DATED HISTORY

One Obstruction, Four Faces: The Barrier to an Emergent Linear Graviton on a Local Lattice

Historical four-part obstruction formulation. Each obstruction applies to the specified local lattice construction and its assumptions.

Manuscript · p. 1, framing

One Obstruction, Three Faces: A Local First-Class Chiral Bosonic Spin-2 Carrier and the Barrier to a Protected Graviton Phase

Evidence · One-Hamiltonian composition requirements

One Obstruction, Four Faces: The Barrier to an Emergent Linear Graviton on a Local Lattice

Joshua Farrow
Independent Research
joshuajamesfarrow@protonmail.com

June 2026

Abstract

What does it take for a local, translation-invariant lattice construction to carry an emergent linear graviton, and why is it hard? We show that four apparently distinct obstacles are one algebraic fact: in three dimensions there is no local, gauge-invariant, symmetric, transverse-traceless (TT) tensor operator of derivative order below three; the minimal one is the Cotton tensor. We prove this by direct enumeration. That single fact (i) locks the dispersion of any exact compact finite-dimensional local realization to Lifshitz $z = 3$; (ii) suppresses the response to any source whose conservation is a derivative identity (k^{+2} for a Beltrami superpotential, k^{+4} for the gauge-invariant Cotton source), while independent local stress variables under a local conservation constraint couple at full $1/k^2$ strength; (iii) shows first-classness alone leaves masslessness unprotected, a positive local first-class mass being relevant; and (iv) forces a lone chiral cone in a local periodic bulk to violate a topological charge-sum rule. To make the map concrete we exhibit a local, positive, first-class chiral bosonic carrier that saturates the two TT branches, realizes them as a single wall-localized cone with finite residue, and—probed by an externally conserved source—reproduces the $1/k^2$ susceptibility with TT Ward identities at machine precision. The carrier evades face (i) by leaving finite local dimension and face (iv) by spending a synthetic dimension; faces (ii) and (iii) survive because they depend on neither resource. A finite-resource law quantifies the cost of each face: local infrared observables converge exponentially in the onsite cutoff, global no-leakage costs grow only logarithmically in system size, and Ward/source exactness is the sole cost that grows linearly. Both halves of the remaining requirement admit effective-level payments by one local object: a constrained local stress sector reproduces the graviton response in Gauss-law fashion, and coupling it through a positive covariant square makes the pole the Goldstone of a noncompact diagonal shift, with the seagull symmetry-fixed rather than tuned (detuning is massive or tachyonic) and protection holding across the full shift-preserving orbit. The irreducible remainder is microscopic: a finite-dimensional gauge structure whose invariance is the shift and whose Gauss law is the constraint, and we state precisely what it must accomplish. Throughout we separate theorems from numerics from calibrated models; in particular, in the concrete realization the linear dispersion is inherited from the input connection, and the construction supplies positivity, locality, single-cone selection, and finite residue—not the exponent.

1 The question and the fact

Getting a relativistic massless spin-2 mode out of a local lattice is famously delicate. Bosonic tensor models on a lattice generically flow to Lifshitz points with dynamical exponent $z \neq 1$ rather

than to a relativistic graviton [5, 6, 7], echoing the way chiral fermions cannot live alone on a lattice [1, 2] and are recovered only by spending an extra dimension [3]. This paper asks what, precisely, obstructs an emergent linear graviton in a local translation-invariant setting, and answers that four standard-looking difficulties are one fact.

Work in linearized 3D gravity: a symmetric metric perturbation h_{ij} (six components), with the diffeomorphism gauge $h_{ij} \mapsto h_{ij} + i(k_i\xi_j + k_j\xi_i)$, transverse condition $k_i h_{ij} = 0$, and traceless condition $h_{ii} = 0$. A *local operator* is a polynomial in the lattice momentum k. The root of everything below is a statement about the minimal degree of a gauge-invariant TT operator.

Lemma 1 (No low-order local gauge-invariant TT operator)**.** *The minimal-degree local (polynomial in k), gauge-invariant, symmetric, transverse, traceless tensor operator $O_{ij}(k) = M_{ij}{}^{kl}(k)\, h_{kl}$ on the linearized 3D metric has degree three, and is unique up to scale; it is the linearized Cotton tensor [4]. Degrees one and two admit only $O \equiv 0$.*

Proof. Impose the three conditions as polynomial identities in k: gauge invariance $M(k)\, g_a(k) = 0$ on the three pure-gauge modes g_a, transversality $k_i O_{ij} = 0$ for all h, and tracelessness $O_{ii} = 0$ for all h. Writing the entries of M as homogeneous polynomials of degree p and evaluating the resulting linear system at generic momenta over-determines each polynomial identity, so its solution space is the null space of the assembled matrix. Enumeration gives solution-space dimensions 0, 0, 1 at $p = 1, 2, 3$: the smallest singular values at $p = 2$ are $O(1)$ (no solution), while at $p = 3$ exactly one is zero. The degree-three generator satisfies all three conditions at fresh momenta to 2×10^{-15}, confirming a genuine identity. Uniqueness of the degree-three TT gauge invariant identifies it with the Cotton tensor [4], the 3D analogue of the Weyl tensor. □

The obstruction is that the transverse projector, which is degree zero, is nonlocal ($\sim k_i k_j/k^2$), and every *local* representative of a gauge-invariant TT object is pushed up to Cotton order. The rest of the paper is four consequences of Lemma 1 and one construction that shows exactly which of them can be bought off and at what price.

2 The carrier that saturates the milestone

We first exhibit a concrete construction that reaches the boundary the obstructions describe: a local, positive, first-class object with exactly two TT branches and the correct conserved-source response. It requires as input a *connection vector* $n(k)$: any local, translation-invariant, $\mathbb{R}^3$-valued symbol with an isolated node at the origin at which it is linear, $n(k) = \Lambda k + O(|k|^2)$ with $\det \Lambda \neq 0$. One explicit choice from a coined tensor walk is described in Appendix A; nothing in this section depends on it beyond the stated properties.

2.1 First-class Stueckelberg quotient

With symmetric tensors A_{ij} (coordinates) and B_{ij} (momenta) and one Stueckelberg pair (ϕ, π) (fourteen real dimensions), define

$$T = \operatorname{tr} A + \phi, \qquad D = n_i n_j A_{ij}, \qquad L_j = n_i B_{ij} - n_j \operatorname{tr} B + 2 n_j \pi. \tag{1}$$

Proposition 1 (Abelian first class)**.** *For every nonzero n the five constraints (1) are first class: all ten Poisson brackets vanish. The quotient reduces fourteen dimensions to four, i.e. exactly two TT canonical pairs, removing helicity 0 and ± 1.*

The maximum bracket is 2.8×10^{-17} at generic n and 3.6×10^{-15} worst over 2000 random connections; the trace–momentum bracket from $\operatorname{tr} A$ cancels against the Stueckelberg $2n_j\pi$, and the D–L_j bracket cancels between n_iB_{ij} and $-n_j \operatorname{tr} B$.

2.2 Chiral single-cone lift

A positive lift that squares a signed carrier doubles the branches to four. The single-cone construction avoids the square with a domain wall in a synthetic coordinate s [3], using a Wilson profile $m_W(k) = \frac{1}{4}(3 - \cos k_x - \cos k_y - \cos k_z)$. Either as a positive scalar stiffness $K_{\text{ch}} = D^\dagger D + \|n\|^2 I$ with $(D\psi)_s = \psi_{s+1} - (0.42{+}2.5m_W)\psi_s$, or as a signed band

$$H_{\text{DW}}(k,s) = n_i\Gamma_i + [M_0 + \beta m_W(k) + 1]\Gamma_5 + \big[-\tfrac{1}{2}\Gamma_5 - \tfrac{i}{2}\Gamma_4\big]_{\langle s,s+1\rangle} + b_{\text{top}}\Gamma_1\,\delta_{s,W-1}, \qquad (2)$$

$(M_0, \beta, b_{\text{top}}, W) = (-1, 10, 2, 24)$, the Γ a Clifford set.

Proposition 2 (Single cone, gapped doublers, wall-localized)**.** *K_{ch} has one soft mode; H_{DW} has one boundary cone. The Wilson profile gaps the lattice doubler even where $\|n\| \to 0$, and the upper-wall mass removes the mirror branch: the low TT branch count is 2 with the wall mass and 4 without.*

At $W = 48$, K_{ch} has one machine-zero eigenvalue with second gap $\sqrt{\lambda_2} = 0.581548$ and right-wall weight $\sim 10^{-19}$; at the coin-doubler momentum $\|n\| \approx 10^{-9}$ yet the lowest frequency is 1.406. An independent H_{DW} build (distinct Clifford representation) gives 2 branches with b_{top} and 4 without, the surviving mode carrying lower-wall weight 1.000000.

2.3 Dispersion (inherited) and conserved-source response

The surviving cone has $\omega(k) = \|n(k)\| + O(|k|^3)$, with fitted exponent $z = 0.999787$; the lowest frequency equals $\|n(k)\|$ to five parts in 10^5. We state the crux plainly: the linear dispersion is *carried by the input connection*, not produced by the construction, which supplies positivity, single-cone selection, doubler removal, and finite residue. For an externally conserved TT source the carrier gives static susceptibility exponent -1.999525, dynamical exponent 0.999762, residue plateau $\lambda_0\chi = 0.823598$, and TT trace and divergence Ward residuals 2.5×10^{-16} and 1.3×10^{-18}. Since $\chi \approx Z_0/\lambda_0$ with $\lambda_0 = \|n\|^2$, the $1/k^2$ certifies a genuine massless pole with finite residue rather than independent new physics.

3 Four faces of Lemma 1

Face 1: finite dimension forces $z = 3$. For an exact compact local-character spin-2 pairing in three dimensions the TT dispersion exponent is $z = 3$, and no bounded-range polynomial canonical change of variables lowers it. By Lemma 1 the minimal local physical characters begin at $O(k)$ (kinetic) and $O(k^2)$ (potential); the symplectic pairing $\Omega(d) = Z(d)X(d)^\top$ is homogeneous of degree three with $\operatorname{rank}\Omega = 2$ everywhere off the origin (its minor-sum-of-squares factors into positive-definite quadratics times a manifestly positive sextic, bounded below by 0.25 on the unit sphere). Random positive Hessians give $z = 3.000000$ over 1200 branches, and the Schur complement $B - C^\top A^{-1}C = O(k^4)$: no auxiliary rescues an $O(k^2)$ stiffness. Independently, $\operatorname{tr}[Q,P] = 0 \neq \operatorname{tr}(iI)$ forbids $[Q,P] = iI$ in any finite dimension [8]. This is the lattice-Lifshitz phenomenon of [5, 6]. The carrier evades this face only by leaving finite local dimension: its boundary oscillator is continuous-variable.

Face 2: first-classness does not protect masslessness; a noncompact shift does. A positive local first-class mass exists and is relevant: adding $m^2 I$ to K_{ch} preserves the constraint ideal (a generic positive copy-space block commutes with the constraints to 2×10^{-15}) yet gaps the infrared, $\omega = \sqrt{m^2 + \|n\|^2} \to m$, with deep-IR exponent $z \to 0$ and crossover $k_* = m/c_g$. The absence of a *low-order local gauge-invariant* protector is Lemma 1 [10]. What does protect the pole, at the effective continuous-variable level, is a *noncompact diagonal shift*. Couple the wall mode to the local Ward stress of Face 3 through the positive covariant square $H_m = \frac{\Omega^2}{2}\|S - \alpha h_{\mathrm{wall}}\|^2$: the term linear in $S\, h_{\mathrm{wall}}$ is the finite source coupling, and the h_{wall}^2 seagull is fixed by the same square rather than tuned—integrating out S cancels it identically, because $\min_S \|S - \alpha h\|^2 = 0$. The coupled model retains an exact zero mode at $k = 0$ (eigenvalue 7×10^{-17}; a hybrid with wall and stress weights each $v_0(0)^2/(1 + v_0(0)^2) = 0.451634$), two linear TT branches, and a dynamical stress response $\chi_S \sim k^{-2.000}$. The zero mode is the Goldstone of the diagonal shift $h \mapsto h + \epsilon v_0$, $S \mapsto S + \alpha v_0(0)\epsilon$, under which the square is invariant. Detuning the seagull upward opens a gap obeying the square-root mass law (fitted exponent 0.499); downward is tachyonic: the symmetric point is the unique stable massless completion. Protection is a symmetry orbit, not a point: random positive covariant squares preserving the shift keep the exact zero and $z = 1$ in every trial, while generic shift-breaking wall terms open $O(10^{-1})$ gaps. This realizes, in Goldstone form, exactly the noncompact protection structure that the compact-character no-go of Face 1 leaves as its escape clause. What remains open is microscopic: the shift is a property of the effective Gaussian square, not yet derived from a finite-dimensional gauge structure.

Face 3: sources built from derivatives couple at Cotton order. The $1/k^2$ response above requires a conserved source, and *how* conservation is achieved fixes the coupling strength. There are three rungs. (i) A source sector of *independent local stress variables* S_{ij} with the local first-order constraint $\partial_i S_{ij} = 0$ couples at full strength: the TT subspace lies inside the divergence-free subspace, a generic constrained stress carries the exact fraction 2/3 of its weight in the TT sector (0.6662 ± 0.0007 measured against $E[\chi_2^2/(\chi_2^2 + \chi_1^2)] = 2/3$), and the per-TT response is $\chi \sim k^{-2.0000}$—the nonlocal TT projector *emerges* from eliminating a local constraint, exactly as the Coulomb interaction emerges from a local Gauss law. (ii) A source whose conservation holds *identically*, as a derivative identity—the Beltrami superpotential $S = \nabla \times \nabla \times \Phi$, degree two—is suppressed to $k^{+2.0000}$. (iii) A source built gauge-invariantly from the metric itself first appears at Cotton order (Lemma 1) and is suppressed to $k^{+3.9999}$; this exponent was verified using the enumerated degree-three generator of the Lemma itself. This is the lattice form of the soft-graviton lesson [9]: universal coupling requires matter whose stress tensor is conserved *on shell* by a constraint, not a superpotential. What the kinematic statement does not supply is the dynamics: a microscopic matter Hamiltonian whose gauge structure enforces $\partial_i S_{ij} = 0$ as a first-class constraint. Universal coupling to local stress-energy is therefore not established by the carrier, but the obstruction is now localized in the enforcement of the constraint, not in locality of the source.

Face 4: a lone chiral cone violates charge conservation. The node charge is the degree of $\hat{n}$ on a small enclosing sphere. The connection has exactly two nodes carrying equal and opposite unit charge (confirmed by two methods, $\operatorname{sign}\det J$ and the solid-angle degree; the absolute sign is an orientation convention), summing to zero: Nielsen–Ninomiya [1, 2], i.e. Poincaré–Hopf for a map from the boundaryless 3-torus [11]. Both cones are isotropic mirror partners ($\det J = \pm 9.95 \times 10^{-4}$, singular values 0.0998). *Within* locality, translation invariance, a finite-range bulk symbol, Hermiticity, and preservation of the origin cone, a single bulk cone cannot arise by smooth local deformation. The domain wall of (2) does not delete the partner; it relocates the opposite

charge to the other synthetic boundary and gaps it from the selected wall ($m_W = 0$ at the cone, 0.83 at the doubler). This is not the unique resolution: the statement is evaded by nonlocal projection, interactions or topological order, cone-destroying annihilation, or—as here —by spending a synthetic dimension [3].

4 The approximation law

Faces 1–4 rule out an exact fixed-dimensional protected graviton; what remains is a controlled approximation whose cost is governed by the soft-mode mean occupation $\bar{n} = 0.171$, i.e. $\eta = \bar{n}/(1+\bar{n}) = 0.145969$. **Local cost is exponential:** the onsite truncation tail is η^d and the canonical-commutator defect is $d(1-\eta)\eta^{d-1}$; $d = 10$ gives tail 4×10^{-9}. A positive onsite guard creates a mass $\Delta = C\sqrt{\lambda\,\delta_{\mathrm{CCR}}}$ (fitted $C = 2.372$; the $\sqrt{}$ is generic, the constant calibrated), buying a finite-size window $L_* = 2\pi c/\Delta$, not a phase. **Global cost is logarithmic:** $E[N_{\mathrm{tail}}] = WL^3\eta^d$ needs $d \sim 3\log L/|\log\eta|$, so $d = 27$ covers $L = 10^6$. **Ward/source cost is linear:** a finite-range approximation to the nonlocal TT compiler, modeled as a Yukawa regulator $m = 1/R$, has residual $m^2/(k_{\mathrm{min}}^2 + m^2)$ at $k_{\mathrm{min}} = 2\pi/L$, so fixed residual requires $R \propto L$, and at $R = L$ the residual is the L-independent constant $1/(4\pi^2+1) = 0.0247$. The exponents are robust; the two prefactors (C, the Yukawa form) are calibrated models. The one non-scalable cost is Ward/source exactness—Faces 2 and 3 as a scaling law—which no cutoff can buy down.

5 What an emergent graviton requires

Lemma 1 is one refusal of locality, and each face is its bill in a different currency. The character and canonical debts (Face 1) are paid by continuous variables. The topological debt (Face 4) is paid by a synthetic dimension, which relocates rather than deletes the partner. The carrier settles these. The *symmetry* debt (Faces 2 and 3)—no low-order local gauge invariant to protect the mass or make a local source conserved—is the linear cost of Section 4, and it is left open. That is exactly why the object is a carrier and not a phase.

The missing ingredient is now sharply localized, and both halves of the symmetry debt admit effective-level payments by a single object. The kinematic half is paid by locality of the constraint: independent local stress variables under $\partial_i S_{ij} = 0$ couple at $O(1)$, the nonlocal TT structure emerging in Gauss-law fashion (Face 3). The dynamical half is paid, at the effective continuous-variable level, by the covariant square of Face 2: the same local positive term supplies the finite source coupling *and* the symmetry-fixed seagull, so that the massless pole is the Goldstone of a noncompact diagonal shift and its protection is an orbit of the symmetry rather than a tuned point. Masslessness and universal coupling are then consequences of one structure, exactly as demanded. What remains is strictly microscopic: the diagonal shift and the Ward constraint are properties of the effective Gaussian sector, and the remaining task is to derive them from a finite-dimensional gauge structure—finite-spin, rotor, or topological order [7]—whose gauge invariance *is* the shift and whose Gauss law *is* the conservation constraint. A first step on that route exists: replacing the noncompact shift by an exact compact $U(1)$ rotor phase—with the wall coupling realized as uniform phase locking $\theta_{s+1} - \theta_s$ rather than a rational lock, which would demand integer-charge hierarchies of order 10^{25}—reproduces the two TT branches, $z \simeq 1$, and $\chi_S \sim k^{-2}$ at the spin-wave level, with the identical orbit-versus-generic protection signature, at the price of a delocalized wall residue (1/17 at synthetic width 16) and a smaller gap. The finite-dimensional cost is then qualitatively different from a hard oscillator truncation: the Holstein–Primakoff defect is $\bar{n}/J$ with $\bar{n} = 0.0952$ (halving with each doubling of the onsite spin), a symmetry-protected expansion around an ordered

sector rather than a violated commutator. What remains for a nonperturbative finite-d phase is the standard list for a compact realization: an ordered phase at finite coupling, control of vortices and anisotropies, a microscopic Gauss law for the Ward stress, survival of the full first-class quotient, and no extra Goldstone branches. The first two are confirmed for the diagonal $U(1)$ sector by direct Monte Carlo on the rotor surrogate: the ordering crossover sits at the known 3D XY critical coupling ($\beta_c \simeq 0.454$), and deep in the ordered phase the vortex density is suppressed by orders of magnitude (2×10^{-3} at $\beta = 0.8$, 10^{-5} by $\beta = 1.2$) with a large finite phase stiffness; the third is confirmed at the symbol/rank/harmonic level: on a finite periodic lattice the stress Gauss law $\partial_i S_{ij} = 0$ has rank exactly $3V - 3$, the deficit being the three identically-vanishing $k = 0$ constraints—so total momentum is the only gauge-invariant charge, precisely the counting structure of lattice electromagnetism ($V - 1$, total charge)—and the carrier coupled to the microscopically constrained stress retains both TT Goldstones, $z = 1$, the k^{-2} response, and the orbit-versus-generic protection signature. The full tensor quotient adds a fifth strike of Lemma 1: the symmetry-fixed h^2 seagull must itself be a degree-zero gauge-invariant TT bilinear, so the only exactly first-class seagulls are the nonlocal TT projector (normalizer residual 5×10^{-16}, real-space kernel with an algebraic $\sim r^{-3}$ tail) and the Cotton-order local operator; every bounded-range candidate fails at $O(1)$ (identity 0.93, trace-free 0.86, divergence-projected 0.52). The escape is the same Gauss-law move that paid Face 3, applied once more: divergence-free *and* trace-free is exactly TT (subspace dimensions $3 \wedge 5 \to 2$; random-stress TT fraction 1.000000000), and both conditions are implementable as local lattice Gauss laws with combined rank exactly $4V - 3$—total momentum again the only charge. Gauging the trace as a second local constraint, in the manner of traceless scalar-charge tensor gauge theory [7], therefore makes the constrained stress exactly TT and the covariant square over it automatically the exact seagull, with the projector's nonlocality generated by constraint elimination rather than inserted by hand. This has been implemented at the harmonic level: promoting the covariant difference $R = S - \alpha A$ to a field with local constraints $\partial_i R_{ij} = 0$ and $R_{ii} = 0$, the constraint null space regenerates the analytic TT projector to machine precision (10^{-15}), the enlarged square and both constraints are exactly invariant under the combined shift, the lattice constraint rank on the covariant field is again $4V - 3$, and the compact $r_0 = 1$ model over the doubly-constrained sector retains both Goldstones, $z = 1$, the k^{-2} stress response, the $1/17$ residue, the square-root detuning law, and full orbit protection—with each polarization's masslessness independently exact under the other's detuning. The assembled model now exists as a single Hamiltonian: full symmetric tensors on a path of $W+1$ compact nodes (matter stress at one end), path-Laplacian stiffness plus the carrier term, and the complete local constraint set (per-slice divergence and trace, covariant R-constraints), with lattice constraint rank exactly $(W+1)(4V - 3)$. In this assembled symbol the no-extra-branches obligation acquires quantitative content: the constrained infrared contains exactly two branches under the full tensor quotient, versus six if only the covariant constraint is imposed—the quotient is essential, and sufficient. Both branches have $z = 1$ to nine digits with the k^{-2} stress response, and the gap above them is the pure path gap $2\sin(\pi/2(W+1))$. A small unfactored compact Metropolis pilot (every tensor component a rotor, constraints as cosine penalties) sits deep in the ordered sector with zero topological defects and constraints locked at the thermal scale. Monte Carlo screens of the unfactored compact variables extend this: across inverse-temperature, constraint-stiffness, and finite-size scans (to $\beta = 1.2$, $\kappa = 32$, $L = 5$, $W = 4$) the ordered sector persists with charge-defect and flux-saturation fractions identically zero, while the harmonic target—now evaluated on the true coined connection with a Wilson term, so that the known doubler is gapped at 0.640—retains $z = 1$, the k^{-2} response, and the $1/17$ residue to nine digits. Flux saturation is then quantified rather than assumed: after removing the single Goldstone, the spin-wave covariance of the assembled chain is bounded ($\max\langle q^2 \rangle = 0.539$) with worst-node occupation $\bar{n} = 0.0208$—an order of magnitude colder than the bare carrier—so the local over-ceiling tail is η^{J+1} with $\eta = \bar{n}/(1 + \bar{n})$, and a $J = 16$ register

on a $24^3 \times 17$ box expects 10^{-23} saturated sites while paying a Holstein–Primakoff defect of only $\bar{n}/J = 1.3 \times 10^{-3}$. Both of those steps have now been screened at their achievable pre-many-body level. With the divergence and trace constraints imposed *exactly*—by null-space elimination in the spectral functions and by hard-quotient updates in the Monte Carlo—the ordered sector strengthens with system size ($|m| = 0.927 \to 0.959$ from $L = 4$ to 8) with vortex fraction identically zero and stable helicity, while the infrared pole residue equals the exact $1/(W+1)$ to seven digits and full Brillouin-zone scans find *no* soft point other than the origin cone (minimum second gap 0.19; the coined doubler held at 0.640). The finite-spin cost is then an error budget with one input: at onsite spin J the Holstein–Primakoff defect is $\bar{n}/J$, the pole retains spectral weight at least $1 - 4\bar{n}/J$ (99.5% at $J = 16$), and the measured residue interval $\frac{1}{W+1}(1 \pm \bar{n}/J)$ collapses onto the exact value. A first exact quantum run of that step now exists at container scale: sparse diagonalization of the finite-spin compact model (rotors hard-truncated to $|n| \leq J$) on a one-dimensional ring times the minimal Ward-square chain, in Hilbert spaces up to 1.95×10^6 states. The Goldstone appears as an Anderson tower obeying the Q^2 law with a stiffness that converges in J; the magnon branch converges to five digits by $J = 4$ to a finite compact value (11% anharmonic softening of the bare harmonic frequency); the matter-operator pole survives at exactly the Debye–Waller strength predicted by the box's own harmonic theory; the finite-spin corrections shrink geometrically in J; and every low-lying state in every charge and momentum sector is accounted for by tower plus magnons, with nothing else soft. The three-dimensional thermodynamic many-body run remains future work, but the pole's survival of genuine compactness, finite spin, and interactions is no longer a conjecture. The carrier, the source ladder, the covariant-square mechanism, and the finite-resource laws together specify what that construction must achieve and what it costs to approximate until it exists.

The companion problem of local-versus-global structure at a quantum boundary—where a globally complicated state is locally simple away from a transition region—has the same Bogoliubov origin and admits an exactly-solvable finite realization [12], which we develop separately.

Methods

All quantitative statements were regenerated from independent re-implementations. Lemma 1 is established by null-space enumeration of the gauge/transverse/traceless conditions imposed as polynomial identities at generic momenta, with the degree-three generator verified at fresh momenta to 2×10^{-15}. The carrier diagnostics (first-class algebra; single-cone spectra of K_{ch} and H_{DW}; conserved-source susceptibility and Ward residuals; mass perturbation) and the obstruction diagnostics (the degree-three symplectic pairing and $z = 3$ lock; the $[Q, P] = iI$ residual bound; the node charges by two methods) were each recomputed from scratch. The covariant-square sector of Face 2 was rebuilt from its defining equations alone (no code accompanies that audit); the inferred parameters were confirmed by six-digit matches on the hybrid weights (0.451634), the first gapped eigenvalue (0.338173), and the wall null weight (0.8236), and the exact zero mode, square-root detuning law, tachyonic branch, and protected-orbit/generic-gap contrast were each reproduced. The compact-rotor variant table was verified against an analytic residue formula $Z_S = (\sum_{s<W} r_0^{2s} + 1)^{-1}$ (all four variants exact), with the charge hierarchy, protection ensembles, and the Brillouin-zone-averaged occupation $\bar{n} = 0.0952$ underlying the $\bar{n}/J$ defect each reproduced independently. The source ladder of Face 3 was verified with independently constructed projectors (closed-form TT projector against SVD null spaces; Div P_{TT} and trace residuals at 5×10^{-17}), an independently built Beltrami map, and the Cotton rung evaluated with the enumerated generator of Lemma 1; the constrained-sector and projected-TT responses coincide identically because the TT

subspace lies inside the divergence-free subspace, so the content of that rung is the $O(1)$ overlap and the contrast with the suppressed rungs, not the equality itself. Two statements are calibrated models, flagged in place: the Cotton-order form factor used as a proxy in the original audit (now superseded by the enumerated-generator evaluation), and the guard/Yukawa prefactors of Section 4. The dispersion exponent is inherited from the input connection and is not a product of the construction.

A A concrete connection

One explicit $n(k)$ satisfying the requirements of Section 2 is the vector part of a coined $SU(2)$ tensor walk, $q(k) = \prod_\mu \big(\cos\alpha,\ \sin\alpha\, h_\mu(k)\big) \cdot (\cos\alpha, -\sin\alpha\, a_\mu)$ with cyclic axes $a = (e_2, e_3, e_1)$, $h_\mu = \cos k_\mu\, a_\mu + \sin k_\mu\, b_\mu$, and coin angle $\alpha = 0.1$; $n(k)$ is its imaginary part. It has an isolated linear node at the origin with $\Lambda = \sin\alpha\, \mathbf{1}$ (isotropic cone, slope $\sin\alpha = 0.0998$) and exactly one mirror partner on the diagonal, as required by Face 4. Any other local symbol with an isolated linear origin node yields the same carrier and the same obstruction map.

References

[1] H. B. Nielsen and M. Ninomiya, *Absence of neutrinos on a lattice. I*, Nucl. Phys. B **185**, 20 (1981).

[2] H. B. Nielsen and M. Ninomiya, *Absence of neutrinos on a lattice. II*, Nucl. Phys. B **193**, 173 (1981).

[3] D. B. Kaplan, *A method for simulating chiral fermions on the lattice*, Phys. Lett. B **288**, 342 (1992).

[4] S. Deser, R. Jackiw, and S. Templeton, *Topologically massive gauge theories*, Ann. Phys. **140**, 372 (1982).

[5] C. Xu and P. Hořava, *Emergent gravity at a Lifshitz point from a Bose liquid on the lattice*, Phys. Rev. D **81**, 104033 (2010).

[6] M. Pretko, *Generalized electromagnetism of subdimensional particles*, Phys. Rev. B **96**, 035119 (2017).

[7] M. Pretko, *Subdimensional particle structure of higher rank $U(1)$ spin liquids*, Phys. Rev. B **95**, 115139 (2017).

[8] A. Wintner, *The unboundedness of quantum-mechanical matrices*, Phys. Rev. **71**, 738 (1947).

[9] S. Weinberg, *Photons and gravitons in S-matrix theory*, Phys. Rev. **135**, B1049 (1964).

[10] K. Hinterbichler, *Theoretical aspects of massive gravity*, Rev. Mod. Phys. **84**, 671 (2012).

[11] J. W. Milnor, *Topology from the Differentiable Viewpoint* (Univ. Press of Virginia, 1965).

[12] I. C. O. Rukan, J. Gulla, and J. Skaar, *A truncated photon*, arXiv:2510.21636 (2025).