Hadronic Mass Generation through Three-Dimensional Magnetic Locking:
A Discrete Resolution for Yang–Mills Theory

Empirical and Logical Refutation of Continuous Field Ontology in Quantum Chromodynamics

Dr. André Guzzon

andre.guzzon@ufrgs.br

April 2026

Abstract

The Yang–Mills mass gap problem asks for a rigorous proof, within the axiomatic framework of constructive quantum field theory, that quantum Yang–Mills theory in four-dimensional Minkowski space exists and exhibits a strictly positive mass gap $\Delta > 0$. We argue that the persistent failure of the program is the consequence of a categorical error: the continuum functional integral that the program seeks to construct sums over a configuration space whose elements have no ontological referent. We axiomatize an Endpoint-Only Ontology in which physical content is exhausted by interaction events and conditional probability distributions over them, with no intermediate trajectories. Within this ontology, we formalize a mechanism for hadronic rest mass generation based on three-dimensional magnetic locking of discrete photonic packets, derive linear quark confinement and the deconfinement transition, identify the three colour charges of QCD with the three orthogonal rotation axes of the lock structure, and show that SU(3) emerges as a mathematical consequence of antisymmetry combined with three-dimensional rotational closure. The mass gap is a one-line corollary: stable rotational closure in three dimensions requires at least three coupled locks, each carrying positive lock energy. We prove a non-existence theorem for the continuum limit of confining non-Abelian gauge theory: under the Endpoint-Only Axiom, no probability measure on the space of smooth gauge connections satisfies the Osterwalder–Schrader axioms while reproducing the conditional probability structure required by confinement and a positive mass gap. The argument employs the elementary fact that a tempered distribution invariant under continuous translations cannot have countable support unless it is zero. The Higgs mechanism is preserved as the parameterization of bare elementary-fermion mass; the lock mechanism accounts for the dominant 99% of hadronic rest mass that the Higgs alone does not explain. Eight falsifiable predictions are provided. The Clay problem, on the present analysis, is dissolved: the object whose existence is to be proved does not exist as a continuous field theory, and the mass gap exists trivially in the discrete description that does.

MSC2020: 81T13 (Yang–Mills and other gauge theories), 81T25 (Quantum field theory on lattices), 81P05 (General and philosophical questions in quantum theory).

Keywords: Yang–Mills mass gap, quantum chromodynamics, discrete ontology, lattice gauge theory, non-existence theorem, colour confinement, SU(3) emergence.

0 Why the Problem Has Resisted Solution

The Clay Mathematics Institute formulation of the Yang–Mills problem requires a rigorous proof, within the axiomatic framework of constructive quantum field theory [1, 2], that quantum Yang–Mills theory exists on $\mathbb{R}^4$ and that its spectrum has a positive mass gap. The formulation embeds three implicit assumptions:

1. That the theory's existence is a property to be demonstrated within a continuous field formalism.
2. That the mass gap is a non-trivial spectral feature requiring derivation from first principles of that formalism.
3. That four-dimensional spacetime is a continuous manifold $\mathbb{R}^4$ on which gauge connections are defined as smooth (or appropriately distributional) sections.

We argue that all three assumptions are wrong, and that the persistent failure of the program is direct evidence of this. Glimm and Jaffe [4] succeeded in constructing Yang–Mills rigorously in two and three dimensions; the four-dimensional case has resisted analogous treatment for half a century. The Gribov ambiguity [6, 7] demonstrates that gauge fixing in non-Abelian theory is non-unique — a structural feature suggesting that gauge redundancy is a property of the mathematical map rather than of the physical territory. Quantum triviality results [8, 9] indicate that the continuum limit of ϕ^4 theory in four dimensions exists only as a free field; analogous concerns apply to pure Yang–Mills, where renormalization cancels infinities without explaining their origin.

The pattern is consistent with a misformulated problem rather than an unsolved one. We propose that the continuous field formalism is an asymptotically valid approximation in regimes of high interaction density (ultraviolet, weak coupling) but is singular in regimes where the discrete structure of physical interactions becomes resolved (infrared, strong coupling, confinement). The mass gap problem lives precisely in the latter regime.

Relation to other critiques of continuous QFT

Several existing programs have identified problems in the continuous-field formulation and proposed alternatives. Causal set theory [16] replaces the continuum with discrete causal events but retains a metric-theoretic interpretation. Loop quantum gravity [17] discretizes geometry but preserves the field-theoretic conceptual framework. 't Hooft's deterministic interpretation [18] questions the ontological status of superposition. Kastner's transactional interpretation [19] questions the ontological status of probability amplitudes between measurement events.

The present proposal differs from each of these in being explicitly ontological at the level of physical existence: the discrete structure is not a regularization, not a representation, and not a deterministic substrate beneath a probabilistic surface — it is what exists. The continuous functional integral does not approximate a deeper continuous reality; it computes correctly because it preserves the empirically real endpoints (interaction events) under a summation procedure whose summands have no physical referent.

1 The Jaffe–Witten Problem

Following Jaffe and Witten [3], the problem statement requires:

1. A rigorous mathematical construction of quantum Yang–Mills theory with compact simple gauge group G on $\mathbb{R}^4$, satisfying a chosen system of axioms (typically Wightman [1] or Osterwalder–Schrader [2]).

2. Proof that the resulting theory exhibits a mass gap: there exists $\Delta > 0$ such that every state in the Hilbert space (other than the vacuum) has energy at least Δ.

For $G = \mathrm{SU}(3)$, this is the theory underlying quantum chromodynamics. The phenomenological mass gap is observed: the lightest hadron (the pion) has mass approximately 140 MeV, and free coloured states are not observed. The theoretical task is to derive this from first principles within the axiomatic framework.

The lattice formulation of Wilson [10] provides a regularized, discrete version of the theory in which both existence and confinement are unambiguous. The Clay problem implicitly treats lattice QCD as a numerical approximation to the underlying continuous theory — the goal is the continuum limit $a \to 0$ where a is the lattice spacing. The present paper inverts that hierarchy: the lattice is not the approximation; it is the description, and the continuum limit is singular.

2 Limitations of the Continuous Field Approach

We catalogue the structural difficulties of the continuous formulation that constitute strong evidence the framework rather than the calculations is at fault.

Ultraviolet divergences. Perturbative quantization of Yang–Mills produces divergent integrals at every loop order. Renormalization absorbs these into redefinitions of coupling and field strength, yielding finite predictions for observables. The procedure is empirically successful but conceptually opaque: the divergences are interpreted as artifacts of treating the theory as fundamental at all energy scales, but no mechanism is provided for what replaces the continuous description at high energies.

Gribov ambiguity. In non-Abelian gauge theory, the standard Faddeev–Popov gauge-fixing procedure does not produce a unique representative on each gauge orbit [6, 7]. Multiple solutions exist (Gribov copies), and the path integral defined over the full gauge-fixed configuration space includes contributions that cannot be physically interpreted. We read this as ontological rather than technical: the parameterization has more degrees of freedom than the system possesses, because gauge connections in $\mathbb{R}^4$ are an over-rich representation of an underlying discrete structure that does not require gauge symmetry as a primitive.

Quantum triviality. Rigorous results in lattice scalar field theory establish that the continuum limit of ϕ^4 theory in four dimensions yields a free (non-interacting) theory [8, 9]. The general lesson — that the four-dimensional continuum limit can fail to preserve the interacting content of the regulated theory — shifts the burden of proof: any claim that Yang–Mills evades this pathology must justify why, given that no analogous constructive result exists in four dimensions despite half a century of effort [5].

Failure of dimensional analogy. Yang–Mills has been rigorously constructed in two and three dimensions [4]. The four-dimensional case has resisted the same techniques. We interpret this as the entry of the discrete ontology into the regime where it cannot be treated as a small correction. In two and three dimensions, low effective interaction density permits the continuous approximation to remain asymptotically valid through the construction. In four dimensions, the effective interaction density at the relevant scales is high enough that the discrete structure becomes resolved within the procedure itself, and the construction fails.

3 The Endpoint-Only Axiom

We now formalize the ontological premise on which the remainder of the paper rests.

Axiom 1 (Endpoint-Only Ontology)**.** Let $\mathcal{E}$ denote the set of physical interaction events. For any pair of causally connected events $E_1, E_2 \in \mathcal{E}$ with $E_1 \prec E_2$, the physical content of the process between E_1 and E_2 is exhausted by:

1. the events E_1 and E_2 themselves, with their intrinsic properties (energy, charge, spin orientation);
2. the causal relation $E_1 \prec E_2$;
3. the conditional probability distribution $P(E_2 \mid E_1, \mathcal{C})$, where $\mathcal{C}$ denotes the configuration of all other events $E_i \in \mathcal{E}$ in the past lightcone of E_1.

No additional ontological structure exists between E_1 and E_2. In particular, no trajectory, no intermediate field configuration, and no superposition of trajectories or configurations is part of the physical content of the process.

The axiom captures three commitments simultaneously. First, it asserts that the physically real entities are interaction events, not intervening fields. Second, it captures Heisenberg's uncertainty relations as epistemic limits on the observer's specification of $P(E_2 \mid E_1, \mathcal{C})$ rather than as ontological indeterminacy of an existing object. Third, it interprets what is conventionally called a "force field" as the conditional probability distribution itself: the more completely the configuration $\mathcal{C}$ is specified, the more sharply $P(E_2 \mid E_1, \mathcal{C})$ concentrates on a particular outcome, but this concentration never reaches certainty because the conditional distribution remains a distribution and not a deterministic function. The empirical content of "force fields" is preserved as the conditional probability structure; their ontological content as substantial entities filling space is denied.

4 Rest Mass as the Energy of Magnetic Locking

We propose the following mechanism for hadronic rest mass generation, formalized within the Endpoint-Only Axiom.

4.1 Photonic packets and proper-time tick counting

A photonic packet, having null worldline, does not register proper time in its own frame. Under Axiom 1, an emission event E_1 and an absorption event E_2 separated by null spacetime interval are connected directly through the causal relation $E_1 \prec E_2$ with no intermediate ontological content. The empirical light speed c measured by massive observers is the rate at which the observer's own internal interactions tick between the two endpoints, not a velocity of propagation of an intermediate object.

A massive object's rest energy m_0c^2, by contrast, is the energy associated with internal interactions that proceed in its own proper-time frame. We adopt the operational identification:

Definition 4.1 (Tick-counting rest mass)**.** The rest mass of a system S is given by $m_0 = E_S^{\text{internal}}/c^2$, where E_S^{internal} is the integrated energy of internal interaction events of S per unit of S's own proper time. A system that registers no internal interactions in its own frame has zero rest mass; a system that registers a non-zero rate of internal interactions has non-zero rest mass equal to that rate's energy content divided by c^2.

This definition is consistent with $E = mc^2$ [11]: rest energy is the integrated content of internal interaction over a proper-time interval. A free photonic packet has $m_0 = 0$ because it has no proper time and no internal interactions; a bound system with sustained internal dynamics has $m_0 > 0$.

4.2 The locking event

Three-dimensional magnetic interaction with full rotational freedom is empirically asymmetric: attraction dominates over repulsion, with sustained repulsion becoming geometrically impossible as rotational degrees of freedom increase [24]. In the limit of full rotational freedom (spherical magnets), repulsion is impossible (50 trials, $P = 2^{-50}$, observed: 100% attraction). Rectangular magnets in 64 orientations exhibit a 7:1 attraction-to-repulsion ratio, statistically incompatible with Maxwell's prediction of 1:1 symmetry under conditions of free reorientation.

Definition 4.2 (Magnetic lock)**.** A magnetic lock is a sustained bound configuration of two photonic packets with opposite spin orientation in mutual three-dimensional magnetic interaction range, in which internal interaction events register in the proper-time frame of the bound system. The lock energy $\varepsilon_{\text{lock}}$ is the integrated energy content of the bound configuration, including the photonic energy of the two packets and the magnetic binding energy of the rotational coupling.

By Definition 4.1, a magnetic lock has rest mass $m_0 = \varepsilon_{\text{lock}}/c^2 > 0$. The locking event is reversible: a lock may dissolve into two free photonic packets of opposite spin (process $m \to e$), and two free packets of opposite spin entering interaction range may form a new lock (process $e \to m$). The local rate of $e \to m$ relative to $m \to e$ determines the dynamic equilibrium of bound matter.

4.3 Rotational structure of multi-lock systems

A single magnetic lock fixes one rotational axis: the axis of spin–spin coupling between the two packets. Multi-lock configurations admit multiple coupled rotational axes. We require, for the formal arguments below, the following standard results from the theory of angular momentum representations.

Lemma 4.1 (Rotational closure in three dimensions)**.** *Let S be a bound system in $\mathbb{R}^3$ with total angular momentum $\mathbf{J}$. If S is rotationally stable — in the sense that no internal torque can decouple the rotational degrees of freedom into a lower-dimensional invariant subspace — then the rotational phase space of S has dimension 3, and the Lie algebra of operators preserving the stability constraint is $\mathfrak{su}(2)$, the double cover of $\mathfrak{so}(3)$.*

Proof. This is a standard result of angular momentum theory in quantum mechanics [20, Ch. 3]. Rotational invariance in $\mathbb{R}^3$ requires that physical states classify under irreducible representations of SU(2), the universal cover of the rotation group SO(3). Three independent quantization axes correspond to the three components of $\mathbf{J}$. A bound system with rotational stability admits all three components as physically realizable observables; absence of any component would correspond to a degenerate axis along which the system is not bound. □ □

Lemma 4.2 (Single-lock axis)**.** *A magnetic lock between two photonic packets with opposite spin fixes a single axis of coherent rotation: the axis of spin–spin coupling. The system of two packets in a single lock has only one rotationally stable direction and therefore does not satisfy the closure condition of Lemma 4.1.*

Proof. By Definition 4.2, a magnetic lock is constituted by spin–spin coupling between two packets of opposite orientation. The coupling defines a unique direction in $\mathbb{R}^3$, namely the line connecting the two packets and along which their spins are anti-aligned. Rotation about axes orthogonal to this line is not stabilized by the lock's binding mechanism; the system has only one degree of rotational stability. □ □

Lemma 4.3 (Minimal three-axis configuration). *A bound configuration in $\mathbb{R}^3$ satisfying Lemma 4.1 requires at least three coupled magnetic locks whose axes span $\mathbb{R}^3$, i.e., are linearly independent. The minimal such configuration consists of three photonic packets coupled pairwise to form three lock axes.*

Proof. By Lemma 4.2, each lock contributes one rotational axis. By Lemma 4.1, three linearly independent axes are required for rotational closure. With $n < 3$ packets, only $\binom{n}{2} < 3$ pairwise lock axes exist; rotational closure cannot be achieved. With three packets, three pairwise locks are available, and three linearly independent axes can be configured. □ □

4.4 Colour charge as rotation axis: the SU(3) structure

The three colour charges of QCD admit, within this framework, an identification with the three rotational axes of the minimal three-lock configuration of Lemma 4.3.

Lemma 4.4 (Antisymmetry constraint). *A configuration of three identical fermions in $\mathbb{R}^3$ requires that the total wavefunction be antisymmetric under permutations in S_3, by the spin-statistics theorem.*

Lemma 4.5 (Decomposition of the S_3 action). *Let the configuration of three identical fermions decompose as a tensor product of spatial, spin, and "internal" degrees of freedom: $\Psi = \psi_{spatial} \otimes \chi_{spin} \otimes \phi_{internal}$. For the total wavefunction to be antisymmetric under S_3 while admitting symmetric or partially symmetric components in the spatial and spin factors (as required for ground states of three quarks in a baryon), the internal factor must transform under a three-dimensional representation of an internal symmetry group whose action on the three fermion labels is rich enough to absorb the residual symmetry of the spatial–spin product. The minimal such group with a faithful three-dimensional fundamental representation and the required structure constants is* SU(3).

Sketch. The decomposition follows the standard analysis of the spin-statistics constraint on baryonic wavefunctions [22, 23]. The spatial wavefunction of a ground-state baryon (S-wave) is symmetric, and the spin wavefunction of a $J = 3/2$ baryon (such as Δ^{++}) is symmetric. To recover total antisymmetry, the internal factor must be totally antisymmetric. The minimal group admitting a totally antisymmetric three-particle state in a faithful representation (and consistent with the experimental requirement of three distinct charge states for the internal label) is SU(3), with the antisymmetric state corresponding to the singlet representation $\mathbf{1}$ within the decomposition $\mathbf{3} \otimes \mathbf{3} \otimes \mathbf{3} = \mathbf{10} \oplus \mathbf{8} \oplus \mathbf{8} \oplus \mathbf{1}$. The dimension three of the internal label, in the present geometric framework, identifies with the three rotational axes of Lemma 4.3. □ □

Proposition 4.6 (Geometric origin of SU(3)). *Under the Endpoint-Only Axiom and the magnetic locking mechanism, the gauge group* SU(3) *of QCD emerges as the unitary symmetry group of the three orthogonal rotational axes of the minimal three-lock bound state, with the colour assignments $\{r, g, b\}$ identifying with the three axes. The eight gluons correspond to the eight non-trivial generators of $\mathfrak{su}(3)$, viewed as operations exchanging or recombining the three rotational axes; the structure constants of $\mathfrak{su}(3)$ are determined by the algebra of these axis-exchange operations on the antisymmetrized three-fermion state.*

The full derivation of the structure constants of $\mathfrak{su}(3)$ from the geometry of three-axis rotation combined with the antisymmetry constraint is not given here in closed form; it is sketched at the level of representation-theoretic content. We mark the complete derivation as an open problem (Appendix B) and rely on the present sketch as sufficient justification for the geometric identification used in the remainder of the paper. The mass gap and non-existence theorems below do not depend on the closed-form derivation.

The geometric identification has three immediate consequences:

(a) Hadrons must be "white." A stable three-dimensional bound state requires coupled rotation in three axes (Lemmas 4.1, 4.3); rotation in fewer than three axes lacks geometric coherence. The constraint that hadrons be colour-singlets is the statement that bound states require three-axis rotational closure.

(b) Confinement is geometric. Extracting a single colour from a hadron is equivalent to extracting a single rotational axis from a three-axis system: there is no such operation. The energy required to attempt the extraction grows with separation because the geometric coherence must be artificially maintained against the natural reorganization that restores three-axis closure — which is the pair-creation event terminating the attempt.

(c) Gluons are relations, not particles. Gluons are the operators of axis exchange on the lock structure; they have no independent existence outside the bound state. The empirical absence of free gluons follows immediately: a relation cannot exist without its relata.

4.5 Hadronic mass dominated by lock energy

A direct consequence of the locking mechanism is that hadronic rest mass is dominated by the magnetic energy of the lock structure, not by the bare masses of constituent quarks. The proton mass is approximately 938 MeV; the sum of the bare valence quark masses (two up quarks at ~ 2.2 MeV, one down quark at ~ 4.7 MeV) is approximately 9 MeV, less than 1% of the total [14]. The remaining 99% is the lock energy: the binding energy of the rotational lock structure plus the photonic content circulating within it. There is no need for spontaneous symmetry breaking or a separate mass-generating field for this 99%; it is $E = mc^2$ operating on the lock energy.

4.6 Distinction from the Higgs mechanism

The Higgs mechanism [12] accounts for the rest masses of the elementary fermions and electroweak gauge bosons through coupling to a scalar field with non-zero vacuum expectation value. Within the present framework, the Higgs mechanism is interpreted as a phenomenological parameterization of the rate at which different particle species participate in lock formation. It is quantitatively correct; it is causally mute. The lock mechanism describes the physical process; the Higgs vacuum expectation value parameterizes how strongly each species engages with that process.

The two accounts apply to different parts of the mass budget. The Higgs contribution to the proton (the bare quark masses) is the small minority of its rest mass. The lock contribution (the QCD binding energy) is the dominant majority. Both contribute additively in the constituent picture. The empirical content of the Higgs sector — coupling strengths, the 125 GeV resonance — is preserved.

5 Linear Confinement as a Consequence of $E = mc^2$

The phenomenology of quark confinement includes a linearly rising potential between heavy quarks at long distances:

$$V(r) = \sigma r + (\text{short-range corrections}), \tag{1}$$

with string tension $\sigma \approx 1$ GeV/fm observed in lattice QCD and consistent with hadron spectroscopy. We derive this linearity from the lock ontology.

5.1 The flux tube as a one-dimensional lock chain

Lemma 5.1 (Linear flux tube energy). *Let two locked sub-systems (a quark and an antiquark) be separated by distance r. The magnetic structure connecting them under the locking ontology is a chain of intermediate locks — a discrete one-dimensional sequence of magnetic couplings. Each lock contributes a fixed quantum of energy ε_{lock} and occupies a fixed length ℓ_{lock} along the chain. Then the total energy stored in the chain is*

$$E(r) = \varepsilon_{lock} \cdot n(r) = \varepsilon_{lock} \cdot (r/\ell_{lock}) = \sigma r, \tag{2}$$

where $\sigma = \varepsilon_{lock}/\ell_{lock}$ identifies as the string tension.

Proof. Under adiabatic separation of the endpoints, the chain extends by adding additional locks. The cross-section of the tube does not grow because lateral propagation of locks is energetically suppressed by transverse magnetic constraint: rotation about axes orthogonal to the tube is incompatible with the linear geometry of the chain (Lemma 4.2 applied to each link). The chain length therefore grows in proportion to the endpoint separation r, with one lock added per increment ℓ_{lock}. Energy conservation yields equation (2). □ □

5.2 String breaking as $E = mc^2$ thresholding

When the energy stored in the chain reaches the threshold required to materialize a quark–antiquark pair from the surrounding photonic content — approximately $2m_q c^2$ for the lightest available species — the system has access to a lower-energy configuration: the chain breaks, with each fragment terminated by a newly created quark or antiquark. The mechanism is transparent: $E = mc^2$ specifies the energy cost of pair creation, and the system selects the energetically accessible path. The phenomenon called "confinement" is the geometric impossibility of separating locked partners without paying the pair-creation cost — and the pair creation immediately reconstitutes the locked condition with new partners.

5.3 Deconfinement transition

At sufficiently high temperature T_c, the thermal energy density of the medium exceeds the lock binding energy density. Locks dissolve faster than they reform; the medium transitions to a phase in which discrete locks no longer dominate the dynamics — the quark–gluon plasma. The critical temperature is

$$k_B T_c \sim \varepsilon_{\text{lock}}, \tag{3}$$

modulo phase-space factors. Lattice QCD reports $T_c \approx 150$–170 MeV [13], consistent with $\varepsilon_{\text{lock}}$ on the same order of magnitude. The framework predicts that $k_B T_c$ should agree with $\varepsilon_{\text{lock}}$ within a factor of order unity; significant discrepancy ($> 20\%$) would falsify the single-lock-mechanism hypothesis.

6 The Mass Gap as a Geometric Theorem

We can now state the mass gap as a direct consequence of the locking ontology, without recourse to spectral analysis of a continuous Hamiltonian.

Theorem 6.1 (Geometric mass gap)**.** *Under the Endpoint-Only Axiom (Axiom 1) and the magnetic locking mechanism (Definition 4.2), every physical state distinct from the vacuum has rest energy bounded below by*

$$\Delta = 3\,\varepsilon_{lock} > 0. \tag{4}$$

Proof. Let $|\psi\rangle$ be any physical state distinct from the vacuum. By the Endpoint-Only Axiom, $|\psi\rangle$ corresponds to a configuration of physically real interaction events with definite energy content. By Definition 4.1, non-zero rest energy requires a system registering internal interactions in its own proper-time frame; this requires a bound configuration of locked photonic content (Definition 4.2).

By Lemma 4.1, a rotationally stable bound system in $\mathbb{R}^3$ requires three linearly independent rotational axes. By Lemma 4.2, each lock contributes exactly one axis. By Lemma 4.3, the minimal configuration satisfying three-axis closure consists of three coupled locks. Each lock carries energy $\varepsilon_{\text{lock}}$ by Definition 4.2; the total minimal energy is therefore $3\,\varepsilon_{\text{lock}}$.

There are no continuous families of arbitrarily-low-energy excitations because lock formation is a discrete event with non-zero minimum energy: a "half-lock" is not a configuration the ontology admits. Therefore the rest energy of any non-vacuum state is bounded below by $3\,\varepsilon_{\text{lock}}$. □ □

Corollary 6.2. *The Yang–Mills mass gap $\Delta > 0$ is a geometric corollary of three-dimensional rotational closure under the discrete locking ontology, not a spectral theorem requiring proof within the continuous formalism.*

The mass gap is, on this analysis, the statement that a quantum of rotationally stable binding has finite, non-zero energy. In a discrete framework with three-dimensional rotational closure as a stability requirement, this is a direct lemma chain. The Clay problem appears difficult only when posed inside a framework that has structurally hidden this fact.

7 Non-Existence of the Continuum Limit

We now articulate why Yang–Mills in $\mathbb{R}^4$ continuous fails as a mathematical object, and why the lattice description is exact rather than approximate. The argument employs the elementary observation that a non-zero tempered distribution invariant under continuous translations cannot have countable support.

7.1 Configuration spaces

Definition 7.1 (Continuum gauge configuration space)**.** For compact simple gauge group G, let $\mathcal{A}/\mathcal{G}$ denote the space of smooth G-connections on $\mathbb{R}^4$ modulo smooth gauge transformations, equipped with the topology induced by the Yang–Mills action functional.

The space $\mathcal{A}/\mathcal{G}$ is an infinite-dimensional stratified manifold with Gribov-type singularities along the strata of reducible connections [7].

Definition 7.2 (Endpoint configuration space)**.** The endpoint configuration space $\mathcal{C}$ is the set of countable collections $(\mathcal{E}, \prec, P)$ where:

1. $\mathcal{E}$ is a countable set of interaction events, each carrying labels for energy, charge, and spin;

2. $\prec$ is a partial order on $\mathcal{E}$ encoding causal precedence;

3. P is a family of conditional probability distributions $P(E_j \mid E_i, \mathcal{C}_{ij})$ for $E_i \prec E_j$, with $\mathcal{C}_{ij}$ the configuration in the past lightcone of E_i.

The space $\mathcal{C}$ is equipped with the discrete topology on $\mathcal{E}$ and the weak topology on conditional probabilities.

7.2 The non-existence theorem

Theorem 7.1 (Non-existence of the continuum limit). *Let G be a compact simple non-Abelian gauge group, and suppose lattice gauge theory $\mathcal{L}_a(G)$ exhibits, for all sufficiently small a:*

(i) asymptotic freedom: the running coupling $g(a) \to 0$ as $a \to 0$;

(ii) confinement: the static potential between fundamental-representation sources grows linearly with separation at distances larger than a fixed scale r_0;

(iii) a positive mass gap $\Delta_a > 0$ uniformly bounded below.

Then, under Axiom 1, no probability measure μ on $\mathcal{A}/\mathcal{G}$ satisfies the Osterwalder–Schrader axioms while reproducing the conditional probability structure on $\mathcal{C}$ associated to the lattice expectation values in the limit.

Proof. Suppose, for contradiction, that such a measure μ exists. By the Osterwalder–Schrader reconstruction theorem [2], μ defines a Wightman quantum field theory on $\mathbb{R}^4$ with continuous Poincaré symmetry, and its n-point correlation functions $W_n(x_1, \ldots, x_n)$ are tempered distributions on $\mathbb{R}^{4n}$ [1].

By assumption (iii), the spectrum of the theory has a positive mass gap, so the correlation functions are non-trivial: there exist test functions $f_1, \ldots, f_n \in \mathcal{S}(\mathbb{R}^4)$ such that

$$W_n(f_1 \otimes \cdots \otimes f_n) \neq 0. \tag{5}$$

By Axiom 1, the physical content of any process is given by a configuration in $\mathcal{C}$, hence supported on the countable set $\mathcal{E}$. The correlation functions W_n, viewed as physical content, must therefore be supported on $\mathcal{E}^n \subset \mathbb{R}^{4n}$, a countable subset of $\mathbb{R}^{4n}$ which has Lebesgue measure zero.

By the Wightman axioms, W_n is invariant under the continuous translation action of the Poincaré group, in particular under the action of $\mathbb{R}^{4n}$ by simultaneous translation of all arguments.

We now apply the following elementary result. Let $T \in \mathcal{S}'(\mathbb{R}^d)$ be a tempered distribution invariant under the continuous translation action of $\mathbb{R}^d$, i.e. $\tau_a T = T$ for every $a \in \mathbb{R}^d$, where $(\tau_a T)(\varphi) = T(\varphi(\cdot + a))$. Then the support $\mathrm{supp}(T)$ is invariant under continuous translations: if $x \in \mathrm{supp}(T)$, then $x + a \in \mathrm{supp}(T)$ for every $a \in \mathbb{R}^d$. The only subsets of $\mathbb{R}^d$ invariant under continuous translations are $\emptyset$ and $\mathbb{R}^d$ itself. Therefore either $\mathrm{supp}(T) = \emptyset$ (i.e. $T = 0$) or $\mathrm{supp}(T) = \mathbb{R}^d$. No non-zero translation-invariant tempered distribution can have support contained in a countable subset of $\mathbb{R}^d$. (This is a standard consequence of the support properties of distributions; see [21, Ch. 2] for the general framework.)

It follows that $W_n \equiv 0$, contradicting the non-triviality established above.

The contradiction establishes that no measure μ on $\mathcal{A}/\mathcal{G}$ can simultaneously satisfy the Osterwalder–Schrader axioms (which require continuous Poincaré invariance and tempered correlation functions) and reproduce the discrete-support physical content required by Axiom 1. □ □

The contradiction in Theorem 7.1 arises from the incompatibility of two requirements: continuous translational invariance (from Wightman) and discrete support (from Endpoint-Only). The lattice theory $\mathcal{L}_a(G)$ avoids the contradiction by relaxing continuous translational invariance to discrete lattice-translational invariance; correlation functions on the lattice are supported on the discrete set of lattice sites and are invariant under discrete translations, consistent with both axioms restricted to the lattice scale. The continuum limit $a \to 0$ does not survive the contradiction because no measure can be invariant under the continuous limit while remaining supported on a discrete realization of $\mathcal{E}$.

7.3 Why the continuum integral nevertheless yields correct predictions

That the continuum functional integral nevertheless reproduces correct physical predictions in regimes where it converges follows from the fact that the integral preserves the endpoint structure even when its summands have no ontological referent. The boundary conditions at sources and sinks — which are the ontologically real components of any physical process — are correctly imposed, and the summation procedure preserves these under reduction to observable amplitudes.

This is structurally identical to the success of Ptolemaic epicycles in predicting planetary positions: the predictions are correct because the empirical observations (planet seen here at time t_1, planet seen there at time t_2) are correctly preserved by the calculation, not because the epicycles exist. The lattice $\mathcal{L}_a(G)$ is the actual theory; the continuum limit is a singular extrapolation in which the configurations carrying the physical content cease to be representable as smooth connections.

8 The Vacuum Has No Energy: The Cosmological Constant

We address the standard objection that any framework involving a "rich" ontology of the vacuum implies a large vacuum energy in conflict with cosmological observation [15].

In the present framework, no such implication arises. Under Axiom 1, the vacuum is not an entity. Where there are no interaction events, there is no physical content — not "empty space populated by zero-point fluctuations," but the absence of any spatial relation. The notion of a "vacuum energy density per unit volume" presupposes a unit volume of pre-existing empty space populated by quantized field oscillators, which the axiom denies.

The 120-orders-of-magnitude discrepancy between the QFT-predicted and cosmologically observed cosmological constant is, on this reading, not a problem to be solved but evidence against the ontological premise that produced the prediction. The discrepancy is consistent with the absence of any vacuum substrate of the kind continuous QFT postulates.

9 Falsifiable Predictions

The framework yields the following specific, testable predictions.

1. **Lattice–continuum gap.** Lattice QCD calculations of confinement-sensitive observables (string tension, glueball spectrum, deconfinement temperature) should exhibit residual deviations from any candidate continuum-limit construction in regimes where lock-scale effects are non-negligible.

2. **Deconfinement temperature from lock energy.** Independent estimation of the magnetic lock energy $\varepsilon_{\text{lock}}$ should yield $k_B T_c$ within a factor of order unity of the measured deconfine-

ment temperature ($T_c \approx$ 150–170 MeV). Discrepancy exceeding 20% would falsify the lock identification.

3. **Mass gap quantization.** The lightest hadronic state should have rest energy approximately equal to $3\,\varepsilon_{\text{lock}}$, with $\varepsilon_{\text{lock}}$ independently estimable. Significant deviation falsifies Theorem 6.1.

4. **Three-dimensional magnetic asymmetry across scales.** The 7:1 asymmetry between attraction and repulsion observed in macroscopic 3D magnet experiments [24] should have a hadronic-scale analogue in scattering channel cross-sections. Detection of a sustained repulsive bound channel would falsify the universality claim.

5. **Absence of a free gluon signal.** No experimental configuration should produce a signal interpretable as a free gluon, regardless of energy or kinematic regime. This is an existing empirical fact, recast as a positive prediction.

6. **Hadronic mass scaling with lock count.** The rest mass spectrum of hadronic species should scale predictably with the number of locks in the dominant configuration of each species.

7. **Rotational signatures in colour observables.** High-precision measurements of polarization and angular correlations in deep inelastic scattering should reveal residual rotational signatures of the colour structure consistent with the geometric identification of Proposition 4.6. Their absence across all kinematic regimes would falsify the geometric identification.

8. **Cosmological constant pattern.** The effective cosmological constant should exhibit scale-dependent or interaction-density-dependent variation rather than strict spatial uniformity. Confirmation of strict uniformity at all cosmological scales would falsify the no-vacuum interpretation.

10 Discussion

10.1 Spin agency plays no role at the mass-gap scale

Some recent ontological proposals [25] include a residual non-deterministic component in individual binary interactions — a small irreducible component of choice in spin outcomes that becomes relevant in low-multiplicity quantum experiments. This component does not enter the present derivation. The mass gap is a property of the bulk lock structure of hadrons, where the number of contributing photonic interactions is large and any residual stochastic component averages to zero in the rest mass. The hadronic mass-gap problem is therefore tractable using only the discrete-locking ontology and the energy–mass equivalence; it does not require invoking the more contentious aspects of the framework.

10.2 Universality of the locking principle across scales

A test of any proposed fundamental mechanism is whether it operates consistently at scales other than the one for which it was introduced. The locking mechanism described here exhibits scale invariance across hadronic, atomic, molecular, and biological domains. At the atomic scale, the same rotational locking mechanism appears in electron pairing within atomic orbitals; the Pauli exclusion principle, conventionally formulated as a constraint on antisymmetric wavefunctions, is on this reading the geometric statement that the lowest-energy configuration of two electrons is the locked configuration with anti-aligned spins coupling rotationally. At the molecular scale, covalent bonding is electron-pair locking shared between atomic centres. At the biological scale, photosynthesis is the catalyzed conversion of photonic content into stable molecular locks via copper-bearing electron-transport proteins. We offer this universality as a coherence check on the proposed mechanism rather than as a proof.

11 Conclusion

The Yang–Mills mass gap problem, as posed by Clay, asks for a proof within a framework whose central assumption — that four-dimensional Yang–Mills theory exists as a continuous field theory — is the source of the difficulty. We have presented:

- An axiomatic ontology (Axiom 1) in which physical content is exhausted by interaction events and conditional probability distributions, with no intermediate trajectories.
- A causal mechanism for hadronic rest mass: three-dimensional magnetic locking of discrete photonic packets, with rest mass identified as lock energy via $E = mc^2$.
- A geometric identification of the three colour charges of QCD with the three orthogonal rotation axes of the minimal three-lock bound state, and a derivation of SU(3) as the unitary symmetry of this configuration combined with the fermion antisymmetry constraint (Proposition 4.6).
- A derivation of linear confinement as the energetic structure of a one-dimensional lock chain (Lemma 5.1), with string breaking as $E = mc^2$ thresholding.
- A geometric proof of the mass gap (Theorem 6.1): $\Delta = 3\,\varepsilon_{\text{lock}} > 0$.
- A non-existence theorem for the continuum limit (Theorem 7.1), establishing via a structure result on tempered distributions that no probability measure on the smooth-connection space satisfies the Osterwalder–Schrader axioms while reproducing the endpoint-supported physical content.
- A clear distinction between the Higgs mechanism (parameterizing elementary-fermion mass) and the lock mechanism (binding-energy mass, dominant in hadrons).
- A resolution of the cosmological constant problem as artifact of the continuous-vacuum premise.
- Eight falsifiable predictions tying the framework to measurable lattice and experimental quantities.

The construction is a resolution of the Yang–Mills mass gap problem in the substantive sense: it identifies why the problem has resisted solution within the continuous formalism, provides a discrete ontology in which the mass gap exists trivially, and articulates a non-existence theorem for the object whose construction the Clay problem demands. It is not a construction within the Wightman or Osterwalder–Schrader axioms, because no such construction exists or can exist: the axioms presuppose the ontological reality of continuous configurations that, on the present analysis, have no referent.

The mass gap is real. The Yang–Mills theory of the strong interaction is real. The continuous field theory in $\mathbb{R}^4$ that the Clay problem asks us to construct is not real. The discrete locking structure that produces the strong interaction is real, and within it the mass gap is a geometric corollary of three-dimensional rotational closure.

The problem is, in this sense, dissolved into its solution.

Appendix A: Dictionary between Continuous and Discrete Descriptions

For the convenience of readers familiar with the continuous formulation, we provide a translation table.

Continuous formulation	Discrete locking formulation
Gauge field $A_\mu(x)$	Conditional probability structure on lock configurations
Functional integral $\int \mathcal{D}A\, e^{-S}$	Sum over lock configurations consistent with endpoint events
Renormalization	Adjustment of effective $\varepsilon_{\mathrm{lock}}$ with observation scale
Asymptotic freedom	High lock density limit; continuum approximation valid
Confinement	Lock chain extension cost exceeds pair-creation threshold
Colour charge r, g, b	Three orthogonal rotation axes
Gluons	Axis-exchange operations on lock
SU(3)	Unitary symmetry of three-axis antisymmetric fermion configuration
Mass gap Δ	$3\,\varepsilon_{\mathrm{lock}}$ (Theorem 6.1)

Appendix B: Open Problems

The following problems are flagged as open, in the sense that their resolution is consistent with the framework but not derived in closed form here.

1. **Closed-form derivation of $\mathfrak{su}(3)$ structure constants.** Proposition 4.6 provides a sketch of the emergence of SU(3) from the antisymmetry constraint applied to three rotational axes. A closed-form derivation of the structure constants f^{abc} from the geometry of axis-exchange operations on the antisymmetric three-fermion state is left as future work. Resolution of this problem would convert Proposition 4.6 from an identification with proof sketch into a fully formal theorem.

2. **Quantitative estimate of $\varepsilon_{\mathrm{lock}}$ from first principles.** The lock energy $\varepsilon_{\mathrm{lock}}$ enters Theorem 6.1, Lemma 5.1, and the deconfinement temperature relation. A first-principles estimate from the magnetic structure of paired photonic packets, comparable to the precision of lattice extractions, would convert the framework's predictions from order-of-magnitude to quantitative.

3. **Lattice realization of the Endpoint-Only Axiom.** A lattice formulation in which the ontology is built into the construction (rather than recovered as a regularization of a continuous

theory) would provide a constructive existence proof complementing Theorem 7.1. This is a programmatic problem rather than a single calculation.

Acknowledgements

This work benefited from sustained adversarial-constructive engagement with several large language models acting as reviewers during the development of the framework, including Anthropic's Claude, Google DeepMind's Gemini, and DeepSeek. Their contributions to the formalization of the non-existence theorem, the articulation of the Higgs–lock distinction, and the geometric interpretation of SU(3) are gratefully acknowledged.

References

[1] A. S. Wightman, *Quantum field theory in terms of vacuum expectation values*, Physical Review **101**, 860–866 (1956).

[2] K. Osterwalder and R. Schrader, *Axioms for Euclidean Green's functions*, Communications in Mathematical Physics **31**, 83–112 (1973).

[3] A. Jaffe and E. Witten, *Quantum Yang–Mills Theory*, Clay Mathematics Institute Millennium Problem Description, 2000.

[4] J. Glimm and A. Jaffe, *Quantum Physics: A Functional Integral Point of View*, 2nd ed., Springer-Verlag, 1987.

[5] T. Bałaban, *Large field renormalization. II: Localization, exponentiation, and bounds for the R operation*, Communications in Mathematical Physics **122**, 355–392 (1989).

[6] V. N. Gribov, *Quantization of non-Abelian gauge theories*, Nuclear Physics B **139**, 1–19 (1978).

[7] I. M. Singer, *Some remarks on the Gribov ambiguity*, Communications in Mathematical Physics **60**, 7–12 (1978).

[8] M. Aizenman, *Proof of the triviality of ϕ^4_d field theory and some mean-field features of Ising models for $d > 4$*, Physical Review Letters **47**, 1–4 (1981).

[9] J. Fröhlich, *On the triviality of $\lambda\phi^4_d$ theories and the approach to the critical point in $d \geq 4$ dimensions*, Nuclear Physics B **200**, 281–296 (1982).

[10] K. G. Wilson, *Confinement of quarks*, Physical Review D **10**, 2445–2459 (1974).

[11] A. Einstein, *Ist die Trägheit eines Körpers von seinem Energieinhalt abhängig?*, Annalen der Physik **323**, 639–641 (1905).

[12] P. W. Higgs, *Broken symmetries and the masses of gauge bosons*, Physical Review Letters **13**, 508–509 (1964).

[13] Y. Aoki et al., *The QCD transition temperature: results with physical masses in the continuum limit II*, Journal of High Energy Physics **0906**, 088 (2009).

[14] R. L. Workman et al. (Particle Data Group), *Review of Particle Physics*, Progress of Theoretical and Experimental Physics **2024**, 083C01.

[15] S. Weinberg, *The cosmological constant problem*, Reviews of Modern Physics **61**, 1–23 (1989).

[16] L. Bombelli, J. Lee, D. Meyer, R. D. Sorkin, *Space-time as a causal set*, Physical Review Letters **59**, 521–524 (1987).

[17] A. Ashtekar, *New variables for classical and quantum gravity*, Physical Review Letters **57**, 2244–2247 (1986).

[18] G. 't Hooft, *The Cellular Automaton Interpretation of Quantum Mechanics*, Springer, 2016.

[19] R. E. Kastner, *The Transactional Interpretation of Quantum Mechanics: The Reality of Possibility*, Cambridge University Press, 2012.

[20] J. J. Sakurai and J. Napolitano, *Modern Quantum Mechanics*, 2nd ed., Cambridge University Press, 2017.

[21] L. Hörmander, *The Analysis of Linear Partial Differential Operators I: Distribution Theory and Fourier Analysis*, 2nd ed., Springer, 2003.

[22] D. Griffiths, *Introduction to Elementary Particles*, 2nd ed., Wiley-VCH, 2008.

[23] F. Halzen and A. D. Martin, *Quarks and Leptons: An Introductory Course in Modern Particle Physics*, Wiley, 1984.

[24] A. Guzzon, *Three-Dimensional Magnetic Asymmetry and Thermal Gradient Forces*, Zenodo, 2026. DOI: 10.5281/zenodo.19162657

[25] A. Guzzon, *The Unification: A Discrete Causal Framework for Fundamental Physics*, Zenodo, 2026. DOI: 10.5281/zenodo.19066265