THREE GAPS, ONE INDEX: THE m-UNIFORM RIGIDITY OF A LOSSY MAGNETIC PAIR IS TOPOLOGICAL, NOT SPECTRAL

JUAN CARLOS PAREDES

Abstract. In one fully specified, machine-checked special case—the charge-k magnetic flux family on an $m \times m$ torus of lossy (rank-deficient) links—we disambiguate three distinct quantities that the almost-commuting-operator and topological-protection literatures all call "the gap," and determine which survive uniformly in the system size m. The **local** gaps—the commutator norm $\|[\mathcal{U}, \mathcal{V}]\|$, the magnetic Hamiltonian's spectral gap, and the spectral-localizer gap they control—all **vanish** like the flux-per-plaquette $2\pi k/m^2 \to 0$. Two other gaps are m-**uniform**: the *core gap* (the rank/isometry structure of the links, independent of the magnetic data) and the Bott branch gap (well-definedness of the index). The consequence is a sharp negative with a constructive moral: the m-uniform operator-norm separation $c(k)$ between charge and flatness—the Plateau quantity—**cannot** be sourced from any local spectral gap, and is therefore irreducibly *topological* (a Voiculescu distance-to-commuting that stays $\approx \sqrt{2}$ while every local gap collapses), transported into the non-invertible setting by the rank-structure core gap, not by a spectral gap. This locates exactly where a proof of the Plateau Conjecture must, and must not, look. We then make the positive locale concrete: reading the rank loss as the ideal of an extension reduces the remaining unitary-to-contraction passage to a single off-diagonal Schur block, which numerics (Exp. E, and an adversarial worst-case Exp. F) show is core-gap-controlled and m-uniform—recasting "extend a dimension-free theorem to contractions" as "bound one block." A direct minimization of the distance to flatness itself (Exp. G) independently fails to collapse the distance, converging up to the *same* $\sqrt{2}$ as the index-jump radius—the conjecture's most direct numerical test, passed.

The $\sqrt{2}$ above is the **codimension-1** value ($d = r + 1$; the tested case $d = 3$, $r = 2$). For codim $= d - r \geq 2$, a *trivial* flat competitor supported on an r-plane orthogonal to the rotating core sits at operator distance **exactly** 1 from the flux family at every m, so distOpFlat $\leq 1 < \sqrt{2}$ there. The separation is therefore codimension-dependent: $c(k) = \sqrt{2}$ at codimension 1, $c(k) = 1$ at codimension ≥ 2—both m- and k-uniform. The **Plateau Conjecture ($c(k) > 0$) holds across** $2 \leq r \leq d$**, with uniform floor** 1; what is codimension-1-specific is the *mechanism*—the dimension-free **unitary** (Voiculescu/Exel–Loring) bound transports only at codimension 1, while for codimension ≥ 2 the binding competitor is the orthogonal one the unitary bound cannot see.

House style. Every claim in this note compiles in Lean ([THEOREM]), cites the literature ([CITED]), carries numerical evidence with its script ([EVIDENCE]), or wears a [CONJECTURE] tag. Numerics (experiments A–G, full detail in §9) are the evidence base and are referenced throughout by letter.

1. The conflation, and the special case that resolves it

"Locally almost-flat data can be globally far from trivial, and a *gap* protects the obstruction." This sentence is true and load-bearing across three literatures—stability of almost-representations, topological protection of indices, and the C*-algebraic theory of almost-commuting matrices [1, 2, 7, 13]. But "the gap" in it names different objects in

Date: July 2, 2026.

different arguments, and in the *lossy* (non-invertible, rank-deficient) regime the differences decide whether a uniform obstruction exists at all.

We resolve the ambiguity in a single instance where every object is explicit and most are machine-checked: the charge-k flux family of the `LossyCocycles` library (`Defs.lean`, `PlateauConjecture.lean`). The instance is not a toy—it is the canonical Harper/Hofstadter magnetic-translation pair, the same object that carries the Chern/Bott index in topological-insulator physics—but rendered as *rank-r partial isometries on each torus edge, with $r < d$*, which is what makes it lossy and pushes it outside the classical invertible theory.

2. Objects (imported, with tags)

On the $m \times m$ torus each directed edge e carries $T(e) \in \mathbb{R}^{d\times d}$ of rank $\le r$.

- **Flat family** [THEOREM-anchored]: $T(e) = V(\operatorname{dst} e)\cdot V(\operatorname{src} e)^{\mathsf{T}}$ for site frames V ($V^{\mathsf{T}}V = I_r$). Pure gauge; every cycle composite is a projection (`Basic.lean`).
- **Flux family** [THEOREM-anchored]: $T(e) = S\cdot R(\varphi_e)\cdot S^{\mathsf{T}}$, $S = \operatorname{stdFrame}(d, r)$, $R(\theta)$ the SO(2) rotation of the core 2-plane, φ the Landau-gauge angle (vertical edges $2\pi kx/m^2$; horizontal wrap edges $-2\pi ky/m$). Total flux $2\pi k$ (`PlaqIndex.lean`, `totalFlux_flux`).
- **Magnetic pair** [THEOREM-anchored] (`MagPair.lean`): $\mathcal{U}$ (horizontal transport and shift), $\mathcal{V}$ (vertical). Flat $\Rightarrow$ $[\mathcal{U}, \mathcal{V}] = 0$ exactly (`magComm_flat`); for any family $\|[\mathcal{U}, \mathcal{V}]\| \le \sup_s \|\operatorname{plaqDisc}(s)\|$, edge-local, m-free (`opNormG_magComm_le`).
- **Distance to flatness**: $\operatorname{distOpFlat}(r, T) = \inf_{T' \text{ flat}} \sup_e \|T(e) - T'(e)\|$. The **Plateau Conjecture** [CONJECTURE]: $\exists\, c(k) > 0$ with $\operatorname{distOpFlat}(r, \text{flux}) \ge c(k)$ for all $m \ge 3$, $2 \le r \le d$ (`PlateauConjecture.lean`). The conjecture is no longer a raw `sorry`: a machine-checked reduction (`plateau_of_rigidity`) derives it, constant preserved and m-free, from the single pair-level estimate `FluxPairRigidity` ($\delta(k) > 0$, independent of m, d, r, separating the flux pair from every commuting contraction pair; elementarily $\delta(k) \le 1$ via the zero pair), which now carries the program's one `sorry`.
- **Bott index** [EVIDENCE/CITED]: for the (complexified) magnetic pair, $\operatorname{bott}(\mathcal{U}, \mathcal{V}) = \frac{1}{2\pi}\sum \arg\operatorname{spec}(\mathcal{U}\mathcal{V}\mathcal{U}^\dagger\mathcal{V}^\dagger) = \operatorname{totalFlux}/2\pi = k$.

3. Three gaps

We separate three quantities all called "gap," and measure their m-dependence (Exp. A–Exp. D).

G1 — the local/spectral gaps. [EVIDENCE: A, D] — all vanish $\propto 1/m^2$. Three apparently different quantities are one family:

- the commutator $\|[\mathcal{U}, \mathcal{V}]\| \le 2|\sin(\pi k/m^2)|$ (Exp. A; matches the machine-checked bound);
- the spectral gap of the magnetic Hamiltonian $H = \mathcal{U} + \mathcal{U}^\dagger + \mathcal{V} + \mathcal{V}^\dagger$ (Exp. D);
- the spectral-localizer gap, which Loring–Schulz-Baldes bound by the bulk gap of H [10, 11].

Each is set by the flux-per-plaquette $2\pi k/m^2$. Measured (Exp. D, $k = 1$): the largest spectral gap of H falls $0.334 \to 0.012$ over $m = 6 \to 32$ ($\approx 2\times$ the flux density), and the gap at the Fermi level $E = 0$ is ≈ 0 throughout. The commutator (Exp. A) falls $0.390 \to 0.006$. **No member of G1 has an m-uniform floor.**

G2 — the core gap (rank/isometry structure). [EVIDENCE: B] — m-uniform, $= 1$. The smallest singular value of the edge maps restricted to the surviving SO(2) core. Because the rotations $R(\varphi)$ are isometries on the core *regardless of the angle*, the core gap is **independent**

of the magnetic data and equals 1 at every m (Exp. B, $\sigma_{\min}(\text{core}) = 1.0$ for $m = 4, \ldots, 16$). This is the bounded-geometry quantity, and it is the one that controls the contraction-to-unitary *polar* transport: the polar (unitary-part) map is m-uniformly Lipschitz, constant ≈ 0.82 flat across m (Exp. B), consistent with the Kahan/Li bound $1/\gamma$ at $\gamma = 1$. The core gap is not a spectral gap; it is the lossy structure itself.

G3 — the Bott branch gap (index well-definedness). [EVIDENCE: A] — m-uniform, $\to \pi$. The distance of $\operatorname{spec}(\mathcal{U}\mathcal{V}\mathcal{U}^\dagger\mathcal{V}^\dagger)$ from the branch cut at -1. The index $\frac{1}{2\pi}\sum \arg$ is well-defined and integer-valued exactly while this gap is positive. Measured (Exp. A): it *grows* $2.75 \to 3.14$ over $m = 4 \to 32$—the index becomes **more** robustly defined as the lattice grows. (Contrast G1, which makes the operators look *more* commuting at the same time.)

4. The disambiguation, and its consequence

> **Observation (special case).** Among the gaps attached to the charge-k flux family, the local/spectral family **G1 vanishes** like $1/m^2$, while **G2 (core gap) and G3 (branch gap) are m-uniform**. [EVIDENCE: A, B, D]

The consequence is the point of the note. A uniform obstruction $c(k)$ is an $O(1)$ quantity; it cannot be lower-bounded by anything in G1, on pain of a bound $\sim 1/m^2$ (this is precisely why the program's Stage A floor $|\sin(2\pi k/m)|/m$ decays—it is a G1-type estimate). So:

> **Corollary.** Any route to an m-uniform $c(k)$ through a **local spectral gap**—the commutator, the Hamiltonian gap, or the spectral-localizer gap—is **impossible** for this family. In particular the Loring–Schulz-Baldes "localizer gap $\Rightarrow$ distance-to-trivial" mechanism [11], dimension-free though it is in principle, yields only a vanishing (useless) bound here, because the gap that feeds it (G1) vanishes. [EVIDENCE: D]

Where, then, does the $O(1)$ come from? Directly: the **topological distance-to-commuting**. Along a homotopy from the flux pair to a trivial (commuting) flat pair, the Bott index stays k until the operator distance reaches $\sqrt{2} \approx 1.4142$ (Exp. C; for the tested codimension-1 case $d = 3$, $r = 2$—a clean $\mathrm{U}(n)$-geodesic re-run shows this radius *rises to* $\sqrt{2}$ *from below* with m rather than being flat, and it is codimension-dependent, §4.1). The separation stays $O(1)$ *while every local gap collapses*—the Voiculescu phenomenon [1]. **Independent corroboration** [EVIDENCE: G]: a *direct* minimization of the conjectured quantity itself—$\mathrm{distOpFlat}(\mathrm{flux}) = \inf_{T' \text{ flat}} \sup_e \|\mathrm{flux}(e) - T'(e)\|$, optimized over the full product of site-frame Stiefel manifolds—does **not** decay with m; the best flat competitor we can find converges *monotonically up to* $\sqrt{2}$ ($1.294, 1.397, 1.410, 1.412, 1.413$ for $m = 4, \ldots, 32$, within 10^{-3} of $\sqrt{2}$ at $m = 32$). Two unrelated computations—the homotopy index-jump radius (C) and the metric infimum (G)—land on the **same** constant $\sqrt{2}$, which is precisely the reframe's claim that the topological radius *equals* the achievable distance. This $O(1)$ is genuinely non-local; it is an abstract Exel–Loring/Voiculescu quantity, not a spectral one. What carries it into the lossy ($r < d$) setting m-freely is **G2**: the core gap makes the polar reduction to genuine unitaries uniformly bi-Lipschitz (Exp. B), after which the dimension-free unitary obstruction applies.

> **One-line moral.** "Locally gapped $\Rightarrow$ rigid" is true here only for the *right* gap: the rigidity is m-uniform because the **rank-structure** core gap (G2) is, and because the index it protects is **topological** (G3 / Voiculescu)—never because a *spectral* gap (G1) is, since none is.

4.1. **Codimension dependence of $c(k)$: $\sqrt{2}$ is special to codimension 1 [EVIDENCE].** Section 4 established the $\sqrt{2}$ only for $d = 3$, $r = 2$ (codimension $d - r = 1$). Sweeping the codimension yields a sharp dichotomy:

- **The core gap and the ideal leak are codimension-invariant.** $\sigma_{\min}(\text{core}) = 1$ and the ideal leak $\|(1-P)[\mathcal{U}, \mathcal{V}]P\| = 0$ at every m for $d = 3, 4, 5, 6$—the SO(2) core is an isometry on its 2-plane regardless of the ambient d. So G2/G4 are the correct (r, d)-uniform objects; the worry that bounded geometry might degrade with codimension is answered *no*.
- **But the separation** distOpFlat **collapses from $\sqrt{2}$ to** 1 **once** codim ≥ 2**.** The magnetic flux rotates only a fixed 2-plane, so a flat competitor needs an r-plane *avoiding that 2-plane*, which fits iff $d - r \geq 2$. The **constant** competitor $T'(e) = WW^{\mathsf{T}}$ (W that orthogonal r-plane) then satisfies $\|\text{flux}(e) - WW^{\mathsf{T}}\| = 1$ *exactly* at every edge and every m (the mismatched pieces—flux on the rotating plane, WW^{T} on W—lie on orthogonal norm-1 subspaces; verified 1.000000 for $r = 2$, $d = 4, 5, 6$ and $r = 3$, $d = 5, 6, 7$). Hence distOpFlat ≤ 1 for all m when codim ≥ 2—a hand-built, optimizer-independent bound. Warm-starting a Riemannian descent *from* this competitor cannot beat 1, so distOpFlat $= 1$ there, and the bound is k-independent (the competitor ignores the flux angle).

  Refined reading. $c(k) = \sqrt{2}$ at **codimension 1** ($d = r + 1$); $c(k) = 1$ at **codimension** ≥ 2—both m- and k-uniform. The **Plateau Conjecture (**$c(k) > 0$**, uniform over** $m \geq 3$**,** $2 \leq r \leq d$**) holds, with uniform floor** 1**.** What is codimension-1-specific is the *mechanism*: the dimension-free **unitary** (Voiculescu/Exel–Loring/Loring) bound transports only at codimension 1. For codim ≥ 2 the binding competitor is the **orthogonal** one, at distance $1 < \sqrt{2}$, which the unitary bound cannot see—equivalently, the Bott index is *not* locally constant on $\sqrt{2}$-balls of the contraction metric there. This realizes the "cheap channel" of §8 in the high-codimension regime: non-invertibility genuinely opens a channel the unitary theorem never sees.

The matching lower bound distOpFlat ≥ 1 for codim ≥ 2 is **proved for** $r = 2$ (and numerically decisive for all r) via the rotating-plane (Π) restriction: $\delta \geq PB := \max_e \|(\text{flux} - T')P_\Pi\|$, whose minimum over flat competitors is 1 (the orthogonal competitor), established by a trace identity and a winding/coboundary contradiction against the dimension-free unitary radius (modulo that standard import). For $r > 2$ the same $\min PB = 1$ is verified numerically, but the proof is **open**: the $r = 2$ argument turns on the Π-shadow connection being a *coboundary* (Chern 0), whereas for $r > 2$ the shadow line bundle can carry nonzero Chern, so closing it needs the higher-r dimension-free obstruction (the same "not-packaged" gap as §8). So for codim ≥ 2, distOpFlat $= 1$ is **proved for** $r = 2$ and numerically decisive for all r (upper bound ≤ 1 from the orthogonal competitor, all r).

(Earlier optimizer runs reporting $\sim \sqrt{2}$ at large m for codim ≥ 2 were artifacts of the descent never finding the orthogonal basin; the hand construction and the Π-restricted bound are the reliable statements, and the companion papers carry the corrected $\delta(k) \in (0, 1]$ bracket.)

4.2. **The codimension-≥ 2 lower bound for $r > 2$: reduction to a magnetic loop-holonomy estimate.** Section 4.1 left distOpFlat ≥ 1 for codim ≥ 2, $r > 2$ as **open**, proved only for $r = 2$. This section records the reduction the program has since obtained: a chain of [VERIFIED]/[PROVED] steps that collapses the open problem to a **single quantitative holonomy inequality of Exel–Loring type**, plus the numerical and structural evidence that pins it. The boundary between proved and conjectural is made unmistakable; the open core is flagged [CONJECTURE].

Setup (from §4.1). $PB := \max_e \|(\mathrm{flux}(e) - T'(e))P_\Pi\| \leq \mathrm{distOpFlat}$, so $PB \geq 1 \Rightarrow \mathrm{distOpFlat} \geq 1$. PB depends only on the complexified Π-shadow $\alpha_s \in \mathbb{C}^r$ ($\|\alpha_s\|_{\mathrm{op}} \leq 1$, the contraction cap). Write $\alpha_s = \nu_s u_s$ (u_s unit, $\nu_s = \|\alpha_s\|$); $p_e := \langle u_t, u_s \rangle$ (Hermitian overlap), $q_e := (u_t, u_s)$ (bilinear), $\omega_s := (u_s, u_s)$, and φ_e the Landau-gauge flux phase.

(R1) Sharp per-edge characterization [VERIFIED, 10^{-15}]. $PB_e < 1 \iff g_e > |G_e|$, with $g_e = \mathrm{Re}[e^{i\varphi_e} P_e] - \frac{1}{2} n_s$ (real), $G_e = e^{-i\varphi_e} Q_e - \frac{1}{2} w_s$ (complex), $\det(H_e) = g_e^2 - |G_e|^2$; cap $n_s + |w_s| \leq 2$.

(R2) Telescoping $\Rightarrow$ Chern $-k$ [VERIFIED]. $PB < 1$ on all edges forces the Π-shadow line $[u_s] \in \mathbb{CP}^{r-1}$ to carry Chern $= -k$ (edge cancellation on the torus). A $PB < 1$ competitor must live in the binding sector.

(R3) Magnitude/line factorization [VERIFIED]. $PB_e < 1 \iff \mathrm{bracket}_e > 0$, where $\mathrm{bracket}_e = \nu_t \,\mathrm{Re}[e^{i\varphi_e} p_e] - \frac{1}{2}\nu_s - |\nu_t e^{-i\varphi_e} q_e - \frac{1}{2}\nu_s \omega_s|$, homogeneous of degree 1 in ν. So feasibility depends only on the line directions u and the ratios $\Delta y_e = \log(\nu_t/\nu_s)$ (a coboundary).

(R4) Empty-edge lemma [PROVED]. If $R_e := \mathrm{Re}[e^{i\varphi_e} p_e] \leq |q_e|$ on some edge, then $PB_e \geq 1$ unconditionally (reverse-triangle plus casework, $|\omega| \leq 1$). So a $PB < 1$ competitor is **all-ray**: $R_e > |q_e|$ on every edge.

(R5) Difference-constraint / homological reformulation [PROVED]. On an all-ray line, $\mathrm{bracket}_e > 0 \iff \Delta y_e > \theta_e$, $\theta_e := \log x_e^0$ (x_e^0 the first positive root of $R_e x - \frac{1}{2} = |x e^{-i\varphi_e} q_e - \frac{1}{2}\omega_s|$). The constraints $y_t - y_s > \theta_e$ admit a single-valued potential y on the torus iff no directed cycle has nonnegative mean weight. The competitor graph has **forward edges only**, so every directed cycle is non-contractible ($H_1(T^2)$). Hence:

> $PB \geq 1$ (the bound holds on an all-ray line) $\iff$ $\mu^* \geq 0$, where $\mu^* :=$ the Karp maximum mean cycle of $\theta =$ the largest mean θ over non-contractible directed loops.

This is the cleanest form of the obstruction: a **single-valued-potential / monodromy** statement. A $PB < 1$ competitor would be a height function climbing by $> \theta_e$ across every edge yet closing on the torus; $\mu^* \geq 0$ says no such function exists—some non-contractible loop has total climb ≥ 0.

(R6) Monotonicity $\Rightarrow$ a Hermitian bound [PROVED]. On a ray edge, $\theta_e \geq \theta_e^{(0)} := -\log(2R_e)$ (the bilinear (q, ω) only raises the threshold; $f_e(1/2R_e) \leq 0$, $f'_e \geq R_e - |q_e| > 0$). Writing $\mu_0^* :=$ the Karp maximum mean cycle of $\theta^{(0)}$, we have $\mu^* \geq \mu_0^*$, so the bound follows from the purely **Hermitian** statement

> **(H)** some non-contractible loop C has $\prod_{e \in C} 2R_e \leq 1$ (i.e. $\mu_0^* \geq 0$), $R_e = \mathrm{Re}[e^{i\varphi_e} p_e]$.

Via the exact identity $2R_e = 2 - \|u_t - e^{i\varphi_e} u_s\|^2$ [VERIFIED, 10^{-15}], $R_e = |p_e| \cos\psi_e$ with $\psi_e = \varphi_e + \arg p_e$, and the combined connection ψ has **total Chern** 0 (flux $+2\pi k$ cancels the Berry curvature $-2\pi k$)—exactly the Exel–Loring setting, with the magnitude $|p_e| \leq 1$ only helping.

(R7) The sector dichotomy [EVIDENCE, decisive]. (H) as a statement for *general* Chern $\neq 0$ is **false**: a coherence attack holding the line all-ray produces all-ray Chern $\approx +1$ lines (slack $R - |q| = +0.015$) with every loop $\prod 2R > 1$ ($\mu_0^* = -0.16$), the bound rescued only by the bilinear ($\mu^* = +0.82$). **But in the binding sector Chern** $-k$ the obstruction survives:

a leashed in-sector continuation pushes coherence to the -2 ceiling (min-mean-$R \approx 0.60$) and there $\mu_0^* \approx +0.45$; the threshold where μ_0^* crosses 0 is $\tau \approx 0.69$, and the $-k$ topology—more winding, more frustration—cannot reach it. So the Hermitian (H) is **reinstated for the only sector the telescoping (R2) makes relevant**.

The open core [CONJECTURE]. Either of the following closes distOpFlat ≥ 1 for codim ≥ 2, $r > 2$ (and hence, with the upper bound of §4.1, distOpFlat $= 1$):

(C-H) *(Hermitian, sector-restricted—primary.)* For an all-ray line $[u_s]$ of Chern $-k$, $\mu_0^* \geq 0$: some non-contractible directed loop C has

$$\prod_{e \in C} 2\,\mathrm{Re}\big[e^{i\varphi_e}\langle u_t, u_s\rangle\big] \leq 1.$$

(C-μ) *(bilinear, unconditional—fallback.)* For any all-ray line of Chern $-k$, $\mu^* \geq 0$.

(C-H) is the bilinear-free, r=2-Loring-flavour target (positively supported by R7). (C-μ) holds with large margin ($\mu^* \approx 1.8$) in every tested case. Both are quantitative U(1)-holonomy estimates for the Chern-0 connection $\psi = \varphi + \arg p$, *stronger than the per-edge $r = 2$ Loring bound* (one bad edge $\neq$ a bad loop). They sit as contraction/tracial generalizations of the Exel–Loring formula [3].

Numerical and structural status.

- $\min PB = 1.00000$ for $r > 2$ to $m = 32$ ($k = 1, 2$), via a converged vectorized minimizer validated to 10^{-15} against the scalar reference; no winding or random start dips below 1. The unwind-to-1 / stay-wound-to-$\sqrt{2}$ mechanism is the honest content.
- **Shortcuts mapped and closed** (so the conjectural core is irreducible): wrap-loop averaging is insufficient (the witness loop is a diagonal staircase, not a row or column); a finite-m reduction is unavailable (the all-ray slack $\to 0$, not negative, so the regime does not terminate); and the θ-blow-up at the boundary forces $\mu^* \geq 0$ only in an exponentially thin layer.

Net (codimension ≥ 2, all r). Upper bound ≤ 1 proved (all r, §4.1). Lower bound ≥ 1 **proved for** $r = 2$ (mod the unitary import) and reduced for $r > 2$ to the single holonomy conjecture (C-H)/(C-μ), numerically decisive to $m = 32$. The reduction (R1)–(R7) is the paper's analytic contribution for the high-codimension case; the holonomy estimate is the isolated open frontier, sharpened in §4.3 to a single branch-cut lemma (R).

4.3. Sharpening the open core: from holonomy to a variance, and a branch-cut rigidity.

Section 4.2 left the open core as the loop-holonomy estimate (C-H). This section reports a chain of *exact, machine-verified* identities that replaces (C-H) by a single, more elementary topological lemma—the contribution being the **localisation**, not yet a proof.

The variance certificate [VERIFIED]. Per non-contractible wrap loop, build the Hermitian magnetic-ring operator $H_{\mathrm{ring}}[s{+}1, s] = z_e := e^{i\varphi_e}\langle u_t, u_s\rangle$. The m-uniform certificate of (C-H)$|_{-k}$ is **not** the localizer signature (blind—a winding integer, dead by R7/Stokes) nor its ground state, but the **uniform-mode second moment** $\mathrm{var}_0 := \langle 0|H_{\mathrm{ring}}^2|0\rangle - \langle 0|H_{\mathrm{ring}}|0\rangle^2$. An exact global **sum rule** (each edge in one loop)

$$\sum_{\text{loops}} \mathrm{var}_0 = \frac{2}{m}\Big(\sum_e |p_e|^2 + C_2\Big) - 4\sum_{\text{loops}} (\mathrm{mean}R_{\text{loop}})^2$$

(C_2 the next-nearest-neighbour bilinear) holds to 3.6×10^{-15}, with $\sum \mathrm{var}_0/m$ m-uniform ($0.70 \to 0.80$). So (C-H)$|_{-k}$ takes the form

> **(V)** *all-ray Chern* $-k \Rightarrow \sum_{\text{loops}} \mathrm{var}_0 \geq a_0 m$, $a_0 > 0$ *m-uniform* ($\Rightarrow$ some loop has $\mathrm{var}_0 \geq a_0/2$).

Reduction of (V) to a branch-cut count [VERIFIED]. (i) Cauchy–Schwarz on each loop ($\mathrm{mean}R^2 \leq \langle R^2 \rangle$, loops partition edges) gives the **static** bound $\sum \mathrm{var}_0 \geq \frac{1}{m}\big[E_L(\operatorname{Im} z) - E_L(\operatorname{Re} z)\big]$, E_L the loop-Dirichlet energy—positive, m-stable, no spectra. (ii) $E_L(\operatorname{Im} z) - E_L(\operatorname{Re} z) = -\operatorname{Re} \sum (\Delta z)^2$ is a **phase-disorder** functional: it vanishes on every non-Nyquist uniform winding and is winding-blind, so positivity is irreducibly 2-dimensional. (iii) The manifestly nonnegative form $\sum \mathrm{var}_0 = \sum_{\text{loops}} \mathrm{Var}_s(\bar{z}_{e_s} + z_{e_{s-1}})$ vanishes **iff** the edge field is a **product** ($z^H_{(x,y)} = h(y)$, $z^V_{(x,y)} = v(x)$, for m odd; plus a Nyquist branch for m even)—**iff** $C_\psi = 0$, where C_ψ is the Chern of the **coboundary** $\psi = \varphi + \arg p$ (raw plaquette flux sums to 0), equal to $-\sum_{\text{plaq}} \mathrm{round}(F_{\text{raw}}/2\pi)$ = the net count of plaquettes whose ψ-holonomy is driven past $\pm\pi$. Thus:

> **(R)** *all-ray* ($R_e > |q_e|$) + *band Chern* $-k \Rightarrow C_\psi \neq 0$—the coboundary ψ is forced to **over-wind** some plaquette past the $\pm\pi$ branch cut. Then $C_\psi \neq 0 \Rightarrow$ not product $\Rightarrow \sum \mathrm{var}_0 > 0$; an integer-jump stability version yields the quantitative (V).

Status. (R) is the high-codimension open core in its sharpest form—a discrete Bott/branch-cut statement sitting exactly in the Exel–Loring regime ("the determinant winds once $\|[U,V]\| < \varepsilon$"), with the all-ray hypothesis $R_e > |q_e|$ playing the role of the $< \varepsilon$ commutator threshold that keeps F_{raw} near, but pushed past, the cut. It is the discrete target for the tracial Exel–Loring import; it is *not* a numerics question (numerics certify $\inf \mathrm{var}_0 > 0$ and verify every identity above; what remains is the forced cut-crossing). Relation to the Lean state: the program's `sorry` is the *pair-level* estimate `FluxPairRigidity` (uniform strict positivity, all codimensions—see §2), of which the machine-checked reduction makes the Plateau Conjecture a corollary; (R) is the *edge-level, sharp-constant* core at codimension ≥ 2, $r > 2$—the two are faces of the same transported-rigidity frontier, (R) delivering the exact constant 1 where `FluxPairRigidity` delivers uniform positivity. **This paper poses (R) as its open problem and stops there.**

5. Why the distinction is easy to miss (the trap)

The three gaps coincide in the *invertible, fixed-size* theory, where everything is $O(1)$ and nobody must choose. They split only in the lossy $m \to \infty$ regime that is the subject here, and they split in opposite directions: G1 $\to 0$ (the pair looks ever more commuting/trivial) while G3 $\to \pi$ and G2 stays 1 (the index is ever more sharply defined and the structure ever intact). An argument that reaches, by reflex, for "the gap" will grab G1—the commutator or the Hamiltonian gap, the most visible one—and conclude either that there is no obstruction (it vanishes) or that the bound decays. Both are artifacts of picking the wrong gap. The program's own history records the reflex: the falsified column-winding route and the decaying Stage A bound are both G1-flavoured; the corrected route is G2+G3.

6. Falsification map

if this fails	then
core gap (G2) decays with m (Exp. B $\sigma_{\min}$ drops)	bounded geometry fails; the polar transport is not m-free; the rank structure cannot carry the index—the corrected route loses its vehicle.
branch gap (G3) $\to 0$ with m (Exp. A)	the index is not robustly defined at large m; the topological reading is unsafe.
distance-to-commuting (Exp. C) decays with m	the separation is *not* $O(1)$; the Plateau is false and $c(k) \to 0$. (Strongest kill.)
direct distOpFlat minimization (Exp. G) decays with m	the *actual conjectured quantity* collapses—its strongest, most direct falsifier. (Exp. G says no: best competitor rises to $\sqrt{2}$, no decay. Caveat: an upper bound, so it can kill but not confirm.)
a member of G1 turns out m-uniform after all	a *constructive* spectral route to $c(k)$ would reopen; the note's central negative is wrong. (Exp. A, Exp. D say no.)
compression leak $\|(1-P)[\mathcal{U}, \mathcal{V}]P\|/\varepsilon$ (G4, Exp. E/Exp. F) grows like a power of m	the cheap channel is real; the support freedom of a flat competitor beats the unitary bound; the compression route of §8 loses its m-uniform Schur estimate. (Random: non-monotone in $[2.1, 4.0]$ out to $m = 32$. **Adversarial** (Exp. F): pinned at 2.0 at every m—no coherent edge accumulation.)

7. What this is, and is not

This note does **not** prove the Plateau Conjecture: the $O(1)$ separation rests on the abstract dimension-free Exel–Loring/Voiculescu bound [1, 2], which we import, not derive. What it contributes is a *localization* result: in a fully explicit special case it shows **which** gap can and cannot deliver an m-uniform obstruction, refutes the natural constructive-via-spectral-gap route with direct evidence, and identifies the rank-structure core gap as the correct bounded-geometry quantity and the index as irreducibly topological. For a proof program this is the difference between searching the right and the wrong half of the problem; for the broader almost-commuting / protection literature it is a cautionary disambiguation with a worked, reproducible example.

8. The compression route: where the wall actually is

Section 4 locates the $O(1)$ separation in an abstract dimension-free Exel–Loring/Voiculescu quantity—a theorem about **unitaries**. Our operators are not unitary: the links lose rank, so the magnetic pair is a pair of **gapped partial isometries**. The standing question—the program's one substantive gap, distinct from a definitional `sorry`—is whether that dimension-free bound *survives the passage to contractions*, or whether non-invertibility opens a "cheap channel" the unitary theorem never has to see. [CONJECTURE]

The reframe is to stop trying to **transport the metric bound** across the polar decomposition (metric bounds are not functorial and do not survive arbitrary maps) and instead read the rank loss as the **ideal of an extension**, with the index living in its boundary map:

$0 \to J$ (rank deficiency / kernel) $\to \mathcal{A} \to \mathcal{A}/J$ (unitary part on the core) $\to 0$,
Bott index $= \partial : K_1(\mathcal{A}/J) \to K_0(J)$.

This is the standard "obstruction = connecting map of a short exact sequence" pattern; the non-invertibility is not a defect to be survived but the structural object that *defines* the class. Two facts already in this note say the class lives **downstairs, on the contraction**, not after any lift: the core gap = 1 m-uniformly (Exp. B) is exactly the statement that the support projection P is a clean, gap-separated genuine projection at every m—i.e. $\mathcal{A}/J$ is well-defined uniformly—and Exp. A computes bott(flux) $= -1$ on the partial isometry *directly*, never on a lifted unitary.

So one does not extend Exel–Loring to contractions. One **compresses to** PH, where the partial isometries restrict to honest unitaries, applies the unitary theorem there unchanged, and is left with a single residue: the off-diagonal block $\|(1-P)[\mathcal{U},\mathcal{V}]P\|$—the cheap channel. This is a finite 2×2 Schur-complement estimate, and the core gap is precisely what bounds the off-diagonal coupling. The wall, on this reading, is **not** "extend a deep theorem" but "bound one off-diagonal block uniformly in m."

G4 — the compression leak. [EVIDENCE: E, F] — m-uniform, core-gap-controlled. Exp. E measures $\|(1-P)[\mathcal{U},\mathcal{V}]P\|$ for the rank-$r<d$ partial-isometry family. For the ideal family the leak is **exactly** 0 at every m: the rank loss is *clean*, P is exactly preserved, the compression is lossless by construction. Letting the support frames wobble site-to-site by $O(\varepsilon)$ (so P is no longer exactly preserved), the leak scales like ε with ratio leak/ε staying in the bounded band $[2.1, 4.0]$ across $m = 4, \dots, 32$—**non-monotone** (it dips at $m = 24$ after rising to $m = 16$), so the apparent early drift is wobble-sampling noise, not a trend: no power-law growth, governed by 1/coreGap = 1 (which is 1.0000 to four digits at every m, on operators up to 3072 dimensions at $m = 32$). The index-carrying *kept* part $\|P[\mathcal{U},\mathcal{V}]P\|$ reproduces the G1 commutator $2|\sin(\pi k/m^2)|$ to all digits, confirming the compression keeps the right object.

Exp. E uses *random* frame wobble, which partially cancels across edges; the sharper test is **adversarial**. Since distOpFlat is a sup over edges, an adversary may spend the ε budget on **every one of the m^2 edges at once**, and coherent choices could *accumulate* into one large singular direction growing with m—the only mechanism by which the cheap channel reopens. Exp. F maximizes the leak by projected gradient ascent over per-edge $\mathfrak{so}(d)$ frame generators ($\|A_e\| \le \varepsilon$, matching the sup-metric; analytic top-singular-value gradient, no m^3 commutator formed). The worst-case ratio **saturates at** 2.0000 **at every** $m = 4, \dots, 32$—not merely bounded but *pinned*, m-independent to four digits. The constant 2 is the local "two-operators-in-a-commutator" factor; its m-independence is the direct statement that the per-edge leaks do **not** add coherently across the lattice. (Gradient ascent yields a lower bound on the true worst case; that even the best adversary we can find is m-flat, and that independent random samples reach 1.99, makes coherent accumulation strongly disfavoured.) The Schur estimate is therefore m-uniform in the data **and against a worst-case competitor**: the cheap channel is controlled by the same rank structure that carries the index.

> **What remains, honestly.** This does not close the conjecture. The qualitative obstruction is downstairs already; the *quantitative* $c(k)$ still rests on the dimension-free unitary bound applied on PH—turning "extend a theorem to contractions" into "bound a single block," now demonstrably bounded and m-uniform even against a coherent competitor (see §6). The one caveat that survives: Exp. F's ascent gives a *lower* bound on the worst case, so "pinned at 2" is strong evidence, not proof, of no accumulation.

The correct literature for the remaining estimate is **quantitative / controlled K-theory** and Loring's lifting-of-contractions results—"K-classes through extensions with norm control." One caution: do **not** reach for the gapped-commutator result of Osborne [6]—its gap is the *commutator* gap, which is G1 and **vanishes**; the gap doing the work is the core gap. That is the §5 trap, one level up.

9. Methods and reproducibility

All core numerics are in `coarse-geometry-numerics.py` (`numpy` 1.26.4; `OMP_NUM_THREADS=1`; seconds of runtime). The port mirrors the Lean definitions and is cross-checked against the machine-checked theorems: $\|[\mathcal{U},\mathcal{V}]\|$ matches $2|\sin(\pi k/m^2)|$, and totalFlux reads $2\pi k$ (flux) and 0 (flat, shear). Seven experiments:

- **A** index identification—bott(flux) $= \pm k$, bott(flat) $= 0$ (10^{-16}), branch gap (G3).
- **B** polar Lipschitz / core gap (G2), rank-$r < d$ partial isometries.
- **C** dimension-free distance-to-commuting (the $O(1)$ / Voiculescu).
- **D** magnetic Hamiltonian spectral gap (G1).
- **E** compression leak $\|(1-P)[\mathcal{U},\mathcal{V}]P\|$ (G4)—the §8 cheap-channel / Schur estimate (random wobble).
- **F** adversarial leak—projected gradient ascent maximizing the G4 leak over per-edge frames (worst-case competitor).
- **G** direct distOpFlat—Riemannian descent minimizing $\sup_e\|\text{flux} - T'\|$ over the flat-competitor site frames; the conjectured quantity itself.

Follow-up experiments underlying §4.1 and the geodesic re-run of §4 are standalone scripts reusing the same primitives: the $\mathrm{U}(n)$-geodesic index-jump radius with an integer-flip detector (the §4 correction to Exp. C); the direct distOpFlat sweep at $k = 1, 2, 3$; the compression Schur identity verified to 10^{-16}; the support-slice probe; the codimension sweep; the orthogonal competitor at distance exactly 1; and the floor probe with the k-independence of the codimension-≥ 2 constant.

The machine-checked companion artifacts (`PlateauConjecture.lean`, `MagPair.lean`, `PlaqIndex.lean`, and the rest of the `LossyCocycles` library), together with all scripts and their recorded outputs, are available at `https://github.com/GoodRoyal/LossyCocycles`; the companion papers present the reduction `plateau_of_rigidity` and the formalization narrative.

Appendix A. Data

Exp. A ($k = 1$): bott(flux) $= -1.000000$ and bott(flat) $\sim 10^{-16}$ at every m. **Exp. D** ($k = 1$): magnetic Hamiltonian $H = \mathcal{U} + \mathcal{U}^\dagger + \mathcal{V} + \mathcal{V}^\dagger$; the two share a grid and are tabulated together.

m	branch gap (G3) (Exp. A)	$2\lvert\sin(\pi k/m^2)\rvert$ (G1, Exp. A)	flux/plaquette $2\pi k/m^2$ (Exp. D)	spectral gap (G1, Exp. D)
4	2.7489	0.390181	—	—
6	2.9671	0.174311	0.17453	0.33405
8	3.0434	0.098135	0.09817	0.19157
12	3.0980	0.043630	0.04363	0.08632
16	3.1170	0.024543	0.02454	0.04879
24	3.1307	0.010908	0.01091	0.02176
32	3.1355	0.006136	0.00614	0.01225

Exp. D's gap at $E = 0$ is ≈ 0 throughout (a state sits at the Fermi level); Exp. A's $m = 4$ point has no Exp. D counterpart (Exp. D's grid starts at $m = 6$).

Exp. B ($d = 3$, $r = 2$): $\sigma_{\min}(\text{core}) = 1.0000$ at every m; polar Lipschitz constant (mean): 0.820 (m=4), 0.811 (m=6), 0.811 (m=8), 0.818 (m=12), 0.817 (m=16).

Exp. C ($k = 1$): index-jump operator distance $\|U_t - U_0\| = 1.4142$ at every $m \in \{4, 6, 8, 12, 16\}$; index $= -1.0000$ up to the jump.

Exp. E ($d = 3$, $r = 2$, $k = 1$): compression leak $\|(1 - P)[\mathcal{U}, \mathcal{V}]P\|$. Ideal family: 0 at every m (support exactly preserved). Wobbled frames ($\varepsilon = 10^{-2}$): leak $\propto \varepsilon$, ratio non-monotone in $[2.1, 4.0]$ (operators up to 3072 dimensions at $m = 32$). **Exp. F** ($\varepsilon = 10^{-2}$): adversarial leak—$\max\|(1 - P)[\mathcal{U}, \mathcal{V}]P\|/\varepsilon$ by projected gradient ascent over per-edge frame generators ($\|A_e\| \le \varepsilon$), warm-started from Exp. E's random baseline so adv $\ge$ rand by construction; the two share a grid and are tabulated together.

m	$\sigma_{\min}$(core)	leak (ideal)	$\|P[\mathcal{U}, \mathcal{V}]P\|$	leak (wobble)	Exp. E leak/ε	Exp. F rand. / **adv.** leak/ε
4	1.0000	0	0.3902	2.10×10^{-2}	2.098	1.806 / **2.0000**
6	1.0000	0	0.1743	2.81×10^{-2}	2.813	1.834 / **1.9999**
8	1.0000	0	0.0981	3.24×10^{-2}	3.240	1.993 / **1.9999**
12	1.0000	0	0.0436	3.25×10^{-2}	3.255	1.889 / **1.9999**
16	1.0000	0	0.0245	3.76×10^{-2}	3.765	1.934 / **1.9999**
24	1.0000	0	0.0109	3.18×10^{-2}	3.175	1.912 / **1.9999**
32	1.0000	0	0.0061	3.98×10^{-2}	3.984	1.985 / **1.9999**

*$\|P[\mathcal{U}, \mathcal{V}]P\|$ reproduces the G1 commutator $2|\sin(\pi k/m^2)|$ (Exp. A)—the kept part is the index-carrying object; the leak is the residue, core-gap-controlled and $O(1)$ in m (the $m = 24$ dip below $m = 16$ shows the early rise was sampling noise, not a trend). Exp. F's worst-case ratio is **pinned at** 2.0 **independent of** m—the per-edge leaks do not add coherently across the m^2 edges. The constant is the local two-operators-in-a-commutator factor; ascent gives a lower bound on the worst case, so this is strong evidence, not proof, of no accumulation. ($m = 24, 32$ run at reduced budget.)*

Exp. G ($d = 3$, $r = 2$, $k = 1$): direct distOpFlat—$\min_{T' \text{ flat}} \sup_e \|\text{flux}(e) - T'(e)\|$ by Riemannian descent over site-frame Stiefel manifolds (300 steps, 3 random restarts plus the $V{=}S$ initialization). The $V{=}S$ baseline is 2.0000 at every m (the worst wrap edge, $\varphi \approx \pi$).

m	$V{=}S$ baseline	best competitor ($\le$ distOpFlat)
4	2.0000	1.2942
6	2.0000	1.3830
8	2.0000	1.3968
12	2.0000	1.4065
16	2.0000	1.4099
24	2.0000	1.4123
32	2.0000	1.4131

Monotone convergence up to $\sqrt{2} = 1.41421$ (within 10^{-3} at $m = 32$); no decay. This is an upper bound on distOpFlat*, so it can falsify but not confirm—yet the best competitor a genuine minimization can find cannot collapse the distance, and lands on the same $\sqrt{2}$ as the independent index-jump radius (Exp. C). The binding $\inf_m$ is at $m = 4$ (≈ 1.29), still $\gg 0$.*

References

[1] D. Voiculescu, *Asymptotically commuting finite rank unitary operators without commuting approximants*, Acta Sci. Math. (Szeged) **45** (1983), no. 1–4, 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] M. B. Hastings, *Making almost commuting matrices commute*, Comm. Math. Phys. **291** (2009), 321–345. arXiv:0808.2474.

[6] T. J. Osborne, *Almost commuting unitaries with spectral gap are near commuting unitaries*, arXiv:0809.0602 (2008). [The gapped case—directly relevant to the contraction question.]

[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. A. Loring, *K-theory and pseudospectra for topological insulators*, Ann. Physics **356** (2015), 383–416. arXiv:1502.03498.

[10] T. A. Loring, H. Schulz-Baldes, *Finite volume calculation of K-theory invariants*, New York J. Math. **23** (2017), 1111–1140.

[11] T. A. Loring, H. Schulz-Baldes, *The spectral localizer for even index pairings*, J. Noncommut. Geom. **14** (2020). arXiv:1802.04517.

[12] E. Lozano Viesca, J. Schober, H. Schulz-Baldes, *Chern numbers as half-signature of the spectral localizer*, J. Math. Phys. **60** (2019), 072101. arXiv:1907.11382.

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

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, drafting and revising prose, and retrieving literature, all under his direction and review. Formal results are machine-checked (Lean 4) where stated; numerical claims are reproducible from the released scripts. Juan Carlos Paredes is the author of this work; the assistance is disclosed in the spirit of the full transparency this venue requires.

Independent researcher
Email address: cpt66778811@gmail.com