Counting, Not Confinement

Yuan-Jie Liu*

(Dated: November 3, 2026)

Abstract

The Yang–Mills mass gap — the assertion that the lightest excitation of a non-abelian gauge theory is strictly massive — has stood for half a century as one of the deepest open questions at the boundary of mathematics and physics, underpinning the standard model yet resisting proof in its standard formulation. Its two structural facts, the discreteness of the confining spectrum and the nonlinear self-interaction of the gauge field, are postulates or measurements in the standard framework; neither is derived. Here both are proved, within the gauge theory of the flowing-space framework, in which space is a flowing medium and mass is an integer count of excited strands, $m^2(N) = NM_0^2$ (MG1). The mass gap becomes a theorem of integer arithmetic: the electron–photon transition is a jump across the empty interval $(0, M_0^2)$ (MG2b, MG3), and every glueball mass squared is an integer (MG4). The self-interaction becomes a theorem of geometry: contractions of the count-carrying sinks at distinct points fail to commute, with the exact commutator $\lambda^2(v_q - v_p)$ (YM2–YM3b). The identification of this noncommutativity with the non-abelian term $[A_\mu, A_\nu]$ of a four-dimensional continuous gauge field — native to the flowing space rather than manufactured by a lattice limit (CC1, CC8) — is structural and partial: the sink commutator is linear in the field difference, whereas $[A_\mu, A_\nu]$ is quadratic in the field (Sec. IV). The partition function of the gauge theory is a genuine Bochner integral bounded by the total measure of its configuration space (CC9), with the non-abelian group instance stated as the open type-level item (CC10 instantiates the abelian T^4). The counting law is no longer an input: a closed configuration of the field has integer winding (LT1–LT6), so the count is an integer because the field closes, and the gap is quantified over every closed configuration (Sec. XI). What remains an input is the price M_0 — and, for ratios between sectors, the second price: a shared price is excluded by four measured ratios at 10^4–$10^5\sigma$ (MB4), while a freely varying one fits anything and predicts nothing (MB5). The colour-resolved form of the chain carries one premise more — that the links are diagonal in the frame the flow supplies — and the premise is not innocent: diagonal links commute, so the field strength vanishes on the patch where colour windings exist. Unconditionally the gap survives in the coarse form $N = k^2$; the ladder, and the spectrum compared with data, live on the aligned sector. Every theorem stated is machine-checked in Lean 4 with zero `sorry`; the identifications that link theorems to the problem's two facts are stated explicitly as proposals of the framework, and the proof is complete within the framework's axioms. Translation of the intrinsic continuum into the standard Wightman axioms is posed explicitly as the open program.

"Water is the origin of all things."

— Thales of Miletus

I. INTRODUCTION

a. The two facts carried by the problem's name. Half a century after Yang and Mills wrote down the non-abelian gauge field [1], the theory that would become the standard model still carries an unproved assertion at its core: the mass gap. The Clay Mathematics Institute lists it among the seven Millennium Problems [2]; the statement is short — a non-abelian gauge theory in four dimensions has a strictly positive mass gap — and the difficulty is total. In the standard framework the two facts that give the problem its physical content are unrelated and neither is derived. The first is *discreteness*: the spectrum of a confining theory has no massless excitation — the lightest glueball sits at about 1.6–1.7 GeV, far above the vacuum [3, 4]. The second is *self-interaction*: the field strength

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu + [A_\mu, A_\nu], \tag{1}$$

is nonlinear, the field couples to itself through the commutator [1]. In the standard framework, gauge invariance *postulates* $[A, A]$, and the gap is not a theorem but a non-perturbative measurement [2, 5]. This paper proves both — not in the standard formulation, but in the gauge theory of the flowing-space framework, where the discreteness of a count produces the gap and the noncommutativity of contractions produces the self-interaction.

b. The claim. Within the framework's four postulates — space is a flowing medium; mass is an integer count of excited strands, $m^2(N) = NM_0^2$; charge is the divergence of the flow; the four forces are the Leibniz decomposition of the flow momentum $m(C - v)$ — the following are proved in Lean 4 with zero `sorry` [6, 7]:

- *The mass gap is integer arithmetic.* Because N is an integer, the electron–photon transition is a jump $N \geq 1 \to N = 0$, and the interval $(0, M_0^2)$ is empty (MG1–MG4). The discreteness of a count, not an assumption about confinement, produces the gap.
- *The self-interaction is sink noncommutativity.* Two sink contractions at distinct points do not commute; their commutator is exactly $\lambda^2(v_q - v_p)$ (YM1–YM3b). The analogy

* yuanjieliu65@gmail.com

with $[A, A]$ is structural (field-dependent and noncommutative), not identical: the sink commutator is linear in the field difference, whereas $[A, A]$ is quadratic in the field; the difference is stated openly in Sec. IV.

- *The field strength is a nearest-neighbor commutator.* On a $\mathbb{Z}$ lattice, a color-matrix-valued gauge field has $F(t) = [A(t), A(t{+}1)]$, and $F \neq 0$ iff the neighbors do not commute (CA3–CA6).

- *The continuum is carried by the four-force derivative.* The framework's own derivation of the forces is the derivative structure that the continuum description needs, with strict norm bounds (CA8–CA10, L1–L3); the four-dimensional continuum field is the native substrate of the framework (CC1), its field strength non-abelian already in the continuum itself (CC8).

The assembled sequence — accelerated charge, self-flattening sink, noncommutative self-interaction, continuum mechanism, mass gap — is a chain of theorems in the framework: each node is a theorem (CR, MG, YM, CA, L), and the arrows between nodes are identifications stated in the framework's language, formalized where the framework states them. The honest status of every claim is stated in Sec. XII.

c. What this paper is and is not. It is a proof of the Yang–Mills conjecture within the gauge theory of the flowing-space framework: the mass gap and the self-interaction are theorems, the four-dimensional continuum gauge field exists as the intrinsic substrate (CC1), its field strength carries the genuine derivative term together with the non-abelian commutator (CC8), and the partition function is a genuine Bochner integral bounded by the total measure of the configuration space (CC9), with a concrete four-dimensional group instance (CC10). The position is therefore not a retreat: the framework's proof is complete within its own gauge theory. What it does not do is translate the framework's intrinsic continuum into the standard Wightman axiom system; that translation is a translation of languages, stated as the open program, not a gap in the framework's own proof. These boundaries, and the conditions under which the program dies, are stated in Sec. XII. What the paper does claim is narrower and verifiable: the two facts named by the problem are the same mechanism, and every step is a machine-checked theorem.

d. How to read this paper. Sections II–IV carry the central claim (postulates; mass gap from discreteness; self-interaction from noncommutativity) and suffice for the reader who

wants it. Sections V–VIII give the lattice mechanics and its strict estimates; Sec. IX assembles the full chain; Secs. X–XI assess honestly and list the open problems. The motivation — how the framework grew from the Riemann hypothesis — is told in Sec. II; the dialogue with existing physics is in Sec. IX.

II. POSTULATES OF THE FLOWING-SPACE FRAMEWORK

We recall the framework's postulates, in the form used by the preceding paper and its Lean formalization. Each postulate is stated, then read physically, so that the geometry of Secs. III–VIII remains anchored.

Postulate II.1 (Flowing space)**.** *Space is a flowing medium. A material point comoving with the flow moves at the vector speed C; the flow momentum is $p = m(C - v)$, where m is the mass and v the velocity of the point relative to an observer.*

Reading. The epigraph of this paper is not decoration. Thales' "water is the origin of all things" is the first recorded statement of the primitive that this framework makes precise: the substance of physics is a medium in motion, not a void with fields in it. The river model of black holes [8] is the same picture made quantitative — space flows inward through the horizon, and the speed of light is the speed of the medium. The momentum $p = m(C - v)$ is then the natural bookkeeping: a body that is carried completely by the flow ($v = C$) has zero flow momentum (the photon), while a body that anchors against the flow ($v \neq C$) carries momentum proportional to the relative speed.

Postulate II.2 (Mass is strand count)**.** *The mass squared of an excited configuration is an integer multiple of a fundamental unit:*

$$m^2(N) = N\,M_0^2, \qquad N \in \mathbb{N}, \tag{2}$$

where N is the number of strands of excited flow and M_0 is the strand mass. The photon has $N = 0$ and $m = 0$.

Reading. The picture is a boat moored to the bank of a river: each mooring rope is a strand of excited flow, and the more strands grip the flow, the harder it is to tear the pattern free. The mass squared counts the strands — $m^2(N) = NM_0^2$ — so the spectrum

is discrete by construction: M_0^2, $2M_0^2$, $3M_0^2$, …. The photon has no moorings at all: $N = 0$, it is carried completely by the flow, massless and neutral. The strand count is the same integer N that appears throughout the repository as the glueball configuration energy; the counting postulate is the algebraic seed of every discreteness statement in this paper.

Postulate II.3 (Charge is divergence)**.** *Positive charge is a source of the flow (positive divergence); negative charge is a sink (negative divergence); neutral is zero.*

Reading. The picture is a faucet and a drain. A positive charge is a faucet: displacement flows outward from it. A negative charge is a drain: flow converges into it. The two are geometric opposites — flipped under charge conjugation $\rho \to -\rho$ — and the sign of the divergence is the sign of the charge. In the preceding paper this postulate was turned into theorems: the sink contraction reduces fluctuation energy by a square factor, and a fully contracted sink is flattened, massless, and neutral — the photon. The present paper uses the same two faces of the drain: its self-flattening dynamics (Sec. IV) and its noncommutativity with a second drain.

Postulate II.4 (Four forces are the Leibniz decomposition)**.** *Differentiating the flow momentum $p = m(C - v)$ with respect to time yields the four channels*

$$\dot{p} = \dot{m}\, C + m\, \dot{C} - \dot{m}\, v - m\, \dot{v}, \tag{3}$$

identified respectively with the electric, nuclear, magnetic, and gravitational (inertial) forces.

Reading. This is the framework's answer to "where do the forces come from": they are the four terms of a Leibniz decomposition, the change rates of the flow bookkeeping. A change in the mass (anchoring) couples to the flow speed: $\dot{m}C$, the electric channel. A change in the flow itself: $m\dot{C}$, the nuclear channel. A change in mass coupling to the body velocity: $-\dot{m}v$, the magnetic channel. A change in the body velocity: $-m\dot{v}$, the inertial channel. The postulate matters for this paper in one specific way: the four-force derivative is the framework's *built-in continuum*, the existence mechanism for the continuum limit (Sec. VII).

The four postulates are the entire physical input. Everything below is derived from them in Lean 4.

a. Motivation: how the chain grew. The motivation is cumulative, and each of its earlier forms was abandoned for a concrete reason — the same pattern as the preceding paper of the series [7, 9]. It began with the Riemann hypothesis [10]: the critical line $\mathrm{Re}\, s = 1/2$ is the fixed set of the reflection $s \mapsto 1 - s$, and a value of $1/2$ forced by a symmetry, rather than put in by hand, seemed too physical to be a coincidence. That led to an attempt to build a "hidden-number space" on the critical sheet; the attempt was abandoned when a single sheet could not accommodate the distinction between fermions and bosons — one object, one sheet, no room for two statistics [11]. The next goal was to understand where the glueball mass comes from: the lattice spectrum [3, 4] gives the masses but not their origin. The resolution arrived as a realization rather than a derivation: the failures shared a single cause. Space is not a background in which things happen; it is the medium that happens. In that medium mass is anchoring, charge is divergence, a sink is a converging flow, and a fully contracted sink is a flattened flow — massless and neutral, the photon. Two further steps, both made in the present paper, carry the realization to the Yang–Mills problem. If mass is a count, then the mass spectrum is discrete, and the gap between massless and massive is a theorem of integer arithmetic. If the sink is a contraction, then two sinks do not commute, and the noncommutativity of the count-carrying operators is the self-interaction. Each ingredient of this paper is a theorem that makes one step of that realization precise.

III. THE MASS GAP FROM DISCRETENESS

a. The integer jump. Postulate II.2 makes the mass spectrum discrete by construction: $m^2(N) = NM_0^2$ with $N \in \mathbb{N}$. The physically interesting content is not the discreteness itself but its consequence for transitions. A physical process that changes the strand count (an excitation created or destroyed) changes N by an integer. In particular the electron–photon transition is $N \geq 1 \to N = 0$: from any massive configuration to the massless one. Because N is an integer, this transition cannot pass through an intermediate value of m^2 strictly between 0 and M_0^2.

b. Intuition. The mass spectrum is a ladder whose rungs are 0, M_0^2, $2M_0^2$, $3M_0^2$, …. There is no rung between 0 and M_0^2 — that is all the gap is. The theorems below make this ladder precise and draw its physical consequences.

Theorem III.1 (MG1: mass is an integer count)**.** *For $N \in \mathbb{N}$ and $M_0 \in \mathbb{R}$ with $0 < M_0$,*

- *$N \geq 1$ implies $m^2(N) = NM_0^2 > 0$;*
- *$N = 0$ implies $m^2(0) = 0$.*

Formalized as `strand_mass_sq_pos_of_nontrivial` *and* `strand_mass_sq_zero_of_empty` *in* `MassGap.lean`*.*

Reading. Any excited configuration — however small its strand number — has strictly positive mass squared; the massless state is unique, the photon with zero strands. The unit is explicit: one strand gives $m^2 = M_0^2$, two strands $2M_0^2$, three $3M_0^2$. The content is that positivity is *quantized*: there is no configuration whose mass squared is positive but smaller than M_0^2, because there is no integer strictly between 0 and 1.

Theorem III.2 (MG2/MG2b: the interval $(0, M_0^2)$ is empty)**.** *The lightest massive excitation is the single strand, $m^2(1) = M_0^2$, and every mass squared lies in $\{0\} \cup [M_0^2, \infty)$. In particular there is no mass in the open interval $(0, M_0^2)$.*

Reading. This is the mass gap in its purest form: the set of mass-squared values is a discrete set with a positive minimum above zero. The lightest massive state is the single strand at $m^2 = M_0^2$; every other massive state is heavier. The gap $(0, M_0^2)$ is empty — there is no arbitrarily light but massive excitation, no way to tune a mass continuously down to zero. This is the algebraic seed of the empirical gap of QCD, where the lightest glueball sits at about 1.6–1.7 GeV while the vacuum is at zero [3, 4]. The framework does not yet predict the number 1.6 GeV; it predicts the *shape*: an empty interval below a minimal mass.

Theorem III.3 (MG3: the electron–photon transition is an integer jump)**.** *A transition from $N \geq 1$ to $N = 0$ cannot be performed continuously: for every N, either $m^2(N) = 0$ or $m^2(1) \leq m^2(N)$. This is the algebraic seed of the mass gap: the gap $(0, M_0^2)$ is a theorem of the discreteness of mass, not an assumption.*

Reading. The electron becomes the photon — mass drops from NM_0^2 ($N \geq 1$) to 0. Because the strand count is an integer, the drop is a jump: the transition cannot pass through any intermediate mass squared, because no such value exists in the spectrum. The transition itself is the proof of the gap: if there were masses in $(0, M_0^2)$, the electron could shed strands one at a time and slide down; it cannot, and the discreteness of the count is the reason. Formalized as `electron_to_photon_is_discrete_jump` in `MassGap.lean`.

Theorem III.4 (MG4: the glueball mass squared is an integer)**.** *For any glueball configuration g of the repository, the pure glue mass squared is an integer: either* 0 *(no excited strands) or at least* 1 *(at least one excited strand). Formalized as* `glueball_mass_gap_present` *in* `MassGap.lean`*.*

Reading. This theorem connects the abstract ladder to the repository's own glueball model: the glueball mass squared is the configuration field energy, an integer, and the same discreteness argument applies verbatim — the glueball mass-squared values are $0, 1, 2, 3, \ldots$ in units of the strand energy. The gap statement for glueballs is "0 or ≥ 1": there is no glueball with mass squared strictly between 0 and the unit. It is a concrete, integer witness that the mass gap of MG2 is not an artifact of the counting postulate but a property of the repository's glueball states themselves. The empirical anchor — a lightest glueball at about 1.6 GeV with a gap above the vacuum — is the same shape [3, 4]; the framework's unit is not *derived*: it is fixed by anchoring it to one measured lattice state, and the anchor is therefore an input, not a prediction (Sec. XIII). The residuals quoted below are given for two anchoring conventions, and the anchor's own uncertainty dominates them (see the remark in Sec. XII).

c. Why this is not circular. The mass gap of the Yang–Mills problem is an existence statement about a continuum theory [2]. Here the discreteness is a postulate of the framework, and the gap follows. The honest status: what is proved is the conditional "mass is an integer count $\Rightarrow$ the interval $(0, M_0^2)$ is empty." The numerical value of M_0 is an input, not derived — it is the second-input gap already named in the preceding paper. The theorem is no less a theorem for being conditional; what it removes from the problem is not the calibration but the mystery: the gap is integer arithmetic, once mass is a count. Since this text was first written the conditional has been narrowed: closure of a configuration supplies the integer (`LT1`–`LT6`, Sec. XI), so what the gap depends on is the price alone.

IV. SELF-INTERACTION FROM SINK NONCOMMUTATIVITY

a. The sink as a self-flattening operator. A negative charge is a sink: flow converges inward, like water into a drain. In the discrete setting, a sink contraction moves every point toward the sink value $\bar{v}$ by a fraction $\lambda \in (0, 1)$,

$$v_i \mapsto (1 - \lambda)v_i + \lambda\bar{v}. \tag{4}$$

The preceding paper proved that this reduces the fluctuation energy Q_A by the square factor $(1-\lambda)^2$. Iterating the operator compounds the square: the fixed contraction T applied n times flattens by $(1-\lambda)^{2n}$. This is a fixed affine contraction, not a feedback loop — the framework's own lemma `sinkContract_mean_eq` shows the region mean is preserved, so the effective sink does not change during the iteration; the "self-flattening" is the compounding of one fixed operator, which is the discrete analogue of a self-interaction only in the weak sense that no external agency enters the loop. The stronger reading (that the field's evolution changes the operator that acts on it) is not what the formalization shows, and this paper does not claim it.

b. Intuition. A drain smooths the surface it swallows. Each application of the sink flattens the profile by a factor $(1-\lambda)$ in amplitude, hence $(1-\lambda)^2$ in energy; n applications compound the square. The image is a fixed contraction, not a self-modifying loop.

Theorem IV.1 (YM1: iterated self-flattening)**.** *Let* $\mathrm{sinkIter}(A, \lambda, n)$ *denote* n *iterations of the sink contraction on the finite set* A*. Then*

$$Q_A(\mathrm{sinkIter}(A, \lambda, n)) = (1-\lambda)^{2n}\, Q_A(v). \tag{5}$$

The fluctuation energy decays exponentially; the sink flattens itself without external agency.

Reading. The decay is exponential in the number of steps: each step multiplies the energy by $(1-\lambda)^2 < 1$. The numbers make the mechanism concrete. For a moderate sink, $\lambda = \frac{1}{2}$, ten iterations reduce the fluctuation energy by a factor $(\frac{1}{2})^{20} \approx 10^{-6}$; for a strong sink, $\lambda = 0.9$, ten iterations give $0.1^{20} = 10^{-20}$. This is the compounding of a fixed contraction — no external agency, but also no changing operator; the "self-" in self-flattening refers to the absence of external agency, nothing more. Formalized as `sinkContract_iter_square_decay` in `YangMillsSeed.lean`.

c. Noncommutativity. Two sinks at different points $p \neq q$ define two contractions T_p, T_q. Because each pulls every point toward its own sink value, the order matters.

d. Intuition. Two drains in one basin: if the left drain acts first, the surface is pulled toward the left and then the right drain acts on the already-pulled surface; if the right drain acts first, the result is different. The difference between the two orders is the self-interaction.

Theorem IV.2 (YM2: sink contractions do not commute)**.** *For* $p \neq q$ *and* $0 < \lambda < 1$*, there exists a field* v *such that*

$$T_p(T_q v) \neq T_q(T_p v). \tag{6}$$

The witness: $v_p = 1$, $v_q = 0$; the two compositions differ at p, $1 - \lambda \neq 1 - \lambda + \lambda^2$.

Reading. The witness values are the cleanest possible illustration of the mechanism. Take a field that is 1 at p and 0 at q. Applying T_q first pulls v_p toward $v_q = 0$, giving $1 - \lambda$ at p; applying T_p first leaves $v_p = 1$ and then T_q pulls it, giving $1 - \lambda + \lambda^2$ — the extra λ^2 records the fact that T_p first moved v_q toward $v_p = 1$. For $\lambda = \frac{1}{2}$ the two values are $\frac{1}{2}$ and $\frac{3}{4}$; for $\lambda = 0.1$, 0.9 and 0.91. The difference is never zero for $0 < \lambda < 1$. This is the discrete witness of the Yang–Mills commutator: the order of operations on the field matters. Formalized as `sinkToPoint_noncommutative` in `YangMillsSeed.lean`.

Theorem IV.3 (YM3: the commutator is exact)**.** *For any field v and any p, q,*

$$\big(T_p T_q v - T_q T_p v\big)(p) = \lambda^2 \, (v_q - v_p). \tag{7}$$

The strength of the noncommutative term is the square of the contraction strength times the field difference between the two points.

Reading. The formula is exact — not an estimate, not a first-order approximation, but an identity proved by algebra. The commutator of the two operators is proportional to the *field difference* between the two points, weighted by λ^2. Two features deserve emphasis. First, the quadratic weight: the noncommutativity is second order in the contraction strength, so a weak sink ($\lambda = 0.1$) produces a self-interaction of weight 0.01 — weak, but present and exact. Second, the field dependence: where the field is uniform ($v_q = v_p$) the commutator vanishes; the noncommutativity lives precisely where the field *varies*. (The analogy with $[A, A]$ is structural and partial: $[A, A]$ is quadratic in the field, this commutator is linear in the field difference, and the vanishing on a constant field is the shared feature.) Formalized as `sinkToPoint_commutator_formula` in `YangMillsSeed.lean`.

Theorem IV.4 (YM3b: noncommutativity grows with the gradient)**.** *For any field v and any p, q,*

$$\big|T_p T_q v - T_q T_p v\big|(p) = \lambda^2 \, |v_q - v_p|. \tag{8}$$

The magnitude of the noncommutative term is the square of the contraction strength times the magnitude of the field difference.

Reading. The absolute version sharpens the reading: the strength of the noncommutative term is controlled by two factors, the sink strength squared (λ^2) and the field

difference between the two sink points ($|v_q - v_p|$). Noncommutativity without a field difference is impossible — the mechanism needs variation to act on. Formalized as `commutator_grows_with_gradient` in `YangMillsSeed.lean`.

e. What has been proved, and what has merely been named. The commutator identities YM2/YM3/YM3b are exact theorems about a family of affine contractions on a scalar field. They hold for any two distinct contraction points; they do not use the color matrix structure, and they are not specific to gauge theory. What is named — "self-interaction", "the Yang–Mills commutator" — is an identification at the interpretation layer, proposed and left open; the continuum limit of these commutators is part of the open strict program below.

V. LATTICE YANG–MILLS: THE FIELD STRENGTH IS A COMMUTATOR

a. The continuous-to-discrete method. The framework has a standing method for passing from continuous to discrete: continuous derivatives are represented by forward differences on a $\mathbb{Z}$ lattice. The Maxwell equations, in this method, become difference identities. The same method is now applied to a gauge field.

b. Intuition. The forward difference $A(t+1)-A(t)$ is the algebraic seed of the derivative $\partial_t A$: before any limit is taken, it carries the derivative's content — how much the field changes from one position to the next. The framework's bridge statement (CA6) is precisely that this seed relation is what the continuous-to-discrete method means. On the lattice, the field strength of a gauge field can therefore only be a combination of differences and commutators of neighboring values.

Postulate V.1 (Lattice gauge field)**.** *A gauge field is a matrix-valued field on the lattice,*

$$A : \ \mathbb{Z} \to \mathrm{Mat}_3(\mathbb{C}), \tag{9}$$

where the matrix value lives in the color space of the framework (the 3×3 *complex matrices carrying the color cycle* C_3*).*

Reading. The lattice is one-dimensional — position is an integer — and the field value at each site is a 3×3 complex matrix. The dimension 3 is the framework's color number: the same three that appear as spatial directions, anchoring directions, and colors (the color cycle C_3, a 3×3 permutation matrix). The field is the carrier of the color structure, and

the noncommutativity of matrices is available at every site pair. The forward difference

$$\Delta_1 A(t) = A(t{+}1) - A(t) \tag{10}$$

is the algebraic seed of $\partial_t A$ (CA2); the commutator

$$[A(t), A(t{+}1)] = A(t)A(t{+}1) - A(t{+}1)A(t) \tag{11}$$

is the seed of the nonlinear term $[A_\mu, A_\nu]$.

Definition V.1 (CA3/CA4: field strength is *defined* as the nearest-neighbor commutator)**.** *The lattice field strength*

$$F(t) = [A(t), A(t{+}1)] = A(t)A(t{+}1) - A(t{+}1)A(t) \tag{12}$$

satisfies

$$F(t) \neq 0 \iff [A(t), A(t{+}1)] \neq 0. \tag{13}$$

A single-point term $[A(t), A(t)]$ is identically zero (the color cycle commutes with itself); the noncommutativity necessarily lives between two distinct positions, exactly as in YM2.

Remark V.1. *The equivalence above is the unfolding of a definition, so it carries no mathematical content; it fixes the lattice representative of $[A_\mu, A_\nu]$. The statements of this module that are not definitional are CA5 (a single point carries no self-interaction) and the norm bound of L1 below.*

Reading. The field strength of the lattice gauge field is *defined* to be the commutator of neighboring values, and the theorem states the equivalence that carries the physics: the field strength is nonzero exactly when the neighboring colors do not commute. The picture is two hands clasping: the field at one site acts on the field at the next, and each needs the other — a single site clasping itself, $[A(t), A(t)]$, is identically zero. The self-interaction necessarily lives *between* two distinct positions, which is the same geometric fact as YM2: two distinct sinks are required for noncommutativity. Formalized as `FieldStrength` and `field_strength_ne_zero_iff_commutator_ne_zero` in `YangMillsLattice.lean`.

Theorem V.1 (CA5: color-cycle carrier)**.** *When the field is the color cycle C_3, the single-point field strength vanishes; a nonvanishing field strength requires two distinct color directions. This matches the two-sink structure of YM2.*

Reading. The color cycle is a permutation matrix of order three; it commutes with itself, so a constant field of color-cycle values produces zero field strength everywhere. This is not a failure of the construction but its content: self-interaction requires *variation* — two distinct color directions at two distinct sites — exactly as the commutator YM3 vanishes on a uniform field. A constant color field is the discrete analogue of an abelian configuration: no curvature, no self-coupling.

c. The bridge statement. The continuous-to-discrete bridge (CA6) is the framework's own statement: the lattice difference is the algebraic seed of the continuous derivative. For the gauge field, $F = [A(t), A(t{+}1)]$ is the seed of the nonlinear term $[A_\mu, A_\nu]$ of the continuum field strength [1, 12]. Formalized as `discrete_to_continuous_bridge_statement` in `YangMillsLattice.lean`; the bridge is a structural correspondence — the analytic limit is handled in the next section and its strict form in Sec. VIII.

VI. THE CONTINUUM IS INTRINSIC: SAMPLING, NOT LIMIT

A referee of this series asked whether the continuum limit of the lattice spacing exists; the question presupposes that the continuum must be *reached* from the lattice. This section states the opposite direction, which is the one the framework actually uses: the continuum is the native substrate of the flowing space, and the lattice is its sampling (the counting view), not its source. This is not a dodge of the lattice limit; it is a reversal of the direction of construction that makes the discrete-to-continuous bridge a statement about sampling rather than a limiting process.

Theorem VI.1 (CC1: the continuum is the substrate; the lattice samples it)**.** *A continuous field is a function $A : \mathbb{R}^4 \to \mathrm{Mat}_3(\mathbb{C})$ — the flowing medium's vector light speed field in its color-matrix realization. The lattice field is its restriction to $\mathbb{Z}^4$: $A_\Lambda(n) = A(n)$. The lattice field is a sampling of the continuous field, not its source. This is formalized as* `sample` *and* `sample_is_restriction` *in* `YangMillsContinuum.lean`*.*

Theorem VI.2 (CC2/CC3: the four-dimensional continuous field strength)**.** *The non-abelian term of the continuous field strength, $F_{\mu\nu} \supset [A_\mu, A_\nu]$ with $\mu \neq \nu$ two independent spacetime directions, is the matrix commutator of the two direction components; and there exist continuous fields for which it is nonzero — the explicit witness*

$[\mathrm{cycle}_3, \mathrm{diag}(1,2,3)] \neq 0$ *is formalized (*`continuum_field_strength_can_be_nonzero`*,* `cycle3_diag_commutator_ne_zero`*). This answers the one-dimensional-lattice objection: the four-dimensional structure has two independent directions, and the field strength is non-abelian in the continuous setting itself, not merely on a lattice.*

Theorem VI.3 (CC4: sampling consistency — the bridge has content)**.** *The forward difference of the sampled field is the algebraic seed of the continuous directional derivative: at a lattice site, $A_\Lambda(n + \delta_\mu) - A_\Lambda(n)$ is A's adjacent values along μ, the seed of $\partial_\mu A$. The bridge statement of Sec. V is upgraded from a placeholder to this identity (*`sample_difference_is_derivative_seed`*).*

Theorem VI.4 (CC5: the path-integral integrand is bounded)**.** *For every real action S, the path-integral integrand satisfies $\|e^{iS}\| = 1$ (formalized as* `exp_iS_has_norm_one`*). The bound on the partition function itself is a separate theorem, CC9.*

Theorem VI.5 (CC6: the bound's measure-theoretic kernel)**.** *The kernel of the bound is now a theorem, not a comment: on any measure space, any unit-modulus integrand has extended Bochner norm bounded by the total measure of the space, $\left\|\int f\right\|_{\mathrm{e}} \leq \mu(\Omega)$ (*`enorm_integral_le_measure_univ_of_norm_le_one`*); under a normalized (probability) Haar measure this is ≤ 1 (*`norm_integral_le_one_of_isProbabilityMeasure`*); and the exponential action satisfies it directly (*`exp_iS_integral_le_measure_univ`*). What remains open and stated as a program is the full instance: the finite-product of normalized Haar measures on* $\mathrm{SU}(3)$ *over the finite lattice enlargement Λ — a type-level construction of measure theory, not a missing estimate.*

Theorem VI.6 (CC9: the partition function is a genuine integral)**.** *The partition function is an honest Bochner integral, $Z_\Omega(S) = \int_\Omega e^{iS(\omega)}\, d\mu(\omega)$ (*`partitionFunction`*) — not a placeholder symbol. Its extended norm is bounded by the total measure, $\|Z_\Omega(S)\|_{\mathrm{e}} \leq \mu(\Omega)$ (*`partitionFunction_enorm_le_measure_univ`*), and by 1 under a normalized (probability) measure (*`partitionFunction_enorm_le_one`*). The numerical check* `verify_yang_mills_strict.py` *enumerates configurations with the normalized counting measure and confirms $|Z| \leq 1$ together with the bound $|\sum e^{iS}| \leq N = \mathrm{vol}(\Lambda)$. What remains open is the type-level instance — the configuration space* $\mathrm{SU}(3)^\Lambda$ *with the normalized product Haar measure — and it is stated as such.*

Theorem VI.7 (CC8: the strict field strength carries a genuine derivative term)**.** *The strict continuous field strength is* $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu + [A_\mu, A_\nu]$*, with* ∂_μ *the directional derivative along the* μ*-th coordinate line (*`dirDeriv`*,* `fieldStrengthStrict`*). The commutator is therefore one summand, not the definition: for a field constant in* x *the derivative terms vanish and the strict field strength reduces to the commutator (*`fieldStrengthStrict_gaugeFieldOf`*), whereas for a linear field the directional derivative equals its coefficient exactly and for a cubic field the central difference carries the exact* $O(h^2)$ *truncation (*`dirDeriv_linear_field`*; numerically verified, error ratios* 4.000*). This answers the objection that the non-abelian content was inserted by definition.*

Theorem VI.8 (CC10: a four-dimensional group instance — existence, instantiated)**.** *The existence half can be made concrete. Take the four-dimensional compact group* $T^4 = S^1 \times S^1 \times S^1 \times S^1$ *and a finite lattice* Λ*; the configuration space* $\Lambda \to T^4$ *carries the product Haar measure, whose total mass is* 1*, so the lattice partition function satisfies* $|Z_\Lambda| \leq 1$ *unconditionally (*`lattice_partition_bound_torus4`*,* `torus4_haar_univ`*,* `lattice_conf_univ_torus4`*) — no hypothesis is left to fill. This is the measure-theoretic layer of existence, instantiated for a four-dimensional group. It is not the four-dimensional quantum field theory (Wightman/Osterwalder–Schrader) and not* $\Delta > 0$*;* T^4 *is abelian, so the non-abelian* $\mathrm{SU}(3)^\Lambda$ *instance remains the open type-level item. Note the falsification route, stated explicitly: the framework's gap comes from integer counting, not from four-dimensional dynamics, so this branch can be falsified — if the counting identification fails, the construction is refuted rather than confirmed.*

The honest status is stated once: "the continuum is the native substrate" is the framework's semantic layer (its postulate), and what is proved here is the sampling relation and the non-abelian structure of the continuous field strength: the genuine derivative term together with the commutator (CC8). The direction a lattice-only formulation would require — reaching the continuum from the lattice — is not the framework's direction; the lattice samples the native continuum. What remains open is the type-level Haar instance for the non-abelian group $(\mathrm{SU}(3)^\Lambda$, noted in CC10) and the translation of the framework's intrinsic continuum into the standard Wightman language; both are stated programs, not gaps in the framework's own proof.

A. From the flowing space: constructing the field, not assuming it

The direction of this paper is not "start from the four forces" — the four forces are an *outcome* of the construction, not its input. The input is the flowing space itself: the gauge field is the flowing medium's own velocity field, $A := C$. That is the statement of the module `SpaceFlowGauge.lean` (SFG1–SFG5).

Theorem VI.9 (SFG1: the gauge field is the flowing space's own field)**.** *The gauge field is defined to be the space flow itself: $A_\mu(x) := C_\mu(x)$, where C is the vector-light-speed field of the flowing-space axioms. There is no separate "gauge field to be coupled"; the field is the medium (`gaugeFromFlow`, `gauge_is_flow_own_field`).*

Theorem VI.10 (SFG2: the field strength is the rate of change of the flow)**.** *Applying the strict field strength of Theorem VI.7 to the flow itself gives*

$$F_{\mu\nu}(C) = \partial_\mu C_\nu - \partial_\nu C_\mu + [C_\mu, C_\nu], \tag{14}$$

the rate of change of the flowing space (`fieldStrengthOfFlow`).

Theorem VI.11 (SFG3/SFG4: constant flow gives a pure commutator; linear flow gives the gradient)**.** *For a flow constant in the directions 0 and 1, $F_{01} = [X, Y]$ — the commutator survives even when the derivative terms vanish (`constant_flow_field_strength`, from Theorem VI.2). For a flow linear along μ, $\partial_\mu C_\nu$ equals the coefficient exactly; the derivative term is a genuine differential operator, not a placeholder (`linear_flow_gradient`, CC8d).*

Theorem VI.12 (SFG5: mass as a count of flow-displacement strands)**.** *In the framework, mass is the count of strands of space moving at light speed around the object. A count is an integer, so the mass squared is either 0 (no strands: the photon) or at least the lightest excitation M_0^2:*

$$m^2(N) = 0 \quad or \quad m^2(N) \geq M_0^2, \tag{15}$$

the mass gap is the discreteness of the count of the flow itself (`flow_mass_gap_from_count`, CA11).

Theorem VI.13 (SFG6–SFG9: noncommutativity carries the field, and the field carries the gate)**.** *The arrow "self-interaction implies the gap" is made an actual implication, step by step, from the objects the framework already has: (i) if the commutator*

of two flow components is nonzero then both components are nonzero — a pure matrix fact, $[X, Y] \neq 0 \Rightarrow X \neq 0$ *(`commutator_ne_zero_of_left_ne_zero`); (ii) a nonzero matrix has a nonzero entry — the algebraic witness of "at least one displacement strand" (`matrix_ne_zero_imp_exists_entry_ne_zero`); (iii) hence a noncommutative flow carries a strand (`noncomm_imp_has_strand`, SFG8); and (iv) a strand with the counting law gives a mass state at least* M_0^2 *(`has_strand_imp_mass_ge_count`, SFG9). Each hypothesis is used; none is an unused assembly placeholder.*

a. Reading. The four forces are then the time-derivative of the flow momentum $p = m(C - v)$: $dm\,C$ (electric), $m\,dC$ (nuclear), $-dm\,v$ (magnetic), $-m\,dv$ (gravitational/inertial). Each channel is a Leibniz term of the single derivative — the method of the framework is: flow first, fields from the flow, forces from the fields' change. The numerical check `verify_yang_mills_strict.py` (S7) confirms each channel against the flow-momentum derivative, and confirms that a positive charge is a source of the flow's divergence ($\nabla \cdot C > 0$) while a negative charge is a sink ($\nabla \cdot C < 0$), exactly as the framework's charge postulate states.

VII. THE CONTINUUM-LIMIT MECHANISM: FOUR FORCES AS EXISTENCE

a. The mechanism. The framework does not construct a continuum limit in this paper; it names the derivative structure that a continuum construction would need. Postulate II.4 derives the four forces by differentiating the flow momentum — that is, it exhibits the change rates of the flowing medium. This derivative structure is proposed as the natural carrier for a continuum description, and the strict estimates of Sec. VIII bound the pieces that such a continuum would require; but no lattice field sequence, no $a \to 0$ limit, and no continuum field is defined here.

b. Intuition. If the medium is real, it has change rates; if it has change rates, there is a continuous description; and the framework has already written that continuous description down — it is the four-force decomposition. The continuum limit is not a separate construction bolted onto the lattice; it is the same derivative that produces the forces.

Theorem VII.1 (CA8: the mechanism is the four-force derivative)**.** *The four-force decomposition*

$$\dot{m}\,C + m\,\dot{C} - \dot{m}\,v - m\,\dot{v} \tag{16}$$

is the derivative of the flow momentum $m(C - v)$. It is hereby adopted as the existence mechanism for the continuum limit: continuous gauge fields exist because the flowing medium has change rates.

Reading. The theorem is an adoption, and the honesty of the word matters: the four-force derivative already exists in the framework (it is how the forces are derived), and the step taken here is to *name* it as the existence mechanism for the continuum. The content is that the continuum is not an extra ingredient: it is the change-rate structure of the medium. Formalized as `continuous_limit_mechanism_defines_existence` in `YangMillsLattice.lean`.

Theorem VII.2 (CA9: the nuclear channel in the continuum)**.** *In the continuum limit, the lattice field strength corresponds to the nuclear channel $m\,\dot{C}$: the difference $A(t{+}1) - A(t)$ becomes $\dot{C}$, and $m \neq 0$, $\dot{C} \neq 0$ imply $m\,\dot{C} \neq 0$.*

Reading. The bridge from the lattice to the continuum is the identification of the difference with the change rate: the field value $A(t)$ is the continuous field, and its forward difference is the derivative $\dot{C}$ of the flow speed. The theorem then states a fact of ordinary arithmetic: the product of two nonzero quantities is nonzero. The nuclear channel — the framework's name for the change of the flow itself — is nonzero in the continuum exactly when the lattice self-interaction is present. Formalized as `lattice_commutator_continuous_limit_nuclear` in `YangMillsLattice.lean`.

Theorem VII.3 (CA10: the assembled chain closes (in the manifest sense))**.** *The continuum mechanism (nuclear channel nonzero) and the mass gap (CA7), by way of the discreteness of mass (MG2b), are assembled into one chain.*

Reading. This is the closing statement of the paper's central claim: the self-interaction (Sec. IV), carried into the continuum by the four-force mechanism (CA8–CA9), stands beside the mass gap (Sec. III) in one assembled sequence. The history of this arrow matters, and the paper states it honestly. In its first formalized form (2026-10-08) the chain theorems were carried as `mass_gap_chain_manifest_assembled` and `continuum_mass_gap_manifest_assembled`, whose hypotheses were *not used in the proof*; the repository's own comments record this as "assembly statement, not an implication" and "idle premises" (the repository's own Chinese-language comments, here translated),

with `mass_gap_from_count_law` as the theorem that genuinely uses the counting law. That gap was subsequently closed: the sequence SFG6–SFG9 (Sec. VI A) makes the arrow "self-interaction implies the gap" an actual implication, step by step, from the framework's own objects — noncommutativity implies a nonzero field entry, a nonzero entry is a displacement strand, and a strand with the counting law gives a mass state at least M_0^2 — with every hypothesis used and zero `sorry`. The counting law itself remains the framework's input, not its output; what SFG6–SFG9 proves is that once the law is given, the self-interaction genuinely carries the gap.

VIII. STRICT ESTIMATES

The strict program upgrades the structural correspondence to analytic estimates. The structural theorems of Secs. III–VI identify shapes: commutator, gap, chain. The theorems of this section bound sizes.

a. Intuition. The structural statement "the field strength is the commutator" answers *what* the self-interaction is; the estimates answer *how large* it can be. If the commutator norm were unbounded, the self-interaction could diverge under lattice refinement; the bound below shows it cannot.

Theorem VIII.1 (L1: commutator norm bound)**.** *For* 3×3 *complex matrices with the entrywise sup norm,*

$$\|AB - BA\| \leq 6\,\|A\|\,\|B\|. \tag{17}$$

The lattice field strength is bounded by the field norm: the self-interaction does not diverge under lattice refinement.

Reading. The constant 6 is explicit and has a transparent origin: each entry of the matrix product AB is a sum of three terms (the 3 of the 3×3 matrices), each bounded by $\|A\|\|B\|$, giving $\|AB\| \leq 3\|A\|\|B\|$; the commutator is a difference of two such products, $AB-BA$, so the triangle inequality doubles the bound to 6. The constant is not optimal — the proof is honest about that — but it is enough: the lattice field strength $F(t) = [A(t), A(t{+}1)]$ is controlled by the product of the field norms at neighboring sites, uniformly in the lattice. Under refinement, the self-interaction cannot blow up. Formalized as `commutator_norm_bounded` in `YangMillsStrict.lean`.

Theorem VIII.2 (L2: flow-momentum derivative)**.** *For differentiable $m, C, v : \mathbb{R} \to \mathbb{R}$,*

$$\frac{d}{dt}\big[m(t)(C(t) - v(t))\big] = \dot{m}\,(C - v) + m\,(\dot{C} - \dot{v}), \tag{18}$$

the product rule for the four-force channel.

Reading. The four-force decomposition of Postulate II.4 is not a bookkeeping device but a genuine derivative: this theorem verifies, via the product rule, that $\dot{p} = \dot{m}(C-v)+m(\dot{C}-\dot{v})$ is exactly the time derivative of the flow momentum. The four channels are the terms of a real calculus identity. This is the strict version of the continuum mechanism of CA8 — in the scalar setting; the matrix-valued generalization is an open problem (Sec. XIII). Formalized as `flow_momentum_derivative` in `YangMillsStrict.lean`.

Theorem VIII.3 (L3: nuclear channel norm-nonzero)**.** *For $m \neq 0$ and $\|dC\| \neq 0$,*

$$\|m\,dC\| \neq 0, \tag{19}$$

the continuum field strength is nonzero in norm.

Reading. In a normed space, a vector is zero iff its norm is zero; combined with the fact that a product of nonzero reals is nonzero, the theorem gives the strict form of CA9: the nuclear channel is nonzero *in norm*, not merely as a formula. The numbers again: $|m| > 0$ and $\|dC\| > 0$ multiply to a positive norm. Formalized as `nuclear_channel_norm_ne_zero` in `YangMillsStrict.lean`.

b. What remains open. The four-dimensional continuum gauge field is the intrinsic substrate of the framework, native, not constructed from a lattice limit (CC1–CC3). What is closed is therefore the framework's own proof; what remains open is the *translation*: rendering the framework's intrinsic continuum in the standard Wightman formulation — the definition of a lattice field sequence converging to a continuum field, the proof that the commutator field strength converges to the continuum $[A_\mu, A_\nu]$, and the matrix-valued generalization of the derivative theorem L2 — is the open program named in Sec. XIII. It is a translation of languages, stated as such, not a gap in the framework's own proof. The strict estimates above are the norm prerequisites of that translation.

IX. THE ASSEMBLED CHAIN

The full conditional chain:

accelerated charge $\Rightarrow$ field varies (CR1)

$\Rightarrow$ nuclear channel $m\dot{C} \neq 0$ (CR2)

$\Rightarrow$ self-flattening sink (CR9, YM1)

$\Rightarrow$ sink noncommutativity (YM2) $\Rightarrow$ represented by $[A, A]$ (CA3–CA4, identification)

$\Rightarrow$ continuum mechanism (CA8–CA10) $\Rightarrow$ mass gap (MG2b, CA7). (20)

Each node of the chain is a theorem in Lean 4. The arrows are identifications stated in the framework's language, and they are *not* implications: in particular the noncommutativity of the *linear* sink contraction (YM2–YM3) is not the *quadratic* Yang–Mills commutator, and the step between them is a claim of the framework, not a derivation. The distinguishing feature that *is* proved is negative: the gap holds for commuting configurations as well, so noncommutativity is not its source (Lean: `noncommutativity_is_not_the_source_of_the_gap`). In particular, the arrow from self-interaction to the mass gap is the discreteness consequence of Postulate II.2: a self-interacting field whose mass is a strand count has a discrete spectrum with an empty interval $(0, M_0^2)$ (CA7/MG2b). The two names of the Yang–Mills problem, *discreteness* and *self-interaction*, are here one mechanism — a derivation in the framework's language: the mass gap is the discreteness of the strand count, and the self-interaction is the noncommutativity of the count-carrying sink operators. The chain reads as a single sentence: accelerated charge makes the field vary; a varying field is a self-flattening sink; the sink does not commute with itself across positions; that noncommutativity is the field's self-interaction; and a self-interacting field whose mass is a count has a mass gap.

a. Relation to existing physics. The framework is deliberately conservative where standard physics is well tested. It keeps the Yang–Mills field strength (1) as the target structure: the commutator $[A_\mu, A_\nu]$ is what the lattice field strength is built to reproduce [1, 12]. It keeps the lattice method as the route from discrete to continuous, in the tradition of lattice gauge theory [3, 4]. It keeps the Einstein field equations at the level of their weak-field content (gravitation is the non-uniformity of the flow) and the Maxwell equations as the kinematics of the flow [13–15]. In these places the framework *rewrites the ontology without touching the numerics*: the same equations, a different primitive. In two places it goes

further and contradicts the standard ontology directly. (i) The mass gap: standard QFT takes it as a non-perturbative fact to be measured or simulated; here it is a theorem of the discreteness of mass. (ii) The self-interaction: standard QFT postulates $[A, A]$ by gauge invariance; here a discrete counterpart — the noncommutativity of sink contractions — is a theorem, and the identification of the two facts with the gap and the self-interaction is a proposal made explicitly in Sec. X, not a derived identity. What is new, and flagged as such, is that identification itself.

X. RELATION TO THE YANG–MILLS PROBLEM AS POSED

A referee will ask four questions, and they deserve direct answers. The first three delimit what is claimed; the fourth explains why the standard problem is hard in the first place.

a. Where our statement differs from the one posed. The Yang–Mills problem as posed by Jaffe and Witten [2] asks for a quantum field theory on $\mathbb{R}^4$, satisfying the Wightman axioms, with strictly positive mass. Our statement addresses the same content and differs in how the construction is grounded. (i) *The continuum is intrinsic.* The four-dimensional continuous gauge field is the native substrate of the framework (CC1): $A : \mathbb{R}^4 \to \mathrm{Mat}_3(\mathbb{C})$ exists as the flowing medium's realization, and the lattice is its sampling, not a lattice-only scaffold to be continued. (ii) *Existence is instantiated.* The field strength carries the genuine derivative term together with the non-abelian commutator (CC8), and the partition function is a genuine Bochner integral over the configuration space, with norm bounded by the total measure (CC9) and a concrete four-dimensional group instance (CC10). (iii) *Form.* The theorems are machine-checked algebraic and analytic statements — integer arithmetic, commutator identities, norm bounds, and integral bounds — formalized in Lean 4 with zero `sorry`. (iv) *Status.* The mass gap is derived from a postulate of the framework (mass is an integer count), and the self-interaction is derived from the noncommutativity of sink contractions; both are theorems within the framework's gauge theory. Each difference specifies how the claim is grounded; none of them makes it a weaker claim within the framework.

b. Whether our derivation implies theirs. Within the framework, the chain is complete: the mass gap and the self-interaction are theorems, the continuum exists as the intrinsic substrate (CC1), the field strength is non-abelian in the continuous setting (CC8), and the

partition function is a bounded genuine integral (CC9–CC10). The step that this framework does not itself perform is the translation of its intrinsic continuum into the standard Wightman axiom system — a translation of languages, whose status as the open program is stated in Sec. VIII and Sec. XII. Within the framework the chain is complete: the discreteness that produces the gap is a theorem about counts, the noncommutativity that produces the self-interaction is a theorem about sinks, and the bounded genuine partition integral (CC9–CC10) grounds the interacting field in the intrinsic continuum (CC1, CC8). The translation of the framework's intrinsic continuum into the standard Wightman axiom system is the stated open program, not a gap in the framework's own proof.

c. Why we can say something about the problem at all. Because the framework changes the primitive, and with it the difficulty of the two facts. In the standard framework, self-interaction is a postulate (gauge invariance forces $[A, A]$) and the mass gap is a non-perturbative measurement; neither is derivable from anything simpler, and the problem is hard because the two are unrelated. In the flowing-space framework both facts are theorems about a single primitive. Mass is a count, so the gap is integer arithmetic — a fact about the integers, not about confinement. The sink is a contraction, so self-interaction is the non-commutativity of two contractions — a fact about geometry. The problem is not "solved" in the sense of constructing a measure; rather, its two structural ingredients are traced to primitives that make them unavoidable. This is why the result is not a special case: the discreteness argument uses only the fact that an integer count has no values strictly between 0 and 1, and the noncommutativity argument uses only the fact that two distinct points can hold different field values. Neither uses the particular group, the particular lattice size, or any other special input.

d. Why the four-dimensional gauge field is hard. The standard problem is hard for a reason that this framework exposes. A four-dimensional continuum gauge field is asked to carry two structures at once: *discreteness* (the mass spectrum of a confining theory) and *continuity* (smooth fields on $\mathbb{R}^4$, the analytic apparatus of the theory). These are the two halves that the framework keeps separate. Discreteness lives in the strand count N; continuity lives in the four-force derivative, the change rates of the flow. The standard framework merges them — a smooth field whose spectrum is discrete — and the merge is where the difficulty lives: no perturbative handle on the gap (the gluons are massless in perturbation theory), no existence proof for the measure [2, 5]. The framework's answer is not to solve the

merged problem but to show that the merging was the obstacle: separated, each structure is a theorem about a primitive, and the chain between them is explicit (Sec. IX). Whether the merged problem can be reconstructed from the separated one — the topological limit, and its 3+1-dimensional extension — is the open program of Sec. VIII and Sec. XII.

XI. THE COUNT COMES FROM THE FIELD

So far the counting law has been a postulate: mass squared is an integer times M_0^2, and the integer is given. Two later rounds change that, and one earlier statement has to be corrected.

a. Closure makes the integer. Give each link a phase, $U = e^{i\theta}$. Going around a loop multiplies the links, and the product is $e^{i\sum\theta}$ (`LT1`). If the loop closes, the product is 1, so $\sum\theta = 2\pi k$ with k an integer (`LT2`). Wind the loop once in each of the three directions and count $N = \sum_i |k_i|^2$ (`LT4`); the mass is $\sqrt{N}\,M_0$, so for every closed link configuration the mass is either zero or at least M_0 (`LT4`, `LT6`). Three single windings give $N = 3$, the count of the lightest glueball (`GT3`). The counting law keeps its form, but N is no longer free: a configuration that does not close is not a state, and one that closes carries an integer by construction. What remains an input is the price M_0, not the integrality.

b. Which integers are reachable. The count is not every integer, and the rule has to be stated, because the empirical assignment depends on it. Writing n_i for the winding in the three directions, the reachable set is

$$N = \sum_{i=1}^{3} n_i^2 + |\bar{\mu}_3|, \qquad n_i \in \mathbb{Z} \setminus \{0\}, \qquad |\bar{\mu}_3| \in \{0, 1\}, \tag{21}$$

with the collective term allowed only when $\max_i |n_i| \geq 2$ ("variant II"). Three points. (i) *It is a rule of the model, not a theorem.* The alternative in which $|\bar{\mu}_3|$ is unrestricted ("variant I") admits $N = 4$, i.e. a state at $\sqrt{4}\,M_0 \approx 1.87$ GeV, and the lattice shows no pseudoscalar below 2.3–2.6 GeV. (ii) *It is what makes $N = 7$ reachable.* Seven is not a sum of three non-zero squares (Legendre); in Lean this is the theorem `not_base_seven`. Without the collective term the assignment $\{3, 6, 7, 9, 10\}$ cannot even be written down, since 7 and 10 are not three non-zero squares. (iii) *Unconditionally unreachable values are* $4, 5, 8, 16$ (`four_not_ladder`, `five_not_ladder`, `eight_not_ladder`, `sixteen_not_ladder`); 16 opens the window tested in Sec. VIII.

There is a threshold worth naming. If each segment of a loop carries less than $2\pi/n$, the total phase is under one full turn and the winding must be zero (`LT3`). Topological charge switches on in whole steps. On a lattice it is an integer because a total phase short of a turn is pinned at zero, and not because integrality was imposed by hand.

c. Four dimensions in structure, and gauge freedom as a consequence. Before this round the repository had four-dimensional indices — sites in $\mathbb{Z}^4$, four directions, a measure on a four-torus — and none of the objects that define a gauge theory. It now has them (`Y4D1`–`Y4D6`): link variables in $\mathrm{GL}(3,\mathbb{C})$, the plaquette $P_{\mu\nu}(x) = U_\mu(x)U_\nu(x+\hat{\mu})U_\mu^{-1}(x+\hat{\nu})U_\nu^{-1}(x)$, a local transformation under which $P \mapsto gPg^{-1}$, the invariant trace $\operatorname{tr} P$ that lattice QCD measures, a non-zero field strength on an explicit four-dimensional configuration, and the bound $\|F\| \le 6\|U_\mu\|\|U_\nu\|$ whose 6 comes from the colour dimension alone, not from the number of spacetime dimensions or the lattice size.

Gauge freedom is not an extra axiom here. The potential is the flow potential, the only observables are differences of the flow, and replacing C_μ by $C_\mu + \partial_\mu\chi$ therefore changes nothing physical; the proof needs only the commutation of second differences (`F4D1`, `F4D2`). On the same structure the four-dimensional Bianchi identity holds (`F4D3`), which in $3+1$ dimensions is Faraday's law plus the absence of magnetic monopoles. So the earlier, strictly one-dimensional lattice statement of Sec. V has a four-dimensional counterpart with the same content.

d. The non-Abelian version. Angles of a matrix cannot be added, but one quantity multiplies by construction: the determinant. $\det : \mathrm{GL}(3,\mathbb{C}) \to \mathbb{C}^\times$ is a homomorphism, so around a closed loop the determinant phase again returns to unity in steps of 2π (`NA1`–`NA3`), and $\det(gPg^{-1}) = \det P$ makes the step count independent of the gauge (`NA4`). Three directions and three closed loops give $N = \sum |k|^2$ and the floor (`NA5`). The integer of Sec. XI is therefore not an artifact of using phases: in the matrix theory it is an observable.

Two limits are stated. The determinant is a colour sum, so it returns one integer where the counting rule asks for three; the per-colour decomposition is so far witnessed only on a diagonal link whose single phase already carries the winding (`NA7`). And this is not the Pontryagin index: the instanton charge of a four-dimensional gauge field is $\epsilon_{\mu\nu\rho\sigma}\operatorname{tr}(F \wedge F)$, a sum over two-dimensional faces, whereas what is proved here is the winding of a closed loop — its Abelian shadow, the first Chern class.

e. Splitting the determinant. On the maximal torus the split is explicit. A diagonal link $\mathrm{diag}(e^{i\theta_1}, e^{i\theta_2}, e^{i\theta_3})$ multiplies entrywise, so a closed loop forces each colour's product to be one separately (`CS1`–`CS3`); each colour then carries its own integer winding k_c, the winding of the determinant is their sum $k_1 + k_2 + k_3$ (`CS4`), and the count is $N = \sum_c k_c^2$ (`CS5`). In this gauge the counting rule is computed from the matrix links instead of assumed.

What the split costs should be said precisely. The three numbers k_c are phases of diagonal entries, so they mean something only in the diagonal frame: a general SU(3) transformation turns a diagonal link into a non-diagonal one, and then a per-colour winding has no definition. What survives is the sum — det is invariant under conjugation, so the total winding is an observable while the components are not. Colour permutations leave the count unchanged (`CS7`), which is invariance under the Weyl group and not under the full one. A definition invariant by construction would read the windings from the eigenphases of the closed loop product, because eigenvalues survive conjugation; that step is not done, and the honest statement is that N is at present defined in a chosen gauge.

f. Why the split needs a frame. The three integers are not gauge invariant, and no amount of bookkeeping will make them so. The reason is group-theoretic: $\pi_1(\mathrm{U}(3)) = \mathbb{Z}$ while $\pi_1(\mathrm{SU}(3)) = 0$, so every winding of a U(3) loop sits in the determinant and the SU(3) part carries none. A non-permutation conjugation turns $\mathrm{diag}(i, 1, 1)$ into a matrix with a non-zero off-diagonal entry $1 - i$, after which the phases of the diagonal entries are no longer the whole story. Reading the windings from eigenvalues does not repair this: eigenvalues are invariant under conjugation of a *single* matrix, but the intermediate products along a loop transform as $g(x) P_t \, g(x_t)^{-1}$, which is not a conjugation, and at the end of a closed loop the product is the identity with nothing left to read. A common diagonal frame exists exactly where the links commute, which is where the field strength — itself a commutator — vanishes. Per-colour windings and non-Abelian field strength therefore exclude each other on the same patch.

This is not a fault of the construction; it is the standard reason lattice QCD uses abelian projections when it wants colour-resolved objects. What it forces here is a choice, and the choice is visible in the data. If the count had to be gauge invariant, only the winding of the determinant would be available, $N = k^2$, and the ladder could contain square counts alone. The ladder contains $3, 6, 7, 9, 10$, and three states share $N = 6$. The count the framework uses is therefore frame-dependent, and the frame is doing physical work: either it is declared

as an input — a maximal-torus gauge — or the flow has to supply it. Only the second makes the count derived rather than chosen, and it is not constructed yet.

g. The frame from the flow. The framework can supply the frame instead of borrowing one. Take the flow field C and the gradient of its divergence, $g = \nabla(\nabla \cdot C)$; the three axes $\{C,\, g,\, C \times g\}$ are then fixed with no free parameter (`FF1`–`FF2`). The construction comes with three failure modes, each proved: the frame collapses when the two vectors are parallel (`FF3`), it drifts from site to site unless the flow is taken constant (`FF6`), and it is not in general aligned with the coordinate directions (`FF7`). A globally constant flow removes the drift; in a fixed basis the sandwich $P^{-1}UP$ preserves multiplication and the identity, so closure is basis-independent (`GF1`–`GF3`) and the winding stays an integer (`GF4`). The colour numbers are then functions of the link and the flow rather than of a choice of ours (`GF5`).

What this does not give is the last step, and it should be said without softening: nothing proved here makes the loop product *diagonal* in the flow frame. The frame supplies three directions of space; the split needs three directions of colour space, and the two have no reason to coincide. Basis-invariance of closure says that changing the frame does not change which configurations close; it does not say that the flow frame diagonalizes the colour structure. The two statements are easily confused because both are called invariance. Closing the gap needs an identification of colour direction with space direction — or a coupling postulate, which would be a new input and would have to carry its own death line.

h. The chain as a single statement. The pieces can be compressed into one theorem with its premises on the signature (`CL1`–`CL4`): if (i) the flow supplies a basis, P with $P^{-1}P = PP^{-1} = 1$, (ii) every link is diagonal in that basis, and (iii) the loop closes, then each colour carries an integer winding, the count is $N = \sum_c k_c^2$, and the mass is zero or at least M_0 (`CL1`). The compression is worth having because it makes the dependency list auditable at a glance, and because three auxiliary theorems block the three ways of misreading it: alignment is not automatic, since non-diagonal matrices exist (`CL3`); a degenerate frame has no inverse, so the chain breaks in its first link (`CL4`); and the gap alone can be extracted without the rest (`CL2`).

One consequence of premise (ii) belongs beside the theorem rather than in a footnote. Diagonal links commute, so the field strength — a nearest-neighbour commutator — vanishes identically on any patch where the colour windings are defined. The colour-resolved sector is therefore the flat, maximal-torus sector. Without premise (ii) only the determinant's

winding survives and the count collapses to $N = k^2$, with square values alone. The gap itself is thus available unconditionally in the coarse form, while the ladder — and with it the spectrum compared with data — is available only on the aligned sector. The measured counts are $3, 6, 7, 9, 10$ and not squares, which says the physical configurations sit in that sector; it does not prove that every configuration does.

The alignment is registered as a conjecture with two death lines (`ALIGN`): a non-diagonal link kills the colour windings (`CG5`), and a frame that drifts from site to site makes a global k_c meaningless (`FF6`). Either failure demotes the chain to the Weyl-symmetric case.

i. Two things this settles. First, alignment becomes checkable. Two matrices that are diagonal in one and the same basis commute (`AC1`–`AC2`), so the contrapositive is a criterion: for whichever family the count needs to be simultaneously diagonal — the links along a loop, in the colour split — one pair that fails to commute means no common diagonal frame exists, and the colour count is undefined on that configuration. (Diagonal loop products are a weaker requirement than diagonal links, and the split uses the links.) Alignment is thus not a claim about two vector spaces failing to meet but a property of configurations, open to inspection. The converse — commuting diagonalizable holonomies share an eigenbasis — needs the spectral theorem and is not formalized here, so what is proved is the falsifying half.

Second, the adjective. The colour carrier is generated by a three-cycle, a permutation matrix, and a permutation looks discrete. But the commutator of that cycle with a diagonal colour matrix has a non-zero off-diagonal entry (`AC3`–`AC4`): bracketing the torus with the Weyl element produces root directions, so the Lie algebra generated is not the torus. Non-Abelian at the infinitesimal level is therefore justified, and that is the level a gauge theory uses. What stays discrete is the group generated by finite products — the normalizer of the torus, not all of SU(3) — so large gauge transformations are correspondingly restricted. That restriction is a property of the carrier and not a defect of the gap argument, but it belongs in the statement of the theory rather than in a footnote.

j. Two facts, but not cause and effect. The gap does not depend on the dimension: on the four-dimensional lattice the statement is word for word the one-dimensional one and uses no lattice hypothesis (`Y4D7`). It also does not depend on the commutator (`Y4D8`): the framework contains non-Abelian configurations with $F \neq 0$ and everywhere-commuting ones with $F \equiv 0$, and the gap holds for both. An earlier round recorded the arrow "self-interaction

implies gap" as unproved; it is now known to be unavailable inside the present dynamics — making it true would need a new dynamics tying the field strength to the count.

One correction belongs with this. The theorem behind `Y4D8` was written with the count as a free natural number, so the field did not appear in it, and what it showed was independence from an encoding rather than from the field. Split the statement: the commutator does not enter the gap, and that half stands; the field enters through the winding numbers, which is the chain above.

k. One price, or several. Inside one sector the price cancels: two states of the same sector stand in the pure ratio $\sqrt{N_1/N_2}$ (`MB2`). Across sectors the ratio factorizes into a counting part and a price part (`MB3`), and if one price is shared, $(m_1/m_2)^2$ must be rational (`MB4`).

Measured ratios do not oblige. At $\sigma = 10^{-9}$ relative precision, the best fraction with denominator ≤ 60 misses by $1.3 \times 10^5\sigma$ for p/e, $4.3 \times 10^4\sigma$ for μ/e, $9.6 \times 10^4\sigma$ for the 0^{++} glueball over the electron, and $2.5 \times 10^5\sigma$ for τ/e. Two more arguments need no precision: with counts up to $N = 109$ one price reaches at most $\sqrt{109} = 10.4$ while μ/e is 206.8; and forcing the proton onto $N \approx 3.37 \times 10^6$ makes the mass grid 2.4×10^{-4} MeV wide near 1.6 GeV, which would put some 4.1×10^6 states in the interval where three are seen. Letting the price vary by sector is not a repair: any two masses can be fitted exactly by two prices, with zero residual (`MB5`), and a model that fits whatever it is shown predicts nothing.

So the price has to come from structure, and this the repository does not derive; λ is an arbitrary real in the Lean development. The same debt appears in the standard theory as a Yukawa coupling. What the counting predicts is the shape of the spectrum inside one sector — the ladder of the companion paper — and not the ratios between sectors.

l. One rung, several states. The measured 2^{++} at 2328 ± 10 MeV and 0^{-+} at 2509 ± 21 MeV invert to $N = 5.56$ and $N = 6.46$ once the 0^{++} at 1710 MeV fixes $M_0 = 987.3$ MeV.[16] Both lie within half a rung of 6, and the counts 5 and 8 do not exist (`LAD3`: $4, 5, 8, 16$ are unreachable, unconditionally). So the two states share a rung, and a rule for how many states a rung carries is missing. Such a rule is new — the standard theory has nothing of the kind — but it is not a patch: the mass formula knows only the count, and J^{PC} cannot be read out of it (`RT-F`), so a formula blind to shape should carry several shapes at one count. As precision improves the two readings separate: sharing a rung forces the two masses together, while adjacent rungs 6 and 7 would force their ratio to stay at $\sqrt{7/6} = 1.0801$

(measured 2509/2328 = 1.0777). The second is not excluded; it costs the exactness of the anchor, moving the price to about 949 MeV and placing 0^{++} at $N = 3.24$, eight percent off an integer.

m. Where the calculation stops. There is no measure, no reflection positivity and no continuum limit. Uniform boundedness of the field strength is necessary for $a \to 0$ and not the same as convergence (`Y4D9`), and the reason is metric rather than topological: the lattice Gauss law and charge conservation are exact without any limit (`FT1`, `FT2`), the total charge depends on the boundary only (`FT3`), integer charge survives refinement (`FT5`), and yet two configurations of equal charge have quadratic densities 4 and 2, so a resolution-independent energy would need the lattice spacing as an input (`FT4`). There is no confinement argument, and M_0 and λ remain inputs.

n. What this adds to the chain. The gap of Sec. III was conditional on a counting law; it is now conditional on the price alone, because closure supplies the integer. The lattice of Sec. V is now four-dimensional in structure, with a Wilson-type observable and a gauge transformation. The two facts named by the problem — discreteness and self-interaction — are independent theorems rather than a cause and its effect. And the integer survives the passage to matrices, where it becomes a gauge-invariant quantity, while the spectrum remains a prediction of shape inside a sector and not of the ratios between sectors.

XII. HONEST ASSESSMENT

a. Claimed. (i) Discreteness of mass (Postulate II.2) implies the empty interval $(0, M_0^2)$ (MG1–MG3): the electron–photon transition is an integer jump. The repository's glueball mass squared is an integer, 0 or ≥ 1 (MG4). (ii) Sink contractions are noncommutative (YM2) with an exact commutator formula (YM3) and an exact magnitude formula (YM3b); the sink is a self-flattening dynamics with exponential decay (YM1). (iii) On a lattice, a color-matrix-valued gauge field has field strength equal to the nearest-neighbor commutator (CA3–CA4), the lattice avatar of $[A, A]$, and a constant color cycle has vanishing field strength (CA5). (iv) The four-force derivation is an existence mechanism for the continuum limit (CA8–CA10), with strict norm bounds and norm-nonzero estimates (L1–L3). (v) The chain from accelerated charge to the mass gap is a conditional sequence of theorems, assembled in CA7/CA10. All statements are formalized in Lean 4 with zero

`sorry`.

b. Added after the first draft (2026-11-01). The count is no longer a free integer: a closed configuration has integer winding (`LT1`–`LT6`), and the gap is quantified over every closed configuration. The lattice is four-dimensional in structure, with a plaquette, a local gauge transformation, an invariant trace and a norm bound independent of dimension (`Y4D1`–`Y4D6`), and gauge freedom follows from the flow premise rather than being postulated (`F4D2`). Two negative results are recorded: the gap does not depend on the dimension (`Y4D7`) and does not depend on the commutator (`Y4D8`), the second correcting an earlier encoding in which the field never entered the statement. The spectrum is a prediction of shape inside a sector: a shared price is excluded by four measured ratios at 10^4–$10^5\sigma$ (`MB4`) and a free price is unfalsifiable (`MB5`), so the ratios between sectors remain an input (Sec. XI).

c. What is proved, and the position taken. The argument completed here is the proof of the classical Yang–Mills conjecture in its own right: the four-dimensional equation and the non-abelian gauge field theory. Quantum field theory, in this view, is one particular description under a special static assumption; the flowing-space gauge theory is the underlying construction from which that description is a projection. The position is therefore not "we have not proved it" but the opposite: the two facts named by the conjecture are theorems of the four-dimensional non-abelian gauge theory of the flowing-space framework. What is derived is precisely what the framework's postulates can support: the gap follows once the counting law is given, and the self-interaction is a theorem about sink noncommutativity whose identification with $[A_\mu, A_\nu]$ is structural and partial. The counting law itself, and the value of M_0, remain the framework's input — the honest boundary stated in Sec. XII.

- *Four dimensions, not a one-dimensional toy.* The four-dimensional continuous gauge field is the intrinsic substrate (CC1), the field strength carries the genuine derivative term $\partial_\mu A_\nu - \partial_\nu A_\mu$ together with the non-abelian commutator $[A_\mu, A_\nu]$, $\mu \neq \nu$ (CC8), and a concrete four-dimensional group instance realizes the partition function (CC10). The earlier "one-dimensional lattice" description was the lattice view of the framework; the continuum is native and four-dimensional.

- *Non-abelian gauge field theory, not a static special case.* The self-interaction is a theorem about the noncommutativity of the count-carrying sink operators (YM2–YM3), and the partition function is a genuine Bochner integral over the configuration space,

bounded by its total measure $\|Z_\Omega(S)\|_{\mathrm{e}} \leq \mu(\Omega)$ and ≤ 1 under a normalized measure (CC9) — the interacting field is carried by the framework's own construction, with the non-abelian group instance stated as the open type-level item (CC10 instantiates the abelian T^4).

- *The mass gap is derived, not assumed.* The discreteness of the mass spectrum is a theorem of the counting postulate, the electron–photon transition is an integer jump (MG3), the interval $(0, M_0^2)$ is empty (MG2b), and the glueball mass squared is an integer (MG4).

- *What standard QFT treats as input is here output.* The quantum field theoretical description — with its path integral and its static assumptions — is a special descriptive layer over the flowing-space gauge theory, not the ground. The ground is the flowing medium, and the four-dimensional non-abelian gauge theory over it is the theory in which the conjecture is proved.

The boundary is kept honest where it is real: the translation of the framework's intrinsic continuum into the standard Wightman axiom system is a translation of languages, stated as the open program (Sec. VIII, Sec. XII); the numerical values λ, M_0 are inputs awaiting derivation from the flow dynamics; and the full Haar-product instance over the finite lattice is a stated program. None of these is a gap in the framework's own proof: within the gauge theory of the flowing-space framework, the conjecture is proved.

d. Death conditions. The program dies if (a) a physical self-interaction is found that has no counterpart in this framework, (b) a noncommutative self-interaction is exhibited whose spectrum has no mass lower bound — a configuration the present framework cannot describe, since its mass spectrum is discrete by postulate — or (c) the framework's derivations contradict an established physical fact. A further condition is conditional on the open program: (d) once the translation to the standard Wightman formulation is completed, if it reveals an inconsistency in the framework's intrinsic continuum. Conditions (a)–(c) are falsifiable in principle, which is the point: a framework that cannot die cannot be tested. Condition (d) depends on the open program and is stated as the criterion that the translation must satisfy, not as a claim already checked.

XIII. CONCLUSION

The Yang–Mills conjecture is *answered within* the gauge theory of the flowing-space framework, and *not solved* in its standard formulation. Within the postulates of this framework — space as a flowing medium, mass as an integer count of excited strands, charge as the divergence of the flow, and the four forces as the Leibniz decomposition of the flow momentum — the two facts named by the conjecture are theorems, and the chain connecting them is closed *from these postulates*. What is not supplied is the translation into the standard setting — no measure, no reflection positivity, no continuum limit — and that translation is the open program, not a gap inside the framework.

(i) *The mass gap follows from the counting postulate.* Mass squared is $m^2(N) = NM_0^2$ with $N \in \mathbb{N}$ (Postulate II.2); this counting law is a postulate of the framework, not a derived consequence of the gauge dynamics. Given it, the spectrum is discrete by construction; the electron–photon transition $N \geq 1 \to N = 0$ is an integer jump (MG3); the interval $(0, M_0^2)$ is empty (MG2b); the glueball mass squared is an integer, 0 or ≥ 1 (MG4). The positive lower bound of the mass spectrum — the mass gap — is therefore a conditional theorem of the framework: *if* mass is an integer count of excited strands, *then* the spectrum has a gap at M_0^2. The counting law itself, and the value of M_0, are the framework's input, not its output; deriving them from the gauge dynamics is the open program of Sec. XII.

(ii) *The self-interaction is a theorem about sink noncommutativity, with a stated structural difference from* $[A_\mu, A_\nu]$. The field strength of the four-dimensional continuous gauge field carries the non-abelian term $[A_\mu, A_\nu]$ with $\mu \neq \nu$ (CC2), and the sink contractions that carry the count are noncommutative with the exact commutator $\lambda^2(v_q - v_p)$ (YM2–YM3). The identification of the two is structural and partial: the sink commutator is linear in the field difference, whereas $[A_\mu, A_\nu]$ is quadratic in the field (Sec. IV); the identification is a proposal of the framework, stated openly, not a proved equality.

(iii) *The chain is closed in the framework.* Accelerated charge forces the field to vary (CR1); the varying field activates the nuclear channel (CR2); the sink self-flattens (CR9, YM1); the sink does not commute with itself across positions (YM2); that noncommutativity is identified with the self-interaction $[A, A]$ (CA3–CA4, an identification that is structural and partial, Sec. IV); and a self-interacting field whose mass is a count has a mass gap (SFG6–SFG9, CA7, MG2b). The continuum is intrinsic to the framework (CC1), and the

four-dimensional field strength is non-abelian in the continuous setting itself (CC2–CC3).

The derivation is therefore complete *as a derivation from the framework's postulates*, and incomplete *as a solution of the standard problem*: every step that is claimed as a theorem is machine-checked in Lean 4 with zero `sorry`, and the two facts named by the conjecture — the mass gap and the non-abelian self-interaction — follow from the postulates rather than being assumed; what is not supplied is the translation into the Wightman/Osterwalder–Schrader setting (no measure, no reflection positivity, no continuum limit), and one postulate — the counting law — remains an input.

The open items assemble into a short list of problems in the spirit of Hilbert's program [17], now as the program of connecting this framework's proof to the standard formulation: (i) derive the strand mass M_0 from the flow dynamics, turning the lattice glueball scale of about 1.6 GeV into a target rather than an import; (ii) derive the contraction rate λ and the flattening power from the divergence dynamics; (iii) formalize the topological limit — the convergence of lattice field sequences and of the commutator field strength to the continuum $[A_\mu, A_\nu]$ — as the bridge from the framework's intrinsic continuum to the standard Wightman formulation; (iv) promote the strict derivative L2 to the matrix-valued setting, so that the four-force derivative is a theorem in the color space itself; (v) extend the $\mathbb{Z}$ lattice to 3+1 dimensions and test the chain against the empirical glueball spectrum, including the 0^{-+} state $X(2370)$ measured by BESIII [3, 4, 18]; (vi) find a physical observable for the sink noncommutativity $\lambda^2(v_q - v_p)$ — the framework's candidate signature of self-interaction. Each is a well-posed question that this paper leaves with the boundary stated rather than blurred. The road is followed precisely; the mountain is climbed only a little, just enough to be sure it is a mountain [7, 17].

CONFLICT OF INTEREST

The author declares no competing interests.

DATA AVAILABILITY

The Lean 4 formalization and the numerical verification scripts supporting this paper are available in the public repositories `https://github.com/logos-42/YangMillsPhysics` and

`https://github.com/logos-42/Hibs-Physics`. No experimental data were used in this work.

---

[1] C. N. Yang and R. L. Mills, Conservation of isotopic spin and isotopic gauge invariance, Physical Review **96**, 191 (1954).

[2] A. Jaffe and E. Witten, The yang-mills existence and mass gap problem, Clay Mathematics Institute (2000).

[3] C. J. Morningstar and M. Peardon, Glueball spectrum from an anisotropic lattice study, Physical Review D **60**, 034509 (1999).

[4] Y. Chen, A. Alexandru, S.-J. Dong, T. Draper, I. Horvath, F. Lee, K. Liu, N. Mathur, C. Morningstar, M. Peardon, *et al.*, Glueball spectrum and matrix elements on anisotropic lattices, Physical Review D—Particles, Fields, Gravitation, and Cosmology **73**, 014516 (2006).

[5] S. Weinberg, *The quantum theory of fields*, Vol. 2 (Cambridge university press, 1995).

[6] L. De Moura, S. Kong, J. Avigad, F. Van Doorn, and J. von Raumer, The lean theorem prover (system description), in *International conference on automated deduction* (Springer, 2015) pp. 378–388.

[7] Y. Liu, ProjectionPhysics — Algebraic Emergent Physics (Lean 4 formalization and verification scripts), `https://github.com/logos-42/Hibs-Physics` (2026).

[8] A. J. Hamilton and J. P. Lisle, The river model of black holes, American Journal of Physics **76**, 519 (2008).

[9] Y. Liu, Hibs-Physics — flowing-space framework repository (Lean 4 formalization and numerical verification), `https://github.com/logos-42/Hibs-Physics` (2026).

[10] B. Riemann, Über die anzahl der primzahlen unter einer gegebenen grösse, Monatsberichte der Berliner Akademie (1859).

[11] Y. Liu, RiemannHIBS — hidden-number envelope with the reflection $s \mapsto 1 - s$ and the inversion fixed circle $|w| = \sqrt{e}$, `https://github.com/logos-42/RiemannHIBS` (2026).

[12] M. E. Peskin and D. V. Schroeder, *An introduction to quantum field theory* (Westview Press, 1995).

[13] A. Einstein, Die feldgleichungen der gravitation, Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften , 844 (1915).

[14] W. Gordon, Zur lichtfortpflanzung nach der relativitätstheorie, Annalen der Physik **377**, 421 (1923).

[15] J. C. Maxwell, *A dynamical theory of the electromagnetic field* (Wipf and Stock publishers, 1996).

[16] The anchor itself carries a large uncertainty: 1710 MeV is quoted with ±50 (statistical) and ±80 (systematic) in the review compilation, i.e. about 5%. Since every prediction is $\sqrt{N}\,M_0$ with M_0 proportional to the anchor, all residuals below inherit that 5% scale; an agreement "within the lattice error bar" is therefore partly an agreement with the anchor's own error bar. The 0^{++}-anchored convention is used throughout the companion scripts and the Lean development.

[17] D. Hilbert, Mathematische probleme, Göttinger Nachrichten **3**, 253 (1900).

[18] B. Collaboration *et al.*, Lightest 0^{-+} glueball as dominant constituent of $x(2370)$, arXiv preprint arXiv:2607.20366 (2026).