Toroidal fermions and Hopf-projected bosons in the square-root-mass geometry: topological criteria, no-go results, and falsifiers

Krzysztof Urbanowicz*
QWID Foundation
(Dated: July 18, 2026)

We ask two questions about the Single Quantum Uncertainty Sphere (SQUS) framework of the companion papers: can a toroidal configuration justify the defining properties of a fermion, and can a boson be a projection of that configuration? Both questions have sharp answers, but not the naive ones. A bare axisymmetric torus cannot be a fermion: its orientation space has trivial fundamental group, so the Finkelstein–Rubinstein mechanism has nothing to act on (Proposition 1). A torus decorated with phase windings (n, k) on its two cycles removes the continuous stabilizer, restores the $\mathbb{Z}_2$ rotation loop, and – through the identification with toroidal Hopf solitons of charge $Q_H = nk$ – admits fermionic quantization exactly when nk is odd. The boson channel is not a metaphorical projection but the Hopf map itself: the spinor bilinear kills the fiber phase carrying the 2π sign, and with it the exchange obstruction, so the projected channel obeys symmetric statistics and may condense. What SQUS adds to this established topology is the mass dictionary: the torus lives on the warped $\xi = \kappa\sqrt{m}$ geometry, which (i) forces spinor antiperiodicity at the horn tip ($U \to -I$ exactly), (ii) decomposes the measured spectral shift as $a = \Phi - \frac{1}{2}$ with a required flux $\Phi = 0.6726(68)$, compatible with 2/3 at 0.9σ, and (iii) predicts two competing radius laws, $f \propto m$ and $\lambda_C \propto 1/m$, whose crossing defines a distinguished mass scale m_*, and (iv) bounds the mass of any coherent mode: a winding must complete a circuit within the vacuum coherence time set by the Gibbons–Hawking fluctuations of the Λ-vacuum, $\Omega \gtrsim H$, giving the seesaw ceiling $m_{\max} = m_*^2/m_L \approx 10^{17}$ GeV and, from sub-Planckianity, forcing $m_* \leq 3.8$ meV – the neutrino window. The sustaining mechanism is a closed fluctuation–dissipation cycle on a vacuum comb: the loop quantizes the vacuum itself at integer winding, the exchange power is capped universally at $P_{\max} = (m_* c^2)^2/\hbar \approx 3\times 10^{-11}$ W, and the ceiling acquires a fourth formulation, $m_{\max} c^2 = P_{\max}\, L/c$, with supercritical modes relaxing onto it as an attractor. We prove internal no-go results (variance–tower incompatibility in the minimal stabilizer class; exclusion of the Hopf-charge mass tower by the $Q^{3/4}$ energy bound; exclusion of the 8–15–17 mixing ansatz at 6.8σ) and give a falsification schedule, including electron pointlikeness bounds and Pauli-violation null tests.

I. INTRODUCTION

The companion papers of this series develop a five-dimensional framework in which the square-root mass variable $\chi \equiv \sqrt{m}$ is a dimensionful fifth coordinate, $\xi = \kappa\chi$, the flavor sector is organized by the Koide/equipartition structure in the variables $\sqrt{m_i}$ [1], and the neutrino spectrum is a holonomy-shifted winding tower on a spatial SQUS loop [2]. The present paper addresses the constituent question: *what kind of object sits on that loop?* The working picture of the framework – a two-phase toroidal resonance for fermions, with bosons as single-phase or bilinear projection channels – is here confronted with the mathematics that already exists for extended objects with spin: the Finkelstein–Rubinstein (FR) quantization of solitons [3], the exchange–rotation theorem of Sorkin [4], the toroidal Hopf solitons of the Faddeev–Niemi model [5, 6], and their fermionic quantization by Krusch and Speight [8].

We use the rigor labels of the companion papers: [A] theorems and exact computations, [B] computations confronted with data, [C] empirical regularities without mechanism, [D] hypotheses and outlook.

* krzysztof.urbanowicz@qwid.org

Two popular intuitions are corrected at the outset, because both fail as stated and both have rigorous replacements. The first intuition is dimensional: "the torus fills three-dimensional space, therefore it is a fermion and a building block of matter." Dimension does not determine statistics – one-dimensional strings quantize as bosons, three-dimensional Skyrmions quantize as bosons for even baryon number [9], and pointlike electrons are fermions. The correct criterion is topological: the 2π-rotation loop must be noncontractible in configuration space [A], and a topological term must select the fermionic sector [3, 10]. The replacement for "fermions build the world because they fill space" is the stability-of-matter theorem [13, 14]: volume filling is a *consequence* of Fermi statistics (bosonic matter collapses with $E \sim -N^{7/5}$ [15]), not its cause. The causal chain this paper assembles runs

$$\text{decorated torus} \Rightarrow \text{FR fermion} \Rightarrow \text{Lieb–Thirring} \Rightarrow \text{extensive ma} \tag{1}$$

with the first implication conditional on odd winding parity (Sec. III) and the rest theorems.

The second intuition is that "a boson fills only two dimensions and can therefore condense." Taken spatially this is backwards: in strictly two dimensions the Mermin–Wagner–Hohenberg theorem [16, 17] forbids true condensation at $T > 0$. The correct version is internal: the boson channel retains *one* phase where the fermion carries two, and the projection that removes the second phase – the Hopf map, Sec. IV – removes precisely the $\mathbb{Z}_2$ loop that generates the exclusion principle. Condensation is then permitted for the projected channel not because of its dimension but because the obstruction has been projected away.

A. What is new here relative to the literature

Because every ingredient above is standard, we state explicitly what this paper adds. (i) A no-go/rescue pair for rigid tori: Proposition 1 (bare axisymmetric torus admits no fermionic quantization) and its resolution by phase decoration, with the parity rule nk odd inherited from [8] (Sec. II–III). (ii) The physical identification of three tori that the literature treats separately: the Clifford torus of the two spinor phases in S^3, the spatial toroidal soliton, and the winding torus of the SQUS spectral tower – all related by the same Hopf fibration, and tied to the mass axis by the warp (Sec. III). (iii) The boson-as-Hopf-pushforward theorem chain: bilinear = Hopf map ⇒ fiber phase (and its $\mathbb{Z}_2$) removed ⇒ symmetric statistics of the projected channel (Sec. IV). (iv) Mass-sector consequences that exist only in SQUS: forced antiperiodicity at the horn tip with a computable deviation angle at the stabilized orbit (Sec. V, App. A); the decomposition $a = \Phi - \frac{1}{2}$ of the measured spectral shift, with $\Phi = 0.6726(68)$ compatible with 2/3; and the two-radius-law crossing scale m_*. (v) Internal no-go results that kill natural readings of the framework's own numbers (Sec. VI), continuing the methodological pattern of the companion papers. Table I summarizes the division of labor between inherited and new content.

II. CONFIGURATION SPACE: THE BARE TORUS FAILS, THE DECORATED TORUS SUCCEEDS

A. Proposition 1: no fermionic quantization for an axisymmetric torus [A]

Let a rigid body have symmetry group $H \subseteq SO(3)$ (the rotations mapping the configuration to itself). Its space of orientations is $Q = SO(3)/H$. FR quantization [3] assigns a $\mathbb{Z}_2$ phase to noncontractible loops of Q; spinoriality requires the 2π-rotation loop to be noncontractible.

Proposition 1. If H contains a continuous subgroup $SO(2)$ (axial symmetry), the 2π-rotation loop about *any* axis is contractible in Q, and no fermionic quantization exists.

Proof. The fibration $H \to SO(3) \to Q$ gives the exact sequence

$$\pi_1(H) \xrightarrow{\iota_*} \pi_1(SO(3)) \to \pi_1(Q) \to \pi_0(H) \to 0. \qquad (2)$$

The 2π rotation about the symmetry axis lies in H and generates $\pi_1(SO(3)) = \mathbb{Z}_2$; since $\pi_1(SO(2)) = \mathbb{Z}$ maps onto it under ι_* (reduction mod 2), the image of $\pi_1(SO(3))$ in $\pi_1(Q)$ is trivial. All 2π-rotation loops about all axes are homotopic in $SO(3)$ (there is only one nontrivial class), hence all are contractible in Q. □

A bare torus – with or without the flip symmetry, i.e. $H = SO(2)$, $O(2)$, or larger – therefore cannot acquire spin $\frac{1}{2}$ from its shape. The popular picture "doughnut ⇒ fermion" is false as stated. This negative result is the reason the two-phase decoration of the SQUS construction is not ornamental but essential.

TABLE I. Division of labor between established results (cited) and the contributions of this paper.

Established	Source
FR quantization; exchange ↔ rotation	[3, 4]
toroidal Hopf solitons; $E \geq cQ^{3/4}$	[5–7]
fermionic iff Q_H odd	[8]
Skyrmion spin from topological term	[9, 10]
stability of matter; bosonic collapse	[13–15]
composite gauge boson no-gos	[18, 19]
New in this paper	**Sec.**
bare-torus no-go; decoration rescue (Prop. 1–2)	II
three-torus identification via Hopf fibration	III
boson = Hopf pushforward; condensation corollary	IV
tip antiperiodicity $U \to -I$; orbit deviation angle	V
$a = \Phi - \frac{1}{2}$, $\Phi = 0.6726(68)$ vs 2/3	V
two-radius crossing scale m_*	V
coherence ceiling $\Omega \gtrsim H$; $m_{\max} = m_*^2/m_L$	V D
vacuum comb; flux cap $P_{\max}$; $m_{\max}c^2 = P_{\max}L/c$	V F
variance–tower incompatibility theorem	VI
$Q^{3/4}$ tower exclusion; 8–15–17 exclusion	VI

B. Proposition 2: decoration restores the $\mathbb{Z}_2$ [A]

Decorate the torus with a phase pattern $e^{i(n\theta_1 + k\theta_2)}$, $nk \neq 0$. A spatial rotation about the symmetry axis now shifts θ_1 and is no longer a symmetry by itself; the residual stabilizer Γ (spatial rotations combined with internal phase shifts that restore the configuration) is at most discrete for the rigid-marking model. Then $\pi_1(\Gamma) = 0$, the exact sequence gives an injection $\mathbb{Z}_2 = \pi_1(SO(3)) \hookrightarrow \pi_1(Q)$, the 2π-rotation loop is noncontractible, and a fermionic FR sector exists. □

Two caveats keep this honest. First, when the internal phase is a genuine field degree of freedom, spatial rotation combined with a compensating *isorotation* is a symmetry; the correct arena is then the field configuration space, treated next. Second, FR quantization makes the fermionic sector *available*, not mandatory; selecting it dynamically requires a topological term in the action, as in the Skyrme model at odd N_c [9] and the Hopf-term mechanism of [10]. The SQUS framework must ultimately supply such a term; we flag this as open [D].

III. THE HOPF BRIDGE: ONE TORUS SEEN THREE WAYS

A. Field realization: toroidal Hopf solitons [A]

The field-theoretic home of a decorated torus is the Faddeev–Niemi model of maps $\boldsymbol{n} : \mathbb{R}^3 \to S^2$ with $\boldsymbol{n} \to \boldsymbol{n}_\infty$ at infinity, classified by the Hopf invariant $Q_H \in \pi_3(S^2) = \mathbb{Z}$ [5]. Minimal-energy configurations are tori and, at higher charge, links and knots [6]; the energy obeys the Vakulenko–Kapitanskii bound $E \geq c\, Q_H^{3/4}$ [7]. For the standard toroidal ansatz with windings n and k on the two cycles,

$$Q_H = nk. \tag{3}$$

The configuration space of the model has $\pi_1 = \mathbb{Z}_2$ in each topological sector, and Krusch and Speight proved, by transporting the question to the Skyrme model along the Hopf fibration, that *Hopf solitons can be quantized as fermions if and only if the* 2π*-rotation loop is noncontractible, which holds precisely for odd Hopf charge* [8]. Combining with Eq. (3):

$$\boxed{\text{fermionic torus} \iff nk \text{ odd,}} \tag{4}$$

i.e. *both* windings odd. The two-phase structure of the SQUS note thus acquires a parity selection rule it did not previously have. The minimal fermion is $(n, k) = (1, 1)$, $Q_H = 1$.

B. The internal side: Clifford torus and the same fibration [A]

A normalized two-component spinor (z_1, z_2), $|z_1|^2 + |z_2|^2 = 1$, lives on S^3. At fixed amplitudes the two remaining phases sweep a Clifford torus $T^2 \subset S^3$ – the fibered tori of the Hopf map $h : S^3 \to S^2$, $h(\Psi) = \Psi^\dagger \boldsymbol{\sigma} \Psi$. The fibers of h are the circles of the overall phase; any two fibers are linked once, and the linking number *is* the Hopf invariant [11]. The spatial soliton of the previous subsection is a map whose target is the base S^2 of this same fibration, and whose preimage circles realize the linking in physical space. The two tori – internal (phases of Ψ) and spatial (soliton core) – are therefore not merely analogous; they are the total space and the fiber structure of one fibration. Figure 1 displays the three views.

C. The SQUS identification [C/D]

SQUS proposes to take this coincidence literally: the spatial SQUS loop of the companion paper's spectral tower and the internal phase torus are the same object, glued by the warp. The supporting observation is geometric: on the canonical metric the physical size of every spatial separation scales as $|\xi|/L$ and collapses at the tip, so the distinction between "configuration-space torus" and "physical-space torus" is itself scale-dependent, degenerating as $\xi \to 0$. We label the identification [C/D]: it is the framework's central structural hypothesis, motivated but not derived. One consequence is immediate and testable within the framework: the tower winding ℓ of [2] must be a *phase* winding (rotor-like), not the Hopf charge – the energetics of Sec. VIB enforce this.

IV. THE BOSON AS A HOPF PUSHFORWARD

A. Projection kills exactly the exclusion loop [A]

The spinor bilinear is the Hopf map: $\boldsymbol{n} = \Psi^\dagger \boldsymbol{\sigma} \Psi$ is invariant under the fiber phase $\Psi \to e^{i\gamma}\Psi$ and in particular under $\Psi \to -\Psi$. The 2π sign of the fermion lives entirely in the fiber that the projection removes. By Sorkin's theorem [4], for identical solitonic objects the exchange loop is homotopic to the 2π-rotation loop of one object; the FR phase of the two loops is therefore the same $\mathbb{Z}_2$ element. Consequently:

Corollary. Any channel built from Hopf-projected (bilinear) data carries trivial FR phase, obeys symmetric statistics, and admits macroscopic occupation of a single state. Condensation of the projected channel is permitted *because the projection removed the topological obstruction*, not because of any dimensional counting. □

The channel taxonomy of the SQUS note follows: the scalar, vector, axial, and tensor bilinears $\bar{\Psi}\Gamma\Psi$ are different components of the pushforward data, and the single-phase combinations $\varphi_{p,q} = p\theta_1 + q\theta_2$ are their toroidal Fourier labels. In this precise sense the boson *is* a projection of the fermionic torus: the pushforward of the two-phase state along the Hopf fibration.

B. The extensivity corollary [A]

Completing the chain (1): once the fermionic sector is selected, the Lieb–Thirring inequality guarantees stability of matter of the second kind – N-body energy bounded below by $-cN$ and volume growing extensively – whereas the same matter with bosonic constituents collapses, $E \sim -N^{7/5}$ [13–15]. The folk statement "fermions are the building blocks because they fill space" is thus recovered as a corollary at the *end* of the chain, with the causal arrow pointing from statistics to volume. Conversely, the projected channels, carrying no exclusion, pile up rather than fill: condensates, coherent fields, and the classical limit. In strictly two spatial dimensions condensation would in fact be obstructed at $T > 0$ [16, 17]; the operative "dimensional reduction" of the boson is the loss of one *internal* phase, not of a spatial dimension.

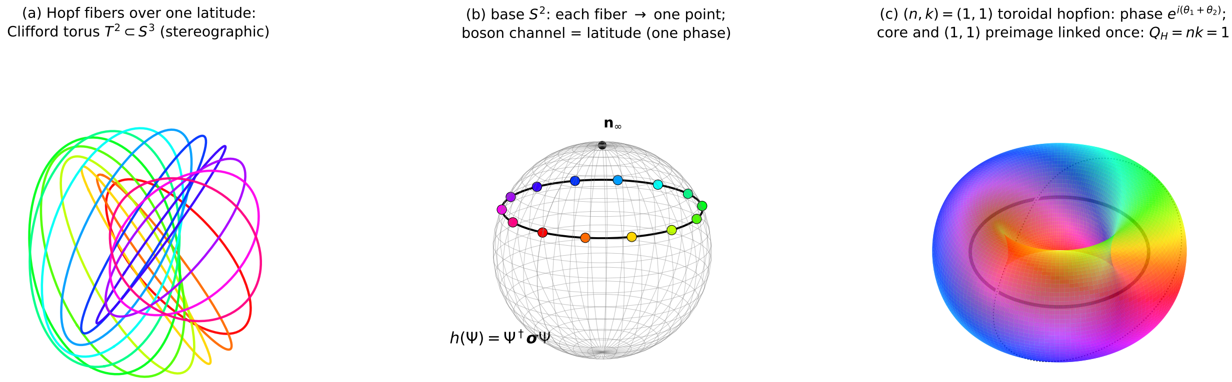

FIG. 1. One torus seen three ways. **(a)** Hopf fibers over a single latitude of the base S^2, stereographically projected from S^3 to $\mathbb{R}^3$: the fibers sweep a Clifford torus, the two-phase configuration space of a normalized spinor at fixed amplitudes. Each fiber is colored by its base point. **(b)** The base S^2 under the Hopf map $h(\Psi) = \Psi^\dagger \boldsymbol{\sigma} \Psi$: every fiber collapses to one point (matching colors), and the Clifford torus projects to a single latitude circle – the boson channel retains one phase where the fermion carries two. The fiber phase, which hosts the 2π sign, is removed by the projection. **(c)** The spatial realization: a $(n, k) = (1, 1)$ toroidal hopfion, its surface colored by the winding phase $e^{i(\theta_1+\theta_2)}$. The core circle (black) and the $(1, 1)$ preimage curve (white) are linked exactly once, realizing $Q_H = nk = 1$ – the minimal fermionic torus of Eq. (4).

V. THE MASS SECTOR: WHAT THE WARP ADDS

A. Forced antiperiodicity at the tip [A]

On the SQUS surface $d\rho^2 + f(\rho)^2 d\varphi^2$ the azimuthal loop is contractible through the tip cap, so the spin structure is unique: a single-valued spinor in the adapted frame is antiperiodic in φ, and the half-integer shift of azimuthal quantum numbers is forced, not chosen. The spinor holonomy is $U = -\exp[i\pi f'(\rho)\sigma_3]$ (App. A); for the quadratic horn $f' \to 0$ at the tip, so

$$U \longrightarrow -I \qquad (\xi \to 0) \text{ exactly}, \tag{5}$$

realizing the 4π structure geometrically in the near-tip regime. At the stabilized orbit ξ_c of the completed profile the deviation from $-I$ is the angle $\pi\xi_c/(2R_s)/\sqrt{B(\xi_c)}$, small for $\xi_c \ll R_s$.

B. Decomposition of the spectral shift [B/C]

The companion paper fits the winding tower $\varepsilon_\ell \propto (\ell - a)^2$ with $a = 0.1726(68)$ [2]. The discrete spin-structure shifts $\alpha \in \{0, \frac{1}{2}\}$ are excluded as the sole origin of a at 25σ and 48σ respectively. But if the tower modes are spinors, the $\frac{1}{2}$ is *mandatory* (Sec. VA), and the measured shift decomposes as

$$a = \Phi - \tfrac{1}{2}, \qquad \Phi = 0.6726 \pm 0.0068, \tag{6}$$

where Φ is the continuous $U(1)$ holonomy (flux or rotation). The required flux is compatible with $\Phi = 2/3$ at $+0.87\sigma$ – equivalently, the known candidate $a = 1/6$ acquires the reading $1/6 = 2/3 - 1/2$: flux of two-thirds of a quantum minus the spinor half. We note, with a look-elsewhere caveat and a [C] label, that $2/3$ is also the equipartition charge $Q = \frac{2}{3}$ of the Koide sector [1]; whether the two occurrences share an origin is open [D], but Eq. (6) converts the question into a number that improved Δm^2 data (JUNO) will sharpen or kill.

C. Two radius laws and the crossing scale [B/D]

SQUS assigns the torus two competing sizes. The warp gives the physical circumferential radius $f = (\kappa^2/R_s)\, m$, growing linearly with mass; quantum mechanics gives the reduced Compton radius $\lambda_C = \hbar/(mc)$, falling with mass. The two laws cross at a single distinguished mass,

$$m_* = \sqrt{\frac{\hbar R_s}{\kappa^2 c}}, \tag{7}$$

the unique mode whose geometric loop and Compton loop coincide. Below m_* the Compton cloud dwarfs the geometric core; above it the geometric radius is the larger. The zitterbewegung identification of the companion paper (circulation at λ_C) and the horn law cannot both hold for all masses; their simultaneous validity at exactly one mass is either an internal inconsistency or a selection principle. We record the second reading as a conjecture [D]: m_* is a candidate for the framework's absolute scale, to be matched against the stabilization sector once κ and R_s are fixed independently.

D. The mass ceiling: three routes to $m_{\max} = m_*^2/m_L$ [A/D]

Does the construction bound the mass of a single mode from above? The geometry alone does not, and this negative statement is exact. The warp supplies an attraction toward the tip that is precisely inverse-square (the conformal, critical case: the A/ξ^2 term of the radial problem), and for the pure horn the Liouville reduction is exactly scale-free – the curvature terms $f'^2/4f^2 - f''/2f$ cancel identically for $f = \xi^2/R_s$, leaving the purely repulsive $\ell^2 R_s^2/\xi^4$ with no bound states. A scale-free potential confines nothing and selects nothing: no maximal mass, and indeed no mass at all, follows from the geometry [A]. On a completed profile the situation at large ξ is if anything softer: beyond ξ_c the profile flattens ($h \to 0$), the circulation gradient that could restore the mode along ξ dies away, and only the marginal inverse-square attraction remains. Large-mass modes are thus barely held – and what actually eliminates them is not a force but a *coherence condition*.

The winding level ℓ is a resonance of a ring of circumference $2\pi f$: it exists as a discrete state only if its width is smaller than the ring's free spectral range, $\Gamma < \hbar c/(2\pi f)$ – the finesse condition of any resonator. The unavoidable floor on Γ is set by the vacuum itself. The Λ-vacuum sections that regularize the tip (companion paper, curvature appendix) carry Gibbons–Hawking fluctuations with correlation time $\sim 1/H$ and correlation length $\sim c/H$ [12]. For $2\pi f \ll c/H$ these fluctuations are common-mode across the loop and dephase nothing; for $2\pi f \gtrsim c/H$ the circulating wave samples uncorrelated patches of the fluctuating background, accumulates $O(1)$ phase variance per circuit, and fails to reconnect with itself: no standing winding, no level. Equivalently, since the null circulation has angular velocity $\Omega = c/f$, the coherence condition is a *minimal angular velocity*,

$$\Omega \gtrsim H : \tag{8}$$

a mode must complete at least of order one circuit per Hubble time, or the universe's own fluctuation background erases it before resonance can form. With the horn law $f = (\kappa^2/R_s)m$ and the bootstrap fixing of κ through m_* (Eq. 7), the three conditions – validity of the tip expansion ($\xi \lesssim L$), the resonator finesse bound, and Eq. (8) – give the same ceiling up to $O(1)$ factors:

$$m_{\max} = \frac{m_*^2}{m_L}, \qquad m_L \equiv \frac{\hbar}{Lc} = 1.19 \times 10^{-33}\ \text{eV}, \tag{9}$$

a seesaw between the self-dual mass and the warp mass; in square-root variables, $\chi_* = \sqrt{\chi_{\max}\chi_L}$ – the self-dual scale is the geometric mean of the ceiling and the Hubble floor. The coherence margin of any mode is the same ratio, $\Omega/H \simeq m_{\max}/m$, and at the ceiling $\Omega/H_0 = 0.83 \approx 1$: the limiting mode circulates exactly once per Hubble time.

Numerically, with m_* identified as the cone neutrino ν_1 (0.364 meV, the one mode for which the companion paper makes the Compton identification, $R_1 = 0.542$ mm),

$$m_{\max} = 1.1 \times 10^{17}\ \text{GeV} \approx 10^{-2}\, m_{\text{Pl}}, \tag{10}$$

a GUT-scale ceiling under which every Standard-Model fermion sits comfortably (the top quark by fifteen orders of magnitude; its coherence margin is $\Omega/H_0 \sim 5 \times 10^{14}$, with a circuit time of ~ 1.5 h on a configurational loop of $\sim 2.6 \times 10^{11}$ m). The logic also runs backwards, and this is the sharpest statement of the subsection: requiring only that the ceiling not exceed the Planck mass – above which gravity intervenes regardless – forces

$$m_* \le \sqrt{m_{\text{Pl}}\, m_L} = 3.8\ \text{meV}, \tag{11}$$

i.e. the framework's self-dual mass *must* lie in the neutrino window. Of the three cone masses only ν_1 satisfies Eq. (11) ($\nu_2 = 8.66$ meV would place the ceiling at $5\, m_{\text{Pl}}$). We label this convergence [C/D] with the usual look-elsewhere caveat; but it is the first internal argument for *why* the self-dual scale should sit at the lightest neutrino rather than anywhere else on fourteen decades of mass axis.

Two costs are incurred and must be stated. First, with this κ the configurational radii are macroscopic – $f(e) \approx \alpha \times 760$ km, $f(t) \approx \alpha \times 1.7$ AU – so the configurational firewall of the companion paper (no spatial charge or current distribution; Sec. VIIA) is not a stylistic precaution but a survival condition, and any channel that couples f to spatial physics falsifies the construction outright. Second, sustaining $E = mc^2$ by circulation above m_* would require winding number $\ell = (m/m_*)^2 - 2\times 10^{18}$ for the electron, 2×10^{29} for the top – so the $\ell = 1$ circulation reading does not extend to the heavy fermions; consistently with the companion paper's admission that the charged spectrum has no mechanism, heavy fermions cannot be windings of the same loop, and the ceiling (9) bounds the domain of the geometry, not the winding tower. A self-consistency check closes the subsection: the smearing of the loop radius by a mode's own width, $\delta f/f = 2\,\Delta\chi/\chi = \Gamma/(mc^2)$, must stay below $\sim 1/(4\ell)$ for reconnection; the widest relevant resonance, the top quark, gives 0.4% against 25% – passed with two orders to spare.

E. Vacuum-driven mass and its convergence condition [B]

If the mass is the stationary variance of vacuum-driven toroidal modes, $m = \sum_{n,k}\langle |A_{nk}|^2\rangle$ with $\langle |A_{nk}|^2\rangle = D_{nk}/(2\gamma_{nk}\omega_{nk}^2)$, then finiteness of m is not automatic: for flat noise the sum diverges logarithmically ($\sum 1/\omega^2$ on T^2), and for quantum zero-point weighting linearly ($\sum 1/\omega$). The geometric filter of the SQUS construction is therefore a *necessary* ingredient, with the quantitative requirement

$$\sum_{n,k} \frac{D_{nk}}{\gamma_{nk}\,\omega_{nk}^2} < \infty, \tag{12}$$

i.e. damping and/or drive spectra must fall faster than the flat case. Equation (12) belongs on the derivation program as a hard constraint, on par with the fluctuation–dissipation relation.

F. Resonant vacuum feeding: the comb, the flux cap, and self-regulation [B/D]

The winding reading has a clean ontology: the field is always present, and the *particle is the circulating phase* – the wave persists, the phase goes around, and the trivial sector $\ell = 0$ (no circulation) is the vacuum. This subsection develops the dynamics of that picture: how the vacuum sustains a circulating phase, at what maximal rate, and what happens beyond the ceiling of Sec. V D.

The comb. On a closed loop, periodicity quantizes the vacuum itself: the loop's field modes exist only at integer winding (shifted to $\ell - a$ by the holonomy), so the on-loop vacuum spectrum is a comb, and *between the teeth there is no vacuum mode at all* – not a depleted vacuum but an absent one, the Casimir structure of a twisted circle. Two consequences follow at once. First, integer-ℓ gating of the feeding is automatic, not imposed: a fractional winding of a single-valued phase is not a configuration, so the only states that exist are exactly the resonant ones. Second, the absent inter-tooth continuum *protects* every level: there is nothing on the loop for a tooth to decay into, so each winding state is stable for free, and transitions $\ell \to \ell'$ require external mixing – which the framework, consistently, does not determine. A pre-emptive exclusion belongs here: if one tried to *derive* the holonomy from the loop's Casimir energy $E_{\mathrm{Cas}} \propto B_2(a) = a^2 - a + \frac{1}{6}$, the minimum ($a = \frac{1}{2}$, 48σ) and the zero ($a = 0.2113$, 5.7σ) are both excluded by the fitted $a = 0.1726(68)$: the comb is load-bearing structure, but it does not set a.

Store versus flux. Could the vacuum feed a mode from its stored zero-point content? No: the ring's Casimir energy is $|E_{\mathrm{Cas}}| = (B_2/2)(m_*/m)\, m_* c^2$, so demanding $mc^2 \leq |E_{\mathrm{Cas}}|$ gives $m \leq \sqrt{B_2/2}\, m_* \approx 0.04$ meV – below even ν_1. The static reservoir starves the entire tower; the feeding must be a *flux.* The flux is the closed fluctuation–dissipation cycle: the bath that damps is the bath that drives ($D = \gamma\hbar\omega$ at zero temperature), dissipated energy returns to the same vacuum, the net exchange vanishes in steady state, and no perpetuum mobile is involved – the mode holds a standing variance in circulation, not a withdrawal. On the horn the quantum balance is exactly marginal: the mass cancels identically from the self-consistency condition, so the vacuum sustains every mass equally – consistent with the log-uniform measure and with the radial flat direction, and leaving discreteness to integer ℓ alone. This channel evades the variance–tower incompatibility theorem of Sec. VIA *legitimately*: it is an open, driven-dissipative mechanism outside the potential class $U(\chi) = \kappa\chi^2/2 + C_\ell/\chi^s$ to which the theorem applies.

The flux cap. The gross exchange rate is $P = \gamma\, mc^2$, and the finesse condition of Sec. V D caps the linewidth at the tooth spacing, $\gamma \leq \mathrm{FSR}/\hbar = c/f$. With the horn law the mass cancels once more and the cap is universal:

$$P_{\max} = \frac{(m_* c^2)^2}{\hbar} = 2.0\times 10^{8}\ \mathrm{eV\,s^{-1}} \approx 3.2\times 10^{-11}\ \mathrm{W}, \tag{13}$$

the same maximal feeding power for every mode from the neutrino to the ceiling – the framework's third constant alongside m_* and m_L. The assembly time $t_{\min} = mc^2/P_{\max}$ runs from 2.5 ms (electron) through ~ 14 min (top) to, at the ceiling, *exactly* L/c: the identity

$$m_{\max} c^2 = P_{\max} \times \frac{L}{c} \tag{14}$$

holds to four digits and is the fourth independent formulation of the ceiling (9) – the largest sustainable mass is the maximal vacuum flux integrated over the Hubble time.

Self-regulation. The dynamics closes into a self-consistent cycle. Above the ceiling the tooth is unresolvable, the resonant intake fails, damping wins, and the mode slides down in χ – but the dissipated energy returns to the bath (the vacuum "recharges"), the shrinking loop *widens* the tooth spacing, and the feeding channel reopens the moment the mode drops below $m_{\max}$: the supercritical mode relaxes onto the ceiling rather than dying, and the ceiling is an attractor from above. Total collapse to the tip is forbidden throughout by topology: ℓ is conserved, so a fed or starved mode alike remains a winding, and the only way to reach $\ell = 0$ is a phase-cutting, topology-changing process – surgery, not starvation. Near the ceiling a genuine two-sided balance is plausible [D]: there $\hbar\omega$ approaches the Gibbons–Hawking temperature of the Λ-vacuum, thermal occupation enhances the intake and pushes upward while coherence loss pushes downward, sharpening the attractor into an equilibrium point. One precision is essential to keep this subsection consistent with the rest of the paper: what self-regulates *everywhere* is the amplitude at a given tooth (the FDT fixed point, restored at rate γ); the *position* in χ remains marginal below the ceiling, so the feeding mechanism sustains the mode wherever it sits, while the radial completion of the companion paper selects where it sits. The two roles – sustenance and selection – must not be conflated; the vacuum feeds the orbit, it does not choose it, except at the ceiling where coherence itself does the choosing.

VI. NO-GO RESULTS

Following the methodological pattern of the companion papers, we kill the most tempting readings of our own construction.

A. Variance–tower incompatibility [A/B]

Consider the minimal stabilizer class $U(\chi) = \frac{\kappa}{2}\chi^2 + C_\ell/\chi^s$ with winding-dependent $C_\ell \propto (\ell-a)^p$ and $s \geq 0$. Minimization gives $E \propto C_\ell^{2/(s+2)}$, so the tower exponent is $2p/(s+2)$. The measured ratios $r_2 = 4.878(44)$, $r_3 = 11.737(182)$ require the exponent 2 with the geometric $p = 2$; this forces $s = 0$ – a rigid rotor, with *no* dependence of the winding energy on the radial variable. Explicitly, $s = 2$ (standard centrifugal) predicts $r_2 = 2.21$, $r_3 = 3.42$ (-61σ, -46σ); the horn centrifugal $s = 4$ is worse (-72σ); the exact quantum isotonic oscillator, scanned over all parameters, has $\chi^2_{\min} = 196$; the variance reading $m = \langle\chi^2\rangle$ has $\chi^2_{\min} = 436$. But at $s = 0$ the ground-state variance is ℓ-independent, while the data demand m_ℓ varying by a factor 138. Hence:

Theorem (incompatibility). Within the minimal class above, the conjunction "m is the variance of the same radial variable whose spectrum forms the tower" and "the tower ratios follow $(\ell-a)^2$" has no solution. □

The escape routes are exactly two: identify mass with the eigenvalue rather than the variance (the spectral postulate of [2]), or let the Sec. VD variance refer to a field distinct from the radial coordinate. Either way, the resonator that stabilizes the loop must be energetically flat: fitting $\varepsilon_\ell = E_0 + c(\ell-a)^2$ to the ratios bounds the radial zero-point at

$$E_0/\varepsilon_1 = +0.04 \pm 0.17, \qquad |E_0| \leq 0.34\,\varepsilon_1 \text{ (95\% CL)}. \tag{15}$$

B. The Hopf-charge mass tower is excluded [B]

If the tower label were the Hopf charge, the energy bound $E \geq cQ^{3/4}$ [7] and the numerical near-saturation [6] would give $\chi_\ell \propto Q^{3/4}$ at most, hence $r_2 \leq 2^{3/4} = 1.68$ against the measured 4.878: excluded without a fit. The division of labor is therefore forced: the Hopf charge classifies the *type* (fermion for odd nk, Eq. (4)), while the spectral tower is a *phase* winding on a frozen loop. Conflating the two roles would falsify the wrong statement.

C. The 8–15–17 mixing ansatz is excluded [B]

The Pythagorean candidate $\sin\theta = 8/17$ gives $\sin^2\theta = 64/289 = 0.22145$. Against the on-shell value computed from the world averages $M_W = 80.3692(133)$ GeV, $M_Z = 91.1880(20)$ GeV one finds $1 - M_W^2/M_Z^2 = 0.22321(26)$, a 6.8σ discrepancy; against the effective leptonic angle 0.23153(16) the discrepancy is 63σ. No renormalization scale rescues it: the running $\hat{s}^2(\mu)$ has its minimum ≈ 0.231 near M_Z and increases in both directions. If a future torus-stability equation were to output the ratio 8:15, that equation would be falsified by these data; conversely, any derived radius ratio must reproduce $\tan^2\theta$ within current errors to survive.

VII. CONSTRAINTS FROM PRECISION PHYSICS

A. Electron pointlikeness [B]

Any *spatially* extended charged structure at the Compton scale is dead on arrival: the electron behaves pointlike in QED tests to far below $\lambda_C = 3.9 \times 10^{-13}$ m. In the Brodsky–Drell parametrization [20], a substructure of size R shifts the anomalous moment by $\delta a \sim m_e R/\hbar c$ (linear) or $\delta a \sim (m_e R/\hbar c)^2$ (chirally protected); the $\sim 10^{-12}$ agreement of a_e with QED then bounds

$$R \lesssim 4\times10^{-25}\,\text{m (linear)}, \qquad R \lesssim 4\times10^{-19}\,\text{m (quadratic)}, \tag{16}$$

and collider contact-interaction limits independently give $R \lesssim 10^{-19}$ m. The SQUS torus therefore *must not* be an electromagnetic charge distribution of size f or λ_C. The framework's defense is the projection itself: electromagnetic probes couple to Hopf-projected (bilinear) data, which are pointlike-blind to the fiber structure; the two-phase structure is visible only in channels sensitive to χ – mass-changing couplings. The sharp version [D]: SQUS-type substructure should appear, if anywhere, as a pattern in Yukawa/Higgs coupling deviations ordered by mass, not in electromagnetic form factors. This is testable at HL-LHC precision on $\kappa_\mu, \kappa_\tau, \kappa_b$, and its absence there would constrain the mass-channel coupling of the torus.

B. Statistics is exact or the topology is wrong [B]

Because the fermionic sign here is topological ($\mathbb{Z}_2$ holonomy), the framework predicts *exactly* zero violation of the Pauli principle – there is no small parameter by which a topological phase can be slightly wrong. Dedicated searches for Pauli-violating atomic transitions (VIP-2/VIP-3 at Gran Sasso) bound the violation probability at the $\lesssim 10^{-31}$ level and any confirmed nonzero signal would falsify the topological origin of statistics assumed here, independently of all other SQUS content. The same rigidity applies to the spin-statistics connection: in local relativistic QFT the connection is a theorem, so the geometric model does not get to choose; it must *reproduce* the theorem through its topological term, and Sec. II's open [D] item (the SQUS analogue of the Wess–Zumino/Hopf term at $\theta = \pi$) is the place where that obligation lands.

C. Composite gauge bosons remain fenced off [A]

The projection reading of Sec. IV deliberately stops short of compositeness. The fence posts are known theorems: a massless spin-1 particle carrying a Lorentz-covariant conserved charge is forbidden [19] – directly blocking composite gluons that carry color – and the

TABLE II. Falsifiers and confirmation routes. Inherited entries (Σm_ν, $m_{\beta\beta}$) are stated in [1, 2].

Observable / result	Kills / confirms	Test
$a \to 1/6$ vs away	confirms/kills Eq. (6)	JUNO Δm^2
Pauli violation > 0	kills topological statistics	VIP-3
e EM form factor	kills spatial-charge reading	a_e, colliders
Yukawa pattern in m	confirms mass-channel structure [D]	HL-LHC
derived ratio $\neq \tan\theta_W$	kills that stability eq.	EW data
tower $\propto Q^{3/4}$	already excluded	Sec. VIB
$m =$ Var + tower	already excluded (min. class)	Sec. VIIA
$\Sigma m_\nu = 59.16$ meV	inherited nearest-term test	cosmology
$m_* > 3.8$ meV established	kills sub-Planck ceiling logic	cone/JUNO
any spatial coupling of f	kills configurational reading	precision EM

photon-as-fermion-pair construction fails on exact Bose statistics [18]. An SQUS gauge sector must therefore arrive as an emergent gauge *redundancy* of the projected description, not as bound states; this constraint should be treated as load-bearing in any future derivation of $SU(2) \times U(1)$ or $SU(3)$ from the torus.

VIII. FALSIFICATION SCHEDULE

Table II lists what would falsify, and what would substantially confirm, the construction. Two remarks. First, the parity rule (4) is presently an internal consistency condition; it becomes a falsifier the moment the framework commits to winding assignments for known fermions. Second, the strongest available *confirmation* route is the flux decomposition (6): a derivation tying $\Phi = 2/3$ to the equipartition charge $Q = \frac{2}{3}$, combined with JUNO-improved Δm^2 shrinking σ_a, would convert a 0.9σ curiosity into a two-sector coincidence with a mechanism – or destroy it.

IX. DISCUSSION

The answers to the title questions are: yes, a torus can justify the "look" of a fermion, but only as a decorated torus with odd winding parity, through the FR mechanism realized by Hopf solitons – shape alone provably cannot (Proposition 1); and yes, a boson can be a projection of the torus, in the exact sense of the Hopf pushforward, which removes the one phase in which the exclusion-generating $\mathbb{Z}_2$ lives – whence condensation. The dimensional intuition that motivated both questions survives as a pair of corollaries with reversed arrows: extensivity of fermionic matter (Lieb–Thirring) and pile-up of projected channels.

What is genuinely SQUS in this paper is the welding of that topology to the mass axis: forced antiperiodicity at the horn tip, the $\Phi - \frac{1}{2}$ decomposition of the measured shift with $\Phi \approx 2/3$, the two-radius crossing scale m_* with its coherence ceiling $m_{\max} = m_*^2/m_L$ and the resulting neutrino-window constraint on m_*, and the convergence condition (12) on the vacuum-driven mass mechanism.

Limitations, stated plainly: the topological term selecting the fermionic sector has not been derived from the SQUS action; the identification of the internal and spatial tori is a hypothesis [C/D], not a theorem; the parity rule has no committed winding assignments yet; m_* is located only conditionally (the ceiling argument fixes it to the neutrino window under sub-Planckianity, and to ν_1 under the Compton identification, both [D]); the coherence bound uses the correlation scale of de Sitter fluctuations at the order-of-magnitude level, not a derived dephasing rate; the feeding channel is an open-system model whose filter is constrained (quantum FDT, convergence, finesse) but not derived, and the near-ceiling thermal balance is a remark, not a computation; the boson channels have not been promoted to gauge fields, and the theorems of Sec. VIIC stand guard over any attempt; the Yukawa prediction is a direction, not a number; and the flux value 2/3 is, at present, a 0.9σ regularity with a look-elsewhere caveat. The construction survives exactly as far as Table II permits.

Appendix A: Spin connection and holonomy on the SQUS surface

For the surface $d\rho^2 + f(\rho)^2 d\varphi^2$ take the zweibein $e^1 = d\rho$, $e^2 = f\,d\varphi$; the Cartan equations give $\omega^{12} = f'(\rho)\,d\varphi$ and Gauss curvature $K = -f''/f$. For the quadratic horn $f = \rho^2/R_s$ (proper-distance form near the tip), $K = -2/\rho^2$: a hyperbolic cusp, with the compensating $+2\pi$ concentrated at the tip so that Gauss–Bonnet closes on the cap. Parallel transport of a spinor once around the loop gives, in the adapted frame,

$$U(\rho) = -\exp[i\pi f'(\rho)\,\sigma_3]\,, \tag{A1}$$

the overall -1 being the frame contribution fixed by the unique spin structure of the contractible loop. Hence $U \to -I$ exactly as $f' \to 0$ (the tip), and at the stabilized orbit ξ_c of the completed profile, where $f'(\xi_c) = \xi_c/(2R_s)$, the deviation from $-I$ is the rotation angle $\pi\,\xi_c/(2R_s)/\sqrt{B(\xi_c)}$. The half-integer shift of azimuthal quantum numbers for single-valued spinor fields follows from the same frame antiperiodicity and is independent of f'.

Appendix B: Exclusion fits

The tower data are the Monte-Carlo-propagated cone ratios $r_2 = 4.878(44)$, $r_3 = 11.737(182)$ of [2]. The stabilizer scans of Sec. VIA use the exact isotonic spectrum $E_n = \omega(2n + 1 + \frac{1}{2}\sqrt{1+8B})$ and the virial identity $\langle\chi^2\rangle = E/\omega$; joint fits minimize over the coupling and the shift a with $a \in [0, 0.99)$. The rotor bound fits $\varepsilon_\ell = E_0 + c(\ell - a)^2$ exactly to the two ratios and propagates errors by Monte Carlo assuming uncorrelated r_2, r_3 (a conservative simplification; the common-Δm^2_{21} correlation slightly tightens the bound). The electroweak comparison uses M_W and M_Z world averages and the effective leptonic angle as quoted in Sec. VIC. All scripts are public (Data Availability).

DATA AVAILABILITY

Analysis scripts reproducing the numerical exclusions and bounds are publicly available at [repository/Zenodo DOI].

[1] K. Urbanowicz, "Equipartition interpretation of the Koide relation," DOI: 10.5281/zenodo.21325642.

[2] K. Urbanowicz, "The square root of mass as a fifth coordinate: geometry, stabilization, and spectral structure," companion paper.

[3] D. Finkelstein and J. Rubinstein, "Connection between spin, statistics, and kinks," J. Math. Phys. **9**, 1762 (1968).

[4] R. D. Sorkin, "A general relation between kink-exchange and kink-rotation," Commun. Math. Phys. **115**, 421 (1988).

[5] L. Faddeev and A. J. Niemi, "Stable knot-like structures in classical field theory," Nature **387**, 58 (1997).

[6] R. A. Battye and P. M. Sutcliffe, "Knots as stable soliton solutions in a three-dimensional classical field theory," Phys. Rev. Lett. **81**, 4798 (1998).

[7] A. F. Vakulenko and L. V. Kapitanskii, "Stability of solitons in S^2 nonlinear σ-model," Sov. Phys. Dokl. **24**, 433 (1979).

[8] S. Krusch and J. M. Speight, "Fermionic quantization of Hopf solitons," Commun. Math. Phys. **264**, 391 (2006), arXiv:hep-th/0503067.

[9] E. Witten, "Current algebra, baryons, and quark confinement," Nucl. Phys. B **223**, 433 (1983).

[10] F. Wilczek and A. Zee, "Linking numbers, spin, and statistics of solitons," Phys. Rev. Lett. **51**, 2250 (1983).

[11] H. Hopf, "Über die Abbildungen der dreidimensionalen Sphäre auf die Kugelfläche," Math. Ann. **104**, 637 (1931).

[12] G. W. Gibbons and S. W. Hawking, "Cosmological event horizons, thermodynamics, and particle creation," Phys. Rev. D **15**, 2738 (1977).

[13] F. J. Dyson and A. Lenard, "Stability of matter I," J. Math. Phys. **8**, 423 (1967).

[14] E. H. Lieb and W. E. Thirring, "Bound for the kinetic energy of fermions which proves the stability of matter," Phys. Rev. Lett. **35**, 687 (1975).

[15] J. G. Conlon, E. H. Lieb and H.-T. Yau, "The $N^{7/5}$ law for charged bosons," Commun. Math. Phys. **116**, 417 (1988).

[16] N. D. Mermin and H. Wagner, "Absence of ferromagnetism or antiferromagnetism in one- or two-dimensional isotropic Heisenberg models," Phys. Rev. Lett. **17**, 1133 (1966).

[17] P. C. Hohenberg, "Existence of long-range order in one and two dimensions," Phys. Rev. **158**, 383 (1967).

[18] M. H. L. Pryce, "On the neutrino theory of light," Proc. Roy. Soc. Lond. A **165**, 247 (1938).

[19] S. Weinberg and E. Witten, "Limits on massless particles," Phys. Lett. B **96**, 59 (1980).

[20] S. J. Brodsky and S. D. Drell, "Anomalous magnetic moment and limits on fermion substructure," Phys. Rev. D **22**, 2236 (1980).