FORMALIZED MATHEMATICS AS A DEBUGGING TOOL: FALSIFYING AN INVARIANT AND DELETING A PHANTOM DIFFICULTY IN LATTICE GAUGE THEORY

JUAN CARLOS PAREDES

Abstract. We present a machine-checked formalization in Lean 4 of a problem in lattice gauge theory: the Plateau Conjecture, which asserts that a charged magnetic flux family on a discrete torus is separated from all flat gauge fields by an operator-norm gap uniform in the system size. During formalization of a natural 1D invariant—the column-winding sum—the proof assistant exposed a fatal flaw: the shear family, a gauge configuration with trivial horizontal edges, carries the same winding as the charged flux family yet lies arbitrarily close to a flat family in the large-system limit. This falsification forced a correction: the invariant must read the full 2D plaquette data. We construct the corrected invariant (total flux), prove it passes the falsification test, and then formalize the reduction of the conjecture itself: a machine-checked theorem derives the full m-uniform statement from a single transported rigidity estimate for the magnetic-translation pair, exposing along the way that a distance-translation step previous drafts treated as half the open problem is not load-bearing at all. The library's only remaining `sorry` is that one estimate, stated over concrete objects with its classical unitary specialization known. The methodological contribution is the use of formal methods not as an after-the-fact certification but as an active research tool that catches subtle errors—twice, in opposite directions: once falsifying a plausible invariant, once deleting a phantom difficulty—in plausible mathematical reasoning.

1. Introduction

1.1. **The physical problem.** In lattice gauge theory, a gauge field assigns a parallel-transport operator to each edge of a discretized space. On a two-dimensional torus, a uniform magnetic field with integer flux k is represented by the Landau-gauge connection: vertical edges carry an x-dependent phase, and a boundary twist on horizontal wrap edges ensures the total flux is $2\pi k$. A gauge field is *flat* if it comes from a globally defined frame field—parallel transport around any closed loop is trivial up to projection.

The Plateau Conjecture (so named for its analogy to Plateau's problem and the quantized plateaus of the quantum Hall effect) asserts:

> For every nonzero integer charge k, there exists a constant $c(k) > 0$ such that for all torus sizes $m \geq 3$ and all ranks $2 \leq r \leq d$, the charge-k flux family is at operator-norm distance at least $c(k)$ from every rank-r flat family.

Informally: topological charge cannot be flattened. The obstruction is uniform in the system size—it survives the continuum limit.

1.2. **The role of formalization.** The conjecture is stated over concrete objects: finite matrices, explicit edge families, elementary operator norms. This makes it amenable to formalization in a proof assistant. We formalized the entire problem in Lean 4: frames, edge families, the Landau-gauge flux family, flat families, operator-norm distance, and two candidate invariants for detecting the obstruction.

Date: July 3, 2026.

The formalization did more than verify—it *debugged*. While attempting to prove the conjecture via a natural 1D invariant (the column-winding sum), we discovered a counterexample: the **shear family**, which shares the flux family's vertical structure but has trivial horizontal edges. The shear family carries the same winding as the flux family yet lies arbitrarily close to a flat family. The invariant was wrong.

This paper tells that story.

1.3. Contributions.

(1) A complete formalization of the Plateau problem: definitions, flux family, flat families, distance-to-flat, and two invariants (winding sum and total flux).
(2) A machine-checked proof that the 1D winding sum fails: the shear family counterexample.
(3) A corrected 2D invariant (total flux) with proofs that it correctly distinguishes flux, flat, and shear.
(4) The magnetic-translation pair and a locality theorem: $\|[\mathcal{U}, \mathcal{V}]\|$ is bounded by the maximum plaquette discrepancy, never amplified by system size.
(5) A machine-checked reduction (`plateau_of_rigidity`): the full conjecture, uniformly in system size with the constant preserved, follows from one scoped estimate (`FluxPairRigidity`—dimension-free rigidity of the flux pair against every commuting contraction pair). The reduction's proof exposed that the converse distance-translation step, listed in earlier drafts as half the remaining problem, is not load-bearing—a second instance, in the opposite direction, of formalization correcting plausible reasoning.
(6) A precise statement of the one remaining open estimate—the library's only `sorry`—with its classical specialization identified (at full rank the pair's complexified core is the Voiculescu pair) and its direct adversarial numerics reported.

2. Formal setting

2.1. Frames and edge families.

A **frame** is a $d \times r$ matrix V with orthonormal columns: $V^\mathsf{T} V = I_r$. This is the condition that makes interior terms cancel in products.

An **edge family** on an $m \times m$ discrete torus assigns a $d \times d$ matrix $T(e)$ to each directed edge e. Edges come in two directions: horizontal ($\to$, stepping $+x$) and vertical ($\uparrow$, stepping $+y$). The torus wraps: coordinates live in $\mathbb{Z}_m$ and addition is modular.

A family is **flat** (of rank r) if there exists an assignment of frames $V(s)$ to sites such that

$$T(e) = V(\mathrm{dst}(e)) \cdot V(\mathrm{src}(e))^\mathsf{T}.$$

This is the discrete analog of a pure-gauge connection: parallel transport factors through a globally defined frame field.

2.2. The flux family.

The **charge-k flux family** in Landau gauge assigns to each edge

$$T((x,y),\uparrow) = S \cdot R(2\pi k x/m^2) \cdot S^\mathsf{T} \qquad \text{(vertical)},$$
$$T((x,y),\to) = S \cdot R(\theta_{\mathrm{horiz}}(x,y)) \cdot S^\mathsf{T} \qquad \text{(horizontal)},$$

where $S = \mathrm{stdFrame}(d,r)$ embeds $\mathbb{R}^r$ into $\mathbb{R}^d$, $R(\theta)$ rotates the first two coordinates of $\mathbb{R}^r$ by θ, and the horizontal angle is zero except at the wrap boundary $x = m-1$, where it carries $-2\pi k y/m$. This boundary term ensures the total flux through the torus is $2\pi k$.

2.3. **Distance to flatness.** The **operator norm** of a matrix is $\|A\| = \sup\{\|Av\| : \|v\| = 1\}$. The **distance to flatness** is

$$\mathrm{distOpFlat}(r, T) = \inf\{\varepsilon : \exists T' \text{ flat of rank } r,\ \forall e,\ \|T(e) - T'(e)\| \le \varepsilon\}.$$

The Plateau Conjecture states: $\mathrm{distOpFlat}(r, \mathrm{fluxFamily}(k)) \ge c(k) > 0$ for all $m \ge 3$.

2.4. **The collapse lemma.** The algebraic foundation is a single identity: for frames V_1, V_2, V_3 with $V_2^{\mathsf{T}} V_2 = I$,

$$(V_3 V_2^{\mathsf{T}})(V_2 V_1^{\mathsf{T}}) = V_3 V_1^{\mathsf{T}}.$$

This extends to chains: a product of flat edge maps along any path collapses to $V_{\mathrm{end}} \cdot V_{\mathrm{start}}^{\mathsf{T}}$. For closed cycles, the composite is $V_{\mathrm{start}} \cdot V_{\mathrm{start}}^{\mathsf{T}}$—an orthogonal projection, idempotent and symmetric.

3. The 1D winding sum and its falsification

3.1. **The winding sum.** The **vertical cycle** at column x is the product of all m vertical edges up that column:

$$\mathrm{vertCycle}(T, x) = T((x, m-1), \uparrow) \cdots T((x, 0), \uparrow).$$

For the flux family, this is $S \cdot R(2\pi kx/m) \cdot S^{\mathsf{T}}$—a rotation by an angle proportional to x.

The **core block** extracts the top-left 2×2 submatrix. The **core vector** of a 2×2 matrix is $(A_{00} + A_{11},\ A_{10} - A_{01}) = (\mathrm{trace}, \mathrm{skew})$. For a rotation $R(\theta)$, this is $(2\cos\theta, 2\sin\theta)$. The **cycle angle** is $\arg(\mathrm{trace} + i\,\mathrm{skew})$.

The **winding sum** is

$$\mathrm{windingSum}(T) = \frac{1}{2\pi} \sum_x \mathrm{angleDiff}\big(\mathrm{cycleAngle}(x),\ \mathrm{cycleAngle}(x+1)\big),$$

where $\mathrm{angleDiff}(a, b)$ is the principal-branch difference in $(-\pi, \pi]$.

3.2. **Anchors: flat $\to 0$, flux $\to k$.** We proved two theorems (machine-checked in Lean).

Theorem 1 (Flat anchor)**.** *If T is flat and satisfies a core gap condition (the cycle's core vector never vanishes), then* $\mathrm{windingSum}(T) = 0$.

Proof. A flat vertical cycle is $V(x, 0) \cdot V(x, 0)^{\mathsf{T}}$, whose core vector is $(\mathrm{trace}, 0)$ with $\mathrm{trace} > 0$ by the gap condition. Every cycle angle is zero, so the sum is zero. □

Theorem 2 (Flux anchor)**.** *For $m > 2|k|$,* $\mathrm{windingSum}(\mathrm{fluxFamily}(k)) = k$.

Proof. The cycle angle at column x is $2\pi kx/m$ (up to a branch integer). The condition $m > 2|k|$ ensures consecutive columns differ by $|2\pi k/m| < \pi$, so the principal-branch difference is exactly $2\pi k/m$ at each step. Summing over m columns gives $2\pi k$; dividing by 2π gives k. The wrap from $x = m-1$ to $x = 0$ is handled by the branch correction lemma. □

3.3. **The shear family.** The **shear family** takes the flux family's vertical edges but replaces all horizontal edges with the identity (the standard projection SS^{T}):

$$\mathrm{shearFamily}(e) = \begin{cases} SS^{\mathsf{T}} & \text{if } e \text{ is horizontal,} \\ \mathrm{fluxFamily}(e) & \text{if } e \text{ is vertical.} \end{cases}$$

This family is "morally flat"—one could gauge-transform the vertical phases away if not for the torus topology. But the winding sum, which only reads vertical edges, cannot distinguish it from the flux family.

3.4. **The falsification.** We proved three properties of the shear family.

Theorem 3 (Shear winds)**.** windingSum(shearFamily(k)) $= k$.

Proof. The winding sum depends only on vertical edges, which are identical to the flux family's. ☐

Theorem 4 (Shear is gapped)**.** *The shear family satisfies the same core gap condition as the flux family (gap ≥ 2).*

Theorem 5 (Shear is close to flat)**.** *For every edge e,* $\|\text{shearFamily}(e) - \text{constFlat}(e)\| \leq 2\pi|k|/m$*, where* constFlat *assigns SS^T to every edge.*

Proof. Horizontal edges are exactly equal (difference zero). Vertical edges differ by $S(R(\theta) - I)S^\mathsf{T}$ where $\theta = 2\pi kx/m^2$. The operator norm of this difference is bounded by $|\theta|$ (using the chord bound $\|R(\theta) - I\| \leq |\theta|$), and $|\theta| \leq 2\pi|k|/m$ since $x < m$. ☐

Corollary 1 (Falsification)**.** *The implication "different winding $\Rightarrow$ uniform separation" is false. The shear family has winding $k \neq 0$ and the constant flat family has winding 0, yet their edge-wise distance is at most $2\pi|k|/m$, which tends to 0 as $m \to \infty$.*

This killed the 1D reduction route to the Plateau Conjecture. The winding sum is a 1D shadow—it sees the phase progression along columns but misses the 2D plaquette structure.

4. The corrected invariant: total flux

4.1. **Plaquette holonomy and total flux.** The **plaquette holonomy** at site $v = (x, y)$ is the product of four edge maps around a unit square:

$$\text{plaqHol}(T, v) = T(\uparrow \text{ at } v)^\mathsf{T} \cdot T(\rightarrow \text{ at top})^\mathsf{T} \cdot T(\uparrow \text{ at right}) \cdot T(\rightarrow \text{ at } v).$$

The **plaquette angle** extracts the rotation angle of the core block. The **total flux** is the sum over all plaquettes:

$$\text{totalFlux}(T) = \sum_v \text{plaqAngle}(v).$$

4.2. **The three anchors.**

Theorem 6 (Flat $\to 0$)**.** *If T is flat with a plaquette gap, then* totalFlux$(T) = 0$.

Proof. A flat plaquette holonomy collapses (via four applications of the collapse lemma) to the projection $V(v)V(v)^\mathsf{T}$, whose core angle is zero. ☐

Theorem 7 (Flux $\to 2\pi k$)**.** *For $m > 2|k|$,* totalFlux(fluxFamily(k)) $= 2\pi k$.

Proof. Every interior plaquette contributes $2\pi k/m^2$. The wrap column and corner plaquettes require branch corrections, but the sum over all m^2 plaquettes telescopes to exactly $2\pi k$. ☐

Theorem 8 (Shear $\to 0$)**.** totalFlux(shearFamily(k)) $= 0$.

Proof. Interior plaquettes contribute $+2\pi k/m^2$, but the wrap-column plaquettes contribute $-2\pi k(m-1)/m^2$. The sum is $(m-1) \cdot m \cdot (2\pi k/m^2) + m \cdot (-2\pi k(m-1)/m^2) = 0$. ☐

This is the decisive test. The corrected invariant correctly distinguishes flux (total flux $2\pi k$) from shear (total flux 0), where the 1D winding sum failed.

5. The magnetic-translation pair and locality

5.1. **Construction.** The **magnetic translations** $\mathcal{U}$ (horizontal shift with horizontal edge twist) and $\mathcal{V}$ (vertical shift with vertical edge twist) act on site-indexed vectors:

$$(\mathcal{U}\psi)(x+1,y) = T((x,y),\rightarrow)\cdot\psi(x,y),$$
$$(\mathcal{V}\psi)(x,y+1) = T((x,y),\uparrow)\cdot\psi(x,y).$$

5.2. **Commutator locality.**

Theorem 9 (Locality)**.** *For any edge family, $[\mathcal{U},\mathcal{V}]$ is block-diagonal in the plaquette index. Specifically, $[\mathcal{U},\mathcal{V}]_{(x+1,y+1),(x,y)} = \mathrm{plaqDisc}(T,(x,y))$ and all other entries vanish.*

Theorem 10 (Locality bound)**.** $\|[\mathcal{U},\mathcal{V}]\| \le \sup_v \|\mathrm{plaqDisc}(T,v)\|$*. The system size m does not amplify the commutator—the bound uses the maximum, not the sum, of plaquette discrepancies.*

5.3. **Flux commutator estimate.**

Theorem 11 (Almost commuting)**.** *For the flux family, $\|[\mathcal{U},\mathcal{V}]\| \le 2|\sin(\pi k/m^2)| \to 0$ as $m\to\infty$. The pair is* almost commuting *in the large-system limit, while carrying topological index k.*

This is the Exel–Loring setup: an almost-commuting pair whose K-theoretic index obstructs exact commutativity [2, 3, 4]. The total flux is the evaluation of this index.

6. What remains: the Plateau Conjecture

6.1. **The open problem.** The Plateau Conjecture is stated as a formal proposition in Lean:

$$\forall k\neq 0,\ \exists\, c(k)>0,\ \forall m\ge 3,\ \forall\, 2\le r\le d: \quad c(k)\le \mathrm{distOpFlat}(r,\mathrm{fluxFamily}(k,d,r)).$$

All definitions are concrete. The conjecture is no longer a monolithic `sorry`: the reduction of §6.2 is now itself machine-checked (`plateau_of_rigidity` in `PlateauConjecture.lean`), and the library's only remaining `sorry` is the single scoped estimate `FluxPairRigidity` of §6.3.

6.2. **The missing step: index stability.** The reduction runs through the **Exel–Loring** obstruction for almost-commuting operators. Two ingredients are classical and, crucially, *dimension-free*:

(i) **Well-definedness.** There is a universal constant $\varepsilon_0>0$ such that whenever $\|[U,V]\|<\varepsilon_0$ for a pair of unitaries, an integer Bott index $\mathrm{bott}(U,V)$ is defined.
(ii) **Dimension-free rigidity.** There is a universal constant $\delta_0>0$, *independent of the matrix size*, such that if $\mathrm{bott}(U,V)\neq 0$ then (U,V) is at operator-norm distance $\ge\delta_0$ from every commuting pair [1, 2, 6, 7].

Granting a version of (i)–(ii) for our pair, the conjecture follows—and follows *uniformly in* m. For every m large enough that $2|\sin(\pi k/m^2)|<\varepsilon_0$, the flux pair has a defined Bott index equal to $k\neq 0$ (the content of totalFlux $=2\pi k$); flat families give *exactly commuting* pairs (`magComm_flat`), hence index 0; so by (ii) the flux pair is at distance $\ge\delta_0$ from every flat pair, with δ_0 independent of m. The finitely many remaining small-m cases are handled separately. **The m-uniformity is therefore not a separate miracle to engineer—it is inherited from the dimension-independence of δ_0, the defining feature of the Exel–Loring obstruction.** Note the direction of the logic: the shrinking commutator $2|\sin(\pi k/m^2)|\to 0$ is the *enabling hypothesis* of (i), not a threat to (ii).

6.3. **What is genuinely open.** Two gaps were thought to separate the classical statement from our setting. Formalizing the reduction settled their status asymmetrically.

(1) **Non-unitarity—open, and now the *entire* remaining problem.** $\mathcal{U}$, $\mathcal{V}$ are rank-r contractions (partial isometries on the surviving subspace), not unitaries—this non-invertibility is the entire "lossy" premise. Ingredients (i)–(ii) must be transported to the gapped contraction setting, plausibly by passing to the *unitary part* on the core subspace (the object the `coreBlock`/`rotAngle` machinery already extracts), but this transport is not automatic and the constants ε_0, δ_0 must be shown to survive it. In the Lean library this is scoped as the Prop `FluxPairRigidity`: for each $k \neq 0$ a $\delta(k) > 0$, independent of m, d, r, separating the flux pair from **every commuting contraction pair** on the same space. It carries the library's only `sorry` (`flux_pair_rigidity`).

(2) **Distance translation—closed; the hard direction is not load-bearing.** Rigidity controls distance in the pair $(\mathcal{U}, \mathcal{V})$; the conjecture is about edge-wise distOpFlat. The formalized reduction shows only the easy direction is needed, and with constant exactly 1: single-translation locality $\|\mathrm{magU}(T) - \mathrm{magU}(T')\| \leq \sup_e \|T(e) - T'(e)\|$ (`opNormG_magU_le`, m-free, the same mechanism as commutator locality), applied in the contrapositive—any edge-wise c-close flat family assembles into a commuting contraction pair c-close to the flux pair (`magComm_flat` + `opNormG_magU_flat_le_one`). The converse translation, flagged in earlier drafts as "the second half," never enters. Hence `plateau_of_rigidity : FluxPairRigidity → PlateauStatement`, machine-checked, constant preserved.

Adapting (i)–(ii) to a gapped non-unitary pair—now the one estimate `FluxPairRigidity`—is the natural question for the operator-algebra and quantitative-topology communities [2, 6, 7, 5, 9, 10]. Two remarks on its scope. First, for $r = d$ the complexified core of the flux pair **is** the classical Voiculescu pair (magnetic translations at flux $2\pi k/m^2$ on m^2 sites) [1], where rigidity is known; the open content is uniformity across rank deficiency $r < d$. Second, quantifying over all $m \geq 3$ (rather than $m \geq m_0(k)$) additionally asserts finitely many small-m cases, a finite check once $m_0(k)$ is explicit.

Empirical status of the estimate. A pre-registered adversarial test gives two facts, in order of reliability. (i) An elementary ceiling: the zero pair commutes, is contractive, and sits at distance $\max(\|\mathcal{U}\|, \|\mathcal{V}\|) = 1$ from the flux pair at every m and every padding—so the constant in `FluxPairRigidity` satisfies $\delta(k) \leq 1$. (ii) Adversarial search for a nearer commuting contraction pair (penalty optimization with exact-commuting Schur certification; `e4_pair_rigidity.py` in the repository's `numerics/` folder; $k = 1$, $m = 3, \ldots, 10$, zero-padding to $z = m^2$) found **no competitor below the trivial distance** 1—the kill condition (distance decaying toward 0, or collapsing under rank-padding) did not fire. The search's own certified constructions all landed *above* 1, i.e., it never reached the trivial basin, so (ii) is a non-firing falsification attempt, not a measured floor. The Bott index of the pair is pinned at k throughout, and $\|[\mathcal{U}, \mathcal{V}]\| = 2|\sin(\pi k/m^2)|$ to machine precision. Net: the open estimate is bracketed $\delta(k) \in (0, 1]$ conjectured, with the strict positivity exactly what remains.

6.4. **Surviving theorems.** The following are proved and available for the final step:

- collapse lemma and chain/cycle algebra;
- flux family has core gap 2 (m-uniform);
- nearby flat families retain core gap ≥ 1 (gap stability, m-uniform);
- total flux distinguishes flux ($2\pi k$), flat (0), and shear (0);
- commutator locality and the almost-commuting estimate;

- Stage A bound: distOpFlat $\geq |\sin(2\pi k/m)|/m$ (positive, but decays like $1/m^2$—the 1D limit).

7. Discussion

7.1. **Formalization as a research tool.** This project exemplifies a workflow where formalization is not an after-the-fact certification but an active participant in mathematical discovery. The shear family counterexample was not an external criticism—it emerged from the formalization process itself. When we attempted to prove the uniform gap lemma (`UnifGapLiftCont`), the proof assistant forced us to consider all edge cases, revealing that the horizontal edges could be made trivial without affecting the 1D invariant.

The "falsification-first" methodology—constructing the simplest possible counterexample to a proposed lemma before investing in its proof—is standard in mathematics. Formalization makes this methodology systematic: the type checker demands that every case be handled, and the cases you cannot handle become counterexamples.

7.2. **Why the shear family is subtle.** The shear family is not an exotic pathology. It is the simplest modification of the flux family: set horizontal edges to identity. A physicist might object that this violates the boundary conditions—but on a lattice torus, the only boundary is the wrap, and the flux family already has a special boundary rule (the Dirac string on horizontal wrap edges). The shear family simply chooses a different boundary rule. Both are legitimate edge families; both satisfy the core gap condition; both have well-defined winding. The fact that they differ in their 2D topology is invisible to any 1D measurement.

7.3. **The Exel–Loring index.** The corrected invariant—the total flux, which equals the index of the magnetic-translation pair—is inherently 2D. It sums over plaquettes, not columns. It reads every edge. It vanishes on the shear family because the positive interior contributions are exactly canceled by negative boundary contributions, reflecting the fact that the shear family is "flat with a gauge twist" rather than genuinely charged.

The remaining mathematical challenge is the index stability theorem: that a sufficiently small perturbation of an almost-commuting pair cannot change its K-theoretic index. This is a known result in operator algebras for unitary pairs [2, 6, 7, 8]. What formalizing the reduction accomplished (§6.3) is a precise accounting of how much of the adaptation is genuinely open: the edge-to-pair assembly, the m-free locality constant, the exact commutativity and contractivity of flat pairs, and the reduction itself are all machine-checked; the classical unitary specialization is known; what remains—the library's only `sorry`—is transporting the dimension-free lower bound from unitary pairs to gapped partial-isometry pairs uniformly in rank deficiency. Direct adversarial optimization against that statement (nearest commuting contraction pair, exact-commuting certification, system sizes $m = 3, \ldots, 10$, rank-deficient padding to twice the dimension) found no competitor below the trivial zero-pair distance of exactly 1, so the estimate survives its falsification attempt with the constant bracketed in $(0, 1]$—the strict positivity being exactly what is formally open.

8. Conclusion

We have presented a machine-checked formalization that:

(1) states the Plateau Conjecture precisely over concrete objects,
(2) constructs and proves properties of a 1D winding invariant,
(3) discovers and verifies a counterexample (the shear family) that falsifies the 1D approach,
(4) constructs a corrected 2D invariant (total flux) that passes the falsification test,
(5) builds the magnetic-translation pair and proves its locality,

(6) reduces the Plateau Conjecture, via the machine-checked `plateau_of_rigidity`, to a single scoped estimate—`FluxPairRigidity`—deleting a phantom second half of the problem in the process,

(7) leaves that one estimate as the open problem, stated over concrete objects with its classical unitary specialization known.

The formalization is available at https://github.com/GoodRoyal/LossyCocycles. All theorems stated in this paper are machine-checked in Lean 4. The library's only `sorry` is the open estimate `FluxPairRigidity`.

Authorship and AI-assistance disclosure

This research program — its conception, its original ideas, its questions, and its direction — is the work of the author, Juan Carlos Paredes. The strategic pivots, the falsification discipline, the decisions about what to pursue, what to test, what to claim, and what not to claim are his throughout. In the course of the program he directed large language models (Anthropic Claude — Opus and Fable model families) as research assistants: implementing and running numerical experiments he commissioned, formalizing proofs in Lean 4 (all machine-verified), drafting and revising prose, and retrieving literature, all under his direction and review. Juan Carlos Paredes is the author of this work; the assistance is disclosed in the spirit of the full transparency this venue requires.

Appendix A. Formalization statistics

File	Content	Lines
`Basic`	collapse lemma, chain/cycle algebra	72
`Defs`	norms, torus, flux family, distance-to-flat	74
`Winding`	core block, core vector, winding sum, flat/flux anchors	569
`Shear`	shear family construction and falsification proof	430
`PlaqIndex`	plaquette holonomy, total flux, three anchors	433
`MagPair`	magnetic translations, commutator locality, flux estimate	373
`StageA`	elementary norm kit, telescoping, Stage A bound	493
`StageB`	gap stability, source-projection toolkit	232
`PlateauConjecture`	statement, assembly locality, the reduction, the open estimate	364

Total: approximately 3,040 lines of Lean 4, all machine-checked except the single `sorry` carrying `FluxPairRigidity`.

References

[1] D. Voiculescu, *Asymptotically commuting finite rank unitary operators without commuting approximants*, Acta Sci. Math. (Szeged) **45** (1983), 429–431.

[2] R. Exel, T. A. Loring, *Almost commuting unitary matrices*, Proc. Amer. Math. Soc. **106** (1989), no. 4, 913–915.

[3] R. Exel, T. A. Loring, *Invariants of almost commuting unitaries*, J. Funct. Anal. **95** (1991), no. 2, 364–376.

[4] R. Exel, *The soft torus and applications to almost commuting matrices*, Pacific J. Math. **160** (1993), no. 2, 207–217.

[5] P. Friis, M. Rørdam, *Almost commuting self-adjoint matrices—a short proof of Huaxin Lin's theorem*, J. Reine Angew. Math. **479** (1996), 121–131.

[6] M. B. Hastings, *Making almost commuting matrices commute*, Comm. Math. Phys. **291** (2009), 321–345. arXiv:0808.2474.

[7] M. B. Hastings, T. A. Loring, *Almost commuting matrices, localized Wannier functions, and the quantum Hall effect*, J. Math. Phys. **51** (2010), 015214. arXiv:0910.5490.

[8] M. B. Hastings, T. A. Loring, *Topological insulators and C^*-algebras: theory and numerical practice*, Ann. Physics **326** (2011), no. 7, 1699–1759. arXiv:1012.1019.

[9] T. J. Osborne, *Almost commuting unitaries with spectral gap are near commuting unitaries*, arXiv:0809.0602 (2008).

[10] T. A. Loring, *A guide to the Bott index and localizer index*, arXiv:1907.11791 (2019).

Independent researcher
Email address: cpt66778811@gmail.com