Bounded Fold Dynamics and the Particle Spectrum in VERSF

A spectral criterion, three threshold mechanisms and two native tests

Keith Taylor

November 2026

General reader summary

In VERSF, a fold is the minimally distinguishable "thing" that can exist in reality. That definition says nothing about shape, parts or energy. This paper asks a narrower question. When such a thing is realized physically, why might it cost a definite, nonzero amount of energy to disturb it?

A guitar string helps. A string of fixed length cannot sound arbitrarily low notes. Its lowest note is set by how long it is, how tightly it is stretched and how much material it carries. Let the string grow without limit, slacken it, or load it down, and the lowest note sinks toward silence. The paper proves the same thing for any connected web of adjustable quantities that are compared with their neighbours: keep the web bounded, keep its links taut and keep its load in check, and there is a lowest note. A bounded web is the simplest guarantee, not the only one: what matters in general is how well connected the web is. The paper also proves a ceiling, so the lowest note can be neither arbitrarily low nor arbitrarily high.

Two other ways of getting a lowest note are identified. One is an anchor: if every part is tied to a fixed reference, there is a lowest note however large the web becomes. The other is a twist: if travelling once around a closed loop brings you back slightly out of step, even moving the whole loop together costs energy. A worked example, six quantities on a ring and a seventh at the centre, shows all three at once. It also shows something practical. A measurement that listens to the ring alone can miss the softest motion entirely.

The example is not arbitrary. VERSF already contains a basic seven-site process, and when the example is given that process's own weights, its notes are exactly the rates at which the process settles: one half per step for the slowest. A second calculation already in the programme is derived again from its starting point and shown to get its lowest note from an anchor, not from comparisons.

The paper does not show that this wheel is the fold, or that any of its notes is a known particle. It does state the remaining gap precisely. Getting a lowest note is the easier part. If a single step of the underlying process is very fine compared with the scale of known particles, then those particles are extraordinarily low notes, and a low note always needs something small or something slack: weak links, a large load, a long chain or a narrow bridge between parts of the web, a weak anchor, or a small twist. The paper lists the routes its own calculations expose, without claiming that the list is complete. The physical dynamics must supply one and show which.

Contents

General reader summary **1**
Abstract **3**
1 The question and the contribution **3**
2 From minimal distinguishability to an excitation spectrum **4**
2.1 The primitive fold and an effective description 4
2.2 Four energies that must remain distinct 5
2.3 Three ways a threshold can arise 5
3 Comparison thresholds **6**
3.1 A floor from bounded comparisons 6
3.2 An operator form for returned matrices 7
3.3 Two ceilings, and the window between 8
4 Reference and holonomy thresholds **9**
4.1 A reference floor 9
4.2 A holonomy floor 9
5 An illustrative seven-coordinate wheel **10**
5.1 The declared action 10
5.2 Exact mode decomposition 11
5.3 What the moving reference changes 12
5.4 A twisted rim 13
6 Coordinates without inertia and mixed terms **14**
6.1 Coordinates without inertia 14
6.2 Mixed terms 14
6.3 Reduction when first-order couplings are present 16
7 Propagation and a sector threshold in space **16**
8 The same modes determine observable response **18**
8.1 Source matching and spectral weight 18
8.2 Two exact probes of the wheel 18
9 Attachment to the VERSF common action **19**
9.1 Return the stiffness, the inertia and the mixed term together 19
9.2 A native comparison instance: the canonical seven-site kernel 20
9.3 A native reference instance: the radial block of the six-defect benchmark 21
9.4 The sector without inertia 22
9.5 Clock, finite step and the scale of a threshold 23
9.6 Stability at quadratic order 24
10 Consequences for the Standard Model programme **24**
11 Conclusion **26**
Appendix A Numerical verification **26**
Appendix B Code for the numerical checks **28**
References **32**

Abstract

In VERSF, a fold is the minimally distinguishable "thing" that can exist in reality. The definition specifies neither a geometry nor an energy cost. We ask under what dynamical conditions bounded manifestations of this primitive acquire nonzero excitation thresholds, and we keep the primitive fold distinct from the effective models used to test the question. Three mechanisms are separated. A comparison threshold arises in any connected network of restoring comparisons: it is controlled from both sides by the algebraic connectivity of the network, for which bounded size is one sufficient guarantee, and it leaves common motion free. A reference threshold arises when coordinates are restored toward a fixed value, and it survives unlimited growth of the carrier. A holonomy threshold arises when transport around a closed loop fails to return a coordinate to itself, and then even common motion costs energy. A seven-coordinate wheel with a moving hub exhibits all three exactly. Its spatial extension has massive internal branches and a gapless collective branch in one quadratic action, and two response functions show how a probe can miss the softest physical mode. Coordinates without inertia and mixed time-odd terms, both of which occur in the native kernels, are treated by exact reductions, including the change of inertia that first-order couplings produce, and by a sharp degradation bound. Two native tests are then made. With the stationary weights and edge conductances of the canonical seven-site history kernel, the wheel reproduces that kernel's relaxation spectrum exactly, with gap one half. The radial block of the six-defect benchmark is derived from its residuals and shown to be a reference threshold. Finally we state what a small threshold requires in step units, conditionally on the step scale, and list the small numbers that the results of the paper expose. The wheel remains an illustrative scalar model. Nothing here derives an energy gap from minimal distinguishability or identifies Standard Model particles.

1 The question and the contribution

A fold is the minimally distinguishable "thing" that can exist in reality. In the VERSF programme this is a primitive definition, prior to a choice of geometry or effective degrees of freedom. The bounded-fold question is whether physical realizations of that distinction can have bounded internal dynamics with a positive excitation threshold. Persistence, localization and a particular internal structure must be established for the realization being studied; they are not inserted into the definition of a fold.

The route to a spectrum has two stages. First, the physical dynamics must select the admitted realizations and their evolution. Second, their excitation energies must be calculated. A restoring energy and a kinetic metric determine the frequencies of a quadratic effective description; quantization then assigns an energy quantum to each mode. Minimal distinguishability does not by itself supply either matrix or a positive minimum over the spectrum.

The VERSF sources provide common-action, record and gauge constructions that are substantially more developed than this intuition alone. Building on the conditional Gaussian constructions [3], the Yang–Mills supplement supplies a native kinetic test, an interacting strong-coupling construction, soft-record obstructions and conditional clock results [1]. The force-strength synthesis supplies the joint-response framework and explains why static response, time evolution and the physical current must belong to one parent [2]. The six-defect master action and its finite benchmark [5], the canonical seven-site history kernel [6], the attachment of opportunity count to the emergent clock [7], the closure-mass spectra [8] and the configuration-dynamics kernel [10] are used here as stated dependencies.

The calculation addresses the second stage and contributes five things.

First, three distinct mechanisms for a nonzero threshold are separated: comparison, reference and holonomy (Sections 2.3, 3 and 4). They differ in what they require and in what they leave free. A statement about "the gap" is not well posed until the mechanism is named.

Second, the comparison mechanism is given two-sided control: a floor and a ceiling set by the algebraic connectivity of the comparison network, with a second, local ceiling (Propositions 1 and 2), together with an operator form that can be applied to matrices returned by a common action (Corollary 1).

Third, a seven-coordinate wheel with a moving hub is solved exactly and shown to realize all three mechanisms, including both limits of the hub inertia and a holonomy that lifts the collective mode (Section 5).

Fourth, two features of the native kernels that a textbook oscillator analysis omits are treated exactly: coordinates with no inertia, mixed time-odd terms, and the change of inertia that arises when the two occur together (Section 6).

Fifth, two native tests are made (Section 9). Evaluated on the stationary weights and conductances of the canonical seven-site kernel, the wheel pair reproduces that kernel's relaxation spectrum exactly. The radial block of the six-defect benchmark is derived from its residuals and classified as a reference threshold.

These results establish explicit spectral mechanisms under stated dynamical conditions. The wheel is not derived here from the primitive fold. Its role is to make the spectral test concrete and to identify the quantities that a physical VERSF realization must supply. Section 9.5 states what a small threshold requires in step units.

Throughout, "proved" refers to the declared mathematical model or to a cited result under its source hypotheses. No measured mass, force strength or preferred particle assignment is used to select a coefficient.

2 From minimal distinguishability to an excitation spectrum

2.1 The primitive fold and an effective description

The primitive definition is minimal distinguishable existence. "Minimal" concerns physical distinguishability; it does not stipulate a smallest spatial length, a polygon, a graph or a collection of subparts. A fold is therefore not defined here as a persistent relational pattern. Such a pattern can be an effective description of a physical realization, provided that relation to the primitive is established.

For the spectral test, an effective model uses internal coordinates q, physical constraints, a connected network of comparisons, and an admitted evolution law. In this comparison model, "bounded" means that the number of retained internal coordinates and the diameter of the comparison network are controlled. These are hypotheses of the spectral estimate, not a replacement for the primitive definition. These bounds do not, by themselves, bound the amplitude of q, prove localization in ordinary space, or establish a stable nonlinear particle.

The connection to energy requires a further result from the physical dynamics. For an internal threshold, the admitted internal excitation energies must have a strictly positive lower bound. For a vacuum mass gap, that bound must cover the entire physical nonvacuum spectrum of the theory in

question. The existence of a minimum distinguishable thing does not specify either energy bound. The following models isolate sufficient dynamical conditions that can be tested in the VERSF construction.

In a regular small-fluctuation chart, write the physical quadratic Lagrangian as

$$L_2 = \tfrac{1}{2}\,(dq/dt)^{\mathrm{T}}\,M\,(dq/dt) - \tfrac{1}{2}\,q^{\mathrm{T}}\,K\,q, \qquad M > 0, \quad K \geq 0. \tag{1}$$

Here M is inertia and K is stiffness. For dimensionless coordinates their units are energy times time squared and energy, respectively. Both must be evaluated on the same physical carrier. Genuine gauge redundancies are removed using their independently established constraints. A surviving zero mode is not removed merely because it would spoil a gap.

The frequencies solve

$$K\,v_\alpha = \Omega_\alpha^2\,M\,v_\alpha, \qquad v_\alpha^{\mathrm{T}}\,M\,v_\beta = \delta_{\alpha\beta}. \tag{2}$$

For a finite collection of positive harmonic modes, the energy of one quantum is $\hbar\,\Omega_\alpha$. A positive internal threshold is therefore a statement about the generalized eigenvalues of the pair (K, M), not about the eigenvalues of K alone.

2.2 Four energies that must remain distinct

An internal excitation spacing measures the cost of changing the state of an already specified system. Formation energy measures the energy difference between a localized configuration and the vacuum. A particle rest energy is obtained from the dispersion of an identified physical excitation. A vacuum mass gap is the infimum of all nonvacuum physical excitation energies in the theory under consideration.

These quantities can be related in a completed dynamical construction, but their equality cannot be assumed. The wheel of Section 5 has a specified quadratic background and no separate nontrivial soliton. It therefore computes normal-mode spacings. Its spatial field extension gives an exact free-particle interpretation for those scalar branches. It does not compute the formation energy of a nonlinear VERSF fold.

For the pure Yang–Mills problem, a positive bound must concern the physical gauge-invariant spectrum of a nontrivial four-dimensional continuum theory [11]. It is not a bound on the energy of an isolated quark. Nor is it a claim that the complete Standard Model is gapped: a physical massless photon supplies excitations of arbitrarily low energy.

2.3 Three ways a threshold can arise

A quadratic restoring energy on coordinates q_i can take three elementary forms, and each produces a threshold in a different way.

Comparison. The energy is a sum of terms $\tfrac{1}{2}\,w_{ij}\,(q_i - q_j)^2$. Only differences are restored. A common displacement costs nothing, so the full carrier has a zero mode and the threshold is a property of the relative carrier. Its size is governed by how well connected the network of comparisons is.

Reference. The energy is a sum of terms $\tfrac{1}{2}\,p_i\,q_i^2$. Each coordinate is restored toward a fixed value. There is no free common displacement, and the threshold does not depend on the size or shape of the carrier.

Holonomy. The energy is a sum of terms $\frac{1}{2}\, w_{ij}\, |q_j - U_{ij}\, q_i|^2$, with transport factors U_{ij} whose product around some closed loop does not act as the identity on the coordinates. The comparisons cannot all be satisfied at once, and the residual disagreement costs energy even for common motion.

The distinction matters for the bounded-fold question because the three require different things of the physical dynamics. A comparison threshold requires connectivity that does not degrade, for which bounded size is one guarantee. A reference threshold requires a reference: something the coordinates are measured against that does not itself move. A holonomy threshold requires an admitted background with nontrivial transport around a loop. Table 1 summarizes the three, with the native instance examined later in the paper.

Table 1. Three threshold mechanisms.

Mechanism	What is restored	Common motion	Effect of growth	Native instance
Comparison	Differences between neighbours	Free	Set by connectivity; falls on growing chains, cycles and grids	Canonical seven-site kernel (Section 9.2)
Reference	Each coordinate toward a fixed value	Has a threshold	None	Radial block of the six-defect benchmark (Section 9.3)
Holonomy	Transported differences	Has a threshold	On a cycle, falls as the loop lengthens	Closure-mass spectra [8]

3 Comparison thresholds

3.1 A floor from bounded comparisons

Let a connected graph G have $n \geq 2$ vertices, an edge set E and graph diameter D. Assign a real coordinate q_i, an inertia $I_i > 0$ and positive edge weights w_{ij}. Consider

$$T = \tfrac{1}{2} \Sigma_i\, I_i\, (dq_i/dt)^2, \qquad V = \tfrac{1}{2} \Sigma_{(ij)}\, w_{ij}\, (q_i - q_j)^2, \tag{3}$$

with the second sum over the edges of G. The common displacement $q_i = s$ costs no potential energy. The internal carrier is its kinetic-orthogonal complement,

$$R_{\mathrm{rel}} = \{\, q : \Sigma_i\, I_i\, q_i = 0 \,\}. \tag{4}$$

This decomposition retains the common coordinate. It is not a declaration that common motion is gauge. Let $a(G)$ denote the algebraic connectivity of G: the second-smallest eigenvalue of its unweighted Laplacian L_G [13].

Proposition 1 (Bounded comparison threshold). Suppose $w_{ij} \geq w_{\min} > 0$ and $I_i \leq I_{\max}$. The smallest squared internal frequency is the minimum over nonzero q in R_{rel} of the quotient $\Sigma_{(ij)}\, w_{ij}\, (q_i - q_j)^2 \,/\, \Sigma_i\, I_i\, q_i^2$, and it satisfies

$$\Omega_{\mathrm{int}}^2 \;\geq\; (w_{\min} \,/\, I_{\max})\; a(G) \;\geq\; 4\, w_{\min} \,/\, (I_{\max}\, n\, D). \tag{5}$$

Consequently, a family with uniform bounds $n \le n_{\max}$, $D \le D_{\max}$, $I_i \le I_{\max}$ and $w_{ij} \ge w_{\min}$ has a uniform positive internal threshold.

Proof. Let $\mu = n^{-1} \Sigma_i q_i$. For q in R_{rel} the weighted mean is zero, so

$$\Sigma_i I_i q_i^2 = \min \text{ over } s \text{ of } \Sigma_i I_i (q_i - s)^2 \ \le \ I_{\max} \Sigma_i (q_i - \mu)^2. \tag{6}$$

The vector with entries $q_i - \mu$ is orthogonal to the constant vector, so by the definition of the algebraic connectivity,

$$\Sigma_{(ij)} (q_i - q_j)^2 \ \ge \ a(G) \Sigma_i (q_i - \mu)^2. \tag{7}$$

Combining (6), (7) and the lower edge-weight bound proves the first inequality of (5). The second is the diameter bound $a(G) \ge 4 / (n D)$ of Mohar [14]. A self-contained argument gives the same statement with the slightly weaker constant $2 / (D (n - 1))$. The variance identity and a shortest path between each pair give

$$\Sigma_i (q_i - \mu)^2 = (1/n) \Sigma_{i<j} (q_i - q_j)^2 \ \le \ \tfrac{1}{2} D (n - 1) \Sigma_{(ij)} (q_i - q_j)^2. \tag{8}$$

Indeed, each path has at most D edges, Cauchy–Schwarz bounds its squared difference by D times the sum of squared edge differences, that path sum is at most the sum over all edges, and there are $n (n - 1)/2$ pairs. This completes the proof.

The estimate shows what boundedness contributes. Together with the weight and inertia bounds it is a sufficient guarantee, because it prevents an arbitrarily long internal comparison path from defeating a fixed local restoring cost. It is not a necessary condition. The first inequality of (5) shows that the controlling quantity is the connectivity: an equal-weight star has $a(G) = 1$ however many leaves it has, so its threshold does not close as it grows. The estimate also exposes what geometry does not supply. Vanishing weights, unbounded inertia or a loss of connectivity can destroy uniformity. A connected tree satisfies the same proposition; a closed cycle is therefore neither necessary nor, without dynamical control, sufficient for this harmonic threshold.

Proposition 1 is the Poincaré inequality of a reversible Markov chain in another notation. If the weights are read as edge conductances and the inertias as stationary weights, the quotient above is the Dirichlet form divided by the variance, Ω_{int}^2 is the spectral gap of the chain, and (8) is the canonical-path estimate [15]. Section 9.2 uses this reading: the canonical seven-site kernel of the programme is such a chain.

3.2 An operator form for returned matrices

The matrices returned by a common action are not presented as a graph Laplacian and a diagonal inertia. The proposition is therefore needed in a form that compares operators.

Definition (coercive comparison block). Let $K \ge 0$ and $M > 0$ be symmetric matrices on coordinates indexed by the vertices of a connected graph G, and suppose K annihilates the constant vector $e = (1, \ldots, 1)$. The pair (K, M) has a coercive comparison block with constants $w_{\min}$ and $I_{\max}$ if, for all q,

$$q^{\mathrm{T}} K q \ \ge \ w_{\min} q^{\mathrm{T}} L_G q \quad \text{and} \quad q^{\mathrm{T}} M q \ \le \ I_{\max} q^{\mathrm{T}} q. \tag{9}$$

Corollary 1 (Operator form). If (K, M) has a coercive comparison block, then every generalized eigenvalue of $K v = \Omega^2 M v$ on the M-orthogonal complement of e satisfies $\Omega^2 \ge (w_{\min} / I_{\max}) a(G) \ge 4 w_{\min} / (I_{\max} n D)$.

Proof. The smallest internal eigenvalue is the minimum, over q not proportional to e, of $q^{\mathrm{T}} K q$ divided by the minimum over s of $(q - s\,e)^{\mathrm{T}} M (q - s\,e)$. By (9) and (7) the numerator is at least $w_{\min}\, a(G)\, \Sigma_i (q_i - \mu)^2$. The denominator is at most its value at $s = \mu$, which is at most $I_{\max}\, \Sigma_i (q_i - \mu)^2$.

Nothing in the corollary requires K to be a Laplacian or M to be diagonal. The graph enters only as a comparison. Any connected graph for which (9) holds will do, and the strongest statement comes from the graph that maximizes $w_{\min}\, a(G)$.

3.3 Two ceilings, and the window between

Proposition 2 (Two ceilings). In the setting of (3), let $d_i = \Sigma_j w_{ij}$ be the total restoring weight at vertex i and let $I_{\mathrm{tot}} = \Sigma_i I_i$. Suppose also that $w_{ij} \le w_{\max}$ and $I_i \ge I_{\min} > 0$. Then

$$\Omega_{\mathrm{int}}^2 \;\le\; \text{min over } i \text{ of } \; d_i\, I_{\mathrm{tot}} \,/\, (\, I_i\, (I_{\mathrm{tot}} - I_i)\,), \tag{10}$$

$$\Omega_{\mathrm{int}}^2 \;\le\; (w_{\max} \,/\, I_{\min})\; a(G). \tag{11}$$

Proof. For (10), displace vertex i alone. The potential energy quotient has numerator d_i. Removing the weighted mean, to bring the trial vector into R_{rel}, leaves the kinetic norm $I_i\, (I_{\mathrm{tot}} - I_i) \,/\, I_{\mathrm{tot}}$. The minimum over R_{rel} cannot exceed this trial value. For (11), take as trial vector an eigenvector f of L_G with eigenvalue $a(G)$. Its entries sum to zero, so the numerator is at most $w_{\max}\, a(G)\, \Sigma_i f_i^2$, and after removal of the weighted mean the kinetic norm is at least $I_{\min}\, \Sigma_i f_i^2$.

When weights and inertias are bounded on both sides, Propositions 1 and 2 say that the comparison threshold is the algebraic connectivity up to fixed factors:

$$(w_{\min} \,/\, I_{\max})\; a(G) \;\le\; \Omega_{\mathrm{int}}^2 \;\le\; (w_{\max} \,/\, I_{\min})\; a(G). \tag{12}$$

A growing family of such networks keeps a uniform positive threshold exactly when its connectivity stays bounded away from zero. Three examples show the range. An equal-weight star has $a(G) = 1$ at every size. A periodic cubic grid of side L, in any dimension, has $a(G) = 4 \sin^2(\pi/L)$, which falls as the grid grows. Two complete graphs joined by a single edge have diameter three at every size, but their connectivity tends to zero as they grow: a bottleneck closes the threshold without any long path.

The local ceiling (10) adds a bound set by one vertex alone. For an equal-weight cycle with n equal inertias the lowest positive squared frequency is $4\, (w/I)\, \sin^2(\pi/n)$. Its floor from (5) is $4\, (w/I) \,/\, (n\, D)$, with D the integer part of $n/2$, and its local ceiling is $2\, (w/I)\, n \,/\, (n - 1)$ (Figure 1, left). Every member of the family is closed, but the threshold tends to zero as n grows. Where a threshold relies on boundedness, that boundedness must survive the limit being taken; it is not enough that it holds at each finite stage.

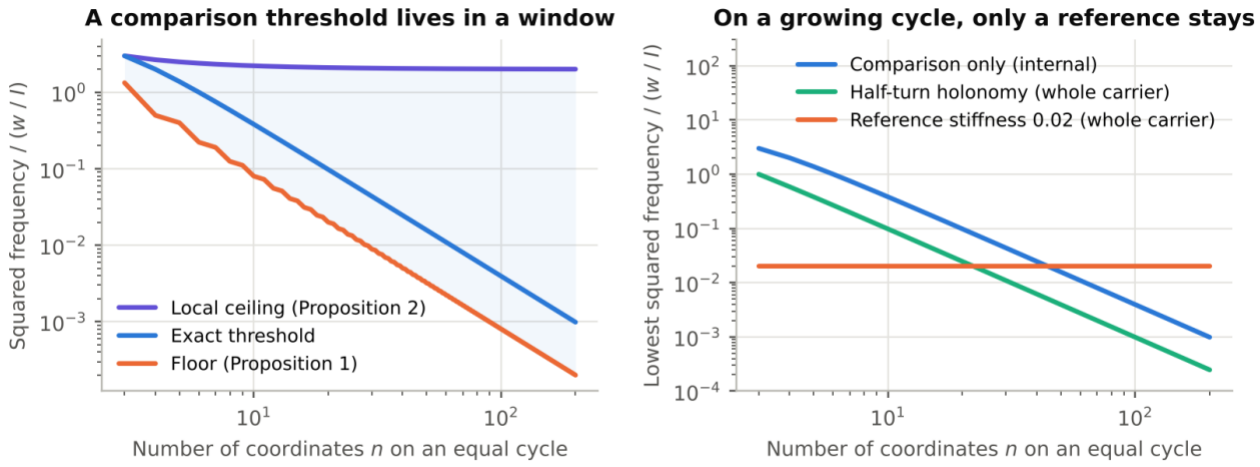


Figure 1. Left: the exact comparison threshold of an equal cycle, with the floor of Proposition 1 and the local ceiling of Proposition 2. Right: the lowest squared frequency of the same cycle under the three mechanisms. The comparison and holonomy values fall as the cycle grows; the reference value does not.

4 Reference and holonomy thresholds

4.1 A reference floor

Proposition 3 (Reference floor). Add to (3) a reference energy $V_{\mathrm{ref}} = \tfrac{1}{2} \Sigma_i\, p_i\, q_i^2$ with $p_i \geq p_{\min} > 0$. Then every squared frequency of the full carrier, common motion included, satisfies

$$\Omega^2 \;\geq\; p_{\min} \,/\, I_{\max}, \tag{13}$$

with no condition on n, on D or on the comparison weights, which may vanish. Moreover, the energy quotient restricted to R_{rel} is at least $p_{\min} / I_{\max} + (w_{\min} / I_{\max})\, a(G)$, and so is the second-smallest squared frequency of the full carrier.

Proof. For every q, $\Sigma_i\, p_i\, q_i^2 \geq p_{\min}\, \Sigma_i\, q_i^2 \geq (p_{\min} / I_{\max})\, \Sigma_i\, I_i\, q_i^2$. The comparison energy is nonnegative, and on R_{rel} it is bounded below by Proposition 1. The statement about the second frequency follows from the min-max principle, because R_{rel} has codimension one.

With unequal reference terms the common displacement is in general no longer a normal mode, and R_{rel} is not carried into itself by the dynamics. The restricted bound is then a statement about the second frequency, not the spectrum of a separate internal system. The separation is exact when the ratio p_i / I_i is the same at every vertex: the common displacement is then a normal mode, with squared frequency equal to that ratio, and all other modes lie in R_{rel}.

A reference threshold therefore needs no boundedness at all. What it needs is the reference: a value toward which each coordinate is restored and which is not itself a dynamical coordinate of the carrier. Section 5.3 shows that a reference is the limit of a comparison partner with unbounded inertia. Section 9.3 shows that the radial block of the six-defect benchmark is of this type, with the vacuum amplitude of the substrate potential as its reference.

4.2 A holonomy floor

Let the coordinates be complex and let each oriented edge carry a phase $U_{ij} = \exp(i\, \theta_{ij})$, with U_{ji} its complex conjugate. Replace the comparison energy by

$$V = \tfrac{1}{2} \Sigma_{(ij)} w_{ij} |q_j - U_{ij} q_i|^2. \quad (14)$$

A configuration costs nothing only if it is carried into itself by transport along every edge. On a connected graph such a configuration exists only if the product of the phases around every closed loop is one. Otherwise (14) has no zero mode at all [8, Theorem 2].

Proposition 4 (Holonomy floor on a cycle). On a cycle of n coordinates with equal weights w, equal inertias I and total holonomy Θ, the sum of the edge phases, the squared frequencies are

$$\Omega_k^2 = (4\, w / I) \sin^2((2\pi\, k - \Theta) / (2n)), \quad k = 0, \ldots, n - 1. \quad (15)$$

The lowest is $(4\, w / I) \sin^2(\Theta_d / (2n))$, where Θ_d is the distance from Θ to the nearest multiple of 2π. It vanishes only for trivial holonomy.

Proof. Replacing each q_j by $\exp(i\, \varphi_j)\, q_j$, a change of phase convention, moves the edge phases without changing their sum, so they may be taken equal to Θ/n. The Fourier modes $q_j = \exp(2\pi\, i\, k\, j / n)$ then diagonalize (14), with $|\exp(2\pi\, i\, k / n) - \exp(i\, \Theta / n)|^2 = 4 \sin^2((2\pi\, k - \Theta) / (2n))$.

This is the phase-cycle spectrum of the closure-mass paper [8, Theorems 3 and 4], restated with an inertia. Two features matter here. First, the floor covers the whole carrier: with nontrivial holonomy there is no free collective mode. This is the sense in which closure can cost energy where an open chain cannot. Second, on a cycle the floor still falls as the loop lengthens, like Θ_d^2 / n^2. Holonomy removes the zero mode; it does not by itself protect the threshold against growth (Figure 1, right). For the half-turn holonomy $\Theta = \pi$ on four coordinates the lowest value is $(w/I)(2 - \sqrt{2})$, the spinorial trial value of [8, Section 6].

5 An illustrative seven-coordinate wheel

5.1 The declared action

We now choose one explicit realization: six rim coordinates q_j, with $j = 0, \ldots, 5$ and $q_6 = q_0$, and a hub coordinate h. Set

$$T = \tfrac{1}{2}\, I\, \Sigma_j\, (dq_j/dt)^2 + \tfrac{1}{2}\, I_h\, (dh/dt)^2, \quad (16)$$

$$V = \tfrac{1}{2}\, \kappa\, \Sigma_j\, (q_{j+1} - q_j)^2 + \tfrac{1}{2}\, \beta\, \Sigma_j\, (q_j - h)^2, \quad \kappa, \beta, I, I_h > 0, \quad (17)$$

with all sums over $j = 0, \ldots, 5$. The rim closes, and the spokes compare it to a reference which is itself allowed to move. No coordinate is pinned to an external fixed value. This is a choice of effective comparison graph, not a decomposition of the primitive fold into seven constituent objects.

The comparison graph is the seven-site wheel that recurs in VERSF: six rim sites and a hub [5, 6, 9]. Its vertices label effective variables and its edges label restoring comparisons. The calculation depends on those connections and their weights, not on embedding the drawing as a regular hexagon. Seven is here the number of sites, the closure number written as K = 7 elsewhere in the programme; in this paper K denotes a stiffness matrix. The graph is the wheel on seven vertices, not the complete graph on seven vertices. Equations (16) and (17) are declared explicitly. We do not claim that they are the complete VERSF six-defect action, its Markov transition operator, or an already derived non-Abelian Hamiltonian. Section 9.2 shows that, at one choice of coefficients, the pair (K, M) of this wheel is nevertheless exactly the pair that governs the canonical seven-site kernel.

The full common shift of every q_j and of h by the same amount s leaves the potential unchanged. The model therefore permits internal stiffness and free collective motion simultaneously.

5.2 Exact mode decomposition

Let $r = (1/6)\,\Sigma_j\, q_j$, $u_j = q_j - r$ and $y = h - r$. Define the inertia-weighted collective coordinate and the relative inertia

$$s = (6\,I\,r + I_h\,h)\,/\,(6\,I + I_h), \qquad I_{\mathrm{rel}} = 6\,I\,I_h\,/\,(6\,I + I_h). \tag{18}$$

Since $\Sigma_j\, u_j = 0$, direct substitution gives

$$T = \tfrac{1}{2}\,(6\,I + I_h)\,(ds/dt)^2 + \tfrac{1}{2}\,I_{\mathrm{rel}}\,(dy/dt)^2 + \tfrac{1}{2}\,I\,\Sigma_j\,(du_j/dt)^2, \tag{19}$$

$$V = 3\,\beta\,y^2 + \tfrac{1}{2}\,\kappa\,\Sigma_j\,(u_{j+1} - u_j)^2 + \tfrac{1}{2}\,\beta\,\Sigma_j\,u_j^2. \tag{20}$$

The coordinates s, y and the five independent rim fluctuations decouple exactly. Fourier modes around the rim have cycle eigenvalues

$$\lambda_m = 4\,\sin^2(\pi\,m\,/\,6), \qquad m = 1, \ldots, 5, \tag{21}$$

which take the values 1, 3, 4, 3, 1.

Proposition 5 (Complete wheel spectrum). The seven frequencies consist of a zero collective mode, five rim modes and one hub–rim mode:

$$\Omega_m^2 = (\kappa\,\lambda_m + \beta)\,/\,I, \qquad \Omega_y^2 = 6\,\beta\,/\,I_{\mathrm{rel}} = \beta\,(1/I + 6/I_h), \qquad \Omega_s = 0. \tag{22}$$

The proof is the diagonal decomposition (19) to (21). In particular,

$$\Omega_{\mathrm{int}}^2 = \min\{\,(\kappa + \beta)\,/\,I,\;\; \beta\,(1/I + 6/I_h)\,\}. \tag{23}$$

For equal inertias and equal rim and spoke stiffness, the squared frequencies in units of κ/I are

$$\{\,0,\;2,\;2,\;4,\;4,\;5,\;7\,\}. \tag{24}$$

Table 2. Modes of the wheel.

Mode	Count	Squared frequency	Value for unit coefficients
Common displacement	1	0	0
First rim pair	2	$(\kappa + \beta)\,/\,I$	2
Second rim pair	2	$(3\kappa + \beta)\,/\,I$	4
Alternating rim	1	$(4\kappa + \beta)\,/\,I$	5
Hub against rim average	1	$\beta\,(1/I + 6/I_h)$	7

With all coefficients equal, (24) is the Laplacian spectrum of the wheel, and its top value 7 belongs to the hub-against-rim vector $(-6, 1, \ldots, 1)$, as recorded in the line-to-loop and closure-mass papers [8, 9]. Proposition 5 extends that static statement to an arbitrary stiffness ratio and to both inertias. The doublets arise from the real sine and cosine components of rim Fourier modes. They do not establish fermion families, colour multiplicities or spin.

5.3 What the moving reference changes

For $I_h = I$, the lowest internal branch changes at $\beta = \kappa/6$. Below that point the hub–rim motion is softer than every nonconstant rim mode. As β tends to zero, the rim remains a closed stiff loop while the internal threshold of the complete cell tends to zero.

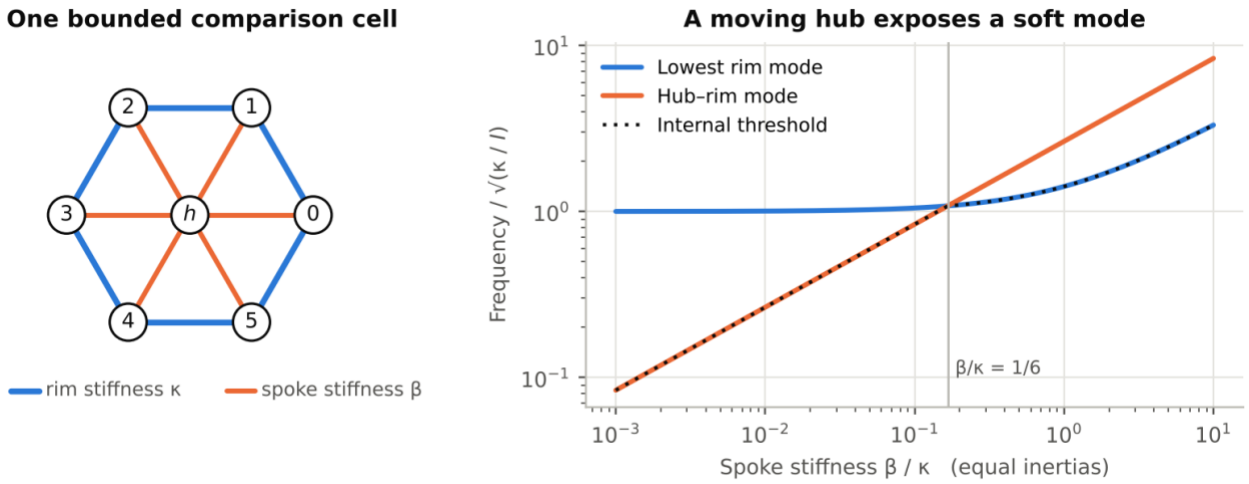


Figure 2. The declared wheel and its exact threshold. The hub is dynamical. Curves use equal inertias; the vertical line marks the branch crossing at β/κ = 1/6.

At $\beta = 0$, the disconnected hub and rim have an additional relative zero mode. This is a failure of positive coupling across the complete carrier, not a failure of the algebraic closure of the rim. The example also shows that a change of inertia can change the ordering of modes without changing the potential. Applying the spectral mechanism to a bounded realization of the fold therefore requires both matrices in (2). Neither is fixed by the drawing or by minimal distinguishability alone.

The inertia of the hub interpolates between two situations that are usually treated as different in kind.

Proposition 6 (Two limits of the hub inertia). Hold κ, β and I fixed.

(a) As I_h tends to zero, the hub is slaved to the rim average, $h = r$. The hub–rim frequency diverges. Six coordinates remain dynamical, and their spectrum is the common zero mode together with the five rim values $(\kappa\,\lambda_m + \beta)\,/\,I$.

(b) As I_h tends to infinity, the six nonzero squared frequencies tend to $(\kappa\,\lambda_m + \beta)\,/\,I$ for $m = 0, \ldots, 5$, with $\lambda_0 = 0$. These are the frequencies of the six-coordinate system in which the hub position is held fixed: a comparison ring with a reference energy of strength β, whose common rim motion has the positive value β/I. The seven-coordinate system keeps its exact zero mode at every finite I_h; in the limit that mode is carried entirely by the hub.

Proof. By (22), $\Omega_y^2 = \beta\,(1/I + 6/I_h)$ diverges in (a) and tends to β/I in (b), while the rim values do not involve I_h. In (a), stationarity in h at a fixed rim gives $h = r$, which removes the y term of (20). In (b), the collective coordinate s of (18) tends to h and keeps zero frequency, while y becomes the displacement of the rim average from the hub, with inertia tending to $6I$ and stiffness 6β.

Proposition 6 says what a reference is in dynamical terms: a comparison partner so inert that it does not respond to the coordinates measured against it. Between the two limits the same mode interpolates continuously (Figure 3, left). Large inertia alone does not remove the free motion of the full system,

however. The pinned system is a restriction. The hub position has to be fixed, by a constraint or by dynamics outside the cell, before the remaining six coordinates have a threshold on their whole carrier. A reference threshold therefore needs something a comparison threshold lacks: a partner whose position the dynamics has fixed.

In the pinned system with $\kappa = \beta$, the rim operator is $3 - S - S^{\mathrm{T}}$, with S the cyclic shift, and its values are 1, 2, 2, 4, 4, 5. This matrix is the face Gram matrix $B_C B_C^{\mathrm{T}}$ of the canonical wheel closure operator $B_C = [\, 1 \mid 1 - S \,]$ of the master action. Its Fourier sectors $k = 2$, 3 and 0 carry the values 4, 5 and 1 that the finite benchmark attaches to its three gauge-channel lifts [5]. The agreement is an identity of matrices on different carriers: vertex scalars here, link cycles there. It carries no particle assignment.

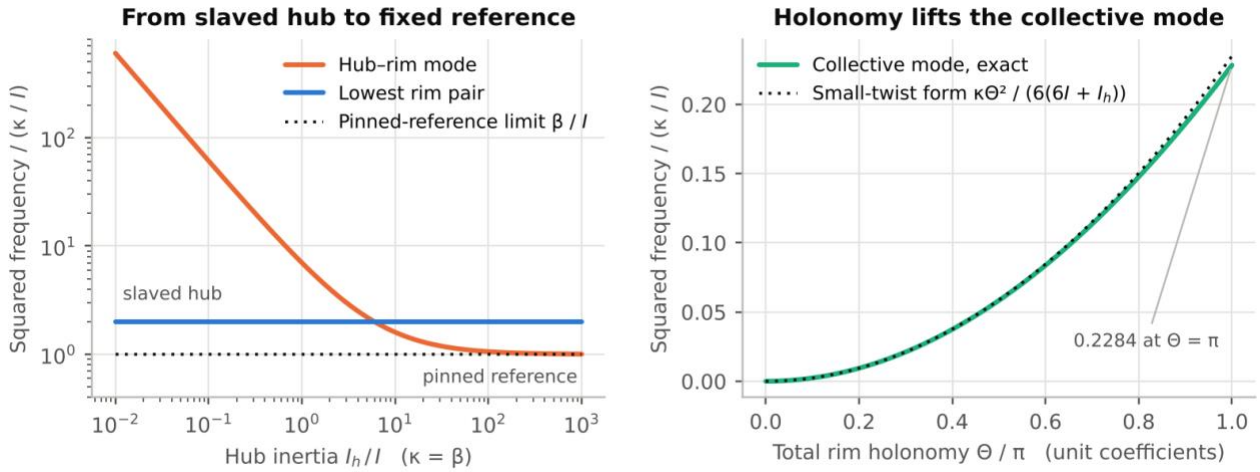


Figure 3. Left: the hub–rim mode as a function of the hub inertia, from the slaved limit to the pinned-reference limit of Proposition 6. Right: the collective mode of the twisted wheel of Proposition 7, with its small-holonomy form.

5.4 A twisted rim

Let the seven coordinates be complex, keep the spokes untwisted, and give each rim edge the phase $\Theta/6$:

$$V = \tfrac{1}{2}\,\kappa \sum_j |q_{j+1} - \exp(i\,\Theta/6)\, q_j|^2 + \tfrac{1}{2}\,\beta \sum_j |q_j - h|^2, \tag{25}$$

with the kinetic energy (16) read with squared moduli. Each of the six triangular faces then carries holonomy $\Theta/6$, and the rim as a whole carries Θ.

Proposition 7 (Holonomy lifts the collective mode). Write $\mu_\Theta = 4 \sin^2(\Theta/12)$. The rim modes $m = 1, \ldots, 5$ have squared frequencies $(4\,\kappa \sin^2((2\pi\, m - \Theta)/12) + \beta) \,/\, I$. The rim-uniform coordinate and the hub mix, and their two squared frequencies are the roots x of

$$I\, I_h\, x^2 - ((\kappa\, \mu_\Theta + \beta)\, I_h + 6\, \beta\, I)\, x + 6\, \beta\, \kappa\, \mu_\Theta = 0. \tag{26}$$

For $\Theta = 0$ the roots are 0 and $\beta\,(1/I + 6/I_h)$. For small holonomy the lower root is

$$\Omega_s^2 = 6\,\kappa\,\mu_\Theta \,/\, (6\, I + I_h) + O(\mu_\Theta^2) \;\approx\; \kappa\,\Theta^2 \,/\, (\,6\,(6\, I + I_h)\,). \tag{27}$$

Proof. Rim Fourier modes with $m \neq 0$ sum to zero and decouple from the hub, as in (20). On the rim-uniform configuration $q_j = r$, each rim term equals $\tfrac{1}{2}\,\kappa\,\mu_\Theta\,|r|^2$. The pair (r, h) therefore has the stiffness matrix with diagonal entries $6\,(\kappa\,\mu_\Theta + \beta)$ and 6β and off-diagonal entries -6β, and inertia $\mathrm{diag}(6I, I_h)$. Its characteristic equation is (26).

The collective mode is exactly free for any untwisted wheel. With a twisted rim it acquires a squared frequency proportional to the square of the holonomy and inversely proportional to the total inertia $6I + I_h$. With unit coefficients and $\Theta = \pi$ the exact value is 0.2284. The lowest rim mode at the same holonomy is 1.2679; its untwisted value is 2 (Figure 3, right). A small holonomy thus gives a small threshold to exactly the mode that would otherwise be free, and leaves the internal modes at the scale of the cell.

Whether an admitted VERSF background carries such a holonomy on the relevant carrier is a separate question. In the six-defect action a nonzero face holonomy is a closure defect, so the common-zero vacuum has none [5]. A holonomy that survives must come from the representation carried by the coordinates, as in the spinorial case of [8], or from a background that is not the common zero.

6 Coordinates without inertia and mixed terms

Two features of kernels returned by a common action fall outside the textbook pair (1). Some coordinates have no kinetic term at all, and the kernel can contain a term of first order in the time derivative. Both occur in the native six-defect benchmark (Section 9). Each can be treated exactly, and so can their combination, which changes the inertia of the coordinates that remain.

6.1 Coordinates without inertia

Proposition 8 (Static reduction). Split the coordinates as $q = (x, c)$. Let the kinetic form act only on x, with inertia $M_x > 0$, let no term of first order in the velocity involve c, and let K have blocks K_{xx}, K_{xc}, K_{cx} and K_{cc} with $K_{cc} > 0$. Then the coordinates c are fixed at each instant by stationarity, $c = -K_{cc}^{-1} K_{cx} x$, and the finite frequencies are the generalized eigenvalues of the pair (K_{eff}, M_x) with

$$K_{\mathrm{eff}} = K_{xx} - K_{xc} K_{cc}^{-1} K_{cx}. \tag{28}$$

The reduction can only soften: $K_{\mathrm{eff}} \le K_{xx}$. If K is the Laplacian of a connected weighted graph, then K_{eff} is the Laplacian of a connected weighted graph on the retained vertices, so Proposition 1 and Corollary 1 apply to the reduced pair. Giving the coordinates c a small inertia η adds frequencies that diverge as η tends to zero and returns the pair (K_{eff}, M_x) in the limit.

Proof. The equation of motion for c contains no time derivative and gives the stated relation; substitution gives (28). The subtracted term is positive semidefinite. The Laplacian statement is the Kron reduction of a weighted graph [16]. The limiting statement is the singular-perturbation property of the pencil; Proposition 6(a) is an explicit case, and Appendix A checks it on random graphs.

Two consequences follow. Under these hypotheses a coordinate without inertia does not contribute a threshold of its own, whatever its static stiffness; it contributes only by softening the coordinates that do move. And a positive static stiffness on such a coordinate cannot be converted into a frequency by supplying an inertia from outside the action. The inertia has to be returned by the same parent, directly or through coupled motion [10]. The hypothesis on first-order terms matters, and Section 6.3 treats the case in which it fails.

6.2 Mixed terms

Allow the quadratic Lagrangian a term of first order in the velocity,

$$L_2 = \tfrac{1}{2} (dq/dt)^{\mathrm{T}} M (dq/dt) + (dq/dt)^{\mathrm{T}} N q - \tfrac{1}{2} q^{\mathrm{T}} K q. \tag{29}$$

The symmetric part of N contributes a total time derivative and has no effect on the equations of motion. The antisymmetric part $A = \tfrac{1}{2}(N - N^{\mathrm{T}})$ is a gyroscopic coupling. The equation of motion is $M\,d^2q/dt^2 + 2A\,dq/dt + Kq = 0$, and the spectrum is no longer determined by the pair (K, M) [17].

Proposition 9 (Gyroscopic bound). Let $M > 0$ and $K > 0$ on the carrier considered, with smallest generalized eigenvalue Ω_0^2, and let γ be the largest eigenvalue modulus of the antisymmetric matrix $M^{-1/2} A M^{-1/2}$. Then every normal-mode frequency ω of (29) is real and satisfies

$$|\omega| \;\geq\; \sqrt{(\gamma^2 + \Omega_0^2)} - \gamma \;=\; \Omega_0^2 \,/\, (\,\gamma + \sqrt{(\gamma^2 + \Omega_0^2)}\,). \tag{30}$$

The bound is attained for two coordinates with $K = \Omega_0^2 M$, for every γ.

Proof. In the coordinates $x = M^{1/2} q$ the equation of motion has unit inertia, stiffness $K_1 = M^{-1/2} K M^{-1/2} \geq \Omega_0^2$ and gyroscopic matrix $A_1 = M^{-1/2} A M^{-1/2}$. For a solution $x = \exp(-i\,\omega\,t)\,v$ with v of unit norm, the inner product with v gives $\omega^2 + 2\,b\,\omega - k = 0$, where $k = v^{\dagger} K_1 v \geq \Omega_0^2$ and $b = v^{\dagger}(i A_1) v$ is real with $|b| \leq \gamma$, because $i A_1$ is Hermitian. Hence $\omega = -b \pm \sqrt{(b^2 + k)}$ is real, and $|\omega| \geq \sqrt{(b^2 + k)} - |b|$. The right-hand side decreases in $|b|$ and increases in k, which gives (30). For two coordinates with $K_1 = \Omega_0^2$ and A_1 the rotation generator of strength γ, the frequencies are $\sqrt{(\gamma^2 + \Omega_0^2)} \pm \gamma$.

A positive threshold therefore survives any finite gyroscopic coupling, but it is degraded, and for strong coupling it falls like $\Omega_0^2 / (2\gamma)$ (Figure 4, left). The static stiffness and the inertia no longer determine the threshold by themselves: the mixed term must be returned with them. When the zero mode of a comparison system decouples from the mixed term, the proposition applies on the relative carrier with $\Omega_0 = \Omega_{\mathrm{int}}$.

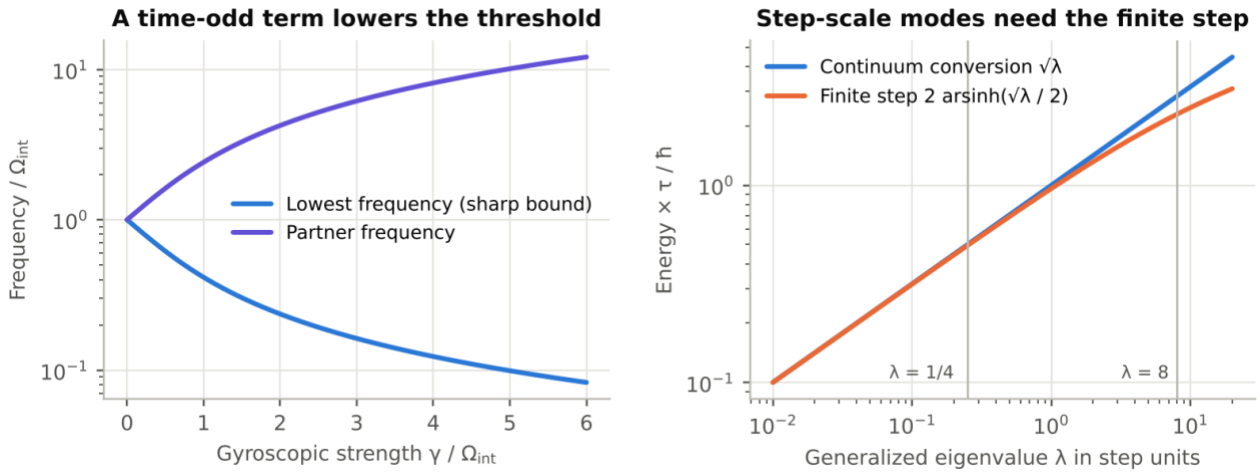


Figure 4. Left: the two frequencies of the sharp case of Proposition 9 as the gyroscopic strength grows. Right: the continuum conversion of a generalized eigenvalue to an energy, compared with the finite-step form of Section 9.5; the marked values are the squared canonical gap and the radial floor of Section 9.

The same term has a different effect when it appears in a positive history weight in place of an oscillatory action. If the quadratic weight along the ordering parameter σ is $\tfrac{1}{2} z'^{\mathrm{T}} M z' + \tfrac{1}{2} z^{\mathrm{T}} K z + z'^{\mathrm{T}} A z$, with all three matrices real and the prime denoting $d/d\sigma$, its stationary histories $\exp(s\,\sigma)$ have decay rates s that are real for $A = 0$ and can become complex otherwise. For two coordinates with unit inertia, stiffnesses k_1 and k_2 and gyroscopic strength c, the rates stay real exactly when $2|c| \leq |\sqrt{k_1} - \sqrt{k_2}|$. In the degenerate case $K = \lambda M$ they are $s = \pm\sqrt{(\lambda - c^2)} \pm i\,c$, complex for every nonzero c. A complex decay rate is not the energy of a state of a positive transfer.

Real rates do not rescue the pair (K, M) either, for two reasons. First, a surviving mixed term changes the rates even when they stay real. With unit inertia, stiffnesses 1 and 4 and $c = 0.1$, the rates are 1.006765 and 1.986561, where the pair alone gives 1 and 2. Second, real rates do not establish a positive transfer. In the same example the covariance of the first coordinate is $1.018267 / (\omega^2 + 1.013576) - 0.018267 / (\omega^2 + 3.946424)$. Its second weight is negative, and a rational scalar covariance is reflection positive only if its poles are simple with non-negative residues [18]. Whenever the antisymmetric part of a mixed history term survives, the rates are therefore to be computed from the full triple, and the positivity or reconstruction conditions that license reading them as physical energies have to be established separately. A threshold may be read from the pair (K, M) alone only when that part vanishes or has been validly removed. The radial block of Section 9.3 is such a case: its mixed term is purely symmetric and drops out.

6.3 Reduction when first-order couplings are present

The absence of a squared-velocity term does not by itself mean the absence of dynamics. If a coordinate without inertia is coupled to the velocity of another coordinate, eliminating it changes the inertia of that coordinate as well as its stiffness.

Proposition 10 (Reduction with first-order couplings). In (29), split $q = (x, c)$, let the inertia act only on x, and let the gyroscopic matrix have blocks A_{xx}, A_{xc} and $A_{cx} = -A_{xc}^{\mathrm{T}}$, with $A_{cc} = 0$ and $K_{cc} > 0$. Eliminating c is exact at every frequency and leaves a system of the same form on x, with K_{eff} as in (28) and

$$M_{\mathrm{eff}} = M_x + 4\, A_{xc}\, K_{cc}^{-1}\, A_{xc}^{\mathrm{T}}, \qquad A_{\mathrm{eff}} = A_{xx} - K_{xc}\, K_{cc}^{-1}\, A_{cx} - A_{xc}\, K_{cc}^{-1}\, K_{cx}. \tag{31}$$

In the positive history form of Section 6.2 the same elimination gives the same K_{eff} and A_{eff}, and

$$M_{\mathrm{eff}} = M_x - 4\, A_{xc}\, K_{cc}^{-1}\, A_{xc}^{\mathrm{T}}. \tag{32}$$

If A_{cc} does not vanish, the coordinates c carry first-order dynamics of their own, and the reduced kernel is a rational function of frequency, not a polynomial.

Proof. The equation of motion for c is $2\, A_{cx}\, dx/dt + K_{cx}\, x + K_{cc}\, c = 0$. Solving it for c and substituting into the equation for x, which contains the term $2\, A_{xc}\, dc/dt$, gives the stated coefficients. In the history form the stiffness enters the equations with the opposite sign, which reverses the sign of the inertia correction.

The simplest case is the history weight $\tfrac{1}{2}\,(x' + c)^2 + \tfrac{1}{2}\, c^2 + \tfrac{1}{2}\, k\, x^2$. It contains no squared derivative of c. Stationarity nevertheless gives $c = -x'/2$, and the reduced weight is $\tfrac{1}{4}\, x'^2 + \tfrac{1}{2}\, k\, x^2$: the inertia of x has fallen from 1 to 1/2, as (32) gives with $A_{xc} = 1/2$ and $K_{cc} = 2$. A static reduction of the stiffness alone would miss this. Proposition 8 is therefore to be used only after its hypothesis has been checked. In general the coordinates without inertia must be eliminated from the complete derivative-dependent quadratic form. The radial block of Section 9.3 is a native instance: eliminating the edge currents lowers the amplitude inertia from w_T to m_ℓ.

7 Propagation and a sector threshold in space

A finite internal oscillator is not yet a relativistic particle. To test the next step, place the same cell at each site x of a periodic cubic spatial lattice with spacing a. Let $Q_x = (h_x, q_{x,0}, \ldots, q_{x,5})$, let $M = \mathrm{diag}(I_h, I, \ldots, I)$, and let K be the wheel stiffness. Choose

$$L_a = \tfrac{1}{2}\, a^3 \Sigma_x \left[(dQ_x/dt)^{\mathrm{T}} M (dQ_x/dt) - Q_x^{\mathrm{T}} K Q_x \right]$$

$$- \tfrac{1}{2}\, a\, c^2 \Sigma_{x,\nu} (Q_{x+a\nu} - Q_x)^{\mathrm{T}} M (Q_{x+a\nu} - Q_x), \qquad (33)$$

where $Q_{x+a\nu}$ is the cell one lattice step from x in the spatial direction ν. In (33), M and K are the corresponding coefficients per spatial volume. Their ratio is unchanged by the common cell-volume factor. The spatial term is proportional to the same kinetic metric. That is a declared common-speed choice, not a consequence of the internal wheel.

Fourier transformation in space and the M-orthonormal internal decomposition give

$$\omega_\alpha(k)^2 = \Omega_\alpha^2 + (4\, c^2 / a^2) \Sigma_\nu \sin^2(k_\nu\, a / 2), \qquad \nu = 1, 2, 3. \qquad (34)$$

Proposition 11 (Uniform internal threshold with a gapless branch). With the positive internal coefficients fixed in physical units, the free theory (33) has an internal excitation threshold $\hbar\, \Omega_{\mathrm{int}}$ independent of external volume and spatial refinement. Its collective branch has $\Omega_s = 0$ and becomes gapless in the infinite-volume limit.

Proof. The spatial term in (34) is nonnegative and vanishes at zero wave number. Each internal branch therefore has its minimum at $\hbar\, \Omega_\alpha$, uniformly in a and in external volume. The collective branch has no constant term, and its frequency approaches $c\, |k|$ at small wave number. Arbitrarily small nonzero wave numbers are available as the volume increases.

For a finite periodic noncompact scalar model, the single exactly uniform collective coordinate is a free particle, so it does not have a normalizable oscillator vacuum. To define the finite-volume oscillator calculation below, we impose zero spatial average on that collective coordinate. This is an explicit infrared condition of the example, not a VERSF gauge identification. All nonzero collective momenta remain, and they still close the full gap as the volume grows. The internal modes are unaffected.

In the spatial continuum limit, (34) gives

$$E_\alpha(k)^2 = \hbar^2\, c^2\, |k|^2 + \hbar^2\, \Omega_\alpha^2, \qquad m_\alpha = \hbar\, \Omega_\alpha / c^2. \qquad (35)$$

Thus the six internal branches are massive scalar particles in this free model. The collective branch is a massless scalar. Its existence demonstrates compatibility between an internal threshold and a gapless physical sector; it does not identify that scalar with the photon. With the twisted rim of Section 5.4 the collective branch is no longer gapless: it becomes a scalar of mass $\hbar\, \Omega_s / c^2$, with Ω_s given by (27), and the free theory then has a positive gap equal to that mass.

The exact decomposition matters. The free Hamiltonian factorizes into collective and internal parts. The subspace containing at least one internal quantum has energy at least $\hbar\, \Omega_{\mathrm{int}}$ above the vacuum, even if it also contains arbitrarily soft collective quanta. The full untwisted theory has no positive gap. An interaction that mixes these sectors would require a new spectral and stability analysis.

For a one-dimensional periodic spatial test with $a = c = \kappa = \beta = I = I_h = 1$, the lowest nonzero collective frequency is $2 \sin(\pi/N)$. Direct diagonalization gives the values of Table 3.

Table 3. One-dimensional spatial test, frequencies in units of $\sqrt{(\kappa/I)}$.

Spatial sites *N*	**Lowest nonzero collective frequency**	**Internal threshold frequency**
8	0.765367	1.414214

Spatial sites N	Lowest nonzero collective frequency	Internal threshold frequency
16	0.390181	1.414214
32	0.196034	1.414214
64	0.098135	1.414214

This explicitly retains the spatial excitations inside the collective band. A local separation between internal modes cannot be substituted for the gap of the entire spatial operator. That distinction is also central to the spectral correction recorded in the corpus review [4].

8 The same modes determine observable response

8.1 Source matching and spectral weight

Let a source $J(t)$ couple through $L_{\text{source}} = J f^{\mathrm{T}} q$. To probe relative motion without driving the free common coordinate, require that the entries of f sum to zero. With $Q = f^{\mathrm{T}} q$, the retarded linear susceptibility is

$$\chi_f(\omega) = \Sigma_\alpha \, (f^{\mathrm{T}} v_\alpha)^2 \, / \, (\, \Omega_\alpha^2 - (\omega + i0)^2 \,), \tag{36}$$

with the sum over internal modes. The frequency locations depend on the generalized spectrum. The residues also depend on how the source couples to each mode. Canonical quantization gives

$$|\langle 0| \, Q \, |1_\alpha\rangle|^2 = (\, \hbar \, / \, (2 \, \Omega_\alpha) \,) \, (f^{\mathrm{T}} v_\alpha)^2, \qquad \chi_f(0) = 2 \, \Sigma_\alpha \, |\langle 0| \, Q \, |1_\alpha\rangle|^2 \, / \, (\, \hbar \, \Omega_\alpha \,). \tag{37}$$

These equations connect mode energy, observable transition strength and static response within one action. They also show why a measured response cannot determine unseen parts of the spectrum. A physical mode with zero overlap with f contributes no pole to this probe.

8.2 Two exact probes of the wheel

Take equal inertias. The adjacent-rim probe $Q_r = q_0 - q_1$ has no overlap with common motion or with the hub–rim mode. With $w = \omega + i0$, its response is

$$\chi_r(\omega) = (1/3) \, / \, (\kappa + \beta - I \, w^2) + 1 \, / \, (3\kappa + \beta - I \, w^2) + (2/3) \, / \, (4\kappa + \beta - I \, w^2). \tag{38}$$

The three numerator weights follow by projecting the adjacent difference onto the two rim doublets and the alternating mode: the weight on rim mode m is $\lambda_m/6$. Their sum is two, the squared Euclidean norm of that source vector.

The second probe is the physical difference $Q_y = h - r = y$. Equation (20) gives

$$\chi_y(\omega) = 1 \, / \, (6\beta - I_{\text{rel}} \, w^2), \qquad \chi_y(0) = 1 \, / \, (6\beta). \tag{39}$$

With $\kappa = \beta = I = I_h = 1$, the exact static values are $\chi_r(0) = 11/20$ and $\chi_y(0) = 1/6$. Their different values and poles are returned by the same Hamiltonian.

As β decreases, the hub–rim mode becomes arbitrarily soft and its susceptibility diverges. The rim probe remains finite, tending to $5/(6\kappa)$. Thus a finite response in one observable does not establish a gap on the whole physical carrier.

Table 4. Thresholds and static responses for $\kappa = I = I_h = 1$.

β	Internal threshold frequency	Adjacent-rim susceptibility	Hub–rim susceptibility
1	1.414214	0.550000	0.166667
0.1	0.836660	0.788213	1.666667
0.01	0.264575	0.828510	16.666667
0.001	0.083666	0.832848	166.666667

The threshold changes branch, while the rim response never sees the hub–rim pole.

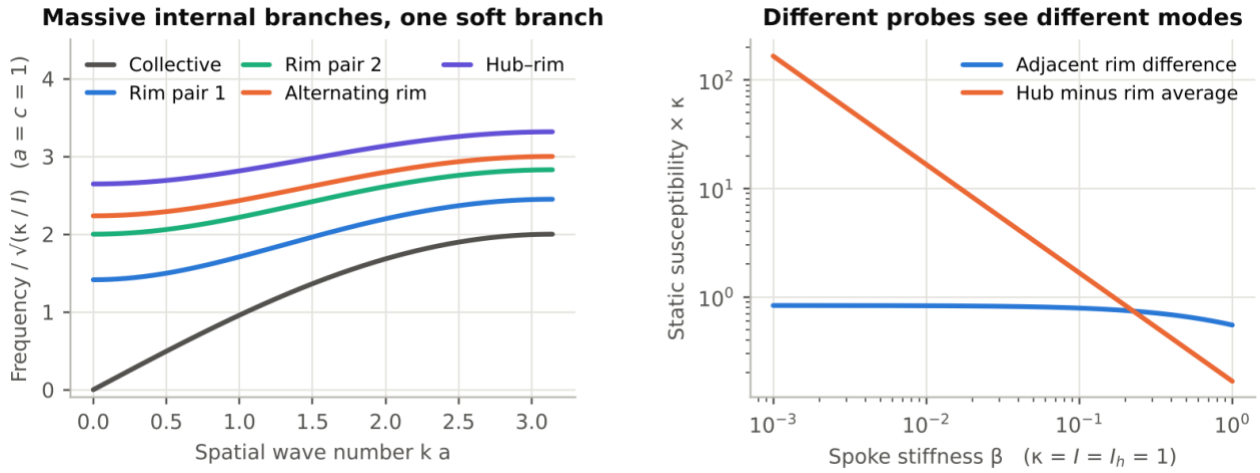


Figure 5. Left: exact spatial dispersion for the equal-weight wheel in one spatial direction. Right: the two static responses as the hub coupling decreases. The adjacent-rim observable misses the mode that closes the internal threshold.

Equations (36) to (39) are a concrete implementation of the same-parent requirement emphasized in the force-strength paper [2]. They are susceptibilities of specified scalar probes. Turning one into a gauge coupling additionally requires the physical current, representation normalization, kinetic residue and matching scale. The mode count or the gap alone supplies none of those identifications.

9 Attachment to the VERSF common action

9.1 Return the stiffness, the inertia and the mixed term together

The bounds of Sections 3 and 4 provide sufficient spectral criteria. The wheel is one example, not a required physical structure or a replacement for the primitive action. On a common-zero residual background, let z be retained physical coordinates, j algebraic auxiliary variables and W the common positive residual metric, and write the linear history residual as $r = P\,j + B\,z + Y\,z'$, with the prime denoting the derivative along the ordering parameter σ. Eliminating j replaces W by the reduced weight

$$W_{\mathrm{red}} = W - W\,P\,(P^{\mathrm{T}}\,W\,P)^{+}\,P^{\mathrm{T}}\,W. \tag{40}$$

The supported inverse and its range conditions are part of the reduction. The elimination is carried out on the complete residual, derivative terms included, so it changes the inertia and the mixed term as well

as the stiffness; it is the general reduction of Proposition 10, not the static reduction of Proposition 8. With any independently present direct terms included, the quadratic history weight is $\tfrac{1}{2}\, z'^{\mathrm{T}} M_\sigma z' + \tfrac{1}{2}\, z^{\mathrm{T}} K_\sigma z + z'^{\mathrm{T}} N_\sigma z$ per unit σ, with

$$K_\sigma = B^{\mathrm{T}} W_{\mathrm{red}} B + K_{\mathrm{direct}}, \qquad M_\sigma = Y^{\mathrm{T}} W_{\mathrm{red}} Y + M_{\mathrm{direct}}, \qquad N_\sigma = Y^{\mathrm{T}} W_{\mathrm{red}} B. \tag{41}$$

In frequency space the kernel is $K_\sigma + \omega^2 M_\sigma + i\,\omega\,(N_\sigma^{\mathrm{T}} - N_\sigma)$. Three points govern the use of (41).

First, all three coefficients are evaluated on the same carrier with the same reduced weight. The symmetric part of N_σ is a total derivative. The antisymmetric part must be retained unless it vanishes or has a justified reduction, and Section 6.2 gives its effect.

Second, (41) is a positive form: stiffness and inertia enter with the same sign. When the antisymmetric part of N_σ vanishes or has been validly removed, its stationary histories decay at rates given by the square roots of the generalized eigenvalues $\lambda(K_\sigma, M_\sigma)$; otherwise the rates are those of the full triple, as in Section 6.2. The oscillatory form (1), with a relative minus sign, is obtained on a justified local Lorentzian branch [1, 7], and the same numbers then appear as frequencies. The generalized eigenvalue is the invariant content. Whether it is read as a squared decay rate or as a squared frequency is part of the clock question of Section 9.5.

Third, retained coordinates on which M_σ vanishes need care. If no mixed term involves them, they are not oscillators and are reduced by Proposition 8. If a mixed term couples them to the velocities of other coordinates, their elimination changes the inertia of those coordinates, by Proposition 10. If a mixed term acts among them, they carry first-order dynamics of their own.

If the background residual is nonzero, second derivatives of the residual can contribute to the Hessian; $B^{\mathrm{T}} W B$ alone is then insufficient. If constraints are nonlinear, the physical Hessian is the constrained second variation, including the appropriate constraint-curvature terms.

The sources already use this common-action projection and explicitly distinguish it from adding a second fluctuation correction to an action that has already integrated those fluctuations [1, 2]. For the present application, the physical construction must first identify how the retained variables describe an admitted realization of the primitive fold. The spectral question is then whether the returned triple, on the complete physical carrier, has a coercive comparison block in the sense of (9), a reference term in the sense of Proposition 3, a holonomy, or none of these.

9.2 A native comparison instance: the canonical seven-site kernel

The history-measure paper supplies a reversible canonical kernel T_0 on the seven-site catalogue. From the hub it moves to each of the seven sites with probability 1/7. From a rim site it moves to the hub, to the site itself and to each rim neighbour with probability 1/4 [6]. Its stationary law is $\pi = (7, 4, 4, 4, 4, 4, 4)/31$. Every edge carries the same stationary flow, $\pi_i\, T_0(i, j) = 1/31$, so its Dirichlet form is $Q_7 = \mathrm{diag}(\pi)(1 - T_0) = (1/31)\, L_W$, with L_W the unweighted Laplacian of the wheel.

Proposition 12 (The canonical kernel is a wheel pair). The relaxation rates of T_0, that is, the eigenvalues of $1 - T_0$, are the generalized eigenvalues of the wheel pair of Section 5 with

$$\kappa = \beta = 1/31, \qquad I = 4/31, \qquad I_h = 7/31. \tag{42}$$

By Proposition 5 they are 0 together with

$$(\kappa + \beta)/I = 1/2 \text{ (twice)}, \quad (3\kappa + \beta)/I = 1 \text{ (twice)}, \quad (4\kappa + \beta)/I = 5/4, \quad \beta\,(1/I + 6/I_h) = 31/28. \tag{43}$$

Proof. The equation $(1 - T_0)\, v = \lambda\, v$ is equivalent to $Q_7\, v = \lambda\, \mathrm{diag}(\pi)\, v$, which is $K\, v = \lambda\, M\, v$ for the stated coefficients.

The values (43) reproduce the spectrum 1, 1/2, 1/2, 0, 0, −3/28, −1/4 of T_0 and its Perron gap 1/2 recorded in [6]. The proposition adds an assignment of each value to a mode. The gap belongs to the first rim pair. The two null directions of T_0 belong to the second rim pair, the value −1/4 to the alternating rim mode, and −3/28 to the hub against the rim average. The hub is the more inert site, in the ratio 7 to 4, and the two stiffnesses are equal. The kernel therefore sits at $\beta/\kappa = 1$, well on the rim-controlled side of the crossing, which for these inertias lies at $\beta/\kappa = 7/24$.

Proposition 1 applied to this pair is the Poincaré inequality of the canonical chain. The wheel has $a(G) = 2$, and here $w_{\min} = w_{\max} = 1/31$, $I_{\min} = 4/31$ and $I_{\max} = 7/31$. The connectivity floor of (5) is 2/7, and its diameter form, with $n = 7$ and $D = 2$, is 2/49. The local ceiling (10) is 31/36. The connectivity ceiling (11) is exactly 1/2, the calculated gap. The ceiling is attained because the lowest mode of the unweighted wheel lives on the rim, where every site has the smallest inertia. The comparison threshold of the wheel is thus not only an analogy: at the coefficients (42) it is the spectral gap of a kernel the programme already uses.

Two cautions keep the result in its place. The first concerns the reading. T_0 is a first-order law, so its generalized eigenvalues are rates, one half per step at the gap, not squared frequencies. The second-order pair with the same rates is $(K\, M^{-1}\, K, M)$, the pair of the squared residual $\tfrac{1}{2}\, (z' + M^{-1}\, K\, z)^{\mathrm{T}}\, M\, (z' + M^{-1}\, K\, z)$, whose generalized eigenvalues are the squares of (43). The one-step kernel also has negative eigenvalues and is not itself a positive transfer; the heat form $\exp(-a_t\, (1 - T_0))$ of the source has the same rates and a positive spectrum [6]. The second caution concerns the grade. The history-measure paper grades T_0 at the canonical-catalogue level, with its magnitudes as properties of the model and not as predictions, and not as the complete physical history law [6]. Proposition 12 is therefore an exact statement about the canonical kernel. It shows that the wheel pair is native to the programme. It does not promote the kernel's rates to particle energies.

9.3 A native reference instance: the radial block of the six-defect benchmark

The supplement evaluates the 226-coordinate six-defect benchmark and obtains a radial inertia, while a 94-dimensional cycle and completion sector remains frequency independent in its continuity-only history lift [1, Theorem 2.1]. For the radial restriction, a wheel Laplacian eigenvalue ℓ gives

$$m_\ell = w_J\, w_T\, /\, (w_J + w_T\, \ell), \qquad k_\ell = 2 + m_\ell\, \ell^2, \qquad \Omega_\ell^2 = \ell^2 + 2\ell\, /\, w_J + 2\, /\, w_T, \tag{44}$$

in units of the ordering parameter, where w_J and w_T are the weights of the current-constitutive and transport-conservation defects of the master action [5]. At the benchmark values $w_J = 1/2$ and $w_T = 1/4$ this is $\ell^2 + 4\ell + 8$.

The formula can be derived in a few lines, and the derivation shows where each term comes from. Let a_i be real fluctuations of the substrate amplitude about its vacuum value on the seven sites, let J collect the currents on the twelve edges, and let d be the incidence matrix of the wheel, so that $d^{\mathrm{T}}\, d = L_W$. With flat transport, three contributions involve the amplitude [5, 10]. The substrate potential has curvature $c_\rho = 2\, \lambda_\rho\, \rho_c^2$ at the vacuum, equal to 2 at the benchmark. The current-constitutive residual is $J - \kappa_J\, d\, a$, with weight w_J. The transport-conservation residual, in the continuity lift, is $a' + d^{\mathrm{T}}\, J$, with weight w_T. Together,

$$\tfrac{1}{2}\, c_\rho\, |a|^2 + \tfrac{1}{2}\, w_J\, |J - \kappa_J\, d\, a|^2 + \tfrac{1}{2}\, w_T\, |a' + d^{\mathrm{T}}\, J|^2. \tag{45}$$

Proposition 13 (Radial block). On a Laplacian eigenvector with eigenvalue ℓ and amplitude coefficient a, eliminating the currents from (45) leaves

$$\tfrac{1}{2}\, m_\ell\, (a' + \kappa_J\, \ell\, a)^2 + \tfrac{1}{2}\, c_\rho\, a^2, \qquad m_\ell = w_J\, w_T \,/\, (w_J + w_T\, \ell). \tag{46}$$

Consequently: (a) the inertia is m_ℓ and the stiffness is $c_\rho + m_\ell\, \kappa_J^2\, \ell^2$, which is (44) at $c_\rho = 2$ and $\kappa_J = 1$; (b) the mixed term $m_\ell\, \kappa_J\, \ell\, a'\, a$ is a total derivative, so no gyroscopic term arises; (c) the generalized eigenvalue is

$$\Omega_\ell^2 = \kappa_J^2\, \ell^2 + c_\rho\, (1/w_T + \ell/w_J). \tag{47}$$

Proof. Let v be the unit eigenvector and take the current along the gradient, $J = b\, d\, v$, where $|d\, v|^2 = \ell$. Then (45) becomes $\tfrac{1}{2}\, c_\rho\, a^2 + \tfrac{1}{2}\, w_J\, \ell\, (b - \kappa_J\, a)^2 + \tfrac{1}{2}\, w_T\, (a' + \ell\, b)^2$. Stationarity in b gives $b = (w_J\, \kappa_J\, a - w_T\, a') \,/\, (w_J + w_T\, \ell)$, and substitution gives (46). Current components orthogonal to $d\, v$ decouple and vanish at stationarity. For $\ell = 0$ there is no gradient, and the last term of (45) gives the inertia w_T directly. The elimination is an instance of Proposition 10 in its history form: with the current normalized along the gradient, $A_{xc} = \tfrac{1}{2}\, w_T \sqrt{\ell}$ and $K_{cc} = w_J + w_T\, \ell$, and (32) gives $w_T - w_T^2\, \ell \,/\, (w_J + w_T\, \ell) = m_\ell$.

Three readings follow.

The floor of this block is a reference threshold. At $\ell = 0$ the comparison part vanishes and (47) gives $c_\rho \,/\, w_T$, which is 8 at the benchmark. The uniform amplitude mode is not free, because the substrate potential restores the amplitude toward its vacuum value. The radial block is the situation of Proposition 3, with the vacuum amplitude as reference, and its floor would persist on an arbitrarily large carrier. It is not an instance of the bounded comparison mechanism of Proposition 1.

The comparison part is the square of a first-order rate. The term $\kappa_J^2\, \ell^2$ in (47) is the square of $\kappa_J\, \ell$, the rate at which the zero-residual law $a' = -\kappa_J\, \ell\, a$ relaxes the mode. This is the relation between first-order and second-order pairs noted after Proposition 12, here with the wheel Laplacian itself as generator.

The values are at step scale. On the unit-weight wheel spectrum of (24), the benchmark values of (47) are 8, 20, 20, 40, 40, 53 and 85, with inertias m_ℓ from 1/4 at $\ell = 0$ down to 1/18 at $\ell = 7$. These are inherited results for declared unit weights. They are not physical mass calculations.

9.4 The sector without inertia

The same source proves that the 94-dimensional cycle and completion block is frequency independent in the continuity lift, with no quadratic temporal term [1, Theorem 2.1]. Static elimination of such a block needs more than a frequency-independent diagonal block: no derivative term may couple it to the coordinates that remain. In the continuity lift the only derivative is the amplitude velocity in the transport-conservation residual, and that residual involves the currents only through their divergence, as in (45). The cycle currents are divergence free, so at the common-zero background no derivative term couples them to the amplitudes. For the remaining coordinates of the block we rely on the cited theorem. Under that condition Proposition 8 applies, and the coordinates of the block are not oscillators of that lift: they are fixed by stationarity at each step, and they contribute to the retained coordinates only through the reduction (28). If a derivative cross-coupling survives, in this lift or in another, the block must be eliminated by Proposition 10 instead, and it then changes the inertia of the coordinates it is coupled to. In either case a positive static cycle stiffness cannot be inserted into (2) with a freely supplied physical inertia and called a native prediction. An inertia for these coordinates must be returned by the same parent: through temporal faces, through the bath block, or through coupled

motion [10]. The appropriate continuation is the source-supported operational generator or temporal-face construction, with its clock and mixed terms attached. This is a specific limitation of the tested history lift, not an absence theorem for all VERSF dynamics.

9.5 Clock, finite step and the scale of a threshold

Let σ count opportunities. On the homogeneous local branch, the attachment of opportunity count to the emergent clock is linear: δ opportunities correspond to the clock interval $\delta\, a_t / c^*$, with a_t and c^* the temporal grid spacing and propagation speed of the emergent description [7, Theorems 3 and 4]. Write $\tau = a_t / c^*$ for this conversion factor. It converts an ordering count into the clock that completed facts generate. It is not a duration that an opportunity occupies against a background time, and it does not fix laboratory seconds [7]. Taking the source action quantum to be $\hbar$ and $t = \tau\, \sigma$, the oscillatory form

$$S_2 = \tfrac{1}{2}\, \hbar \int d\sigma\, [\, z'^{\mathrm{T}}\, M_\sigma\, z' - z^{\mathrm{T}}\, K_\sigma\, z\,] \tag{48}$$

implies

$$M_{\mathrm{phys}} = \hbar\, \tau\, M_\sigma, \qquad K_{\mathrm{phys}} = (\hbar / \tau)\, K_\sigma, \qquad \Omega_{\mathrm{phys}}^2 = \lambda(K_\sigma, M_\sigma) / \tau^2. \tag{49}$$

This is the coefficient conversion used in the supplement [1, Section 2.2]. It must be applied on the admitted domain of the clock. A change of units alone does not control every physical mode through a continuum limit.

Equation (49) treats σ as continuous. A finite-step kernel gives a different relation, and the relation depends on the step rule. For the symmetric second-difference step, in which a mode obeys $z_{n+1} - 2\, z_n + z_{n-1} = \lambda\, z_n$ in the history form, histories decay by the factor $\exp(-\varepsilon)$ per step, with

$$\lambda = 4\, \sinh^2(\varepsilon/2), \qquad E = \hbar\, \varepsilon / \tau = (2\, \hbar / \tau)\, \operatorname{arsinh}(\, \sqrt{\lambda}\, / \, 2\,). \tag{50}$$

This is the form of the finite-step energy formula of [7]. It agrees with $\hbar\, \sqrt{\lambda}\, / \tau$ up to the factor $1 - \lambda/24$ for small λ, and departs from it at step scale: at $\lambda = 8$ it gives $2.292\, \hbar/\tau$ against $2.828\, \hbar/\tau$ (Figure 4, right). The relation is not universal. In a first-order kernel a mode with rate λ is multiplied by $1 - \lambda$ at each step, so $\exp(-\varepsilon) = 1 - \lambda$ defines a real, finite decay exponent only when $1 - \lambda$ is positive. In the canonical kernel of Section 9.2 that covers the two gap modes alone. They are halved at each step, which as a continuous rate is ln 2 per step, against 1/2 per step for the heat form of the same generator. The modes for which $1 - \lambda$ is zero or negative have no such exponent in the one-step kernel and need the separate treatment noted in Section 9.2.

Corollary 2 (Small numbers behind a small threshold). Let $\varepsilon = m\, c^2\, \tau\, /\, \hbar$ be the mass of an excitation in units of the step scale $\hbar\, /\, (\tau\, c^2)$, so that $\lambda = \varepsilon^2$ in the continuum conversion (49) and $\lambda = 4\, \sinh^2(\varepsilon/2)$ in the step form (50). The following hold in the continuum conversion.

(a) A comparison threshold has $\varepsilon^2 \geq (w_{\min} / I_{\max})\, a(G)$. A small ε requires a small ratio of restoring weight to inertia, a small connectivity, or both. A small connectivity requires $n\, D \geq 4 / a(G)$, but it can come from a bottleneck as well as from long paths. For an equal cycle, $n \approx 2\pi\, \sqrt{(w/I)}\, /\, \varepsilon$.

(b) A reference threshold has $\varepsilon^2 \geq p_{\min} / I_{\max}$. A small ε requires a reference stiffness that is small in units of the inertia.

(c) The collective mode of the twisted wheel has $\varepsilon^2 \approx \kappa\, \Theta^2\, /\, (6\, (6I + I_h))$. A small ε requires a small holonomy, a small rim stiffness or a large total inertia.

(d) A mixed term lowers any of these. In the sharp case of Proposition 9, $\varepsilon \approx \varepsilon_0^2 / (2\,\gamma\,\tau)$ for strong gyroscopic coupling, where ε_0 is the value without it.

Proof. Equations (5), (13), (27) and (30), with (49). The statement on connectivity is the second inequality of (5).

The corollary states the scale question without resolving it. A nonzero threshold is the easy direction: Propositions 1, 3 and 4 each give one, and at declared unit weights the native values are of order one in step units, as (43) and (47) show. Whether a small mass is a demanding requirement depends on the step scale, which the clock attachment leaves open. If the step scale lies far above the mass of an observed particle, ε is a very small number, and some small number must produce it. The corollary lists the ones that the results of this paper expose: a weak coupling or a large inertia, as in the hub–rim mode of the wheel at small β; a poorly connected or very long comparison network; a weak reference; a small holonomy; a strong mixed term. For a cycle, (a) says that the number of coordinates around the loop is about $2\pi\sqrt{(w/I)}$ times the reduced Compton wavelength of the excitation in units of $c\,\tau$. The list classifies the routes found here and is not exhaustive. Reductions change inertias (Proposition 10), and interactions can do what no quadratic form does. None of these small numbers is supplied by minimal distinguishability, and none may be selected by a measured mass. Each must be returned by the admitted dynamics with its physical weights.

The uniformity in Proposition 11 concerns external spatial refinement while a fixed internal model is held in physical units. It does not prove that a primitive VERSF refinement preserves seven effective internal coordinates or their coefficients. Holding an internal threshold fixed in physical units while the step conversion is refined requires λ to scale like τ^2, which is the small number of Corollary 2 in another form. If additional internal coordinates survive refinement, they must enter the spectral test; if they are eliminated, their frequency-dependent feedback and observable weight must be controlled.

9.6 Stability at quadratic order

On a fixed physical relative carrier, suppose a perturbed pair (K_1, M_1) has the same null vector as (K, M) and satisfies

$$K_1 \geq (1 - \delta_K)\,K, \qquad 0 < M_1 \leq (1 + \delta_M)\,M, \qquad 0 \leq \delta_K < 1. \tag{51}$$

The Rayleigh quotient gives

$$\Omega_{\mathrm{int},1}^2 \;\geq\; \left((1 - \delta_K) / (1 + \delta_M) \right) \Omega_{\mathrm{int}}^2. \tag{52}$$

If the perturbation also introduces a gyroscopic term of strength γ, Proposition 9 applies to the perturbed pair, and the threshold is at least $\sqrt{(\gamma^2 + \Omega_{\mathrm{int},1}^2)} - \gamma$. Thus the harmonic threshold is stable under controlled relative changes of stiffness, inertia and mixed term. This is not a stability theorem for an interacting quantum vacuum. Nonlinear interactions can change the vacuum, couple internal excitations to a gapless sector, and turn stable free particles into resonances. Those effects require the actual admitted Hamiltonian and observables.

10 Consequences for the Standard Model programme

The bounded-fold proposal starts from minimal distinguishable existence. Its relevance to the Standard Model depends on the physical dynamics connecting that primitive to the admitted matter and gauge states and their spectrum. This work should build on the existing representation and gauge architecture.

The corpus review and the supplement credit conditional non-Abelian gauge structure, linked curvature vertices, local perturbative quantum completion, a Higgs-line metric relation, a conditional clock, source-functional fermion kinetics and later mixed-tensor compatibility [1, 4].

The models of this paper test spectral mechanisms: internal comparisons, references or holonomies produce restoring modes; the physical kinetic metric and any mixed term set their time scale; quantization supplies their excitation energies; and current operators determine how those excitations appear in response. Each step is explicit in the wheel. Selection of that effective description from the primitive, and identification of its excitations with physical particles, remain separate requirements.

The classification has a direct use. A claim that a sector of the programme is gapped should name its mechanism, because the three behave differently under growth of the carrier and require different inputs. On present evidence the canonical kernel supplies a comparison threshold with a free collective mode, the radial block supplies a reference threshold, and the closure-mass spectra supply holonomy thresholds. Corollary 2 then prices each in step units.

For quarks and leptons, the scalar wheel coordinates must be replaced by, or embedded in, the admitted fermionic carrier. Spin, chirality, statistics, gauge representation and flavour mixing have to be inherited from that construction. The wheel doublets alone carry none of those properties. In particular, its six internal real modes are not evidence for the six quark flavours.

Electroweak fermion masses and the W and Z masses require the appropriate Higgs and Yukawa structure; the Standard Model relations depend on those couplings and the symmetry-breaking vacuum [12]. A bounded-fold account would need to calculate the corresponding source-normalized operators or their consequences. Counting closed modes cannot replace that calculation.

The colour problem is different again. The electroweak completion retained by the VERSF source chain is colour singlet, so it does not directly supply a gluon mass matrix [1, 4]. A Yang–Mills gap concerns the physical colour-singlet spectrum of the pure gauge theory, including neutral excitations. The scalar internal thresholds proved here are examples of spectral mechanisms, not such a proof.

Finally, the collective scalar of (33) illustrates why a gapped internal sector can coexist with massless propagation, and Proposition 7 shows what it takes to give that scalar a mass. Identifying a photon requires the correct gauge field, constraints, transverse polarizations and current coupling. The collective scalar is not assigned that role.

A concrete next calculation is now specified. First identify an admitted physical realization of the minimal distinguishable fold and the effective variables that describe it, without prescribing a wheel topology. On one admitted VERSF background and physical carrier, evaluate the full constrained triple (41), any memory terms and the source vertices. Eliminate the coordinates without inertia from the complete form, by Proposition 8 or Proposition 10 as their couplings require. Identify the actual relative subspace and its reference dynamics, and classify each threshold as comparison, reference or holonomy. Test the generalized spectrum in the common physical clock, with the finite-step form where the values are at step scale, and then evaluate the same modes in the fermion and gauge observables. Finally, identify which small number of Corollary 2, if any, produces the lowest admitted mode, and say where the dynamics supplies it. That execution would establish whether these sufficient criteria apply to a realization selected by VERSF dynamics. A different native structure must be assessed through its own full spectrum, not forced into the wheel model.

11 Conclusion

A fold is the minimally distinguishable "thing" that can exist in reality. This paper preserves that primitive definition and studies conditional routes from an effective dynamical description to an excitation threshold. Three mechanisms are separated. A comparison threshold is controlled by the connectivity of the network and lies between the floor of Proposition 1 and the ceilings of Proposition 2. Bounded size with controlled weights and inertias guarantees it, and no hexagonal structure is required. A reference threshold needs a partner whose position is fixed, and survives growth. A holonomy threshold needs nontrivial transport around a loop and reaches the collective mode. The wheel shows all three exactly: six internal branches have thresholds while the common branch remains free, the hub inertia connects the comparison system to the pinned one, and a twisted rim lifts the collective mode. Its spatial extension yields massive scalar branches alongside gapless collective propagation, and its response functions expose a soft mode invisible to a selected probe.

Two native tests give the wheel a place in the programme that the drawing alone could not. With the canonical kernel's stationary weights and conductances, the wheel pair reproduces that kernel's relaxation spectrum exactly and assigns its gap of one half to the first rim pair. The radial block of the six-defect benchmark follows from its residuals in a few lines, and its floor is a reference threshold supplied by the substrate potential.

The remaining physical connection is precise. The admitted VERSF dynamics must relate the primitive distinction to its realizations, supply their full dynamics and clock, including coordinates without inertia and any mixed term, and determine their excitation energies and observable identities. Minimal distinguishability is the starting point; a spectral bound is a consequence only when its dynamical hypotheses are met. Beyond that, if the step scale lies far above observed masses, the dynamics must supply a small number. A threshold is nonzero under mild conditions. A mass far below the step scale needs a weak coupling or a large inertia, a poorly connected or long comparison network, a weak reference, a small holonomy, a strong mixed term, or a route outside this list, and the dynamics must say which. The calculations provide a test of that connection, not a definition of the fold in terms of an oscillator or a hexagon. Establishing the Standard Model spectrum requires applying the test on the admitted fermion and gauge carriers.

Appendix A Numerical verification

All numerical values in this paper are dimensionless model evaluations. The analytic results do not depend on the numerical checks. Every random test uses one generator with seed 20261115, in the order of the listing in Appendix B. Random graphs have a number of vertices drawn uniformly from 2 to 12 and an edge probability drawn uniformly between 0.15 and 0.9, and disconnected samples are discarded. Weights and inertias are independent and uniform between 0.1 and 5, and reference terms between 0.2 and 3, unless the listing states otherwise. Table 5 gives each check with its scope and worst case.

Table 5. Numerical checks.

Result	Check	Worst case
Propositions 1 and 2	4,000 random connected weighted graphs with random inertias	exact / floor ≥ 1.002; local ceiling / exact ≥ 1.000 (attained for two vertices); connectivity ceiling / exact ≥ 1.002
Section 3.3, families	stars of 5, 50 and 200 vertices; two complete graphs of 5, 20 and 40 vertices joined by one edge; a 6 by 6 periodic grid	connectivity 1 for every star; 0.298, 0.091, 0.048 at diameter 3; grid value 1
Proposition 3	1,500 random graphs with random reference terms; lowest and second frequencies	exact / floor ≥ 1.04; common motion a normal mode for constant p_i / I_i, residual 2×10^{-14}
Proposition 4	cycles of 3, 4, 6, 7 and 12 coordinates at five holonomies	residual 3×10^{-15}
Proposition 5	36 combinations: κ in {0.2, 1, 3}, β in {0.001, 0.1, 1, 5}, three inertia pairs	residual 3×10^{-14}
Proposition 6	hub inertia 10^{-6} and 10^{6} at unit coefficients; hub position held fixed	spectra {0, 2, 2, 4, 4, 5}, {0, 1, 2, 2, 4, 4, 5} and {1, 2, 2, 4, 4, 5}
Proposition 7	three coefficient sets at four holonomies; values at Θ = π	residual 8×10^{-15}; collective 0.228380, lowest rim 1.267949
Proposition 8	253 random graphs, 1 to $n - 2$ coordinates without inertia, limit $\eta = 10^{-9}$	residual 4×10^{-6}; reduced matrix a Laplacian to 8×10^{-15}
Proposition 9	3,000 random systems of 2 to 6 coordinates	all frequencies real; lowest / bound ≥ 1.0002
Section 6.2, history form	2,000 random two-coordinate systems; the example with stiffnesses 1 and 4	criterion for real rates agrees in every case; example rates and weights as quoted, covariance residual 4×10^{-16}
Proposition 10	400 random systems in each form; finite roots of the full and reduced pencils	residual 2×10^{-14} (oscillatory), 2×10^{-13} (history)
Proposition 11 and Table 3	full 7N spectrum against the branch formula, N = 8, 16, 32, 64	residual 4×10^{-14}
Section 8 and Table 4	pseudoinverse of the seven-coordinate stiffness at four values of β	agreement to the digits shown
Proposition 12	spectrum of $1 - T_0$ against the wheel pair; squared-residual pair; connectivity bounds	residual 9×10^{-16}; floor 2/7 and ceiling 1/2
Proposition 13	twelve edge currents eliminated at four frequencies, two weight sets	residual 2×10^{-14}
Section 9.5	second-difference step relation at four values	residual 5×10^{-16}

The wheel stiffness is assembled as a sum of twelve edge matrices, six rim and six spoke, and passed to a symmetric generalized eigensolver. The two static susceptibilities are computed by Moore–Penrose inversion; the source vectors annihilate the common mode, so the pseudoinverse gives the unique source response despite the arbitrary common coordinate. The spatial stiffness on N sites is the Kronecker sum of the wheel stiffness K and the cyclic second-difference matrix $L_N = 2 - S_N - S_N^{\mathrm{T}}$, with all coefficients equal to one, and its complete spectrum is compared with every branch of (34), so that the internal threshold of Table 3 is tested directly. The exact uniform collective coordinate is excluded only as described in Section 7. Systems with a singular inertia are solved as a linear pencil in the coordinates and velocities, and the finite roots of the full and reduced systems are matched one to one.

These checks establish the implementation of the declared models. They do not select coefficients or attach coordinates to a physical VERSF particle.

Appendix B Code for the numerical checks

The listing below reproduces every entry of Table 5. It requires numpy, scipy and networkx.

```
"""Numerical checks for "Bounded Fold Dynamics and the Particle Spectrum in VERSF".

Run:  python3 bounded_fold_checks.py        (numpy, scipy, networkx; a few seconds)
Every random test uses the generator numpy.random.default_rng(SEED) below, in the order written.
Results are printed and written to bounded_fold_check_results.json.

Sampling rules
  random graph    : n uniform in 2..12; Erdos-Renyi edge probability uniform in [0.15, 0.9];
                    disconnected samples are discarded
  weights, inertia: independent uniform in [0.1, 5] unless stated
  reference terms : independent uniform in [0.2, 3]
"""
import json
import numpy as np
import networkx as nx
from scipy.linalg import eig, eigh

SEED = 20261115
rng = np.random.default_rng(SEED)
R = {}


def edge_matrix(n, i, j, phase=1.0):
    e = np.zeros(n, complex if np.iscomplexobj(phase) else float)
    e[j], e[i] = 1.0, -phase
    return np.outer(e, e.conj())


def random_graph():
    while True:
        n = int(rng.integers(2, 13))
        g = nx.gnp_random_graph(n, rng.uniform(0.15, 0.9), seed=int(rng.integers(1e9)))
        if nx.is_connected(g):
            return n, g


def laplacians(n, g, lo=0.1, hi=5.0):
    K, L, w = np.zeros((n, n)), np.zeros((n, n)), []
    for i, j in g.edges:
        w.append(rng.uniform(lo, hi))
        K += w[-1] * edge_matrix(n, i, j)
        L += edge_matrix(n, i, j)
    return K, L, np.array(w)


def wheel(kappa, beta, theta=0.0):
    """Hub is coordinate 0, rim coordinates 1..6; theta is the total rim holonomy."""
    K = np.zeros((7, 7), complex if theta else float)
```

```
    for j in range(6):
        K += kappa * edge_matrix(7, 1 + j, 1 + (j + 1) % 6, np.exp(1j * theta / 6) if theta else 1.0)
        K += beta * edge_matrix(7, 1 + j, 0)
    return K


def gen_eigs(K, M):
    return np.sort(eigh(K, M, eigvals_only=True))


# ---- Propositions 1 and 2: floor, local ceiling, connectivity ceiling
worst = dict(floor_connectivity=np.inf, floor_diameter=np.inf, ceiling_local=np.inf,
             ceiling_connectivity=np.inf)
N = 4000
for _ in range(N):
    n, g = random_graph()
    K, L, w = laplacians(n, g)
    I = rng.uniform(0.1, 5, n)
    lam, a, D = gen_eigs(K, np.diag(I))[1], np.sort(np.linalg.eigvalsh(L))[1], nx.diameter(g)
    worst['floor_connectivity'] = min(worst['floor_connectivity'], lam / (w.min() * a / I.max()))
    worst['floor_diameter'] = min(worst['floor_diameter'], lam / (4 * w.min() / (I.max() * n * D)))
    local = np.min(np.diag(K) * I.sum() / (I * (I.sum() - I)))
    worst['ceiling_local'] = min(worst['ceiling_local'], local / lam)
    worst['ceiling_connectivity'] = min(worst['ceiling_connectivity'], (w.max() * a / I.min()) / lam)
R['prop1_2'] = dict(graphs=N, **worst)

# families: connectivity, not size, controls the threshold
star = [np.sort(nx.laplacian_spectrum(nx.star_graph(n - 1)))[1] for n in (5, 50, 200)]
barbell = [(nx.diameter(g), np.sort(nx.laplacian_spectrum(g))[1])
           for g in (nx.barbell_graph(m, 0) for m in (5, 20, 40))]
torus = np.sort(nx.laplacian_spectrum(nx.grid_2d_graph(6, 6, periodic=True)))[1]
R['families'] = dict(star_connectivity=star, barbell_diameter_and_connectivity=barbell,
                     torus_6x6=torus, torus_formula=4 * np.sin(np.pi / 6) ** 2)

# ---- Proposition 3: reference floor
worst, N, eq_resid, uneq = np.inf, 1500, 0.0, 0.0
for _ in range(N):
    n, g = random_graph()
    K, L, w = laplacians(n, g)
    I, p = rng.uniform(0.1, 5, n), rng.uniform(0.2, 3, n)
    ev = gen_eigs(K + np.diag(p), np.diag(I))
    a = np.sort(np.linalg.eigvalsh(L))[1]
    worst = min(worst, ev[0] / (p.min() / I.max()))
    if n > 1:
        worst = min(worst, ev[1] / (p.min() / I.max() + w.min() * a / I.max()))
    # common motion is a normal mode exactly when p_i / I_i is constant
    e = np.ones(n)
    eq_resid = max(eq_resid, np.abs((K + np.diag(0.7 * I)) @ e - 0.7 * I * e).max())
    v = np.linalg.solve(np.diag(I), (K + np.diag(p)) @ e)
    uneq = max(uneq, np.ptp(v))
R['prop3'] = dict(graphs=N, min_ratio=worst, common_mode_residual_equal_ratio=eq_resid,
                  common_mode_spread_unequal=uneq)

# ---- Proposition 4: holonomy on a cycle
resid = 0.0
for n in (3, 4, 6, 7, 12):
    for theta in (0.0, 0.3, np.pi, 2.0, 5.5):
        H = sum(edge_matrix(n, j, (j + 1) % n, np.exp(1j * theta / n)) for j in range(n))
        exact = np.sort([4 * np.sin((2 * np.pi * k - theta) / (2 * n)) ** 2 for k in range(n)])
        resid = max(resid, np.abs(np.sort(np.linalg.eigvalsh(H)) - exact).max())
R['prop4'] = dict(residual=resid, half_turn_four=4 * np.sin(np.pi / 8) ** 2,
                  two_minus_root2=2 - np.sqrt(2))

# ---- Proposition 5: wheel spectrum;  Proposition 6: hub inertia
resid, lam_m = 0.0, [4 * np.sin(np.pi * m / 6) ** 2 for m in range(1, 6)]
for kappa in (0.2, 1, 3):
    for beta in (0.001, 0.1, 1, 5):
        for I, Ih in ((1, 1), (0.3, 2), (2, 0.7)):
            exact = np.sort([0] + [(kappa * l + beta) / I for l in lam_m] + [beta * (1 / I + 6 / Ih)])
            ev = gen_eigs(wheel(kappa, beta), np.diag([Ih] + [I] * 6))
            resid = max(resid, np.abs(ev - exact).max())
R['prop5'] = dict(combinations=36, residual=resid)
slaved = gen_eigs(wheel(1, 1), np.diag([1e-6] + [1.0] * 6))
```

```
inert = gen_eigs(wheel(1, 1), np.diag([1e6] + [1.0] * 6))
pinned = np.sort(np.linalg.eigvalsh(wheel(1, 1)[1:, 1:]))          # hub position held fixed
R['prop6'] = dict(slaved_hub=slaved[:6].round(6).tolist(), inert_hub=inert.round(6).tolist(),
                  hub_fixed=pinned.round(6).tolist(), inert_hub_zero_mode=float(abs(inert[0])))

# ---- Proposition 7: twisted wheel
resid = 0.0
for kappa, beta, I, Ih in ((1, 1, 1, 1), (0.7, 2.0, 0.4, 1.9), (3, 0.2, 1.5, 0.6)):
    for theta in (0.2, 1.0, np.pi, 4.0):
        mu = 4 * np.sin(theta / 12) ** 2
        b = (kappa * mu + beta) * Ih + 6 * beta * I
        disc = np.sqrt(b * b - 24 * I * Ih * beta * kappa * mu)
        exact = [(b - disc) / (2 * I * Ih), (b + disc) / (2 * I * Ih)]
        exact += [(4 * kappa * np.sin((2 * np.pi * m - theta) / 12) ** 2 + beta) / I
                  for m in range(1, 6)]
        ev = gen_eigs(wheel(kappa, beta, theta), np.diag([Ih] + [I] * 6).astype(complex))
        resid = max(resid, np.abs(ev - np.sort(exact)).max())
ev = gen_eigs(wheel(1, 1, np.pi), np.eye(7, dtype=complex))
small = gen_eigs(wheel(1, 1, 0.01), np.eye(7, dtype=complex))[0]
R['prop7'] = dict(residual=resid, collective_at_pi=ev[0], lowest_rim_at_pi=ev[1],
                  small_twist_ratio=small / (0.01 ** 2 / 42))

# ---- Proposition 8: static reduction (no first-order coupling)
resid, kron, N = 0.0, 0.0, 0
for _ in range(300):
    n, g = random_graph()
    if n < 4:
        continue
    N += 1
    K, _, _ = laplacians(n, g, 0.2, 3.0)
    nc = int(rng.integers(1, n - 1))
    x, c = slice(0, n - nc), slice(n - nc, n)
    Keff = K[x, x] - K[x, c] @ np.linalg.solve(K[c, c], K[c, x])
    I = rng.uniform(0.3, 3, n - nc)
    limit = gen_eigs(K, np.diag(list(I) + [1e-9] * nc))[:n - nc]
    resid = max(resid, np.abs(limit - gen_eigs(Keff, np.diag(I))).max())
    kron = max(kron, np.abs(Keff.sum(1)).max(), (Keff - np.diag(np.diag(Keff))).max())
R['prop8'] = dict(graphs=N, residual=resid, laplacian_violation=kron)


def pencil_roots(M, A, K, sign):
    """Finite roots s of  M s^2 + 2 A s + sign K = 0  (sign +1 oscillatory, -1 history form)."""
    n = len(K)
    left = np.block([[np.zeros((n, n)), np.eye(n)], [-sign * K, -2 * A]])
    right = np.block([[np.eye(n), np.zeros((n, n))], [np.zeros((n, n)), M]])
    s = eig(left, right, right=False)
    return s[np.isfinite(s) & (np.abs(s) < 1e7)]


def same_set(a, b):
    a, b = np.array(a), list(np.array(b))
    if len(a) != len(b):
        return np.inf
    worst = 0.0
    for z in a:
        k = int(np.argmin([abs(z - y) for y in b]))
        worst = max(worst, abs(z - b.pop(k)))
    return worst


def spd(n, floor=0.3):
    B = rng.normal(size=(n, n))
    return B @ B.T + floor * np.eye(n)


# ---- Proposition 9: gyroscopic bound
worst, N, real = np.inf, 3000, 0.0
for _ in range(N):
    n = int(rng.integers(2, 7))
    M, K = spd(n), spd(n)
    A = rng.normal(size=(n, n)) * rng.uniform(0, 3)
    A = (A - A.T) / 2
    s = pencil_roots(M, A, K, +1)
```

```
    real = max(real, np.abs(s.real).max())
    Lc = np.linalg.cholesky(M)
    gamma = np.linalg.norm(np.linalg.solve(Lc, np.linalg.solve(Lc, A).T).T, 2)
    Om2 = gen_eigs(K, M)[0]
    worst = min(worst, np.abs(s.imag).min() / (np.sqrt(gamma ** 2 + Om2) - gamma))
R['prop9'] = dict(systems=N, min_ratio=worst, largest_real_part=real)

# history form, two coordinates, unit inertia: rates real exactly when 2|c| <= |sqrt(k1) - sqrt(k2)|
agree, N = 0, 2000
for _ in range(N):
    k1, k2, c = rng.uniform(0.1, 4), rng.uniform(0.1, 4), rng.uniform(0, 1.2)
    s = pencil_roots(np.eye(2), c * np.array([[0, 1.0], [-1.0, 0]]), np.diag([k1, k2]), -1)
    agree += (np.abs(s.imag).max() < 1e-7) == (2 * c <= abs(np.sqrt(k1) - np.sqrt(k2)))
R['history_form'] = dict(trials=N, criterion_agrees=int(agree))

# the same form with real rates: the mixed term still moves them, and one spectral weight is negative
k1, k2, c = 1.0, 4.0, 0.1
roots = pencil_roots(np.eye(2), c * np.array([[0, 1.0], [-1.0, 0]]), np.diag([k1, k2]), -1)
ex_rates = np.sort(np.abs(roots.real))[::2]
pole_a, pole_b = ex_rates ** 2
weights = [(k2 - pole_a) / (pole_b - pole_a), (k2 - pole_b) / (pole_a - pole_b)]
om = 0.7
D = np.array([[om ** 2 + k1, 2j * om * c], [-2j * om * c, om ** 2 + k2]])
fit = weights[0] / (om ** 2 + pole_a) + weights[1] / (om ** 2 + pole_b)
R['history_example'] = dict(rates=ex_rates.tolist(), poles=[pole_a, pole_b], weights=weights,
                            residual=float(abs(np.linalg.inv(D)[0, 0] - fit)))

# ---- Proposition 10: reduction with first-order couplings
resid_osc, resid_hist, N = 0.0, 0.0, 400
for _ in range(N):
    nx_, nc = int(rng.integers(2, 5)), int(rng.integers(1, 4))
    n, x, c = nx_ + nc, slice(0, nx_), slice(nx_, nx_ + nc)
    # oscillatory form: inertia on x only, gyroscopic block A_cc = 0
    Mx, K = spd(nx_), spd(n)
    A = rng.normal(size=(n, n))
    A = (A - A.T) / 2
    A[c, c] = 0
    M = np.zeros((n, n))
    M[x, x] = Mx
    Kcc_inv = np.linalg.inv(K[c, c])
    Keff = K[x, x] - K[x, c] @ Kcc_inv @ K[c, x]
    Aeff = A[x, x] - K[x, c] @ Kcc_inv @ A[c, x] - A[x, c] @ Kcc_inv @ K[c, x]
    Meff = Mx + 4 * A[x, c] @ Kcc_inv @ A[x, c].T
    resid_osc = max(resid_osc, same_set(pencil_roots(M, A, K, +1), pencil_roots(Meff, Aeff, Keff, +1)))
    # history form built as a sum of squares  |B z + Y z'|^2 / 2  with no z' on the c coordinates
    B, Y = rng.normal(size=(n + 3, n)), rng.normal(size=(n + 3, n))
    Y[:, c] = 0
    K, M, Nm = B.T @ B, Y.T @ Y, Y.T @ B
    A = (Nm - Nm.T) / 2
    Kcc_inv = np.linalg.inv(K[c, c])
    Keff = K[x, x] - K[x, c] @ Kcc_inv @ K[c, x]
    Aeff = A[x, x] - K[x, c] @ Kcc_inv @ A[c, x] - A[x, c] @ Kcc_inv @ K[c, x]
    Meff = M[x, x] - 4 * A[x, c] @ Kcc_inv @ A[x, c].T
    assert np.linalg.eigvalsh(Meff).min() > -1e-9
    reduced = pencil_roots(Meff, Aeff, Keff, -1)
    resid_hist = max(resid_hist, same_set(pencil_roots(M, A, K, -1), reduced))
R['prop10'] = dict(systems=N, residual_oscillatory=resid_osc, residual_history=resid_hist,
                   example_inertia=1 - 4 * 0.5 * (1 / 2) * 0.5)

# ---- Proposition 11 and Section 8: spatial test and responses
resid, table3 = 0.0, []
for Ns in (8, 16, 32, 64):
    S = np.roll(np.eye(Ns), 1, axis=0)
    full = np.kron(np.eye(Ns), wheel(1, 1)) + np.kron(2 * np.eye(Ns) - S - S.T, np.eye(7))
    branches = np.sort([O + 4 * np.sin(np.pi * k / Ns) ** 2
                        for O in (0, 2, 2, 4, 4, 5, 7) for k in range(Ns)])
    ev = np.sort(np.linalg.eigvalsh(full))
    resid = max(resid, np.abs(ev - branches).max())
    table3.append((Ns, round(float(np.sqrt(ev[1])), 6), round(float(np.sqrt(2.0)), 6)))
rim, hub = np.zeros(7), np.zeros(7)
rim[1], rim[2] = 1, -1
hub[0], hub[1:] = 1, -1 / 6
table4 = []
```

```
for beta in (1, 0.1, 0.01, 0.001):
    P = np.linalg.pinv(wheel(1, beta), rcond=1e-12)
    threshold = float(np.sqrt(min(1 + beta, 7 * beta)))
    table4.append((beta, round(threshold, 6), round(float(rim @ P @ rim), 6),
                   round(float(hub @ P @ hub), 6)))
R['prop11'] = dict(residual=resid, table3=table3, table4=table4)

# ---- Proposition 12: canonical seven-site kernel
T0 = np.zeros((7, 7))
T0[0, :] = 1 / 7
for j in range(1, 7):
    for k in (0, j, 1 + j % 6, 1 + (j - 2) % 6):
        T0[j, k] = 1 / 4
pi = np.array([7] + [4] * 6) / 31
K, M = wheel(1 / 31, 1 / 31), np.diag(pi)
rates = np.sort(np.linalg.eigvals(np.eye(7) - T0).real)
squared = gen_eigs(K @ np.linalg.inv(M) @ K, M)
connectivity = np.sort(np.linalg.eigvalsh(wheel(1, 1)))[1]
R['prop12'] = dict(residual=float(np.abs(rates - gen_eigs(K, M)).max()), rates=rates.round(8).tolist(),
                   connectivity_floor=float(connectivity / 7),
                   connectivity_ceiling=float(connectivity / 4),
                   stationary_residual=float(np.abs(pi @ T0 - pi).max()),
                   squared_residual_pair=float(np.abs(squared - rates ** 2).max()))

# ---- Proposition 13: radial block, twelve edge currents eliminated
edges = [(1 + j, 1 + (j + 1) % 6) for j in range(6)] + [(0, 1 + j) for j in range(6)]
d = np.zeros((12, 7))
for k, (i, j) in enumerate(edges):
    d[k, j], d[k, i] = 1, -1
ell, V = np.linalg.eigh(d.T @ d)
resid = 0.0
for wJ, wT, c_rho, kJ in ((0.5, 0.25, 2.0, 1.0), (0.7, 1.3, 0.9, 0.6)):
    for om in (0.0, 0.3, 1.1, 2.7):
        rows = np.vstack([np.hstack([-kJ * np.sqrt(wJ) * d, np.sqrt(wJ) * np.eye(12)]),
                          np.hstack([1j * om * np.sqrt(wT) * np.eye(7), np.sqrt(wT) * d.T]),
                          np.hstack([np.sqrt(c_rho) * np.eye(7), np.zeros((7, 12))])])
        H = rows.conj().T @ rows
        reduced = H[:7, :7] - H[:7, 7:] @ np.linalg.solve(H[7:, 7:], H[7:, :7])
        for l, v in zip(ell, V.T):
            m = wJ * wT / (wJ + wT * l)
            resid = max(resid, abs(v @ reduced @ v - (c_rho + m * (om ** 2 + kJ ** 2 * l ** 2))))
R['prop13'] = dict(residual=float(resid), benchmark=[int(round(l * l + 4 * l + 8)) for l in ell])

# ---- Section 9.5: symmetric second-difference step
resid = max(abs(np.exp(-e) - 2 + np.exp(e) - 4 * np.sinh(e / 2) ** 2) for e in (0.05, 0.5, 1.5, 2.3))
R['step'] = dict(residual=float(resid), energy_at_8=float(2 * np.arcsinh(np.sqrt(8) / 2)),
                 continuum_at_8=float(np.sqrt(8)))


def plain(o):
    if isinstance(o, dict):
        return {k: plain(v) for k, v in o.items()}
    if isinstance(o, (list, tuple)):
        return [plain(v) for v in o]
    return float(o) if isinstance(o, (np.floating, np.integer)) else o


R = plain(R)
R['seed'] = SEED
json.dump(R, open('bounded_fold_check_results.json', 'w'), indent=1)
print(json.dumps(R, indent=1))
```

References

Papers of the VERSF Theoretical Physics Programme, references [1] to [10], are available at versf-eos.com.

[1] K. Taylor. *Physical Dynamics and the Yang–Mills Mass Gap in VERSF*. 12 November 2026. Sections 1 to 3, 6, 8 to 10 and Appendix A.

[2] K. Taylor. *From Primitive Dynamics to Physical Force Strengths in VERSF*. 9 November 2026. Sections 3, 5 and 7 to 9.

[3] K. Taylor. *Commitment Constraints and the Yang–Mills Mass Gap in VERSF*. 10 November 2026. Conditional Gaussian comparison results developed further in [1].

[4] *VERSF Corpus Review for the Yang–Mills Programme*. Research review for Keith Taylor, 10 November 2026. Sections 2 to 6. Used for the stated dependency summary and the corrected bulk-band calculation; the underlying corpus theorems retain their own hypotheses.

[5] K. Taylor. *The VERSF Master Substrate Action and Constitutive-Descent Theorem* (MASD-1). The six-defect action, the canonical wheel closure operator and the CFB-1 finite benchmark.

[6] K. Taylor. *The VERSF History-Measure, Generation-Geometry and Branch-Closure Theorem* (HMGBC-1). The canonical seven-site kernel, its stationary law, spectrum and heat form.

[7] K. Taylor. *The Sequential Opportunity-to-Lorentzian Commitment-Time Attaching Map in VERSF*. Theorems 3 and 4 and the finite-step energy formula.

[8] K. Taylor. *Why Matter Has Mass and Space Does Not: The Price of Closure in VERSF*. Theorems 2 to 4 and Section 6.

[9] K. Taylor. *The Refined Line-to-Loop, Commitment-Handoff and Axle-Loop Theorem in VERSF*. The wheel Laplacian spectrum and the hub-against-rim vector.

[10] K. Taylor. *Configuration Dynamics, Refinement Geometry and Record-Conditioned Curvature Response in VERSF*. Parts VIII and XIII: temporal faces, the transport-conservation residual, the exact reduction rule and induced inertia.

[11] A. Jaffe and E. Witten. *Quantum Yang–Mills Theory*. Clay Mathematics Institute Millennium Problem description. claymath.org.

[12] M. Cepeda, L. Reina and P. Savard. *Status of Higgs Boson Physics*, revised August 2025, Sections 11.2.1 and 11.2.2. In F. Takahashi et al. (Particle Data Group), *Review of Particle Physics*, International Journal of Modern Physics A 41, 2630011 (2026). pdg.lbl.gov.

[13] M. Fiedler. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal* 23, 298–305 (1973).

[14] B. Mohar. Eigenvalues, diameter, and mean distance in graphs. *Graphs and Combinatorics* 7, 53–64 (1991).

[15] P. Diaconis and D. Stroock. Geometric bounds for eigenvalues of Markov chains. *Annals of Applied Probability* 1, 36–61 (1991).

[16] F. Dörfler and F. Bullo. Kron reduction of graphs with applications to electrical networks. *IEEE Transactions on Circuits and Systems I* 60, 150–163 (2013).

[17] P. Lancaster. *Lambda-matrices and Vibrating Systems*. Pergamon Press, Oxford (1966).

[18] F. Arici, D. Becker, C. Ripken, F. Saueressig and W. D. van Suijlekom. Reflection positivity in higher derivative scalar theories. arXiv:1712.04308 (2017), Theorem 3.7.