cutoff $R \sim a$.

A falling term must be quantum, and its sign is not secured. A genuine R^* requires a falling ($\sim 1/R$) term, which can only be the quantum (Casimir) vacuum energy of the CFT [23]. A direct computation clarifies, but does not close, the situation.

- The *bulk* two-dimensional Casimir energy on the closed Oloid is governed by the conformal anomaly, which is topological ($\propto c\chi$, with $\chi = 2$ for the sphere-like Oloid) and hence *R-independent*: it gives a constant, not a $1/R$ term, and does not stabilize. (An earlier one-dimensional-circle estimate $c_{\text{eff}} = \theta^2 - \frac{1}{12} = \frac{1}{36}$ was an artifact of treating the seam as a 1D circle; it is not the bulk 2D value.)
- A genuine $1/R$ can come only from the one-dimensional modes on the seam loop—the Callan–Harvey chiral modes already present in the framework (Sections 3.3, 7). For a

chiral fermion twisted by θ , the Casimir coefficient is

$$c_{\text{eff}}(\theta) = \frac{1}{2} \left(\theta - \frac{1}{2} \right)^2 - \frac{1}{48}, \quad (29)$$

which is *negative* (attractive, no stabilization) for the structural twists $\theta = \frac{1}{2}$ ($\mathbb{Z}_2$) and $\theta = \frac{1}{3}$ ($\mathbb{Z}_3$), and positive only in other sectors (e.g. $\theta = 0$). The sign of the stabilizing term is thus not secured; it depends on the exact seam-CFT spectrum and on which twisted sector the modulus occupies.

- That sector selection is constrained by the WZW/coset data ($\theta \in \{0, \frac{1}{2}\}$ for $\text{SU}(2)_1$; $\{0, \frac{1}{3}, \frac{2}{3}\}$ for the $\mathbb{Z}_3$ orbifold) but is tied to the unresolved Standard-Model $\rightarrow$ seam operator identification (Section 7.1).

Status. The intuition is correct that the modulus must be stabilized and that doing so would generate the scale, and the mechanism would be a one-dimensional seam-mode Casimir balancing the tension σR . But neither the sign nor the coefficient of the required term is secured with present ingredients, and even were a minimum to form, the resulting value $\sim \sqrt{\sigma/c_{\text{eff}}}$ would still ride on the single dimensionful anchor σ (dimensional transmutation). The *spectrum* of masses is a separate matter, governed by Paper I’s δ -folding rather than by this mechanism—the naive winding scaling $A_n \propto n^2$ would give $m_\mu/m_e = \frac{1}{2}$, off by a factor of ~ 400 . We therefore record size-modulus stabilization as a genuine open problem (XI-G8): a well-posed mechanism with the right structure, not a solved derivation.

7 Structural By-Products

Beyond α , the seam geometry yields three dimensionless structural results. They are robust in the sense of Section 6—they concern counting and representation theory, not scales—and we grade each by how firmly it is established: a forced topological identity (Section 7.1), rigorous character sums (Section 7.2), and a unifying observation (Section 7.3).

7.1 Colour-generation unification: one $\mathbb{Z}_3$

The Standard Model leaves unexplained why there are *both* three colours and three generations. In IGPS these are one topological fact: the seam’s single $\mathbb{Z}_3$, shared between its knot topology and the $\text{SU}(3)$ gauge centre.

The knot relation forces the centre charge. The seam is the trefoil, the $(2,3)$ -torus knot, with knot group

$$\pi_1(S^3 \setminus \text{trefoil}) = \langle x, y \mid x^2 = y^3 \rangle. \quad (30)$$

A flat $\text{SU}(3)$ connection over the knot complement is a homomorphism ρ with $\rho(x)^2 = \rho(y)^3$. Projecting to the centre $Z(\text{SU}(3)) = \mathbb{Z}_3$, the relation becomes $2[x] = 3[y] = 0 \pmod{3}$; since $\gcd(2,3) = 1$ this forces $[x] = 0$ and leaves $[y]$ as the free $\mathbb{Z}_3$ generator. Because Paper VII [5] selects $\text{SU}(3)$ precisely *via* a non-trivial $\mathbb{Z}_3$, the existence of $\text{SU}(3)$ on the seam forces $[y] \neq 0$: the $\text{SU}(3)$ centre- $\mathbb{Z}_3$ is the trefoil’s y - $\mathbb{Z}_3$ —forced, not merely isomorphic. We note that this $\mathbb{Z}_3$ sector structure is robust to the precise identity of the flavour vertex algebra: Paper VIII [6] proves that the true coset $\mathcal{V}_{\text{IGPS}} = \text{SU}(3)_1^{\otimes 3}/\text{SU}(3)_3$ carries the same $\mathbb{Z}_3$ triality charges, fusion ring, and sector count as the $\text{SU}(3)_1$ approximation, so the counting used here survives the replacement of the $k = 1$ approximation by the exact coset.

The zero-mode count is forced by the Alexander polynomial. The trefoil’s Alexander polynomial is

$$\Delta(t) = t^2 - t + 1, \quad (31)$$

of degree 2, with roots the primitive sixth roots of unity $e^{\pm i\pi/3}$ [24]. The $\mathbb{Z}_3$ generator (the squared monodromy) acts on H_1 (rank 2 = $\deg \Delta$) with eigenvalues ω, ω^2 —the two non-trivial $\mathbb{Z}_3$ characters, one each—while the trivial character sits in $H_0 = \mathbb{Z}$. Hence

$$N_{\text{gen}} = \underbrace{\deg \Delta}_{H_1: \omega \oplus \omega^2 = 2} + \underbrace{1}_{H_0: \text{trivial}} = 3 = |\mathbb{Z}_3|, \quad (32)$$

neither 0 (H_0 always contributes the constant mode), 2 (the trivial mode is genuinely present), nor 6 (H_1 has rank exactly 2). For a general $(2, q)$ -torus knot $\deg \Delta = q - 1$, so $N_{\text{gen}} = q$; only the $(2, 3)$ -trefoil yields $3 = |\mathbb{Z}_3|$ matching $\text{SU}(3)$. Together with the three centre charges of the $\text{SU}(3)$ fundamental ($N_c = 3$),

$$N_c = N_{\text{gen}} = 3 = |\mathbb{Z}_3|, \quad (33)$$

a single topological origin.

Status. Both decisive steps are established: the centre- $\mathbb{Z}_3$ is forced to equal the trefoil- $\mathbb{Z}_3$, and the three zero modes populate its three sectors, one each, forced by Eq. (31) together with $H_0 = \mathbb{Z}$ —forced topological facts with no tuned coefficient. The one standing assumption is the operator-level identification of the seam Dirac zero modes with $H_0 \oplus H_1$ of the trefoil structure. Appendix E writes the seam Dirac operator explicitly (a flat bulk operator with an Atiyah–Patodi–Singer boundary condition at the seam) and reduces this assumption to a single precise geometric lemma: that the Oloid surface functions as a *Seifert surface* for the trefoil, so that its twisted homology equals the Alexander module rather than merely matching Betti numbers (the rank check $b_1 = 2g = 2 = \deg \Delta$ holds, giving $N_{\text{gen}} = 3$). The homological counting is rigorous and the operator is now explicit; this Seifert identification for the specific Oloid geometry is the remaining obligation (XI-G7). This is the strongest structural connection of the paper—the Standard Model’s colour–generation coincidence explained as one forced topological identity.

7.2 Unified spin-dependent symmetry: D_2 representation theory

Three symmetry factors used separately across the IGPS papers are one representation-theoretic fact about how the Oloid symmetry group D_2 (order 4, $\{e, R_x, R_y, R_z\}$) acts on fields of different spin. The effective multiplicity of a field is the coherent character sum over the symmetry it sees:

- **Spin-0 (scalar):** the trivial representation A , with $\sum_g \chi_A(g) = |D_2| = 4$.
- **Spin-1 (vector):** the three components span $B_1 \oplus B_2 \oplus B_3$, each with $\sum \chi = 0$, so a vector has no D_2 -invariant component; it survives only on the largest subgroup leaving a component invariant, $\mathbb{Z}_2 = \{e, R_x\}$, where $\chi = 2$. Factor $|\mathbb{Z}_2| = 2$.
- **Spin- $\frac{1}{2}$ (spinor):** transforms under the double cover Q_8 ; needing 4π to return to itself ($\pi_1(\text{SO}(3)) = \mathbb{Z}_2$), the doubling gives $4\pi \times 2 = 8\pi = \Phi_{\text{total}}$.

The shared linchpin is $\mathbb{Z}_2 \subset D_2$: the vector lives on it ($\times 2$), the spinor doubles through it ($4\pi \rightarrow 8\pi$), and the scalar contains it within the full D_2 ($\times 4$). This unifies the proved $\Phi_{\text{total}} = 8\pi$ (the foundation of the α result, Section 2.1), the photon $\times 2$, and the scalar $\times 4$ into one statement about D_2 representations and their double cover. The character sums are rigorous; the one caveat is that the spinor case combines a *phase* (4π) with the $\mathbb{Z}_2$ *counting* factor, whereas scalar and vector are pure counting factors—so the unification is of common origin (D_2 structure), not of identical algebraic form.

7.3 One seam operator, APS data by chirality

The generation count (Section 7.1) and the Dyson correction (Section 3) invoke the APS index of the *same* seam Dirac operator D_{seam} . Its index data $\text{Ind}(D_{\text{seam}}) = \frac{1}{2\pi} \oint A + \frac{1}{2} \eta(0)$ split according to the chirality of the matter coupling:

- **Chiral coupling** (the Standard-Model fermions, for the generations): the spectral asymmetry $\eta(0) \neq 0$ is essential, and the integer value is fixed to be exactly 3 by the Alexander-polynomial homology of Section 7.1 (not by rounding a near-integer analytic index; Paper III’s [3] $\text{Ind} \approx 2.98 \rightarrow 3$ estimate is heuristic, whereas the homological count is exact).
- **Vectorlike coupling** (the photon, for the self-energy): the spectrum is charge-conjugation symmetric, so $\eta(0) = 0$, leaving only the inflow term $\frac{1}{2\pi} \oint A$; on the photon holonomy with the Gauss–Bonnet factor $\oint \kappa_g ds = \pi$ (Section 2.3) this gives $\alpha/2$, the $\frac{1}{2\pi}$ piece of the Dyson coefficient.

One operator, one APS structure; the chirality of the coupling selects which piece appears. This is the softest of the three results: the factor $\frac{1}{2\pi}$ is also the universal holonomy normalization, so its appearance in both places is partly automatic. The non-trivial content is that the *same* D_{seam} underlies both calculations (Papers III and the present work), with chirality as the discriminator—a unification of source, not an identity. Appendix E writes D_{seam} explicitly—a flat bulk operator with an Atiyah–Patodi–Singer seam boundary condition—and shows that the $\eta(0) = 0$ (vectorlike) versus $\eta(0) \neq 0$ (chiral) split is fixed by whether the boundary holonomy is real or a complex cube root of unity.

The seam representation theory. Taken together, the three connections are facets of one object, the seam: its D_2 symmetry seen differently by each spin (Section 7.2), its $\mathbb{Z}_3$ topology giving both colour and generations (Section 7.1), and its Dirac operator underlying both the generation count and the α -correction, split by chirality (Section 7.3).

8 Open Problems

We collect the outstanding obligations in one place, each with its current status. We distinguish three kinds: closed results that retain a standard, stated technical assumption (marked *); partial results whose dimensionless content is established but whose absolute scale awaits the one anchor; and one genuinely open structural problem. Numbering follows the working record (there is no item G5; that slot held a numerical coincidence, since dismissed, Section 4.3).

XI-G1 (closed, modulo the dimensionless-count premise). This obligation, the linear-versus-quadratic accumulation, is discharged *constructively* (Section 2.6, Appendix C, especially the node-operator construction of Appendix C.5). The value $\sqrt{3}$ is the isotropy L^2 norm and, equivalently, the verified Oloid contact-line modulus $\sqrt{3}r$ [33, 34] (Section 2.4); the *power* to which it accumulates is forced as follows. The seam carries a $\widehat{\mathfrak{u}}(1)$ current algebra (the $c = 1$ free boson, Section 6). The node-count operator is the Bures arc-length of the charge-coherent loop,

$$N = \oint \sqrt{4 \langle j_0^2 \rangle_c} d\phi,$$

built entirely from the algebra; and because a count of *distinguishable* states is a state-space *distance* (Helstrom; Braunstein–Caves [29, 28]), it must be the Riemannian length $\oint \sqrt{F}$, not the quadratic Fisher information $\oint F$. The linear power is thus dictated by the metric axioms, and the quadratic (237) reading is excluded twice over: it is the Fisher information rather

	Obligation	Status
XI-G1	Derive (not posit) that β^{-1} accumulates the strain density $\sqrt{3}$ <i>linearly</i> rather than quadratically, giving $E_{\text{node}} = \sqrt{3}\pi$ and $\alpha^{-1} \approx 137$ rather than $24\pi^2 \approx 237$.	Closed*
XI-G2	Establish that $8\sqrt{3}\pi^2$ matches the infrared $\alpha(m_e)$ rather than a high-scale value.	Partially closed
XI-G3	Derive the hierarchy v/M_{Pl} .	Partial
XI-G4	Fix the SU(3) kink scale in GeV to predict $\alpha_s(m_Z)$.	Partial
XI-G6	Show the APS term $\frac{1}{2\pi} \oint A$ enters the photon self-energy, giving $c = 12 + \frac{1}{2\pi}$.	Closed**
XI-G7	Show the SU(3) centre- $\mathbb{Z}_3$ is forced to equal the trefoil- $\mathbb{Z}_3$, unifying $N_c = N_{\text{gen}} = 3$.	Closed*
XI-G8	Stabilize the size modulus to generate the absolute scale.	Open

Table 3: Outstanding obligations. */** denote closure modulo a standard, stated technical assumption (detailed below).

than a length, and in the gapless theory it is not even a finite pure number. The asterisk on “closed” marks the single physical axiom, stated openly as the Distinguishability Coupling Principle (Appendix C.4): that the strength of the electromagnetic interaction *is* the operational distinguishability of the coupled charge sectors. The Fisher form, the Bures metric, and the length measure are derived beneath it (Chentsov–Petz, an optimal-measurement selection, and the counting no-go), not posited. We have argued the principle itself is irreducible: no unitarity, gauge, positivity, or quantum-speed-limit principle forces the identification of the coupling with distinguishability, since each acts only within the distinguishability calculus. Granting this one postulate—together with the minimal charge-coherent accumulation family—the node operator, the linear measure, and hence the choice of $\alpha^{-1} \approx 137$ rather than 237 are forced, not selected among candidates. We distinguish carefully what this does and does not settle. The *linear form* (the no-go for 237, Theorem C.1) is a theorem and uses none of the numerical inputs. The *value* $8\sqrt{3}\pi^2$, by contrast, rests on three stated primitives recorded in the ledger of Appendix C.5: the entry of $\sqrt{3}$ and 8π as verified Oloid moduli (C), the discretization convention $ds_{\min} = 1$ that turns the Bures length into a count (D), and the identification $\sqrt{F} = 1/\beta = \sqrt{3}\pi$ of the Fisher root with the geometric stiffness (E)—the last importing $\sqrt{3}$ from the embedding geometry rather than from the scalar charge variance. Closing the value (deriving $\sqrt{F} = \sqrt{3}\pi$ from the algebra, and fixing $ds_{\min}$) remains open; the linearity does not wait on it.

XI-G2 (partially closed). The tree value is a topological (scale-free) quantity and therefore naturally matches the infrared, Thomson-limit $\alpha(m_e)$ rather than a running value at a high scale; the conformal structure (Section 6.1) makes this expectation precise. A full treatment of the matching remains to be written out.

XI-G3, XI-G4 (partial: dimensionless done, scale open). For the electroweak hierarchy the ratio m_H/v and the Planck anchor are in hand, but the hierarchy v/M_{Pl} itself is open; a Goldberger–Wise route was investigated and fails with the correct $4/\epsilon^2$ prefactor (Section 4.3). For the strong sector the dimensionless fixed-point value $\alpha_{s,\text{kink}}^{-1} = 4$ is derived, but the absolute GeV scale is open (Section 5.3). Both are the same issue—a conformal theory needs one

dimensionful input—and both would be addressed by a resolution of XI-G8.

XI-G6 (closed).** Both pieces of the Dyson coefficient $c = 12 + \frac{1}{2\pi}$ are grounded in proved IGPS structure: the vertex $12 = \alpha_{\text{kink}}^{-1}$ (Paper X [7]) and the APS index term $\frac{1}{2\pi}$ (Paper III [3]), as set out in Section 3. The one remaining technical step is the half-Oloid loop integral supplying the Gauss–Bonnet phase π , a geometric computation rather than new physics (**).

XI-G7 (closed*). The colour-generation identity $N_c = N_{\text{gen}} = 3$ follows from two forced topological facts (Section 7.1): the centre- $\mathbb{Z}_3$ equals the trefoil- $\mathbb{Z}_3$ (from $x^2 = y^3$ with $\gcd(2, 3) = 1$ and $\text{SU}(3)$ -existence), and the three zero modes populate its three sectors, one each, forced by $\Delta(t) = t^2 - t + 1$ with $H_0 = \mathbb{Z}$. The counting is rigorous. Appendix E now writes the seam Dirac operator explicitly (flat bulk plus an APS seam boundary condition), grounds the index formula and the $\eta = 0/\eta \neq 0$ chirality split in that operator, and reduces the standing assumption (*) to a single geometric lemma: that the Oloid surface is a Seifert surface for the trefoil seam, so that its twisted homology is the Alexander module (the rank check $b_1 = 2g = 2 = \deg \Delta$ holds). Establishing that Seifert identification for the specific Oloid geometry is the remaining step.

XI-G8 (open). Stabilizing the conformal size modulus would generate the absolute scale (Section 6.3). The mechanism would balance the seam tension σR against a quantum (Casimir) falling term, but the bulk two-dimensional Casimir is topological and R -independent, while the seam-mode Casimir has a sector-dependent sign that is negative for the structural $\mathbb{Z}_2$ and $\mathbb{Z}_3$ twists [Eq. (29)]. Neither the sign nor the coefficient is secured with present ingredients, and any generated value would still ride on the single anchor σ . This is the one genuinely open structural problem, tied to the Standard-Model $\rightarrow$ seam operator identification.

The character of the open items. None of the above is a load-bearing free parameter of the α result. XI-G1, the one foundational step of the tree-level value, has its *linear form* closed constructively—the node-count operator is the Bures arc-length of the seam $\widehat{u}(1)$ charge loop and the linear measure is forced by quantum distinguishability (Appendix C.5, Theorem C.1), modulo the framework’s own definition of α^{-1} as a node count; its *value* $8\sqrt{3}\pi^2$ rests on the stated geometric and discretization primitives (the Oloid moduli, $ds_{\min} = 1$, and the Fisher-root identification), whose algebraic derivation remains open; XI-G2, G6, and G7 are completions of results the evidence already supports; XI-G3, G4, and G8 concern the single dimensionful anchor that any conformal framework requires, exactly as the Standard Model takes Λ_{QCD} and v as inputs. The zero-parameter status of the dimensionless prediction (13) is therefore unaffected by what remains open.

9 Conclusion

We have shown that the fine-structure constant follows from the seam geometry of the single Oloid, under one explicitly stated physical postulate and with no parameters fitted to α ,

$$\frac{1}{\alpha} = 8\sqrt{3}\pi^2 + \left(12 + \frac{1}{2\pi}\right)\pi\alpha \approx 137.036,$$

agreeing with the measured value to a relative 3×10^{-6} . The tree-level $8\sqrt{3}\pi^2$ is fixed by the phase area $\Phi_{\text{total}} = 8\pi$ and the stiffness $\beta = 1/(\sqrt{3}\pi)$; the one-loop coefficient $c = 12 + \frac{1}{2\pi}$ combines the Wess–Zumino–Witten vertex with the Atiyah–Patodi–Singer index. Each ingredient is derived independently in Papers I–X [1, 3, 7], so that, given the postulate, nothing is tuned to α , and the knife-edge sensitivity of the formula makes the agreement evidence for the geometry rather than a coincidence permitted by a free parameter. We do not claim to derive α^{-1} from the seam

dynamics with no further input; we claim the sharper and more defensible thing—that a single, named, falsifiable physical principle, together with the independently-fixed Oloid geometry, yields the value. That principle, the Distinguishability Coupling Principle (Appendix C.4), is that the strength of an interaction *is* the operational distinguishability of the coupled sectors; from it the Fisher–Rao form, the Bures metric, and the linear (length) accumulation of the strain density follow as theorems (Chentsov–Petz, an optimal-measurement selection, and the counting no-go of Theorem C.1), rather than being separately assumed. The factor $\sqrt{3}$ is the three-dimensional isotropy norm and, equivalently, the verified Oloid contact-line modulus.

We have also identified the structural reason the framework fixes dimensionless quantities but not absolute scales: the tree-level seam theory is a two-dimensional free-boson conformal field theory. A CFT fixes dimensionless data; the intrinsic $c = 1$ data supply $\lambda_H = 31/240$ and $\alpha_{s,\text{kink}}^{-1} = 4$ directly, while α^{-1} follows from the conformal current normalization together with the Oloid geometric moduli *under the Distinguishability Coupling Principle* (with the tree-level $\sin^2 \theta_W = \frac{1}{4}$ fixed in form but 8% from experiment, flagged as open in Section 5.4)—while carrying no intrinsic mass scale, so one dimensionful anchor must remain, exactly as Λ_{QCD} and v are inputs in the Standard Model. This places the framework’s successes and its open scale questions on the same footing rather than treating the latter as defects.

Two open problems stand out, at opposite ends of difficulty. The sharpest concerns the *value* of the tree term: granting the Distinguishability Coupling Principle, the linear (length) form of the count is a theorem (Theorem C.1), but the identification of the Fisher root with the geometric stiffness, $\sqrt{F} = \sqrt{3}\pi$, and the discretization scale $ds_{\min}$ enter as stated primitives rather than as outputs of the seam algebra; deriving them—bridging the seam current algebra and the Oloid moduli—remains the immediate open obligation. The deepest is the Standard-Model $\rightarrow$ seam operator identification that underlies both the colour-generation result (the $*$ in XI-G7) and the size-modulus stabilization (XI-G8); resolving it would turn the strongest structural connection into a fully closed theorem and would determine whether the conformal size modulus can be stabilized to emit the one remaining scale.

Finally, the colour-generation unification— $N_c = N_{\text{gen}} = 3$ as one forced $\mathbb{Z}_3$ of the trefoil seam, with the count fixed by the Alexander polynomial $\Delta(t) = t^2 - t + 1$ —stands somewhat apart from the α derivation and is arguably the most striking single result here: the Standard Model’s twin appearance of the number three explained as a single topological identity. It may warrant a dedicated treatment. Taken together, the results of this paper are facets of one object—the seam—seen through its D_2 symmetry, its $\mathbb{Z}_3$ topology, and its Dirac operator: a seam representation theory from which the dimensionless constants of the Standard Model appear to follow, with a single dimensionful scale the one thing that geometry, here as elsewhere, does not supply.

A Numerical Summary

This appendix collects the quantities appearing in the paper in one place. The results are dimensionless and derived (zero-parameter); the absolute scales are not predicted but require one dimensionful anchor (Section 6), as discussed at the end of this appendix. The measured reference value is $\alpha^{-1} = 137.035\,999\,177(21)$ [35].

Geometric inputs and the fine-structure constant

Other dimensionless results (derived)

Absolute scales

The absolute mass scales are not predicted here. As established in Section 6, the conformal seam theory fixes dimensionless data but carries no intrinsic scale, so one dimensionful anchor

Quantity	Expression	Value	Note
Phase area	$\Phi_{\text{total}} = 8\pi$	25.133	topological
Stiffness	$\beta = 1/(\sqrt{3}\pi)$	0.183776	geometric
Tree value	$8\sqrt{3}\pi^2$	136.757	-2.0×10^{-3}
Closed form	$\sqrt{3}(8\pi^2 + \frac{1}{2\pi})$	137.033	-2.3×10^{-5}
Full (Dyson)	$8\sqrt{3}\pi^2 + (12 + \frac{1}{2\pi})\pi\alpha$	137.036	$+3 \times 10^{-6}$

Table 4: The fine-structure constant. Given the Distinguishability Coupling Principle (Appendix C.4), the full result is a zero-parameter consequence; the relative error is quoted conservatively (the bare arithmetic gives $+2.5 \times 10^{-8}$).

Quantity	Expression	Value	Note
Higgs quartic	$\lambda_H = 31/240$	0.12917	coset value
Higgs/VEV ratio	$m_H/v = \sqrt{2\lambda_H}$	0.508	-0.08% vs measured
Weak mixing	$\sin^2 \theta_W$	1/4	tree; $+8\%$ vs 0.231, open (§5.4)
Strong coupling	$\alpha_{s,\text{kink}}^{-1} = k + h^\vee$	4	IR fixed point
Coupling ratio	$\alpha_s/\alpha _{\text{kink}}$	$3 = N_c$	colours
Colour/generation	$N_c = N_{\text{gen}} = \mathbb{Z}_3 $	3	trefoil $\mathbb{Z}_3$

Table 5: Dimensionless quantities fixed by the seam CFT, each derived without a parameter tuned to it (Sections 5, 7).

must be supplied externally (as Λ_{QCD} and v are in the Standard Model). Any dimensionful quantity in the framework—for instance a five-dimensional Planck mass or warp factor in the gravity extension—is therefore *computed from* such an anchor (the direction is $v \rightarrow \dots$), not a parameter-free output, and the hierarchy v/M_{Pl} itself remains open (XI-G3, Section 8). We do not tabulate these here, to avoid any suggestion that the absolute scales are predicted; the dimensionless results of Tables 4 and 5 are the content of this paper.

B Gauss–Bonnet Gives $E_{\text{topo}} = \pi$

This appendix supplies the computation behind Eq. (7). Work on the $\mathbb{Z}_2$ fundamental domain Σ_+ , the half-Oloid, which is a developable surface ($K_G = 0$) with the topology of a disk, $\chi(\Sigma_+) = 1$, whose boundary is the half-seam. The Gauss–Bonnet theorem for a surface with boundary and corners reads

$$\int_{\Sigma_+} K_G dA + \oint_{\partial\Sigma_+} \kappa_g ds + \sum_i \theta_i = 2\pi \chi(\Sigma_+) = 2\pi. \quad (34)$$

Since $K_G = 0$, the area term drops, leaving the geodesic and corner terms; we compute each *independently* from the Oloid geometry, so that Eq. (34) becomes a check rather than a definition.

Geodesic-curvature term. The Oloid is generated by two unit circles in perpendicular planes, each through the other’s centre; the seam is the curve along which the convex hull’s two circular arcs meet. On a developable surface the geodesic curvature is preserved under unrolling to the plane, where the seam is a circular arc of radius R , so $\kappa_g = 1/R$. Over the $\mathbb{Z}_2$ half-domain

the arc is a semicircle (subtended angle π , length πR), giving

$$\oint_{\partial\Sigma_+} \kappa_g ds = \frac{1}{R} \cdot \pi R = \pi. \quad (35)$$

Corner term. The half-seam domain closes at the two cusps where the generating circles cross. Because those circles lie in perpendicular planes, they meet transversally at right angles, so the exterior angle at each cusp is $\pi/2$; with two cusps,

$$\sum_i \theta_i = 2 \cdot \frac{\pi}{2} = \pi. \quad (36)$$

Equations (35) and (36) are computed independently, and their sum $\pi + \pi = 2\pi = 2\pi\chi(\Sigma_+)$ satisfies Eq. (34): Gauss–Bonnet is verified, not assumed.

Energy calibration. The minimal energy of one $\mathbb{Z}_2$ winding on Σ_+ is identified with the magnitude of the geodesic term,

$$E_{\text{topo}} = \left| \oint_{\partial\Sigma_+} \kappa_g ds \right| = \pi, \quad (37)$$

which is Eq. (7). This identification of the winding energy with the geodesic-curvature integral is the seam-action normalization of Papers I and III [1, 3]—a physical input, not a corollary of Gauss–Bonnet; the geometric content of this appendix is the independent evaluation of Eqs. (35) and (36). $\square$

C Why the Accumulation is Linear, Not Quadratic

This appendix establishes the single nontrivial step of the tree-level result of Section 2.6: that the strain density $\sqrt{3}$ enters the stiffness β^{-1} to the first power,

$$\beta^{-1} = \int_0^\pi \sqrt{3} d\phi = \sqrt{3} \pi, \quad (38)$$

giving $\alpha^{-1} = 8\pi/\beta = 8\sqrt{3}\pi^2 \approx 136.76$, rather than quadratically,

$$\int_0^\pi (\sqrt{3})^2 d\phi = 3\pi \implies \beta = \frac{1}{3\pi}, \quad \alpha^{-1} = 8\pi/\beta = 24\pi^2 \approx 237. \quad (39)$$

Since the action density $\frac{1}{2}\beta|Q|^2$ is quadratic in the field, the quadratic reading (39) is the one a naive energetic accounting would select, and the burden is to show that it is *excluded*, not merely disfavoured. We give an argument that turns on a single structural fact—the conformal (scale-free) character of the tree theory, proved in Section 6—and we are explicit about the one interpretive hinge it rests on.

C.1 The node count is a dimensionless functional

The quantity α^{-1} is a pure number; in the present framework it is the count of distinguishable informational nodes that the phase budget $\Phi_{\text{total}} = 8\pi$ supports. Two features of this statement do the work. First, α^{-1} is *dimensionless*: whatever functional of the seam configuration computes it must return a pure number. Second, the tree-level seam theory is *scale-free*: it is a two-dimensional conformal field theory (the massless phase Φ with $\Delta\Phi = 0$, a $c = 1$ free boson; Section 6), and so possesses no intrinsic mass scale, length, or energy gap.

These two facts are jointly restrictive. A dimensionless functional of a scale-free configuration can depend only on dimensionless data: the phase winding and the (conformally invariant) geometry. It cannot depend on any energy or length, because none exists to form a dimensionless ratio.

C.2 The two scale-free count functionals are linear

There are two natural functionals that count configurations in a scale-free way, and both are linear in the field amplitude $|Q| = |\partial\Phi|$.

Topological winding. The homotopy class of the phase map is counted by the winding

$$W = \frac{1}{2\pi} \oint d\Phi, \quad (40)$$

an integer (an element of $\pi_1(\text{U}(1)) = \mathbb{Z}$) and manifestly scale-free. It is linear in $|Q|$ by construction: doubling the phase gradient along the loop doubles the enclosed winding. A count of winding sectors is therefore a first-power functional of the field.

Information-geometric length. This is the reading proper to an information-geometric framework. The action density is the squared line element of the seam's information metric,

$$\frac{1}{2}\beta|Q|^2 = \frac{1}{2}\beta g_{ab} d\Phi^a d\Phi^b = \frac{1}{2} ds^2, \quad ds = \sqrt{\beta}|Q|, \quad (41)$$

with β the metric coefficient. By the Fisher–Rao (classical) and Fubini–Study (quantum) distinguishability theorems [26, 27], the number of mutually distinguishable states along a path equals the path's *length* in this metric, $\int ds$. Crucially, a length is the square root of the quadratic form,

$$\int ds = \int \sqrt{\beta}|Q| d\phi, \quad (42)$$

hence *one* power of $|Q|$, not two. The distinguishability count is well-defined for a gapless theory: it requires only the metric, not a spectrum or a gap.

Both functionals thus yield the linear law (38). The coefficient $\sqrt{3}$ in the line element has a concrete geometric origin that we can name precisely. It is the L^2 norm of the three equal components of the isotropically oriented phase current,

$$\sqrt{1^2 + 1^2 + 1^2} = \sqrt{3}, \quad (43)$$

the metric (Riemannian, hence L^2) combination—not the arithmetic mean $\langle \cos \theta \rangle = \frac{1}{2}$ (an L^1 quantity, which would give $\alpha^{-1} = 16\pi^2 \approx 158$) and not the variance $\langle \cos^2 \theta \rangle^{-1} = 3$ (the squared, energetic combination of Eq. (39)). This same $\sqrt{3}$ is, moreover, a *measured modulus of the seam surface itself*: the Oloid's developable generator (its rolling-contact segment) has the exact constant length $\sqrt{3}r$ for circle radius r [33, 34], a fact we have verified directly from the convex-hull geometry. The isotropy L^2 norm and the Oloid contact length are thus the same number arrived at two ways, and the stiffness $\beta = 1/(\sqrt{3}\pi)$ inherits its $\sqrt{3}$ from this geometric datum rather than from any fitted constant.

C.3 The quadratic reading is excluded by scale-invariance

The reading that would give 237 takes the node cost to be the energy $|Q|^2$ itself. To convert an energy into a dimensionless count one must divide by an energy scale—a gap Δ , so that $N = E/\Delta$, or a cutoff. The conformal tree theory is gapless ($\Delta \rightarrow 0$) and supplies no cutoff, so $N = E/\Delta$ is not a finite number: the energy reading does not define a node count at tree level at all. The value 237 is therefore not a competing branch that the framework happens not to select; it is ill-defined in the scale-free theory, presupposing a scale the theory does not possess.

We can package these observations into a single statement that excludes the quadratic reading outright, rather than refuting it case by case.

Theorem C.1 (No-go for the quadratic count). *Let $\gamma : \phi \mapsto \rho(\phi)$, $\phi \in [0, \Phi_{\text{total}}]$, be the charge-accumulation path of the tree-level seam (Eq. (47)), with the loop homogeneous (the tree theory is translation-invariant along the seam, so no point of γ is distinguished). Let $N[\gamma]$ be a functional that returns a dimensionless count of mutually distinguishable seam states, and require that N satisfy*

- (i) reparametrization invariance: N is unchanged under any orientation-preserving reparametrization $\phi \mapsto f(\phi)$ of the accumulation (a count cannot depend on the bookkeeping rate at which charge is tallied);
- (ii) scale-free dimensionlessness: N is a pure number built from the conformal data of the gapless $c = 1$ seam, with no intrinsic gap or cutoff available;
- (iii) the metric axioms: as a count of distinguishable states, N is a distance in state space (non-negative, symmetric, satisfying the triangle inequality), hence the Riemannian length of γ in the quantum Fisher metric F .

Then N is the Fisher–Rao (Bures) length

$$N[\gamma] = \oint \sqrt{F} d\phi \propto \oint |Q| d\phi, \quad (44)$$

linear in $|Q|$, and the quadratic functional $\oint F d\phi \propto \oint |Q|^2 d\phi$ (the reading that yields 237) satisfies none of (i)–(iii). In particular the quadratic count is excluded.

Proof. We show the quadratic functional violates each axiom and that the linear length is the unique survivor.

(i) *Reparametrization.* Along a fixed path the charge is $Q = d\Phi/d\phi$, so under $\phi \mapsto f(\phi)$ the one-form $|Q| d\phi = |d\Phi|$ is invariant, giving $\oint |Q| d\phi$ invariant. The quadratic integrand $|Q|^2 d\phi = (d\Phi)^2/d\phi$ carries an extra inverse differential and is not a one-form; explicitly, under a constant rescaling $\phi = \lambda\tau$ one finds $\oint |Q|^2 d\phi \mapsto \lambda \oint |Q|^2 d\phi$, and under a nonlinear reparametrization it changes by a path-dependent factor. (Both were verified directly, including for non-constant Q .) More generally, for a power-law measure $\oint |Q|^p d\phi$ the constant rescaling produces a factor λ^{1-p} , so $p = 1$ is the unique reparametrization-invariant power.

(ii) *Scale-freeness.* $|Q|^2$ has dimension of energy; to render $\oint |Q|^2$ a pure number one must divide by an energy scale, a gap Δ or cutoff. The gapless $c = 1$ theory has $\Delta \rightarrow 0$ and provides no cutoff, so the quadratic functional is not a finite dimensionless count. The length element $|Q| d\phi = |d\Phi|$ is already dimensionless and needs no such scale.

(iii) *Metric axioms.* By the Helstrom–Braunstein–Caves identification (Subsection C.5), a count of distinguishable states is the Bures distance, i.e. the Riemannian length $\oint \sqrt{F}$. The quadratic $\oint F$ is the integrated Fisher information, a curvature of the statistical manifold, not a length: it fails the triangle inequality and, as shown in (i), reparametrization invariance. Hence it is not a distance and cannot be a distinguishability count.

Uniqueness. Axioms (i) and (iii) force N to be a Riemannian length $\oint w(\phi) \sqrt{F} d\phi$ with a scalar weight w . Homogeneity of the loop (no distinguished point) together with scale-freeness (ii) forbids any non-constant w , so w is constant; absorbing it fixes $N = \sqrt{F} \Phi_{\text{total}}/\sqrt{F_0}$ -normalized to the Bures length (44). Thus the linear length is the unique functional satisfying (i)–(iii). $\square$

Remark C.2 (Why topological invariance is *not* among the axioms). *One might expect “topological invariance” on the list. It must be omitted: the Bures length (44) depends on the seam metric (it changes if the stiffness β changes) and is therefore not a topological invariant. Only the winding $\oint d\Phi/2\pi$ is topological; requiring topological invariance would exclude the Fisher length we are keeping, not just the quadratic reading. Topology enters the count not as a constraint on the measure but solely as the integration range $\Phi_{\text{total}} = 8\pi$ (the $\mathbb{Z}_2$ spin winding of Section 2.1).*

Likewise, plain additivity of the integral does not discriminate—both $\oint \sqrt{F}$ and $\oint F$ are additive under concatenation—so it is subsumed into the metric axioms (iii) as distance additivity, which only the length obeys.

Remark C.3 (What the theorem does and does not fix). *Theorem C.1 fixes the form of the count—a linear Bures length—and thereby the choice between 137 and 237. It does not fix the numerical value: the factors $\sqrt{3}$ (Oloid contact line), π ($\widehat{u}(1)$ normalization) and 8π (Oloid area range) are independent inputs entering through F and Φ_{total} (Subsection C.5). The premise “a count is a distinguishability distance” is thus not circular: it constrains the power, not the number, and $\alpha^{-1} = 8\sqrt{3}\pi^2$ is a separate computation downstream.*

Spectral (Weyl) cross-check. One might attempt to count nodes spectrally, via the Weyl mode-counting function $N(E) \sim (\text{Area}) E$ of the seam Laplacian. This route fails for the same reason and, instructively, fails *towards* the quadratic value: $N(E)$ is linear in the cutoff energy $E \sim |Q|^2$, so it would scale as $|Q|^2$, but it requires the cutoff E as external input, which the conformal theory does not provide. The spectral count is thus inapplicable at tree level; to the extent it could be forced, it reproduces the excluded energetic (39), not the linear (38). It provides no independent support for the linear law—a point we record rather than overstate. Likewise, a tunnelling/instanton action $\int \sqrt{2V} d\phi$ would be linear in the field, but it requires a potential $V(\Phi)$ with separated vacua; the massless phase has only gradient energy and no such potential, so this reading is inapplicable and is not invoked.

The honest content is therefore this: among functionals that are well-defined in the scale-free tree theory, the two available ones (winding and Fisher–Rao length) are both linear; the quadratic (energetic and spectral) readings all require a scale the conformal theory lacks. As Theorem C.1 makes precise, the requirements that the count be reparametrization-invariant, scale-free, and a genuine distance single out the linear Bures length and exclude the quadratic functional outright. Scale-invariance does not merely prefer the linear law—it removes the quadratic alternative from the space of well-defined tree-level counts.

C.4 The Distinguishability Coupling Principle

Before constructing the node operator we state, as plainly as possible, the one physical premise on which the entire identification of α^{-1} rests. We state it at the level of what an interaction *is*, rather than in terms of a particular metric: the specific information-geometric machinery (Fisher, Bures, the square-root length) is then *derived* beneath it by established classification theorems, as we show. We prefer to name the premise as an axiom rather than let it act silently as an “interpretation.”

Distinguishability Coupling Principle. *The strength of the electromagnetic interaction is the operational distinguishability of the coupled charge sectors: it counts how many charge states of the seam are discriminable, by optimal measurement, as charge accumulates around the loop.*

The principle mentions only charge sectors, measurement, and discriminability—no metric is presupposed. The Fisher–Rao form, the Bures metric, and the square-root (length) measure are consequences, each supplied by a standard theorem under a stated hypothesis. We make the chain and its riders explicit.

From distinguishability to the Fisher form (Chentsov–Petz). Take the seam charge sectors to be the smooth family $\rho(\phi)$ (Eq. (47)), a statistical manifold. *We assume that seam coarse-graining—the blocking or integrating-out of short-wavelength seam modes—acts as a charge-preserving completely positive, trace-preserving (CPTP) map.* This is physically natural

for the exact level-one $\hat{u}(1)$, but it is an assumption about the coarse-graining maps, not an automatic fact, and we flag it as such. Under it, the distinguishability metric must be monotone (non-increasing) under coarse-graining, and the classification theorems of Chentsov (classical) and Petz (quantum) [30, 31, 32] then fix the metric: any monotone Riemannian distinguishability measure is, up to an overall scale, a member of the Fisher–Rao family. The Fisher form $F = 4\langle j_0^2 \rangle_c$ is therefore *not* posited; it is forced, given the CPTP hypothesis, by monotonicity alone.

Two riders, stated openly. Two qualifications accompany this. First, Chentsov–Petz fix the metric only *up to an overall constant*; the constant is precisely the resolution $ds_{\min}$ of Eq. (51) (primitive (D)), which the classification does not determine. Second, Petz yields a whole *family* of monotone metrics; singling out the *Bures* (symmetric logarithmic derivative) metric requires the additional principle that the distinguishability is the one *achieved by optimal measurement*—the quantum Cramér–Rao/Helstrom bound [29, 28]. This optimal-measurement selection is a natural operational reading of “distinguishability,” but it is a rider, not a consequence of Petz’s theorem itself. With these two riders the metric is the Bures metric.

From a count to a length (the no-go). Finally, the Distinguishability Coupling Principle asks for a *count* of discriminable sectors, and a count of resolvable states along a path is the geodesic *length* in that metric, $\oint \sqrt{F} d\phi$, by Theorem C.1. Assembling the three steps,

$$\alpha^{-1} = \oint \sqrt{F} d\phi, \quad F = 4\langle j_0^2 \rangle_c, \quad (45)$$

is now a derived consequence of the principle, not an independent posit: the Fisher form from Chentsov–Petz (given CPTP coarse-graining), the Bures selection from optimal measurement, and the length from the counting no-go, with only the overall resolution $ds_{\min}$ left free.

Why the principle itself is irreducible. What remains genuinely axiomatic is the principle’s identification of the *coupling* with distinguishability—the statement that interaction strength simply *is* how discriminable the coupled sectors are. We examined whether a dynamical principle forces this: unitarity and the optical theorem, positivity, gauge invariance and the Ward identity, reparametrization invariance, and the quantum speed limit (Mandelstam–Tamm/Margolus–Levitin). Each constrains the distinguishability calculus itself; none forces the prior identification of the coupling with distinguishability. That identification is the irreducible physical content, and it sits at the level of what an interaction is—in the spirit of operational and relational reconstructions of physics, where a force is what its capacity to discriminate states amounts to. We adopt it as the single axiom and exhibit its consequences.

Relation to the standard vacuum polarization. The principle is sharp because it departs, in a definite and quantifiable way, from the textbook treatment. The quantum Fisher information is, up to normalization, the current–current two-point function at zero momentum—the vacuum polarization, $\Pi(0) = \langle jj \rangle_c = \frac{1}{4}F$. In standard quantum field theory the coefficient renormalizing F^2 , and hence α^{-1} , is the integrated polarization $\oint \Pi \propto \oint F$. The Distinguishability Coupling Principle instead yields the geodesic length $\oint \sqrt{F}$. The two are not a matter of convention: they are different functionals on the same statistical manifold and give numerically different answers, $8\sqrt{3}\pi^2 \approx 137$ for the length against a substantially larger value for the integral (Subsection C.3). The principle is thus *falsifiable*: it selects the length, and the measured $\alpha^{-1} \simeq 137.036$ is what discriminates. We do not present this agreement as a proof of the principle, but we record that it makes a definite, checkable choice that experiment favours.

Resolution, not capacity. We are careful about *which* information-theoretic quantity the principle names. The distinguishability count is a quantity of *resolution*: it counts how many charge sectors can be told apart by measurement around the seam, and as a count of resolvable states it is a geodesic length. This is logically distinct from the information *capacity* of the seam regarded as a channel—how much charge information it could transmit—which is a different functional of the same statistical data, accumulated in a different way and answering a different question. The Distinguishability Coupling Principle is stated at the level of resolution deliberately: it is the discriminability of charge states, not the carrying capacity of a channel, that the counting construction of Subsection C.5 requires. We do not conflate the two; the principle concerns resolution, and the length measure follows from it alone.

Consistency with the loop sector. This places the tree and loop sectors in a single language. The one-loop corrections of Appendix E are computed as genuine contributions to the photon two-point function Π (the Dyson coefficient $c = 12 + \frac{1}{2\pi}$ is literally read from the seam self-energy); the tree term is, by Eq. (45), the Bures length built from the same current data. Tree and loop are then two facets of one vacuum polarization on the seam—the tree as its square-root susceptibility, the loop as its standard corrections—rather than a geometric posit glued to a field-theoretic calculation.

C.5 The node operator from the seam current algebra

The preceding subsections establish that, among the functionals well-defined on the scale-free tree theory, the count is a length and is linear. Granting the Distinguishability Coupling Principle (Subsection C.4), we now strengthen this from a classification of admissible functionals into a *construction*: we exhibit the node-count operator explicitly, built from the seam current algebra, and show that the linear (length) reading is not merely the admissible one but the only one that is a count of distinguishable states. This removes the last sense in which the Fisher–Rao reading might be called a “best candidate” rather than a derived object.

The seam algebra. The tree-level seam is the $c = 1$ free boson of Section 6, whose chiral current $j(z) = i\partial\Phi(z)$ generates the $\widehat{u}(1)$ current algebra at level one,

$$[j_m, j_n] = m \delta_{m+n,0}, \quad Q = j_0, \quad (46)$$

with $Q = j_0$ the seam $U(1)$ charge and the vertex operators $V_q = :e^{iq\Phi}$: carrying charge q . This level-one algebra is the abelian $\widehat{u}(1)$ of the electromagnetic charge sector, and is distinct from the non-abelian flavour vertex algebra $\mathcal{V}_{\text{IGPS}}$ of Paper VIII [6] (the $c = 2$ coset $SU(3)_1^{\otimes 3}/SU(3)_3$); the present construction uses only the $c = 1$ free boson and is unaffected by the coset identity of the flavour sector. As charge accumulates around the seam loop, the physical configuration is a one-parameter family of seam states

$$\rho(\phi) = \frac{e^{i\phi Q} \rho_0 e^{-i\phi Q}}{\text{tr}(e^{i\phi Q} \rho_0 e^{-i\phi Q})}, \quad \phi \in [0, \Phi_{\text{total}}], \quad (47)$$

the charge-coherent family generated by Q , with ϕ the accumulated holonomy (the phase budget) and $\Phi_{\text{total}} = 8\pi$ its total range. Equation (47) is the natural and minimal choice: it is the orbit of the charge generator itself, requiring no datum beyond the algebra (46) and a reference state ρ_0 .

The count is forced to be a Bures arc-length. The framework defines α^{-1} as a count of *distinguishable* seam states (Subsection C.1). The operational content of “distinguishable” is not free: by the quantum detection theory of Helstrom [29] and the statistical-distance theorem

of Braunstein and Caves [28], the number of states that can be told apart by measurement along the family (47) is the Riemannian length of the curve $\rho(\phi)$ in the quantum Fisher–Rao (Bures) metric. Explicitly, the quantum Fisher information of the charge-coherent family is the connected charge variance read directly off the algebra,

$$F(\phi) = 4(\langle Q^2 \rangle_{\rho(\phi)} - \langle Q \rangle_{\rho(\phi)}^2) = 4 \langle j_0^2 \rangle_c, \quad (48)$$

and the distinguishable-state count is the Bures arc-length of the loop,

$$N = \oint ds_{\text{Bures}} = \oint \sqrt{F(\phi)} d\phi = \oint \sqrt{4 \langle j_0^2 \rangle_c} d\phi. \quad (49)$$

Equation (49) is the node-count operator. It is not posited: every ingredient—the generator $Q = j_0$, the variance $\langle j_0^2 \rangle_c$, the family (47)—is furnished by the seam $\hat{u}(1)$ algebra, and the square-root-of-variance form is dictated by the Helstrom–Braunstein–Caves identification of distinguishability with Bures length.

Why linear and not quadratic, at the operator level. The same curve $\rho(\phi)$ admits two distinct functionals,

$$\underbrace{N = \oint \sqrt{F} d\phi}_{\text{Bures length}} \quad \text{versus} \quad \underbrace{I = \oint F d\phi}_{\text{Fisher information}}, \quad (50)$$

the first linear in the accumulated charge, the second quadratic. A *count of distinguishable states* is a distance in state space: it must satisfy the metric axioms—non-negativity, the triangle inequality, and invariance under reparametrization of ϕ . Only the length $\oint \sqrt{F}$ does; the integrated information $\oint F$ is not a distance (it is not additive along concatenated paths and is not reparametrization-invariant), and so cannot be a count of states. This is the operator-level resolution of the “137 versus 237” question promised in Subsection C.3: the linear reading is forced because it is the unique functional of the family (47) that is a distinguishability count, while the quadratic reading $|Q|^2$ is the Fisher information, a curvature of the statistical manifold rather than a length along it. Topological winding (Subsection C.2) is likewise excluded as the *primary* count: a winding is a homotopy invariant, not a measurement-distinguishability count, and so does not answer “how many states can be told apart.” (It re-enters only as the integration *range* $\Phi_{\text{total}} = 8\pi$, below.)

Assembly of the value. Taking the accumulation to be homogeneous around the seam (the tree theory is translation-invariant along the loop), F is constant, so the count is

$$N = \frac{1}{ds_{\min}} \sqrt{F} \Phi_{\text{total}}, \quad (51)$$

where $ds_{\min}$ is the Bures separation between two just-resolvable states. Two inputs are needed to turn (51) into a number, and we are explicit that neither is delivered by the $\hat{u}(1)$ algebra alone:

- *Discretization.* The Helstrom–Braunstein–Caves theorems fix the Bures *length* $\sqrt{F} \Phi_{\text{total}}$ but not the resolution $ds_{\min}$ at which one node is counted per distinguishable step. We set $ds_{\min} = 1$, and we are precise about what this means. In the normalization used here, $F = 4 \langle j_0^2 \rangle_c$ is the standard symmetric-logarithmic-derivative quantum Fisher information of the unitary family (47), for which $F = 4 \text{Var}(Q)$ [28]; the line element $ds = \sqrt{F} d\phi$ then fixes one node at $d\phi = 1/\sqrt{F}$, which is exactly one standard deviation of the optimal single-shot estimator of ϕ —the quantum Cramér–Rao step $(\Delta\phi) \sqrt{F} = 1$ [29]. Thus $ds_{\min} = 1$ is the choice “one node per standard deviation of optimal distinguishability,” the

operational unit in which the count is expressed. We stress that this is a *unit of the count*, not a derived parameter: just as Shannon entropy fixes its unit by declaring one bit, and the Fisher–Rao length by declaring one nat, the node count fixes its unit by declaring one Cramér–Rao deviation. No equation of the theory derives why one deviation, rather than some multiple, individuates a node; we adopt the one-deviation unit and say so plainly. This freedom is not peculiar to the present construction: it is exactly the overall scale that the Chentsov–Petz classification leaves undetermined (the Fisher–Rao metric is unique only up to a constant), carried by every information-geometric quantity. It affects the *value* but not the linear *power* (Theorem C.1 never uses $ds_{\min}$). We list it as primitive (D) rather than dress it as derived.

- *The Fisher root.* We set

$$\sqrt{F} = \frac{1}{\beta} = \sqrt{3}\pi. \quad (52)$$

Here π is the $\hat{u}(1)$ current normalization, an intrinsic $c = 1$ datum, and $\Phi_{\text{total}} = 8\pi$ is twice the verified Oloid area $4\pi r^2$ (Section 2.3). The factor $\sqrt{3}$, however, is the Oloid contact-line / isotropy modulus of Eq. (43), and it enters (52) through the geometric stiffness β , *not* as an output of the charge variance $\langle j_0^2 \rangle_c$. We flag this distinction sharply, because it is easy to misstate: the stiffness factor $\sqrt{3}$ is a property of the *current* $Q_\mu = \partial_\mu \Phi$ (a spatial three-vector, whose isotropic L^2 norm is $\sqrt{3}$; Section 2.4), whereas the Fisher generator is the *scalar* charge j_0 , which has no spatial components and whose variance does not by itself produce a $\sqrt{3}$. Equation (52) is therefore an *identification* of the Fisher root with the geometric stiffness $1/\beta$, motivated and dimensionally consistent, but *not* a derivation of the value from the seam algebra: the $\sqrt{3}$ is imported from the embedding geometry (item (C) below), exactly as in Section 2.4.

With these two inputs,

$$N = \sqrt{F} \Phi_{\text{total}} = (\sqrt{3}\pi)(8\pi) = 8\sqrt{3}\pi^2 = \alpha_{\text{tree}}^{-1}, \quad (53)$$

the tree value of Section 2.3. The construction thus delivers the node operator and the linear *form* from the algebra, while the numerical *value* rests on the discretization convention $ds_{\min} = 1$ and the geometric identification (52); we record both plainly rather than presenting the value as algebra-derived.

Honest ledger. We separate what this construction derives from what it still assumes, so that no step is hidden:

Derived. (i) The node-count operator (49) exists and is built from the seam algebra; it is the Bures arc-length of the charge-coherent loop, with $F = 4\langle j_0^2 \rangle_c$. (ii) “Node = Bures-distinguishable state” is forced, not assumed: given that α^{-1} is a state *count*, the count must be a state-space distance, and quantum distinguishability *is* Bures distance (Helstrom; Braunstein–Caves). (iii) The linear power is forced over the quadratic: only $\oint \sqrt{F}$ is a distance, Eq. (50).

Assumed. (A) The Distinguishability Coupling Principle (Subsection C.4): that the interaction strength *is* the operational distinguishability of the coupled charge sectors. This is the single physical axiom; the Fisher form, the Bures metric, and the square-root (length) measure are *derived* beneath it—by Chentsov–Petz (given charge-preserving CPTP coarse-graining), by the optimal-measurement selection, and by the counting no-go (Theorem C.1) respectively—rather than posited. We have argued the principle itself is irreducible: no unitarity, gauge, positivity, or speed-limit principle forces the identification of the coupling with distinguishability, since all of those act within the distinguishability calculus, not on

what the coupling is. (B) That the physical accumulation is the charge-coherent family (47) generated by Q ; this is the minimal natural choice, but a different family would give a different F . (C) That the geometric moduli $\sqrt{3}$ and 8π are those of the Oloid. The seam is not an abstract $c = 1$ boson but a sigma-model whose worldsheet is the Oloid surface (Paper I [1]); j_0 , F , and the accumulation range are defined using that surface and its induced normalization. The phase budget $\Phi_{\text{total}} = 8\pi$ (twice the solid angle / area; Sec. 2.1) and the generator length $\sqrt{3}$ (the contact-line modulus / stiffness; Sec. 2.4) both enter F and ds through the geometry and the definition of the charge—not as outputs of the $c = 1$ current algebra. We are deliberate that the two stand on the same footing: although 8π is topological ($2 \times 4\pi$) and $\sqrt{3}$ metric, *neither* is produced by the worldsheet algebra, and the identification of 8π with the area, like that of $\sqrt{3}$ with the stiffness, is a feature of the Oloid sigma-model, forced only relative to the Paper I posit that the seam is the Oloid in $\mathbb{R}^3$. There is no theorem deriving either modulus from the $c = 1$ structure alone; this is the single residual geometric primitive, shared by both. (D) *Discretization*: the count equals the Bures length only under the convention $ds_{\text{min}} = 1$ (Eq. (51)); the theorems fix the length, not the one-node-per-unit-distance resolution, so this normalization sets the overall scale of the value. (E) *The Fisher root*: the identification $\sqrt{F} = 1/\beta = \sqrt{3}\pi$ (Eq. (52)) is not derived from the charge variance; the $\sqrt{3}$ is the geometric stiffness (a property of the current Q_μ , not of the scalar charge j_0 that generates F), entering through the charge normalization as in (C). The algebra fixes the *form* $F = 4\langle j_0^2 \rangle_c$, not its value.

Granting the framework’s definition (A) and the minimal state family (B), the node operator, the identity node = distinguishable state, and the linear measure are genuine consequences of the seam algebra together with quantum measurement theory. The numerical *value*, by contrast, rests on the further inputs (C)–(E): the geometric moduli $\sqrt{3}$ and 8π , the discretization $ds_{\text{min}} = 1$, and the identification of the Fisher root with the stiffness. This is the precise and honest sense in which $\alpha^{-1} = N_{\text{Fisher}}$ is established: a *theorem for its form*—a forced Bures arc-length built from $\hat{u}(1)$, linear by Theorem C.1—and a *geometric identification for its value*, in which $\sqrt{3}$, 8π , and the resolution enter as stated primitives rather than algebra outputs. The linearity, and hence the exclusion of 237, does not depend on any of (C)–(E); only the value does.

C.6 A correction to earlier framing

Previous drafts of this work described the linear accumulation as a “BPS versus Dirichlet” distinction, or attributed it to a stress–strain (Hooke) modelling choice in which the stiffness carries one power of the strain by fiat. Both descriptions are superseded. The field $Q_\mu = \partial_\mu \Phi$ is a U(1) phase gradient, not a self-dual soliton, and the Bogomolny construction saturates at an amplitude-independent winding number, a different statement from linear accumulation. And the linearity is not a posited Hooke response: it is forced, given scale-invariance, by the requirement that a dimensionless node count be a length or winding rather than an energy.

C.7 Status and the interpretive hinge

The argument closes the question of the power (linear versus quadratic) *conditional on a single interpretive premise*, which we state plainly: that α^{-1} is a dimensionless count of distinguishable informational nodes, rather than a raw energy expectation. This is the defining interpretation of α^{-1} in the present framework and is consistent with α^{-1} being, as a matter of fact, a pure number and the coefficient (not the value) of the gauge kinetic term; it is not an additional assumption imported from outside. Granting it, the linear law (38) is not merely the admissible reading but the forced one: as shown constructively in Subsection C.5, the count is the Bures arc-length $\oint \sqrt{4\langle j_0^2 \rangle_c} d\phi$ of the seam $\hat{u}(1)$ charge loop (49), the linear power is dictated by the metric axioms (Eq. (50)), the isotropy factor enters as the Oloid contact-line modulus $\sqrt{3}$,

and the energetic value 237 is excluded as both ill-defined in the gapless theory and not a distinguishability count.

This is a genuine advance over a purely eliminative reading. We are not merely observing that the Fisher–Rao length is the best-behaved candidate among several; we exhibit the node operator, derive its square-root (linear) form from quantum measurement theory, and identify each numerical factor. The honest ledger of Subsection C.5 records what remains genuinely primitive—the framework’s definition of α^{-1} as a state count (A), the choice of the charge-coherent accumulation family (B), and the entry of $\sqrt{3}$ and 8π as verified Oloid target moduli (C). None of these is a hidden energetic scale; together they place the linearity on a constructive rather than a comparative footing.

We record the residual honestly. If one rejected the premise and insisted that α^{-1} be computed as an energy expectation, the linear argument would not bind—but that reading does not yield a finite dimensionless count in the scale-free theory either, so it does not furnish a viable route to 237. Within the framework’s own definition of what α^{-1} is, the “magic switch” between 137 and 237 is resolved: the switch position is fixed by scale-invariance, and only the linear setting produces a well-defined tree-level count.

D The Dyson Self-Consistent Equation

The full result of Section 3 solves the quadratic (Dyson) equation

$$\frac{1}{\alpha} = 8\sqrt{3}\pi^2 + c\pi\alpha \iff c\pi\alpha^2 + 8\sqrt{3}\pi^2\alpha - 1 = 0, \quad (54)$$

whose physical (positive) root is

$$\alpha = \frac{-8\sqrt{3}\pi^2 + \sqrt{192\pi^4 + 4c\pi}}{2c\pi}. \quad (55)$$

This is a standard Dyson resummation, $\alpha_{\text{full}} = \alpha_{\text{tree}} + (\text{self-energy involving } \alpha_{\text{full}})$, with the self-energy coefficient c fixed externally (Section 3). Table 6 gives the solution of Eq. (55) as c takes its successive values.

c	α^{-1}	Relative error
0 (tree)	136.757	-2.0×10^{-3}
$12 = \alpha_{\text{kink}}^{-1}$	137.032	-2.7×10^{-5}
$12 + \frac{1}{2\pi}$	137.036	$+3 \times 10^{-6}$

Table 6: Solution of Eq. (55) for successive values of the self-energy coefficient, against $\alpha^{-1} = 137.035\,999\,177(21)$ [35]. The final row is the full prediction; its relative error is quoted conservatively (the bare arithmetic gives $+2.5 \times 10^{-8}$). Each value of c is fixed by independently derived data, not adjusted to α .

Sensitivity to the tree coefficient

The knife-edge claim can be made quantitative. Writing $I \equiv \alpha^{-1}$ and $A \equiv 8\sqrt{3}\pi^2$, Eq. (54) is $I^2 = AI + c\pi$. Differentiating at fixed c ,

$$2I \frac{dI}{dA} = I + A \frac{dI}{dA} \implies \frac{dI}{dA} = \frac{I}{2I - A}, \quad (56)$$

so the fractional sensitivity is

$$\frac{dI/I}{dA/A} = \frac{A}{2I - A} = \frac{8\sqrt{3}\pi^2}{2(137.036) - 8\sqrt{3}\pi^2} \approx 0.996. \quad (57)$$

A fractional change in the stiffness β (hence in A) thus produces an almost equal fractional change in α^{-1} . The 3×10^{-6} agreement of the full prediction therefore requires β to be correct to the same order, $\sim 3 \times 10^{-6}$ (i.e. $\sim 3 \times 10^{-4}\%$)—the knife-edge sensitivity invoked in Section 4.1, now as an explicit derivative rather than an estimate.

E APS Anomaly Inflow and the Dyson Coefficient

This appendix derives the Dyson coefficient $c = 12 + \frac{1}{2\pi}$ of Section 3 from the seam data of Paper III, and—because the point is easy to get wrong—states carefully which spectral object each piece comes from. Paper III establishes the APS index formula with a $\frac{1}{2\pi}$ boundary term,

$$\text{Ind}(\mathcal{D}) \approx \frac{1}{2\pi} \oint_{\Sigma} A ds + \frac{1}{2} \eta(0), \quad (58)$$

together with Callan–Harvey anomaly inflow (the bulk anomaly cancelled by the boundary η -invariant) and the explicit $U(1)$ inflow $S_{\text{bulk}} = \frac{g^2}{8\pi} \int_{\mathcal{M}} A \wedge F$ [3, 12, 13].

The two pieces of c

Linear (APS) piece. The boundary term $\frac{1}{2\pi} \oint A$, evaluated on the photon holonomy $a = \frac{1}{2\pi} \oint A$ and combined with the half-Oloid Gauss–Bonnet phase π (Appendix B), gives

$$\Delta\left(\frac{1}{\alpha}\right)_{\text{APS}} = \frac{1}{2\pi} \cdot \pi \cdot \alpha = \frac{\alpha}{2}, \quad (59)$$

so the APS contribution to c is $\frac{1}{2\pi}$.

Quadratic (WZW) piece. The photon couples to the GKO-coset electromagnetic current. In the Paper X [7] normalization the EM two-point coefficient is $K_{\text{EM}} = \sin^2 \theta_W k_{\text{SU}(2)} + \cos^2 \theta_W k_Y = \frac{5}{8}$; this is the current *normalization*, not the Dyson coefficient. The coefficient $c = 12$ is instead the *vertex* factor, the inverse Schwinger–Dyson coupling at the seam fixed point proved in Paper X. It is fixed by the conformal weight of the doublet primary, not by the weak mixing angle:

$$\frac{1}{\alpha_{\text{kink}}} = \frac{k_{\text{SU}(2)} + h_{\text{SU}(2)}^{\vee}}{h_{1/2}^{(1)}} = \frac{1 + 2}{1/4} = 12, \quad \Pi_{\text{quad}} = \frac{1}{\alpha_{\text{kink}}} \cdot \pi \cdot \alpha = 12\pi\alpha, \quad (60)$$

where $h_{1/2}^{(1)} = j(j+1)/(k+2)|_{j=1/2, k=1} = \frac{1}{4}$ is the Knizhnik–Zamolodchikov conformal weight of the $\text{SU}(2)_1$ doublet primary [6]. This weight coincides numerically with the tree-level $\sin^2 \theta_W = \frac{1}{4}$ but is an independent quantity (a $k = 1$ conformal weight, not a level ratio; see Section 5.4), so the vertex—and the α^{-1} result—does not depend on the weak-mixing derivation. The one remaining technical step is showing that the two-dimensional loop integral on the half-Oloid yields exactly the Gauss–Bonnet phase π —a geometric integral, not a new physical input.

The seam determinant: linearity only

The chiral zero modes on the one-dimensional seam give the exact Dirac determinant in the holonomy background,

$$W[a] = -\log[2 \sin(\pi a)] = -\log(2\pi a) + \frac{\pi^2}{6}a^2 + \dots \quad (61)$$

The holonomy is sourced by the photon: with $a = \frac{1}{2\pi} \oint A$ and the gauge field normalized by the electromagnetic coupling $e = \sqrt{4\pi\alpha}$, the holonomy carries exactly one power of the coupling,

$$a \propto e = \sqrt{4\pi\alpha} \implies a^2 \propto \alpha. \quad (62)$$

Hence the leading $\frac{\pi^2}{6}a^2$ term in Eq. (61) is *linear* in α —which is what makes the correction a Dyson term rather than a constant shift. This is the determinant’s entire role: it establishes linearity, it does *not* set the coefficient.

Index versus determinant (not to be conflated). The coefficient $c = 12 + \frac{1}{2\pi}$ comes from two *distinct* spectral quantities of the seam operator: the WZW vertex (12) and the APS *index* boundary term ($\frac{1}{2\pi}$). Neither is the determinant’s Casimir $\pi^2/6$. The determinant (a product of eigenvalues) and the index (a count of zero modes, or spectral asymmetry) are different objects of the same operator and must not be identified. In particular, the constant alternative $\Delta(\alpha^{-1}) = \pi^2/36$ is ruled out: it would require a projection factor $\frac{2}{3}$, whereas a one-dimensional curve projects with factor $\frac{1}{3}$, giving $\pi^2/72$, which undershoots.

An invalid higher-order attempt (recorded to prevent its re-derivation). One might try to gain precision from the determinant’s next term, $\frac{1}{2}\zeta(4)a^4$. Fixing $a(\alpha)$ by matching the a^2 term to the $\frac{1}{2\pi}$ coefficient produces a clean-looking $(1/20)\alpha^2$. This is invalid on two counts. First, the $\frac{1}{2\pi}$ is the APS *index*, not the determinant *Casimir*, so the normalization is matched to the wrong object. Second, and decisively, the a^4 term is a *four*-photon amplitude (light-by-light scattering), not a correction to the photon two-point function, and therefore does not renormalise α^{-1} at all. The genuine $O(\alpha^2)$ correction to the coupling comes from a true two-loop two-point function, and it *vanishes identically* (shown in the two-loop exactness subsection below): the free-boson seam gives a one-loop-exact vacuum polarisation. The residual $+3 \times 10^{-6}$ is thus not an artifact of a dropped next term—there is no such term—but it is also not pinned by the data at that precision (Section 4).

Structural parallel

The same seam operator supplies both the generation count and the α -correction, but through *different* pieces of its APS data, split by the chirality of the coupling (Section 7.3). The discriminator is the spectral asymmetry $\eta(0)$. For the photon the coupling is vectorlike (charge-conjugation symmetric), so the Dirac spectrum is symmetric under $E \rightarrow -E$; the spectral asymmetry then cancels in pairs and

$$\eta(0) = \lim_{s \rightarrow 0} \sum_{\lambda \neq 0} \text{sgn}(\lambda) |\lambda|^{-s} = 0, \quad (63)$$

leaving only the inflow term $\frac{1}{2\pi} \oint A$, which on the photon holonomy yields the $\frac{1}{2\pi}$ piece of the Dyson coefficient [Eq. (59)]. For the chiral Standard-Model fermions the spectrum is *not* $E \rightarrow -E$ symmetric, $\eta(0) \neq 0$, and the spectral asymmetry is essential to the generation count.

The generation count is an exact integer from topology. We emphasize that the generation count is fixed to be *exactly* 3 by homology, with no rounding of a near-integer analytic index. The trefoil’s Alexander polynomial $\Delta(t) = t^2 - t + 1$ has degree 2, so the $\mathbb{Z}_3$ monodromy acts on H_1 (rank 2) with the two non-trivial characters ω, ω^2 , while $H_0 = \mathbb{Z}$ carries the trivial character; the total is $N_{\text{gen}} = \deg \Delta + 1 = 3 = |\mathbb{Z}_3|$ (Section 7.1). This is an integer count, not an analytic index requiring an η -correction. We note that Paper III’s analytic-index estimate, $\text{Ind}(\mathcal{D}) \approx 2.98 + (\text{correction}) \rightarrow 3$, is heuristic—the value 2.98 is a placeholder, not a derived quantity, and the standard η -invariant of the stated trefoil spectrum $\lambda_n \propto (n + \frac{1}{3})$ is $\eta(0) = 1 - \frac{2}{3} = \frac{1}{3}$, which does not reproduce it. The rigorous statement is the homological one: the count is exactly 3.

Quantity	APS datum used	Result
Generation count	chiral, $\eta(0) \neq 0$ (homology fixes integer)	3
α^{-1} correction	vectorlike, $\eta(0) = 0$; vertex 12 + inflow $\frac{1}{2\pi}$	$c = 12 + \frac{1}{2\pi}$

Table 7: Both results draw on the same seam operator, through different pieces of its APS data: the generation count uses the chiral sector ($\eta(0) \neq 0$), with the integer value fixed rigorously by the Alexander-polynomial homology (Section 7.1); the α -correction uses the vectorlike sector ($\eta(0) = 0$), giving the WZW vertex 12 plus the inflow $\frac{1}{2\pi}$. The shared structure is the operator, not a single identical term.

Explicit construction of the seam Dirac operator

The preceding parallel invokes a single operator D_{seam} in two roles; a reader is entitled to see it written down rather than inferred from its index alone. We give the explicit operator and identify precisely what is rigorous and what remains a single geometric lemma (XI-G7).

The operator on the surface. The Oloid carrier is a developable surface: its Gaussian curvature vanishes away from the seam, so it is locally isometric to the plane and unrolls onto a flat patch with Cartesian coordinates (x, y) . On such a patch the zweibein is $e^a = dx^a$, the spin connection vanishes, and the Dirac operator is exactly the flat operator

$$D_{\text{seam}} = i \sigma^a \partial_a = i(\sigma^1 \partial_x + \sigma^2 \partial_y), \quad \text{Hermitian symbol } H(k) = \sigma^1 k_x + \sigma^2 k_y, \quad (64)$$

with $\sigma^{1,2}$ the Pauli matrices. In the chiral basis the off-diagonal blocks are $\partial_z, \partial_{\bar{z}}$, the Cauchy–Riemann operator. All nontrivial data live on the *seam*, where two ingredients enter: (i) the holonomy $U \in \text{Spin}(2)$ acquired when the unrolled edges are re-glued, and (ii) a conical curvature defect, since Gauss–Bonnet forces the total curvature of the closed developable to be concentrated on the seam. The operator (64) together with the spectral (Atiyah–Patodi–Singer) boundary condition set by U is a well-posed APS problem, and its index is exactly the formula used above,

$$\text{Ind}(D_{\text{seam}}) = \frac{1}{2\pi} \oint_{\Sigma} A ds + \frac{1}{2} \eta(0), \quad (65)$$

with A the boundary connection induced by U and η the asymmetry of the induced one-dimensional boundary operator. The “one-dimensional seam Dirac determinant” used in Section 3 is the determinant of precisely this boundary operator: the two-dimensional bulk and the one-dimensional boundary are the two natural strata of the *same* D_{seam} , not competing descriptions.

The chirality split, explicitly. Charge conjugation acts antiunitarily as $\mathcal{C} = \sigma^1 K$ (K complex conjugation). On the symbol it sends $H(k) \mapsto \sigma^1 H(-k)^* \sigma^1 = -H(k)$, so $\{\mathcal{C}, D_{\text{seam}}\} = 0$: the bulk operator is charge-conjugation symmetric. Whether this survives at the seam is fixed by the holonomy:

- **Vectorlike (photon).** The boundary holonomy is real ($U \in \{\pm 1\}$), commuting with the real structure $\mathcal{C}$. The spectrum is symmetric under $E \rightarrow -E$, so

$$\eta(0) = \lim_{s \rightarrow 0} \sum_{\lambda \neq 0} \text{sgn}(\lambda) |\lambda|^{-s} = 0,$$

leaving only the inflow term $\frac{1}{2\pi} \oint A$ of Eq. (65). This is the $\frac{1}{2\pi}$ piece of the Dyson coefficient.

- **Chiral (matter).** The boundary holonomy is a nontrivial cube root of unity $\omega = e^{2\pi i/3}$. Since $\mathcal{C}$ is antiunitary and $\bar{\omega} \neq \omega$, the boundary condition breaks the $E \rightarrow -E$ pairing; $\eta(0) \neq 0$ and the index acquires the $\frac{1}{2}\eta(0)$ term.

The single operator therefore produces $\eta(0) = 0$ or $\eta(0) \neq 0$ according only to whether the boundary holonomy is real or a complex cube root—the explicit mechanism behind the chirality split.

What remains: one geometric lemma (XI-G7). For the chiral sector the integer value of the count is the rank of the relevant twisted homology, which by the classical theorems of Milnor and Turaev [25, 24] is computed by the trefoil’s Alexander polynomial: $\dim H_1 = \deg \Delta = 2$, and with the constant mode $H_0 = \mathbb{Z}$ this gives $N_{\text{gen}} = 3$ (Section 7.1). By the Hodge theorem $\ker D_{\text{seam}}$ is isomorphic to this cohomology *once* D_{seam} is identified with the twisted Hodge–Dirac operator of the flat $\mathbb{Z}_3$ bundle. On a flat surface that identification is automatic; the one step that is *not* automatic, and that we state rather than prove, is that the Oloid surface functions as a *Seifert surface* for the trefoil seam—so that its twisted homology equals the knot’s Alexander module rather than merely some surface cohomology. This is consistent at the level of ranks: the trefoil has Seifert genus 1, hence a Seifert surface with $b_1 = 2g = 2 = \deg \Delta$, reproducing $N_{\text{gen}} = b_1 + 1 = 3$. Establishing the Seifert identification for the specific Oloid geometry—not merely matching Betti numbers—is the single remaining lemma recorded as XI-G7. The geometric operator (64), the index (65), and the η -split are explicit; the chiral count rests on this one geometric identification together with the (framework-level) input that the seam is the trefoil.

Two-loop exactness of the coupling correction

The Dyson correction $c\pi\alpha$ is a one-loop quantity. For the relation $\alpha^{-1} = 8\sqrt{3}\pi^2 + c\pi\alpha$ to be a *prediction* rather than the leading term of an uncontrolled series, the genuine $O(\alpha^2)$ (two-loop) correction to the coupling must be shown to vanish or to be negligibly small. We show here that it vanishes identically, by a direct computation of the seam fermion determinant.

The relevant theory. The photon couples to the *vectorlike* seam current—the $\eta(0) = 0$ sector of D_{seam} identified in Appendix E, not the chiral sector that fixes the generation count. With the seam carrying the $c = 1$ free boson of Section 6.1, the photon–seam system is the (vector) Schwinger model: a two-dimensional massless Dirac fermion minimally coupled to a U(1) gauge field,

$$\mathcal{L} = \bar{\psi} i \not{D} \psi, \quad \not{D} = \gamma^\mu (\partial_\mu - ieA_\mu), \quad (66)$$

equivalently, after bosonisation, a free boson whose conserved current $j^\mu = \frac{1}{\sqrt{\pi}} \epsilon^{\mu\nu} \partial_\nu \phi$ couples linearly to A_μ .

The determinant is exactly quadratic in A . Hodge-decompose the gauge field, $A_\mu = \partial_\mu \lambda + \epsilon_{\mu\nu} \partial^\nu \varphi$; the pure-gauge part λ decouples. The physical part φ is removed from the Dirac operator by a *local chiral rotation* $\psi \rightarrow e^{ie\gamma^5 \varphi} \psi$, which is possible because of the two-dimensional chiral duality

$$\gamma^\mu \gamma^5 = -i \epsilon^{\mu\nu} \gamma_\nu. \quad (67)$$

The rotation is anomalous: the Fujikawa Jacobian [16] is the two-dimensional axial anomaly, which is *linear* in A ,

$$\partial_\mu j_5^\mu = \frac{e}{\pi} \epsilon^{\mu\nu} \partial_\mu A_\nu. \quad (68)$$

Accumulating the Jacobian over the rotation parameter $s \in [0, 1]$ introduces one factor of the (linear) anomaly and one integration $\int_0^1 s ds = \frac{1}{2}$, producing an effective action that is *exactly quadratic* in A and terminates there,

$$\ln \det(i\mathcal{D}[A]) = \ln \det(i\mathcal{D}) + \frac{e^2}{2\pi} \int A_\mu \left(\delta^{\mu\nu} - \frac{\partial^\mu \partial^\nu}{\partial^2} \right) A_\nu. \quad (69)$$

This is Schwinger’s exact solution [14]: the photon acquires mass e^2/π and the vacuum polarisation $\Pi^{\mu\nu} = (e^2/\pi)(\delta^{\mu\nu} - q^\mu q^\nu/q^2)$ is the *complete* result.

Consequence: the two-loop correction vanishes. Equation (69) contains no term of order e^3 , e^4 , or higher. Hence every photon n -point function with $n \geq 3$ vanishes (in particular the a^4 four-photon term of Eq. (61) is light-by-light scattering and does not renormalise the coupling), and the photon two-point function—the correction to α^{-1} —receives *no* $O(\alpha^2)$ contribution:

$$\Delta(\alpha^{-1})_{2\text{-loop}} = 0. \quad (70)$$

The exactness has a precise origin: the Adler–Bardeen theorem [15] guarantees the anomaly coefficient in Eq. (68) is one-loop exact, and in two dimensions the entire determinant is reconstructed from the anomaly, so its one-loop-exactness propagates to the full effective action. The candidate $O(e^4)$ one-particle-irreducible graphs (the internal-photon rainbow and the vertex correction) therefore cancel by the Ward identity; Eq. (69) is the non-perturbative statement of that cancellation.

Status. The relation $\alpha^{-1} = 8\sqrt{3}\pi^2 + c\pi\alpha$ is consequently exact at the level of the photon two-point function: the α -series of the coupling correction terminates at the one-loop term, with no dropped $O(\alpha^2)$. The result rests on two structural inputs, stated plainly: (C1) the seam is exactly the $c = 1$ free boson of Section 6.1 *non-perturbatively* (a marginal or relevant self-coupling would spoil the Gaussianity and reintroduce higher-order terms; Section 6.1 establishes the free-boson structure at tree level, and the present result requires it to persist); and (C2) the photon couples to the vectorlike current, which is the $\eta(0) = 0$ sector of Appendix E. Neither is a tuning. We emphasise that this exactness concerns the *definedness* of the prediction (no truncation), and is logically separate from the *sensitivity* with which the data constrains the coefficient c , discussed in Section 4.

References

- [1] P. Ninsook, “Geometry and Spin Structure of the Single Oloid Manifold as the Origin of Leptonic Mass (Information-Geometric Physics System I),” <https://doi.org/10.5281/zenodo.18617881> (2026).
- [2] P. Ninsook, “Multi-Seam Configuration and the Topological Scaling of Baryonic Mass (Information-Geometric Physics System II),” <https://doi.org/10.5281/zenodo.18618316> (2026).
- [3] P. Ninsook, “Gauge Forces, Higgs Mechanism, and Particle Spectrum from Seam Topology (Information-Geometric Physics System III),” <https://doi.org/10.5281/zenodo.18655801> (2026).

- [4] P. Ninsook, “Topological Origin of Flavor Mixing and the Cabibbo Scale via Exact Path Integrals (Information-Geometric Physics System IV),” <https://doi.org/10.5281/zenodo.20131097> (2026).
- [5] P. Ninsook, “Quark Mass Hierarchies, CKM Mixing, and CP Violation from Spectral Seam Geometry (Information-Geometric Physics System VII),” <https://doi.org/10.5281/zenodo.20538222> (2026).
- [6] P. Ninsook, “The True Coset Structure of IGPS and the $\mathbb{Z}_3$ Replacement of the $k=1$ Approximation (Information-Geometric Physics System VIII),” <https://doi.org/10.5281/zenodo.20538369> (2026).
- [7] P. Ninsook, “Electroweak Symmetry Breaking and the Higgs from Spectral Worldvolume Geometry (Information-Geometric Physics System X),” <https://doi.org/10.5281/zenodo.20670826> (2026).
- [8] E. Witten, *Non-abelian bosonization in two dimensions*, Commun. Math. Phys. **92** (1984) 455–472.
- [9] P. Goddard, A. Kent, and D. Olive, *Virasoro algebras and coset space models*, Phys. Lett. B **152** (1985) 88–92.
- [10] P. Goddard, A. Kent, and D. Olive, *Unitary representations of the Virasoro and super-Virasoro algebras*, Commun. Math. Phys. **103** (1986) 105–119.
- [11] P. Di Francesco, P. Mathieu, and D. Sénéchal, *Conformal Field Theory*, Graduate Texts in Contemporary Physics, Springer-Verlag, New York, 1997.
- [12] M. F. Atiyah, V. K. Patodi, and I. M. Singer, *Spectral asymmetry and Riemannian geometry. I*, Math. Proc. Cambridge Philos. Soc. **77** (1975) 43–69.
- [13] C. G. Callan, Jr. and J. A. Harvey, *Anomalies and fermion zero modes on strings and domain walls*, Nucl. Phys. B **250** (1985) 427–436.
- [14] J. Schwinger, *Gauge invariance and mass. II*, Phys. Rev. **128** (1962) 2425–2429.
- [15] S. L. Adler and W. A. Bardeen, *Absence of higher-order corrections in the anomalous axial-vector divergence equation*, Phys. Rev. **182** (1969) 1517–1536.
- [16] K. Fujikawa, *Path-integral measure for gauge-invariant fermion theories*, Phys. Rev. Lett. **42** (1979) 1195–1198.
- [17] A. Connes, *Noncommutative Geometry*, Academic Press, San Diego, 1994.
- [18] L. Randall and R. Sundrum, *A large mass hierarchy from a small extra dimension*, Phys. Rev. Lett. **83** (1999) 3370–3373, [hep-ph/9905221](https://arxiv.org/abs/hep-ph/9905221).
- [19] W. D. Goldberger and M. B. Wise, *Modulus stabilization with bulk fields*, Phys. Rev. Lett. **83** (1999) 4922–4925, [hep-ph/9907447](https://arxiv.org/abs/hep-ph/9907447).
- [20] S. Coleman and E. Weinberg, *Radiative corrections as the origin of spontaneous symmetry breaking*, Phys. Rev. D **7** (1973) 1888–1910.
- [21] M. Born and L. Infeld, *Foundations of the new field theory*, Proc. R. Soc. Lond. A **144** (1934) 425–451.
- [22] J. Wess and B. Zumino, *Consequences of anomalous Ward identities*, Phys. Lett. B **37** (1971) 95–97.

- [23] H. B. G. Casimir, *On the attraction between two perfectly conducting plates*, Proc. K. Ned. Akad. Wet. **51** (1948) 793–795.
- [24] J. W. Alexander, *Topological invariants of knots and links*, Trans. Amer. Math. Soc. **30** (1928) 275–306.
- [25] J. Milnor, *A duality theorem for Reidemeister torsion*, Ann. of Math. **76** (1962) 137–147.
- [26] S. Amari, *Information Geometry and Its Applications*, Applied Mathematical Sciences **194**, Springer, Tokyo, 2016.
- [27] W. K. Wootters, *Statistical distance and Hilbert space*, Phys. Rev. D **23** (1981) 357–362.
- [28] S. L. Braunstein and C. M. Caves, *Statistical distance and the geometry of quantum states*, Phys. Rev. Lett. **72** (1994) 3439–3443.
- [29] C. W. Helstrom, *Quantum Detection and Estimation Theory*, Academic Press, New York, 1976.
- [30] N. N. Chentsov, *Statistical Decision Rules and Optimal Inference*, Translations of Mathematical Monographs **53**, American Mathematical Society, Providence, 1982.
- [31] E. A. Morozova and N. N. Chentsov, *Markov invariant geometry on state manifolds*, J. Soviet Math. **56** (1991) 2648–2669.
- [32] D. Petz, *Monotone metrics on matrix spaces*, Linear Algebra Appl. **244** (1996) 81–96.
- [33] H. Dirnböck and H. Stachel, *The development of the oloid*, J. Geom. Graph. **1** (1997), no. 2, 105–118.
- [34] S. R. Finch, *Convex hull of two orthogonal disks*, preprint (2012); see also U. Bäsel, *The mean width of the oloid and integral geometric applications of it*, arXiv:1604.07245 (2016).
- [35] E. Tiesinga, P. J. Mohr, D. B. Newell, and B. N. Taylor, *CODATA recommended values of the fundamental physical constants: 2022*, Rev. Mod. Phys. **97** (2025) 025002.