The Vacuum Has Steps

A constructive, geometric route to the Yang–Mills mass gap, told from the inside

Emmanouil Karolos Čižek

13 December 2025

Abstract

A monograph-length programme proposes a fully rigorous construction of four-dimensional Yang–Mills theory and a quantitative mass gap, using a blend of Euclidean probability, multiscale renormalisation, and Einstein–Cartan torsion geometry. This article explains the architecture and motivation in plain language, highlighting the physical meaning of each step and the role each mathematical tool plays in the proof strategy.[1;2]

When you work long enough on a famous open problem, you learn two uncomfortable lessons. First: most people's mental picture of the problem is not the problem. Second: the moment you introduce any tool that smells unfamiliar, somebody will assume you are trying to sell them a miracle.[2]

The Clay Yang–Mills mass gap problem sits right at that crossroads. It sounds like a physics question, but it is actually a demand for a particular kind of mathematical honesty. It asks you to take one of the most important interacting quantum field theories in nature and do what we routinely do for Brownian motion: **construct it cleanly, prove it exists, and prove a robust, quantitative property of its spectrum**.[2]

In my monograph I pursue a constructive route to that goal.[1] The architecture is classical in spirit: build a Euclidean measure with regulators, establish Osterwalder–Schrader (OS) axioms including reflection positivity, remove regulators in a controlled way, reconstruct the physical Hilbert space and Hamiltonian, then prove a positive spectral gap.[4;5] What is atypical is the language in which I organise the gauge degrees of freedom: an Einstein–Cartan setting in which a **torsion one-form** serves as the primary gauge-like field, supported by a flow (the Einstein–Cartan Ricci–torsion flow, or ECRT flow) that provides geometric structure and stability bookkeeping.[21–23]

Because torsion has not been detected experimentally at QCD scales, the natural suspicion is that this is theatre: an attempt to hitch a ride on geometric celebrity. I understand that suspicion. So I want to explain the programme the way a careful magazine feature would, but in a tone that does not ask for faith. I will describe what each tool is doing, what physical intuition it corresponds to, and where the approach lives or dies: on uniform estimates, not on grand vocabulary.

The mass gap is not a buzzword

Physicists often say "mass gap" the way we say "friction": we mean a real thing, but we rarely pause to define it with enough precision that a proof can grab hold of it.[2]

Here is the version that matters.

A quantum field theory has a Hamiltonian H whose spectrum begins at the vacuum energy. A **positive mass gap** means there is a strict jump: the vacuum sits alone, and the next energy level is at least $m > 0$ above it. In practice, this shows up in correlation functions: if you measure a gauge-invariant observable here and another far away, their correlation decays like $e^{-m\,\text{distance}}$.[2;4;5] No long-range whispering. The vacuum is stiff.

This is *not* the same question as "why is the proton heavy?" The proton's mass is mostly interaction energy in the strong field, yes, but the Clay problem is posed for **pure** Yang–Mills theory, without quarks.[2] It is about the gauge field itself generating an intrinsic energy scale even though the classical gauge bosons are massless.

The "particles" associated to the gap in this setting are not free gluons (which are not in the physical spectrum) but **gauge-invariant excitations**. In a pure gauge theory those are often discussed as glueball-like states. The key is: the gap is a property of the **physical, gauge-invariant sector**.[2]

That phrase, "physical sector", is where much of the difficulty hides.

Why the proof begins in imaginary time

If you want a proof you can take to court, you do not begin with the Minkowski path integral weight e^{iS}. That object oscillates; it is not a probability measure; it is not even clearly defined without extra structure. You can do brilliant physics with it, but the Clay problem is not grading brilliance. It is grading definition.[2]

So we Wick rotate to Euclidean signature, where the weight looks like e^{-S_E} and becomes eligible to define a measure.[4;5] Once you have a measure, you can ask the sort of questions analysts like: are moments finite? do correlations decay? can you take limits?

But Euclideanisation comes with a non-negotiable constraint. The Euclidean object must not be an arbitrary statistical model. It must be one that can be **reconstructed** into a real-time quantum theory with positive norms and a positive-energy Hamiltonian. That is the content of the **Osterwalder–Schrader framework**.[4;5]

At the centre of OS theory is a condition that sounds abstract and is absolutely not: **reflection positivity**. It is a statement that if you reflect an observable across a time-slice and pair it with its original, the result is nonnegative.[4;5] This is the Euclidean shadow of the positivity of inner products in Hilbert space. Without it, you can build a "theory" that looks fine in Euclidean variables and then reconstruct something unphysical with negative-norm states.[6]

In other words: Euclidean measures are powerful, but only certain Euclidean measures correspond to physical quantum theories. The OS axioms are the filter.[4;5]

A minimal schematic of this constructive pipeline is shown in Fig. 1.

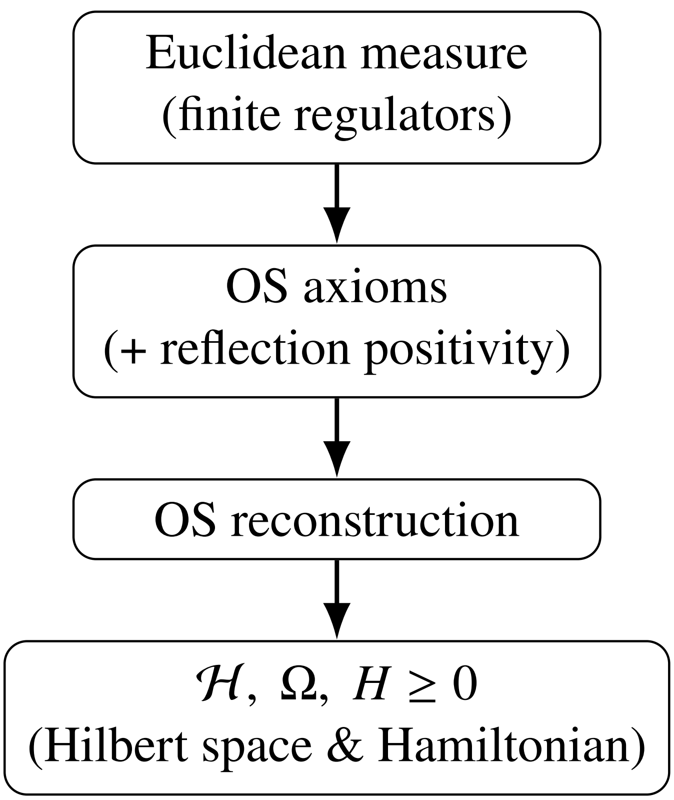

Figure 1: A minimal schematic of the constructive pipeline: Euclidean measure → OS axioms → OS reconstruction → Hilbert space and Hamiltonian.

Slabs, because time is where the physics is

Most people picture Euclidean field theory built on all of $\mathbb{R}^4$ at once. Constructively, that is a good way to lose control and never get it back. I organise the construction on **Euclidean slabs**: finite thickness in Euclidean time, and finite (then growing) spatial volume.[4;5]

The slab is a deceptively useful object. It does three jobs at once.

First, it makes reflection positivity natural: a slab has a canonical reflection across its mid-plane. You can build "mirror" constructions around that plane rather than inventing positivity out of thin air.[4;5]

Second, slabs can be **concatenated**. Glue two slabs of thickness t end to end and you get a slab of thickness $2t$. That gluing operation is not cosmetic: it is the Euclidean precursor of time evolution. It is how the **transfer operator** appears, the object whose logarithm becomes the Hamiltonian after reconstruction.[4;5]

This is sketched in Fig. 2.

Third, slabs allow a clean separation of bulk and boundary. That separation becomes crucial later when the spectral gap is proved not by assuming compactness or spectral input, but by studying **mixing** of a boundary dynamics.[13–15]

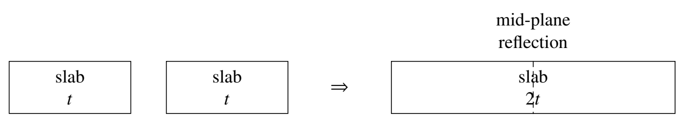

Figure 2: Slab concatenation: two thickness-t slabs glue to a thickness-$2t$ slab, compatible with reflection about the mid-plane and underpinning the transfer-operator (semigroup) picture.

Torsion from the first page, not the epilogue

Here is where my approach diverges in style from many standard presentations.

In a conventional gauge theory formulation, you take a connection A valued in the Lie algebra of $SU(N)$, discretise it as group elements on lattice edges, build actions from plaquette holonomies, and so on.

In my monograph I work in a Cartan geometric setting in which a **torsion one-form** τ is used as the primary gauge-like variable.[1;21] This is not a late reinterpretation. It is how the theory is written from early on, including how holonomy observables are encoded.

What does that mean physically?

A gauge field is, in its most primitive sense, a rule for how internal "colour" rotates as you move through space. Its measurable content, in a non-Abelian theory, is encoded in **holonomy**:

parallel transport around loops. That is what Wilson loops probe.[7]

Torsion, in Cartan geometry, measures a "twist" in how frames are transported. It is also a one-form. It can be treated as part of a connection-like structure.[21] The framework I use is arranged so that the holonomy of τ reproduces the Wilson loop content one normally associates to A.[1]

Now, the obvious objection is also the one I hear most often:

> "Torsion has never been observed in experiments. Doesn't that make this unphysical?"

Only if I were claiming torsion is a new observable field at QCD scales. I am not. The torsion sector is introduced as a **structural device** that improves the constructive control problem, and then a decoupling theorem is used to prove that it drops out of gauge-invariant correlators. In the language of BRST, torsion fields form a **BRST doublet**: they live in the cohomologically trivial sector and cannot affect physical observables.[17;20]

So torsion is not being smuggled in as extra physics. It is scaffolding that is designed to disappear from the physical sector, but is useful in the construction.

Whether you accept that depends on whether the decoupling is proved, and whether torsion genuinely helps control the estimates that matter. Those are legitimate technical questions. I do not ask readers to accept torsion because it sounds profound. I ask them to inspect what it is doing.

Building the Euclidean measure: making "the path integral" real

At finite cutoff, the first goal is simply: define the theory as an honest measure. Not a slogan.

In the slab setting, and at finite regulators, the interacting Euclidean measure is built out of three conceptual ingredients:[1;4;5]

1) A regularised Gaussian free part. The free part is the measure you would have if the theory were linear. It is Gaussian. But raw Gaussian measures in field theory are typically singular at short distances, so we regularise. One disciplined regularisation uses the **heat kernel**, which naturally smooths short-distance behaviour and respects symmetry. In group-valued contexts, heat kernels are canonical objects. Physically: you are preventing ultraviolet fluctuations from becoming infinite without breaking the symmetries you need later.

2) An interacting Boltzmann factor. This is where the Yang–Mills interaction content lives. It is the part that makes the theory nontrivial and the part that tries hardest to destroy your nice bounds.

3) A mirror coupling enforcing reflection positivity. Rather than gauge fixing and hoping positivity survives, the construction uses a mirror measure architecture: fields on the "positive time" side are coupled to reflected fields on the "negative time" side in a way that makes reflection positivity provable. Conceptually: you are building the Euclidean theory so that it can be rotated back to real time without producing negative probabilities or negative norms.[4;5]

At this stage, everything is finite. The theory is a fully defined probability law on a finite-dimensional configuration space (lattice, finite volume, finite cutoff). The work is to prove uniform bounds that survive when the regulators are removed.

OS axioms at finite cutoff: the first non-trivial checkpoint

The monograph's first major checkpoint is that for each fixed cutoff λ the lattice-level theory satisfies the OS axioms (OS0–OS5 in the numbering used there).[1;4;5] In plain language, this says:

- the measure exists and moments are finite;
- symmetries are correct;
- reflection positivity holds;
- regularity/locality properties needed for reconstruction are satisfied;
- a transfer-operator picture is compatible with slab concatenation.

Proving this is not glamorous. It is the constructive equivalent of checking that your aircraft wings are attached before trying to cross the Atlantic.

Three tools are central here:

Mirror symmetry arguments (to build positivity into the measure).
Chessboard estimates (to amplify local inequalities into global control by repeated reflections and tilings).
Moment bounds (to keep the interaction from producing heavy tails or singular spikes).[4;5]

Chessboard estimates deserve a word because they sound like a metaphor and they are not. One partitions space into blocks and uses reflection positivity iteratively to compare the weight of complicated configurations to products of weights of simpler, reflected configurations. The "chessboard" name is not poetic: you literally reflect and replicate blocks like a chessboard pattern.[4;5]

The renormalisation group, but with a leash: finite-range multiscale control

At this point you have a well-defined lattice theory. You do *not* yet have the continuum theory. To get there you must remove the regulators. That is the hardest part in four dimensions.

The renormalisation group is the natural framework for passing from microscopic definitions to macroscopic limits. In most physics usage it is an organising principle. In constructive work it must become an **algorithm with bounds**. [11;12]

I use a Balaban-type **finite-range renormalisation group**. [11;12] The guiding idea is to decompose the covariance into a sum of scale slices, each with finite spatial range. Schematically,

$$C = \sum_{k=0}^{K} C_k,$$

where each C_k only couples fields within a controlled distance at that scale. This finite-range property is not a technical flourish. It prevents correlations from becoming "everywhere at once" at a single integration step, which is exactly what destroys convergence in nonperturbative settings.

Then one performs **inductive block-spin integration**: integrate out the shortest-range fluctuation layer, obtain an effective theory for the remaining modes, repeat scale by scale.

Two kinds of estimates become essential:

Determinant/Gram bounds (Gram–Hadamard-type controls). When you integrate Gaussian degrees of freedom, determinants appear, and determinants are where combinatorial growth hides. Gram–Hadamard bounds turn determinants into products of norms, providing a ceiling that does not explode factorially.

Polymer/cluster expansion convergence (Kotecký–Preiss-type criteria). Interacting partition functions can be reorganised as sums over connected clusters ("polymers"). The Kotecký–Preiss criterion gives a quantitative condition under which the expansion converges absolutely. [10] Physically: it ensures that clusters do not percolate uncontrollably and that the interaction remains tame enough, at each scale, to support a convergent description. [10]

A further demand in the monograph is that these constants be controlled **uniformly in** $SU(N)$, not just for a fixed N. This matters because the Clay statement is formulated for general compact simple gauge groups. [2] If your constants degrade badly as N grows, your "proof" may be a proof only for a particular case.

The continuum object μ_∞: what it is and what it isn't

After multiscale control, the programme takes the double limit: volume $L \to \infty$ and cutoff $\lambda \to \infty$, alongside the refinement of lattice spacing. The target is a continuum Euclidean measure

$$\mu_\infty$$

defined on **connections modulo gauge transformations**. [1]

The phrase "modulo gauge transformations" is not pedantry. Gauge theory contains redundancy. Two configurations related by a gauge transformation represent the same physics. Any honest construction must ensure that physical observables do not depend on arbitrary gauge choices. [2]

In the monograph's framing, μ_∞ is required to be:

- gauge invariant;
- reflection positive on the physical observable algebra;
- compatible with slab concatenation (so a transfer semigroup exists).

Mathematically, the limiting process can be organised as a **projective system** of finite-scale measures whose consistency under coarse-graining and restriction allows one to define a limit. [1] Conceptually: you are building an infinite object as the coherent limit of finite ones, rather than writing down an infinite expression and hoping it behaves.

OS reconstruction: where probability becomes quantum mechanics

Once μ_∞ is in place with OS axioms, OS reconstruction yields the familiar physical structures: [4;5]

- a Hilbert space $\mathcal{H}$;
- a vacuum vector Ω (cyclic for the observable algebra);
- a self-adjoint Hamiltonian $H \geq 0$;
- local fields and, via analytic continuation, Minkowski-space Wightman fields.

This is the point where a Euclidean measure stops being "a nice probability distribution" and becomes "a quantum field theory". [4–6] It is also where the problem becomes a spectral problem: does H have a gap? [2]

A mass gap from mixing: Harris theory on boundary dynamics

Many constructions rely on compactness or spectral assumptions at finite cutoff that effectively smuggle in what you want to prove. I try to avoid that by deriving exponential clustering and a first mass gap via **Harris mixing** for boundary Langevin dynamics on the slab. [13–15]

Here is the intuition.

You can define a Markov process on boundary configurations whose invariant measure is the boundary law induced by the bulk. [1] If that Markov process mixes rapidly (forgets its initial state exponentially fast), then boundary correlations decay exponentially, and the OS transfer semigroup inherits a spectral gap. [4;5;13;14]

Harris theory gives a robust way to prove exponential mixing without assuming compactness: [13–15]

- A **Lyapunov drift** condition prevents escape to infinity: there is a function V such that expected V decreases under the dynamics outside a controlled region.[14]
- A **minorisation/small-set** condition ensures there is enough randomness to mix distributions on some set.[14]

Contraction is measured using a **Kantorovich/Wasserstein distance** between probability measures.[16] Uniform contraction implies exponential mixing, which implies exponential correlation decay, which feeds into OS clustering conditions, and then into a spectral gap for the Hamiltonian.[4;5;13–15]

This route is attractive because it replaces spectral assumptions with explicit dynamical estimates. It is also not decorative: if you cannot prove uniform contraction, you do not get to claim a mass gap.[2]

Confinement as an inequality: Wilson loops, corridors, and string tension

The other pillar is the Wilson loop area law.[7]

A Wilson loop is the gauge-invariant observable obtained by parallel transport around a closed curve and taking the trace.[7] Physically it probes the energy cost of separating colour sources. An area law,

$$\langle W(C)\rangle \approx e^{-\sigma\,\mathrm{Area}(C)},$$

means the expectation decays exponentially with the minimal spanning area of the loop. The constant $\sigma > 0$ is the **string tension**. In physical language: the flux tube has tension; pulling sources apart costs energy proportional to distance.[7]

The monograph's route is adapted from statistical mechanics:[1;10–12]

- polymer and cluster expansions for plaquette variables;
- finite-range RG to keep the expansion controlled across scales;
- a "corridor" construction around minimal spanning surfaces, localising the dominant contributions to a neighbourhood of the surface and controlling fluctuations outside.

An important organisational refinement in later appendices is to remove dependence on a particular corridor implementation from the continuum limit.[1] This is part of a general theme: constructions used to prove an inequality should not become hidden assumptions baked into the final object.

From string tension to quantitative mass gap: the Glimm–Jaffe bridge

Once a strict area law is established with $\sigma > 0$, one can use a Glimm–Jaffe-type argument to derive a **quantitative** spectral gap bound.[9] In the framing used in the monograph, one obtains a lower bound of the form

$$m \;\geq\; \tfrac{1}{2}\,\sigma^{1/2}.$$

I do not ask readers to fixate on the factor $1/2$. The significance is that the mass gap is not merely existential; it is tethered to a physically meaningful confinement scale.[2]

This is the point where the story becomes satisfying in a physicist's sense: a Wilson loop inequality produces an energy gap inequality.[7;9] The vacuum's stiffness and the flux tube's tension are quantitatively linked.

BRST–BFV: isolating the physical Hilbert space constructively

Even after OS reconstruction, one still has to face the most annoying fact about gauge theory: the degrees of freedom we write are not the degrees of freedom we measure. A Hilbert space built naïvely from gauge potentials contains unphysical states.[20]

The monograph uses a non-perturbative BRST–BFV framework to define the physical Hilbert space:[17–19]

- Construct a BRST charge Q that is **nilpotent** ($Q^2 = 0$) and **closed** on the OS Hilbert space.[17;20]
- Identify physical states with BRST cohomology at ghost number 0:
$$\mathcal{H}_{\mathrm{phys}} \cong \ker Q/\operatorname{im} Q.$$
[17;20]
- Implement Slavnov–Taylor identities at the constructive level, not merely as perturbative formalities.[17]

A substantial technical burden lies in making this compatible with the already-constructed OS theory: closed-range arguments, Hilbert complex structure, and a "gap for Δ" analysis (where Δ is the BRST Laplacian-type operator).[1;20] This is not garnish. It is how you ensure the quotient is a Hilbert space rather than a symbolic expression.

The essential output is: the Hamiltonian gap lifts to the BRST-reduced physical space.[1;17–19] Otherwise, you might prove a gap in a space contaminated by gauge redundancy, which would not answer the Clay problem as stated.[2]

Torsion again: why a geometric layer is included if the constructive chain already closes

This is the question I cannot dodge, and I do not want to.

If the constructive OS–RG–Harris–Wilson–BRST chain already yields existence, area law, and gap, why add a chapter about torsion one-forms and ECRT flow?[1]

My answer is restrained: the geometric layer is not required for the *logical existence* of the constructive chain, but it serves two purposes that I believe are legitimate, provided they are kept honest.

First, it provides a **geometric encoding** of gauge information through τ, giving a unified language in which Wilson loops, transfer structure, and evolution can be discussed.[1;21]

Second, it provides a **structured stability narrative**: geometric flows come with monotonic quantities (entropy-like functionals) and non-collapse principles that can function as conceptual bookkeeping for uniformity under scaling.[22;23] In a constructive proof, "uniformity under limits" is the entire game. If a flow helps organise and motivate the stability mechanisms while remaining subordinate to explicit estimates, it can add coherence without adding circularity.

And then comes the critical safeguard: the torsion sector is shown to decouple from the gauge-invariant sector by a BRST doublet mechanism.[17;20] This is what prevents torsion from becoming an untestable new physics claim. In the framework, torsion is intrinsic to the construction but **not** intrinsic to the physical predictions.[1]

A sceptical reader should still ask: does torsion actually improve the estimates, or is it just a change of costume? That is a fair question. It can be answered only by reading the coercivity corrections, the RG appendices controlling the torsion quartic sector, and the decoupling argument itself.[1]

What is testable, and what is not

One final point matters for readers who think like experimentalists.

You cannot test "ECRT flow" directly. You test gauge theory by testing gauge-invariant correlators, spectra, and scaling behaviours.[2] In my framework, if torsion decouples, then torsion-specific signatures should not appear in ordinary low-energy observables. That is the point of decoupling.[1;17]

So what would count as empirical relevance?

- Consistency with known nonperturbative behaviour of $SU(N)$ Yang–Mills (for example, scaling relations, confinement indicators).[7]
- Agreement of gauge-invariant predictions with lattice computations in regimes where both are applicable.[7]
- Robustness of universal ratios and correlation decay structure.[2]

The geometric layer is not a new particle proposal. It is a structural method for building the theory. The experimental contact remains through the usual gauge-invariant quantities.

Why I'm cautious in my claims

Constructive field theory is unforgiving. A proof is not a slogan about "taking limits"; it is the ability to take limits while constants stay bounded and positivity survives.[2] A single non-uniform estimate can silently ruin the whole continuum claim.

So the most honest summary I can give is this:

I have tried to build a complete chain in which each link is explicit, regulator-controlled, and compatible with OS reconstruction and BRST reduction:[1]

- a finite-cutoff Euclidean measure with reflection positivity by design;[4;5]
- a finite-range RG that controls interactions across scales;[11;12]
- a continuum measure μ_∞ obtained as a coherent limit;[1]
- OS reconstruction producing a Hilbert space and Hamiltonian;[4;5]
- a mass gap derived via Harris mixing and/or via a string tension bound;[7;9;13–15]
- confinement as a strict area law, not an assumption;[7]
- BRST–BFV reduction producing the physical Hilbert space without gauge fixing;[17–20]
- torsion as a geometric variable present from the start, and decoupled at the end.[1;17;21]

If this works, then the torsion one-form is not a publicity stunt. It is a scaffold that helps hold the constructive building process together and then provably disappears from the physical façade.[1]

And if it does not work, the failure will not be philosophical. It will be somewhere very specific: a constant that fails to remain uniform, a convergence criterion that is not met, a positivity step that leaks, or a functional-analytic compatibility condition that collapses.[2] That is how constructive mathematics keeps you honest, whether you enjoy it or not.

In the end, the Clay problem is asking for a theory that has always been treated as "real" in physics to become real in the strictest sense: a defined object, a reconstructed Hilbert space, and a provable gap.[2;4;5] My aim in the monograph is to make that transition as transparent as possible: each tool justified by a job it must do, each job tied to either physical meaning (positivity, locality, decay) or to the hard necessities of taking limits in four dimensions.[1]

If it reads like a long walk rather than a clever shortcut, that is because it is. Four-dimensional interacting quantum gauge theory does not forgive shortcuts, and it certainly does not reward anyone for sounding impressive.

Outlook

If this programme is correct, the payoff is not a new phenomenological party trick, but something far less glamorous and far more valuable: a four-dimensional Yang–Mills theory that actually *exists* as a mathematical object, together with a positive spectral gap proved without importing the gap as a hidden assumption.[2] In other words, the vacuum is not a continuous ramp; it has steps. And the steps stay put as the regulators are removed, which is the part that usually breaks when one tries to do this seriously.[1]

The entire construction can be read as a long attempt to make two principles coexist without hypocrisy: *(i)* Euclidean methods should behave like probability theory (so we can estimate things), and *(ii)* the result should still reconstruct to a physical Hilbert space with a positive Hamiltonian (so we are not merely solving a different problem).[4;5] Reflection positivity is the strict parent in the room: it enforces the second principle while constantly threatening to veto the first. Slabs, mirror couplings, and the OS transfer semigroup are essentially the diplomatic treaty that keeps them from declaring war.[4;5]

On the spectral side, I have tried to avoid the classic circular move: assuming a gap, proving clustering, and then announcing that the gap exists.[2] The Harris-mixing route is appealing precisely because it starts from something auditable: a drift condition, a minorisation condition, and a contraction estimate in a Kantorovich (Wasserstein-type) distance.[13–16] If those estimates hold uniformly in the regulators, then exponential mixing forces exponential clustering, and exponential clustering forces a gap for the transfer semigroup.[4;5;13;14] It is not poetic, but it is honest: you can point at the constants and ask whether they are real or imaginary.

Similarly, the Wilson loop technology is not meant as a decorative reference to confinement folklore.[7] A strict area law with $\sigma > 0$ is an inequality with teeth. It says that large loops are exponentially suppressed by area, not by wishful thinking, and it supplies a quantitative scale.[7] The Glimm–Jaffe-style bridge from string tension to spectral gap (schematically $m \gtrsim \sigma^{1/2}$) is, at bottom, a reminder that in a local quantum theory the vacuum cannot simultaneously sustain a genuine flux-tube tension and remain arbitrarily easy to excite.[9] One can debate constants; one cannot debate the direction of the implication if the hypotheses truly hold.

Now to torsion, because nobody is going to let me get away without addressing it.[21] The torsion one-form τ is introduced early because it reorganises the control problem in a geometric language where stability and coercivity estimates can be stated cleanly and tracked across scales.[1;21] The ECRT flow is then a bookkeeping device for "uniformity under evolution" in the same spirit that finite-range RG is bookkeeping for "uniformity under coarse-graining".[11;12;22;23] If that sounds suspiciously like I am trying to make Ricci flow do my homework, I assure the reader that Ricci flow is not, in fact, available to sit exams in my place. The geometric layer should be judged by a simple standard: does it strengthen regulator-uniform bounds, and does it *provably* decouple from gauge-invariant observables?[17;20] The BRST-doublet decoupling theorem is meant to prevent torsion from becoming a new low-energy force claim. If decoupling is airtight, torsion is scaffolding, not an extra building.[1]

The most honest thing I can say about the remaining work is also the least exciting: everything depends on constants. Constructive field theory is not impressed by eloquence. It wants inequalities that survive limits, uniform determinant bounds that do not quietly deteriorate with $SU(N)$, convergence criteria that are verified rather than announced, and functional-analytic compatibility arguments (OS reconstruction, transfer operator, BRST reduction) that do not smuggle in the conclusion.[2;4;5;17–19] If you are looking for a dramatic ending, I regret to inform you that the true climax of this story is a stack of estimates that continue to hold when $a \to 0$, $L \to \infty$, and $\lambda \to \infty$. I realise this is not the sort of plot twist that sells movie tickets.

If the programme is wrong, it will not be wrong in an interesting metaphysical way. It will be wrong at a specific line where a bound is not uniform, a positivity argument leaks, a cluster expansion fails its convergence condition, or a closed-range argument refuses to close.[2;10] That is, oddly, comforting: it means the problem is not being decided by charisma. It is being decided by whether the vacuum really has steps, and whether the proof has the nerve to show where they come from.

References

[1] E. K. Čižek. *A Constructive Einstein–Cartan–Yang–Mills Theory with Positive Mass Gap in Four Dimensions* (version 1.1.0), Zenodo (2025). DOI: `https://doi.org/10.5281/zenodo.17246443`.

[2] A. Jaffe and E. Witten. *Quantum Yang–Mills Theory*. Clay Mathematics Institute, Millennium Prize Problems (problem statement and background). Available as PDF: `https://www.claymath.org/wp-content/uploads/2022/06/yangmills.pdf`.

[3] Clay Mathematics Institute. *Yang–Mills Existence and Mass Gap (Millennium Prize Problem statement).* `https://www.claymath.org/wp-content/uploads/2022/06/yangmills.pdf`.

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

[5] K. Osterwalder and R. Schrader. Axioms for Euclidean Green's functions II. *Communications in Mathematical Physics* **42** (1975), 281–305.

[6] R. F. Streater and A. S. Wightman. *PCT, Spin and Statistics, and All That*. Benjamin (1964).

[7] K. G. Wilson. Confinement of quarks. *Physical Review D* **10** (1974), 2445. DOI: `https://doi.org/10.1103/PhysRevD.10.2445`.

[8] Y. M. Makeenko and A. A. Migdal. Exact equation for the loop average in multicolor QCD. *Physics Letters B*

88 (1979), 135–137. DOI: `https://doi.org/10.1016/0370-2693(79)90131-X`.

[9] J. Glimm and A. Jaffe. *Quantum Physics: A Functional Integral Point of View*. Springer (1987).

[10] R. Kotecký and D. Preiss. Cluster expansion for abstract polymer models. *Communications in Mathematical Physics* **103** (1986), 491–498.

[11] T. Bałaban. Renormalization group approach to lattice gauge field theories I. Generation of effective actions in a small field approximation and a coupling constant renormalization in four dimensions. *Communications in Mathematical Physics* **109** (1987), 249–301.

[12] T. Balaban. Renormalization group approach to lattice gauge field theories II. Cluster expansions. *Communications in Mathematical Physics* **116** (1988), 1–22. DOI: `https://doi.org/10.1007/BF01239022`.

[13] T. E. Harris. The existence of stationary measures for certain Markov processes. In *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability, 1954–1955*, Vol. II, 113–124. University of California Press (1956).

[14] S. P. Meyn and R. L. Tweedie. *Markov Chains and Stochastic Stability*. Springer (2nd ed., 2009).

[15] M. Hairer. Yet another look at Harris' ergodic theorem for Markov chains. In *Seminar on Stochastic Analysis, Random Fields and Applications VI*. Birkhäuser (2011).

[16] C. Villani. *Optimal Transport: Old and New*. Springer (2009).

[17] C. Becchi, A. Rouet, and R. Stora. Renormalization of gauge theories. *Annals of Physics* **98** (1976), 287–321. DOI: `https://doi.org/10.1016/0003-4916(76)90156-1`.

[18] E. S. Fradkin and G. A. Vilkovisky. Quantization of relativistic systems with constraints. *Physics Letters B* **55** (1975), 224–226.

[19] I. A. Batalin and E. S. Fradkin. A generalized canonical formalism and quantization of reducible gauge theories. *Physics Letters B* **122** (1983), 157–164.

[20] M. Henneaux and C. Teitelboim. *Quantization of Gauge Systems*. Princeton University Press (1992).

[21] F. W. Hehl, P. von der Heyde, G. D. Kerlick, and J. M. Nester. General relativity with spin and torsion: Foundations and prospects. *Reviews of Modern Physics* **48** (1976), 393–416. DOI: `https://doi.org/10.1103/RevModPhys.48.393`.

[22] R. S. Hamilton. Three-manifolds with positive Ricci curvature. *Journal of Differential Geometry* **17** (1982), 255–306.

[23] G. Perelman. The entropy formula for the Ricci flow and its geometric applications. arXiv:math/0211159 (2002).