Universal Even–Odd Coherent-Error Parity Dichotomy, Swap-Point Criticality, and Exact-$\mathbb{Q}$ Decoding Certificates for CSS Quantum Codes

Lluis Eriksson*

November 2026

Abstract

Uniform coherent Z-rotations $U(\theta) = \bigotimes_{j=1}^{n} e^{i\theta Z_j}$ induce non-Pauli logical channels whose optimal recovery and finite-size scaling differ sharply from Pauli-twirled stochastic phase-flip noise. Extending the recent surface-code parity analysis of Dai, Wan, Dai, and Zhu (`arXiv:2611.13098`, 2026), we establish three general structural theorems for arbitrary $[[n, 1, d]]$ Calderbank–Shor–Steane (CSS) codes, formalize their algebraic core in zero-dependency Lean 4 (`v4.35.0`, 18 theorems, 0 `sorry`, 0 `Classical.choice`), and construct an exact-$\mathbb{Q}$ certificate verifier (`COHERENT-CSS-CERT-1.0`) free of floating-point trust.

First, for *every* $[[n, 1, d]]$ CSS code whose Z-stabilizers have even Hamming weight, we prove a **Universal Even–Odd Parity Dichotomy**: when $|L_Z|$ is odd, every syndrome Kraus operator is proportional to a pure logical Z_L-rotation, so syndrome-conditioned unitary recovery achieves exact unit entanglement fidelity $F_e^{\mathrm{CPTP}}(\theta) \equiv 1$ for all $\theta \in [0, \pi/2]$; when $|L_Z|$ is even, every syndrome Kraus operator is Hermitian up to phase, and the exact trace-norm identity $(|a+b| + |a-b|)^2 = 4\max(a^2, b^2)$ forces maximum-likelihood Pauli decoding to attain the exact global CPTP optimum $F_e^{\mathrm{Pauli}}(\theta) \equiv F_e^{\mathrm{CPTP}}(\theta)$ over all quantum channels. Second, using exact $\mathbb{F}_2$ harmonic analysis on the self-orthogonal subcode $\mathcal{K}_Z = \mathcal{S}_Z \cap \mathcal{S}_Z^{\perp}$, we prove a **Complete** 0–1 **Swap-Point Binary Law**: at $\theta = \pi/4$, $F_e^{\mathrm{Pauli}}(\pi/4)$ can take *only* two values across the entire universe of $[[n, 1, d]]$ CSS codes—$F_e^{\mathrm{Pauli}}(\pi/4) = 1/2$ if and only if $L_Z \perp \mathcal{K}_Z$ (equivalently $(L_X + \mathcal{S}_X) \cap \mathcal{S}_Z = \emptyset$), and $F_e^{\mathrm{Pauli}}(\pi/4) = 1$ if and only if there exists $g^* \in \mathcal{K}_Z$ with $|L_Z \cap g^*| \equiv 1 \pmod 2$. Third, as an immediate topological corollary, rotating an *even-d* rotated surface code patch by 90° (X-majority vs. Z-majority check orientation) flips its $\theta = \pi/4$ logical fidelity from complete collapse (1/2) to complete coherent immunity (1), witnessed explicitly by the left-boundary column stabilizer $g^* = \prod_{j=0}^{d/2-1} Z_{(2j,0)} Z_{(2j+1,0)} \in \mathcal{K}_Z$. Exact-$\mathbb{Q}$ polynomial evaluation across 9 certified benchmark codes and a 129-code census (matching complex 2×2 SVD to 1.45×10^{-14} and rejecting 37/37 adversarial forgery attacks) uncovers exact all-angle coherent–twirled cancellation on X-majority $d = 2$ surface codes, 113.26× coherent infidelity suppression on Shor $[[9, 1, 3]]$ at $p = \sin^2\theta = 1/5$ (reaching exact zero infidelity $F_e^{\mathrm{Pauli}}(1/4) = 1$ at $\theta = \pi/6$), and a strict coherent *penalty* ($F_e^{\mathrm{Pauli}} < F_e^{\mathrm{twirl}}$) on doubly-even Quantum Reed–Muller $[[15, 1, 3]]$.

1 Introduction

Coherent single-qubit over-rotations $U(\theta) = \bigotimes_{j=1}^{n} \exp(i\theta Z_j)$ arise ubiquitously in trapped-ion, neutral-atom, and superconducting quantum processors from systematic control miscalibration and static crosstalk. Unlike stochastic Pauli-twirled dephasing, coherent rotations interfere across stabilizer-equivalent error cosets inside each syndrome sector [1, 2, 3, 4, 5]. While Majorana-fermion and statistical-mechanics mappings have illuminated the large-distance coherent threshold of planar

*Independent researcher (Sweden). Prepared with AI assistance (Google Antigravity). Code, Lean 4 proofs, and exact-$\mathbb{Q}$ certificates: https://github.com/lluiseriksson/coherent-css-dichotomy (MIT license).

surface codes [6, 7, 8], a striking recent advance by Dai, Wan, Dai, and Zhu [9] revealed that finite-size rotated surface codes exhibit a sharp *even–odd distance parity dichotomy*: for odd distance d, non-Pauli syndrome-conditioned Z_L-rotations strictly outperform Pauli decoding (indeed achieving unit fidelity under pure uniform Z-rotations), whereas for even distance d, maximum-likelihood Pauli decoding is already optimal over all completely positive trace-preserving (CPTP) recovery channels.

This discovery immediately raises four fundamental mathematical questions:

1. **Universality beyond surface codes:** Is the even–odd CPTP-vs.-Pauli optimality dichotomy specific to planar geometry, or is it an exact algebraic consequence of binary weight parity in *any* $[[n, 1, d]]$ CSS code [10, 11]?
2. **Exact classification at maximal coherence ($\theta = \pi/4$):** At the swap point $\theta = \pi/4$ ($p = \sin^2\theta = 1/2$), Pauli-twirled decoding collapses to $F_e^{\text{twirl}}(\pi/4) = 1/2$. Can coherent Pauli decoding achieve intermediate values in $(1/2, 1)$, or is it rigidly quantized?
3. **Boundary orientation of even-d surface codes:** For even distance d, $r_X + r_Z = d^2 - 1$ is odd, breaking $X \leftrightarrow Z$ self-duality. How does a 90° patch rotation ($r_X = d^2/2$ vs. $r_Z = d^2/2$) affect coherent decoding?
4. **Exact machine-verified certification:** Can coherent and CPTP entanglement fidelities be evaluated and audited in exact rational arithmetic $\mathbb{Q}$ with zero floating-point trust and a machine-checked Lean 4 proof kernel [12]?

We resolve all four questions completely. Our main contributions are:

- **Theorem 2.1 (Universal Even–Odd Parity Dichotomy):** Every $[[n, 1, d]]$ CSS code with even-weight Z-stabilizers satisfies $F_e^{\text{CPTP}}(\theta) \equiv 1$ for all θ when $|L_Z|$ is odd, and $F_e^{\text{CPTP}}(\theta) \equiv F_e^{\text{Pauli}}(\theta)$ for all θ when $|L_Z|$ is even.
- **Theorem 2.2 (Exact 0–1 Swap-Point Binary Law):** Via $\mathbb{F}_2$ harmonic analysis on $\mathcal{K}_Z = \mathcal{S}_Z \cap \mathcal{S}_Z^{\perp}$, *every* $[[n, 1, d]]$ CSS code satisfies $F_e^{\text{Pauli}}(\pi/4) \in \{1/2, 1\}$ with zero intermediate values possible, governed by the exact orthogonality criterion $L_Z \perp \mathcal{K}_Z \iff (L_X + \mathcal{S}_X) \cap \mathcal{S}_Z = \emptyset$.
- **Corollary 2.3 (Even-d Surface-Code 90° Orientation Phase Transition):** On any even-d rotated surface code, rotating the patch by 90° switches between X-majority collapse ($F_e^{\text{Pauli}}(\pi/4) = F_e^{\text{CPTP}}(\pi/4) = 1/2$) and Z-majority coherent immunity ($F_e^{\text{Pauli}}(\pi/4) = F_e^{\text{CPTP}}(\pi/4) = 1$), certified by an explicit weight-d left-boundary column stabilizer $g^* \in \mathcal{K}_Z$.
- **Theorem 2.4, Proposition 2.5, and Formal Verification:** For every rational $p = \sin^2\theta \in \mathbb{Q} \cap [0, 1]$, all three fidelities lie in $\mathbb{Q}$ and are independently audited by our zero-float `COHERENT-CSS-CERT-1.0` verifier alongside 18 zero-dependency Lean 4 theorems (0 `sorry`, 0 `Classical.choice`). Exact computation across 9 benchmark codes and 120 random even-$\mathcal{S}_Z$ CSS codes confirms every law, uncovers exact all-angle coherent–twirled cancellation on X-majority $d = 2$ surface codes, quantifies 113.26× coherent infidelity suppression on Shor $[[9, 1, 3]]$ at $p = 1/5$, and reveals a strict coherent *penalty* on doubly-even Quantum Reed–Muller $[[15, 1, 3]]$.

2 Mathematical Framework and Main Theorems

2.1 CSS Syndrome Kraus Decomposition

Let $\mathcal{C} = [[n, 1, d]]$ be a binary CSS quantum code defined by full-row-rank parity-check matrices $H_X \in \mathbb{F}_2^{r_X \times n}$ and $H_Z \in \mathbb{F}_2^{r_Z \times n}$ satisfying $H_X H_Z^{\top} = 0 \pmod 2$ and $r_X + r_Z = n - 1$, with stabilizer subgroups $\mathcal{S}_X = \text{rowspace}_{\mathbb{F}_2}(H_X)$ and $\mathcal{S}_Z = \text{rowspace}_{\mathbb{F}_2}(H_Z)$, and logical operators $L_X \in \mathcal{S}_Z^{\perp} \setminus \mathcal{S}_X$, $L_Z \in \mathcal{S}_X^{\perp} \setminus \mathcal{S}_Z$ obeying $L_X \cdot L_Z \equiv 1 \pmod 2$.

Under the uniform coherent Z-rotation

$$U(\theta) = \bigotimes_{j=1}^{n} (\cos\theta\, I + i \sin\theta\, Z_j) = \sum_{e \in \mathbb{F}_2^n} (\cos\theta)^{n-|e|} (i \sin\theta)^{|e|} Z^e, \tag{1}$$

measuring all r_X X-stabilizer generators yields a syndrome $s \in \mathbb{F}_2^{r_X}$ with canonical minimum-integer representative $e_s \in \mathbb{F}_2^n$ ($H_X e_s = s$). The syndrome fiber partitions into two $\mathcal{S}_Z$-cosets $\mathcal{C}_s^{(0)} = e_s \oplus \mathcal{S}_Z$ and $\mathcal{C}_s^{(1)} = e_s \oplus L_Z \oplus \mathcal{S}_Z$ with weight enumerators $A_s(w) = |\{e \in \mathcal{C}_s^{(0)} : \mathrm{wt}(e) = w\}|$ and $B_s(w) = |\{e \in \mathcal{C}_s^{(1)} : \mathrm{wt}(e) = w\}|$. Applying the pure-$Z$ destabilizer Z^{e_s} projects the post-measurement logical evolution onto the single-qubit Kraus operator

$$K_s(\theta) = \alpha_s(\theta)\, I_L + \beta_s(\theta)\, Z_L, \qquad \alpha_s(\theta) = \sum_{w=0}^{n} A_s(w)\, c^{n-w} (it)^w, \quad \beta_s(\theta) = \sum_{w=0}^{n} B_s(w)\, c^{n-w} (it)^w, \tag{2}$$

where $c = \cos\theta$, $t = \sin\theta$, and $\sum_{s \in \mathbb{F}_2^{r_X}} (|\alpha_s(\theta)|^2 + |\beta_s(\theta)|^2) = 1$.

For any syndrome-conditioned recovery channel $\mathcal{R} = \{\mathcal{R}_s\}_{s \in \mathbb{F}_2^{r_X}}$, the Schumacher entanglement fidelity [13, 14] of $\sum_s \mathcal{R}_s(K_s \cdot K_s^\dagger)$ decomposes additively over syndromes. Maximizing over Pauli recoveries $\{I_L, Z_L\}$ yields

$$F_e^{\mathrm{Pauli}}(\theta) = \sum_{s \in \mathbb{F}_2^{r_X}} \max(|\alpha_s(\theta)|^2,\, |\beta_s(\theta)|^2), \tag{3}$$

while maximizing over *all* trace-preserving quantum channels $\mathcal{R}_s$ on a single logical qubit yields the exact trace-norm optimum (attained by the polar-decomposition unitary $V_s^\dagger$ of K_s):

$$F_e^{\mathrm{CPTP}}(\theta) = \frac{1}{4} \sum_{s \in \mathbb{F}_2^{r_X}} \|K_s(\theta)\|_1^2 = \frac{1}{4} \sum_{s \in \mathbb{F}_2^{r_X}} \Big(|\alpha_s(\theta) + \beta_s(\theta)| + |\alpha_s(\theta) - \beta_s(\theta)| \Big)^2. \tag{4}$$

For comparison, Pauli-twirling $U(\theta)$ produces i.i.d. phase-flip noise at rate $p = \sin^2\theta$ ($q = 1 - p$), with coset probabilities $p_s^{(0),\mathrm{twirl}}(p) = \sum_w A_s(w) q^{n-w} p^w$, $p_s^{(1),\mathrm{twirl}}(p) = \sum_w B_s(w) q^{n-w} p^w$, and optimal Pauli-twirled fidelity $F_e^{\mathrm{twirl}}(p) = \sum_s \max(p_s^{(0),\mathrm{twirl}}(p), p_s^{(1),\mathrm{twirl}}(p))$.

2.2 Theorem 2.1: Universal Even–Odd Coherent-Error Parity Dichotomy

Theorem 2.1 (Universal Even–Odd Coherent-Error Parity Dichotomy)**.** *Let $\mathcal{C} = [[n, 1, d]]$ be any CSS code such that every generator of H_Z (hence every $g \in \mathcal{S}_Z$) has even Hamming weight. Under uniform coherent Z-rotations $U(\theta) = \bigotimes_{j=1}^n e^{i\theta Z_j}$, for every syndrome $s \in \mathbb{F}_2^{r_X}$ and every $\theta \in [0, \pi/2]$:*

(a) ***Odd-$|L_Z|$ Unitary Branch:*** *If $\mathrm{wt}(L_Z) \equiv 1 \pmod 2$, then $i^{-|e_s|}\alpha_s(\theta) \in \mathbb{R}$ and $i^{-|e_s|}\beta_s(\theta) \in i\mathbb{R}$. Consequently $K_s(\theta) = i^{|e_s|}\sqrt{p_s(\theta)}\, e^{i\phi_s(\theta) Z_L}$ is proportional to a pure logical Z_L-rotation, so syndrome-conditioned unitary recovery achieves exact unit fidelity*

$$F_e^{\mathrm{CPTP}}(\theta) = F_e^{\mathrm{unitary}}(\theta) \equiv 1 \quad \forall \theta \in [0, \pi/2], \tag{5}$$

while Pauli-only decoding suffers exact infidelity $1 - F_e^{\mathrm{Pauli}}(\theta) = \sum_s \min(|\alpha_s|^2, |\beta_s|^2)$.

(b) ***Even-$|L_Z|$ Hermitian Branch:*** *If $\mathrm{wt}(L_Z) \equiv 0 \pmod 2$, then both $\tilde{\alpha}_s(\theta) = i^{-|e_s|}\alpha_s(\theta) \in \mathbb{R}$ and $\tilde{\beta}_s(\theta) = i^{-|e_s|}\beta_s(\theta) \in \mathbb{R}$ are real, so $K_s(\theta) = i^{|e_s|}(\tilde{\alpha}_s I_L + \tilde{\beta}_s Z_L)$ is Hermitian up to scalar phase. Consequently maximum-likelihood Pauli decoding achieves the exact global CPTP optimum over all quantum channels:*

$$F_e^{\mathrm{Pauli}}(\theta) \equiv F_e^{\mathrm{CPTP}}(\theta) = \sum_{s \in \mathbb{F}_2^{r_X}} \max(|\alpha_s(\theta)|^2,\, |\beta_s(\theta)|^2) \quad \forall \theta \in [0, \pi/2]. \tag{6}$$

Proof. By the exact Hamming-weight XOR identity $\mathrm{wt}(e\oplus g)+2\,\mathrm{wt}(e\cap g)=\mathrm{wt}(e)+\mathrm{wt}(g)$ (Lean 4 Theorem 1), whenever $\mathrm{wt}(g)\equiv 0 \pmod 2$ we have $\mathrm{wt}(e\oplus g)\equiv \mathrm{wt}(e) \pmod 2$ (Lean Theorem 3). Thus every $e\in\mathcal{C}_s^{(0)}=e_s\oplus\mathcal{S}_Z$ has $\mathrm{wt}(e)-|e_s|\in 2\mathbb{Z}$, so $i^{\mathrm{wt}(e)-|e_s|}=(-1)^{(\mathrm{wt}(e)-|e_s|)/2}\in\{\pm 1\}$ and $\tilde\alpha_s(\theta)=i^{-|e_s|}\alpha_s(\theta)\in\mathbb{R}$.

In case (a), $\mathrm{wt}(L_Z)\equiv 1 \pmod 2$ implies $\mathrm{wt}(e_s\oplus L_Z\oplus g)\equiv |e_s|+1 \pmod 2$ for every $g\in\mathcal{S}_Z$ (Lean Theorem 5), so $\tilde\beta_s(\theta)=ib_s(\theta)$ with $b_s(\theta)\in\mathbb{R}$. The two diagonal entries of $i^{-|e_s|}K_s(\theta)=\mathrm{diag}(\tilde\alpha_s+ib_s,\,\tilde\alpha_s-ib_s)$ have identical moduli $\sqrt{\tilde\alpha_s^2+b_s^2}=\sqrt{|\alpha_s|^2+|\beta_s|^2}$ (Lean Theorem 7), whence $\frac14\|K_s\|_1^2=|\alpha_s|^2+|\beta_s|^2=p_s(\theta)$ and summing over s gives $F_e^{\mathrm{CPTP}}(\theta)=1$. The Pauli infidelity formula follows from $\max(x,y)+\min(x,y)=x+y$ (Lean Theorem 8).

In case (b), $\mathrm{wt}(L_Z)\equiv 0 \pmod 2$ implies $\mathrm{wt}(e_s\oplus L_Z\oplus g)\equiv |e_s| \pmod 2$ (Lean Theorem 4), so both $a=\tilde\alpha_s\in\mathbb{R}$ and $b=\tilde\beta_s\in\mathbb{R}$ are real. Expanding the four sign quadrants of $(a+b,a-b)$ gives the exact identity

$$(|a+b|+|a-b|)^2=4\max(a^2,b^2)\qquad\forall a,b\in\mathbb{R} \tag{7}$$

(machine-checked over $\mathbb{Z}$ in Lean Theorem 6, `hermitian_trace_norm_eq_max`, and valid over $\mathbb{R}$ by homogeneity). Substituting (7) into (4) yields $\frac14\|K_s(\theta)\|_1^2=\max(\tilde\alpha_s^2,\tilde\beta_s^2)=\max(|\alpha_s|^2,|\beta_s|^2)$ on every syndrome s, proving $F_e^{\mathrm{CPTP}}(\theta)=F_e^{\mathrm{Pauli}}(\theta)$ identically. □

2.3 Theorem 2.2: Exact 0–1 Swap-Point Binary Law at $\theta=\pi/4$

Let $\mathcal{K}_Z:=\mathcal{S}_Z\cap\mathcal{S}_Z^\perp=\{g\in\mathcal{S}_Z : g\cdot h\equiv 0 \pmod 2\ \forall h\in\mathcal{S}_Z\}$ denote the self-orthogonal subcode of $\mathcal{S}_Z$. Every $g\in\mathcal{K}_Z$ satisfies $\mathrm{wt}(g)=g\cdot g\equiv 0 \pmod 2$.

Theorem 2.2 (Complete 0–1 Swap-Point Binary Law at $\theta=\pi/4$)**.** *Let $\mathcal{C}=[[n,1,d]]$ be any CSS code (not necessarily even-weight $\mathcal{S}_Z$). At the maximal-coherence angle $\theta=\pi/4$ ($p=1/2$), the optimal Pauli entanglement fidelity $F_e^{\mathrm{Pauli}}(\pi/4)$ takes* only *two possible values, $\{1/2,1\}$, with no intermediate values in $(1/2,1)$:*

(a) ***Swap-Point Collapse*** *($F_e^{\mathrm{Pauli}}(\pi/4)=1/2$): Holds if and only if $L_Z\perp\mathcal{K}_Z$ (equivalently $L_Z\in\mathcal{S}_Z+\mathcal{S}_Z^\perp$, equivalently $(L_X+\mathcal{S}_X)\cap\mathcal{S}_Z=\emptyset$). In this regime, $|\alpha_s(\pi/4)|^2=|\beta_s(\pi/4)|^2$ on 100% of syndromes $s\in\mathbb{F}_2^{r_X}$, forcing $F_e^{\mathrm{Pauli}}(\pi/4)=1/2$ (and if $\mathcal{S}_Z$ is even-weight and $|L_Z|$ is even, $F_e^{\mathrm{CPTP}}(\pi/4)=1/2$).*

(b) ***Complete Coherent Immunity*** *($F_e^{\mathrm{Pauli}}(\pi/4)=1$): Holds if and only if $L_Z\not\perp\mathcal{K}_Z$, i.e., there exists a witness $g^*\in\mathcal{K}_Z$ with $|L_Z\cap g^*|\equiv 1 \pmod 2$ (equivalently $g^*\in(L_X+\mathcal{S}_X)\cap\mathcal{S}_Z$). In this regime, $\alpha_s(\pi/4)\,\beta_s(\pi/4)=0$ on 100% of syndromes $s\in\mathbb{F}_2^{r_X}$, forcing $F_e^{\mathrm{Pauli}}(\pi/4)=F_e^{\mathrm{CPTP}}(\pi/4)=1$.*

Proof. At $\theta=\pi/4$, $\alpha_s(\pi/4)=2^{-n/2}\sum_{e\in\mathcal{C}_s^{(0)}}i^{\mathrm{wt}(e)}$. Substituting $e_1=e\oplus g$ ($g\in\mathcal{S}_Z$) and using $\mathrm{wt}(e\oplus g)-\mathrm{wt}(e)=\mathrm{wt}(g)-2\,\mathrm{wt}(e\cap g)$ (Lean Theorem 1) gives $i^{\mathrm{wt}(e\oplus g)-\mathrm{wt}(e)}=i^{\mathrm{wt}(g)}(-1)^{\mathrm{wt}(e\cap g)}$. Writing $e=e_s\oplus h$ ($h\in\mathcal{S}_Z$) yields

$$2^n|\alpha_s(\pi/4)|^2=\sum_{g\in\mathcal{S}_Z}i^{\mathrm{wt}(g)}(-1)^{\mathrm{wt}(e_s\cap g)}\sum_{h\in\mathcal{S}_Z}(-1)^{\mathrm{wt}(h\cap g)}=|\mathcal{S}_Z|\sum_{g\in\mathcal{K}_Z}(-1)^{\mathrm{wt}(g)/2+\mathrm{wt}(e_s\cap g)}, \tag{8}$$

because $\sum_{h\in\mathcal{S}_Z}(-1)^{h\cdot g}=|\mathcal{S}_Z|\,\mathbf{1}_{g\in\mathcal{S}_Z^\perp}$ by $\mathbb{F}_2$-character orthogonality, and $g\in\mathcal{K}_Z\implies\mathrm{wt}(g)$ is even so $i^{\mathrm{wt}(g)}=(-1)^{\mathrm{wt}(g)/2}$. Replacing e_s by $e_s\oplus L_Z$ yields identically

$$2^n|\beta_s(\pi/4)|^2=|\mathcal{S}_Z|\sum_{g\in\mathcal{K}_Z}(-1)^{\mathrm{wt}(g)/2+\mathrm{wt}(e_s\cap g)+\mathrm{wt}(L_Z\cap g)}. \tag{9}$$

Now observe the key linearity lemma: for any $g_1,g_2\in\mathcal{K}_Z$, $\mathrm{wt}(g_1\cap g_2)\equiv 0 \pmod 2$, so $\mathrm{wt}(g_1\oplus g_2)\equiv\mathrm{wt}(g_1)+\mathrm{wt}(g_2) \pmod 4$ (Lean Theorem 10, `self_orthog_wt_mod4_additive`). Therefore $\chi_s(g):=(-1)^{\mathrm{wt}(g)/2+\mathrm{wt}(e_s\cap g)}$ and $\psi(g):=(-1)^{\mathrm{wt}(L_Z\cap g)}$ are both *linear* $\mathbb{F}_2$-characters on the group $\mathcal{K}_Z$!

If $L_Z \perp \mathcal{K}_Z$, then $\psi \equiv 1$ on $\mathcal{K}_Z$, so (8) and (9) coincide identically on every syndrome s, and Lean Theorem 11 (`equal_coset_gives_half_fidelity`) forces $2F_e^{\text{Pauli}}(\pi/4) = \sum_s(|\alpha_s|^2 + |\beta_s|^2) = 1$, i.e. $F_e^{\text{Pauli}}(\pi/4) = 1/2$.

Conversely, if $\exists g^* \in \mathcal{K}_Z$ with $|L_Z \cap g^*| \equiv 1 \pmod 2$, then ψ is a non-trivial character on $\mathcal{K}_Z$. Consequently the two characters χ_s and $\chi_s \cdot \psi$ are distinct for every syndrome s, so they can *never* both be the trivial character on $\mathcal{K}_Z$. Since the sum of any non-trivial character over $\mathcal{K}_Z$ vanishes identically, at least one of $|\alpha_s(\pi/4)|^2$ and $|\beta_s(\pi/4)|^2$ is zero on every syndrome s. Lean Theorem 12 (`vanishing_coset_gives_unit_fidelity`) then forces $F_e^{\text{Pauli}}(\pi/4) = \sum_s(|\alpha_s|^2 + |\beta_s|^2) = 1$, whence $1 = F_e^{\text{Pauli}}(\pi/4) \le F_e^{\text{CPTP}}(\pi/4) \le 1$. Finally, $g^* \in \mathcal{K}_Z$ with $g^* \cdot L_Z = 1$ means X^{g^*} commutes with $\mathcal{S}_Z$ and anticommutes with L_Z, so $g^* \in (L_X + \mathcal{S}_X) \cap \mathcal{S}_Z$, and conversely. □

2.4 Corollary 2.3: Even-d Rotated Surface Code 90° Orientation Phase Transition

For an odd-distance rotated surface code $[[d^2, 1, d]]$, $r_X = r_Z = (d^2-1)/2$ and 90° patch rotation is an $X \leftrightarrow Z$ symmetry; since $|L_Z| = d$ is odd and $\mathcal{K}_Z = \{0\}$ (or $L_Z \perp \mathcal{K}_Z$), $F_e^{\text{CPTP}}(\theta) \equiv 1$ while $F_e^{\text{Pauli}}(\pi/4) = 1/2$. For *even* distance $d \ge 2$, however, $r_X + r_Z = d^2 - 1$ is *odd*, splitting the rotated surface code into two dual orientations under pure coherent Z-rotations:

- **X-majority orientation** ($r_X = d^2/2$, $r_Z = d^2/2 - 1$): $d/2 - 1$ weight-2 Z-checks lie along each vertical boundary.
- **Z-majority orientation** ($r_Z = d^2/2$, $r_X = d^2/2 - 1$): $d/2$ disjoint weight-2 Z-checks tile the entire left boundary column $\{(r, 0) : 0 \le r < d\}$ (and right boundary column), obtained from the X-majority patch by a 90° rotation (swapping $H_X \leftrightarrow H_Z$, $L_X \leftrightarrow L_Z$).

Corollary 2.3 (Even-d Rotated Surface Code 90° Orientation Phase Transition)**.** *For every even $d \ge 2$, both orientations of the $[[d^2, 1, d]]$ rotated surface code satisfy $F_e^{\text{CPTP}}(\theta) \equiv F_e^{\text{Pauli}}(\theta)$ for all θ (Theorem 2.1(b)). At $\theta = \pi/4$, however, rotating the patch by 90° induces a sharp $1/2 \leftrightarrow 1$ phase transition:*

*(a) In the **Z-majority orientation** ($r_Z = d^2/2$), the product of the $d/2$ disjoint left-boundary weight-2 Z-checks is the full left column $g^* = \prod_{j=0}^{d/2-1} Z_{(2j,0)} Z_{(2j+1,0)} = Z^{\text{supp}(X_L)} \in (L_X + \mathcal{S}_X) \cap \mathcal{S}_Z$. Hence $g^* \in \mathcal{K}_Z$ with $|L_Z \cap g^*| = |\{(0,0)\}| = 1$, forcing **complete coherent immunity** $F_e^{\text{Pauli}}(\pi/4) = F_e^{\text{CPTP}}(\pi/4) = 1$.*

*(b) In the **X-majority orientation** ($r_X = d^2/2$), every Z-stabilizer generator has even intersection with the left-column logical-Z string $L_Z = \{(r, 0) : 0 \le r < d\}$, so $L_Z \perp \mathcal{S}_Z$ (hence $L_Z \perp \mathcal{K}_Z$) and the code suffers **complete logical collapse** $F_e^{\text{Pauli}}(\pi/4) = F_e^{\text{CPTP}}(\pi/4) = 1/2$.*

Proof. In the Z-majority orientation on the $d \times d$ grid (d even), the left boundary carries weight-2 Z-stabilizers on pairs $\{(2j, 0), (2j+1, 0)\}$ for $j = 0, \ldots, d/2 - 1$. Their product is the weight-d left column $g^* = \{(r, 0) : 0 \le r < d\} \in \mathcal{S}_Z$. Every bulk Z-plaquette meets column 0 in either 0 or 2 qubits, every left-boundary Z-check meets it in 2 qubits, and every right-boundary Z-check meets it in 0 qubits; hence $g^* \in \mathcal{S}_Z^\perp$, so $g^* \in \mathcal{K}_Z = \mathcal{S}_Z \cap \mathcal{S}_Z^\perp$. The top-row logical-$Z$ string $L_Z = \{(0, c) : 0 \le c < d\}$ intersects g^* exclusively at corner $(0, 0)$, giving $|L_Z \cap g^*| = 1 \equiv 1 \pmod 2$. Theorem 2.2(b) immediately yields $F_e^{\text{Pauli}}(\pi/4) = F_e^{\text{CPTP}}(\pi/4) = 1$.

Conversely, in the X-majority orientation (obtained by swapping $H_X \leftrightarrow H_Z$ and $L_X \leftrightarrow L_Z$), L_Z is the left column $\{(r, 0) : 0 \le r < d\}$, bulk Z-plaquettes sit on (r, c) with $r + c$ even, and top/bottom boundary weight-2 Z-checks sit on $\{(0, c), (0, c+1)\}$ and $\{(d-1, c), (d-1, c+1)\}$ exclusively at *odd* column indices $c \in \{1, 3, \ldots, d-3\}$ (since both $-1 + c$ and $(d-1) + c$ are even iff c is odd for even d). Consequently, every top and bottom boundary Z-check is disjoint from column 0 (0 overlap), while every bulk Z-plaquette intersects column 0 in either 2 qubits ($c = 0$) or 0 qubits ($c \ge 1$). Thus every generator of H_Z—and hence every $g \in \mathcal{S}_Z \supseteq \mathcal{K}_Z$—has even overlap with L_Z, proving $L_Z \perp \mathcal{K}_Z$ and forcing $F_e^{\text{Pauli}}(\pi/4) = F_e^{\text{CPTP}}(\pi/4) = 1/2$ by Theorem 2.2(a) for every even $d \ge 2$. □

2.5 Exact-$\mathbb{Q}$ Polynomial Representation and Special Exact Phenomena

Theorem 2.4 (Exact-$\mathbb{Q}$ Polynomial Evaluation over $\mathbb{Q}[p]$)**.** *Let $\mathcal{C} = [[n, 1, d]]$ be any even-Z-stabilizer CSS code and let $p = \sin^2\theta \in \mathbb{Q} \cap [0, 1]$, $q = 1 - p \in \mathbb{Q}$. Then for every syndrome $s \in \mathbb{F}_2^{r_X}$, all weights in $\mathrm{supp}(A_s)$ share the same parity (likewise for B_s), and the squared coset amplitudes are exact rational polynomials in (p, q):*

$$|\alpha_s(p)|^2 = \sum_{w_1,w_2} A_s(w_1)A_s(w_2)\,(-1)^{(w_1-w_2)/2}\,q^{\,n-(w_1+w_2)/2}\,p^{\,(w_1+w_2)/2} \in \mathbb{Q}, \tag{10}$$

with the analogous formula for $|\beta_s(p)|^2$ via B_s. Consequently $F_e^{\mathrm{Pauli}}(p)$, $F_e^{\mathrm{CPTP}}(p)$, and $F_e^{\mathrm{twirl}}(p)$ are exact rational numbers computable without floating-point arithmetic, trigonometry, or square roots.

Proof. Multiplying $\alpha_s(\theta) = \sum_{w_1} A_s(w_1)c^{n-w_1}(it)^{w_1}$ by $\overline{\alpha_s(\theta)} = \sum_{w_2} A_s(w_2)c^{n-w_2}(-it)^{w_2}$ gives $i^{w_1}(-i)^{w_2} = i^{w_1-w_2} = (-1)^{(w_1-w_2)/2}$ and $c^{2n-w_1-w_2}t^{w_1+w_2} = q^{n-(w_1+w_2)/2}p^{(w_1+w_2)/2}$, where $(w_1 \pm w_2)/2 \in \mathbb{Z}$ because all weights in $\mathcal{C}_s^{(0)}$ share parity by Theorem 2.1. Rationality of $F_e^{\mathrm{Pauli}}(p)$ and $F_e^{\mathrm{CPTP}}(p)$ then follows from (3) and Theorem 2.1. □

Proposition 2.5 (Doubly-Even Constructive Interference and Exact $d = 2$ Cancellation)**.**

(a) ***Steane*** *$[[7, 1, 3]]* ***Unit Trivial-Syndrome Law:*** *Since $\mathcal{S}_Z$ is doubly-even with weight enumerator $\{0^1, 4^7\}$ and $L_Z \oplus \mathcal{S}_Z$ has $\{3^7, 7^1\}$, all cross-terms carry phase $+1$ (Lean Theorem 14), giving $p_0(\pi/4) = 2^{-7}((1+7)^2 + (7+1)^2) = 1$ identically (Lean Theorem 15; transversal logical $S^\dagger$ gate).*

(b) ***Exact*** *$d = 2$ X**-Majority Coherent–Twirled Cancellation:*** *For `rotsurf_d2_Xmaj` ($[[4, 1, 2]]$, $H_Z = [1111]$, $H_X = [1010, 0101]$), the four syndromes yield coherent vs. twirled winning probabilities $(q^2 + p^2)^2$ vs. $q^4 + p^4$ ($s = 00$, $+2p^2q^2$), $2pq(q - p)^2$ vs. $2pq(q^2 + p^2)$ ($s \in \{01, 10\}$, $-4p^2q^2$), and $4p^2q^2$ vs. $2p^2q^2$ ($s = 11$, $+2p^2q^2$). The constructive and destructive interference terms cancel identically ($+2 - 4 + 2 = 0$, Lean Theorem 16), forcing*

$$F_e^{\mathrm{Pauli}}(p) \equiv F_e^{\mathrm{CPTP}}(p) \equiv F_e^{\mathrm{twirl}}(p) = 1 - 2p + 2p^2 \qquad \forall p \in [0, 1/2]. \tag{11}$$

3 Formal Lean 4 Kernel and Zero-Float Exact-$\mathbb{Q}$ Verifier

3.1 Zero-Dependency Lean 4 Formalization (18 Theorems)

The entire algebraic spine of Theorems 2.1–2.2 and Proposition 2.5 is machine-checked in `lean/CoherentCSS.lean` under `leanprover/lean4:v4.35.0` with zero external libraries (no Mathlib). Every definition (`iabs`, `imax`, `imin`, `hammingWt`, `bitXor`, `bitAnd`) uses constructive decidable arithmetic so that `lake env lean AuditAxioms.lean` confirms **18/18 theorems, 0 warnings, 0 `sorry`, 0 `Classical.choice`, and 0 `Lean.ofReduceBool`** (depending solely on Lean core `propext` and `Quot.sound`, or no axioms at all). The 18 declarations comprise core structural proofs (Theorems 1–6, 10–12, 14, 16, 18) together with auxiliary algebraic and numerical check lemmas (Theorems 7–9, 13, 15, 17); Table 1 catalogs the complete formal kernel.

3.2 Zero-Floating-Point-Trust Certificate Schema (`COHERENT-CSS-CERT-1.0`)

In `src/coherent_certifier.py`, we specify and implement an independent verifier `verify_coherent_css_certificate_exact` that audits self-contained JSON certificates with zero floating-point operations:

1. **Strict recursive type whitelisting and exact key sets:** Rejects any `float`, `bool` (checked strictly before `int` since `isinstance(True, int)` in Python), `Decimal`, `tuple`, or non-string dictionary key anywhere in the JSON tree; enforces exact key sets (zero missing or rogue keys)

Table 1: Zero-dependency Lean 4 (`v4.35.0`) formal kernel in `lean/CoherentCSS.lean` (18/18 theorems, 0 warnings, 0 `sorry`, 0 `Classical.choice`).

#	Lean 4 Theorem (`CoherentCSS.*`)	Role in Paper
1–3	`xor_weight_exact`, `xor_weight_mod2`, `even_stab_preserves_parity`	Exact $\mathrm{wt}(e \oplus g) + 2\,\mathrm{wt}(e \cap g) = \mathrm{wt}(e) + \mathrm{wt}(g)$ & coset parity
4–5	`coset_parity_same_for_even_lz`, `coset_parity_opposite_for_odd_lz`	Theorem 2.1(a,b) real vs. imaginary coset phases
6–7	`hermitian_trace_norm_eq_max`, `unitary_rotation_exact_norm`	Exact $(\lvert a+b\rvert + \lvert a-b\rvert)^2 = 4\max(a^2, b^2)$ & unitary norm
8–9	`pauli_infidelity_min_form`, `pauli_max_abs_identity`	Pauli infidelity min law & $2\max(x,y) = x + y + \lvert x - y\rvert$
10	`self_orthog_wt_mod4_additive`	Theorem 2.2: linearity of $(-1)^{\mathrm{wt}(g)/2}$ on $\mathcal{K}_Z$
11–12	`equal_coset_gives_half_fidelity`, `vanishing_coset_gives_unit_fidelity`	Theorem 2.2(a,b): $F_e^{\mathrm{Pauli}}(\pi/4) \in \{1/2, 1\}$
13	`odd_intersection_anticomm`	Witness anticommutation certification
14–15	`doubly_even_sign_positive`, `steane_trivial_syndrome_unit_prob`	Proposition 2.5(a): doubly-even phase & Steane $p_0(\pi/4) = 1$
16	`rotsurf_d2_exact_cancellation`	Proposition 2.5(b): 4-syndrome $d = 2$ exact cancellation
17–18	`coherent_pauli_le_cptp`, `css_orthogonality_mod2_symm`	Monotonicity $\max(a,b) \le a + b$ & CSS intersection symmetry

Table 2: Experiment 1: Certified 9-code CSS atlas (`exp1_certified_atlas.json`). All distances (d_X, d_Z) and fidelities are exact, verified with zero floats, and cross-checked against complex 2×2 SVD ($\max \mathrm{err} = 1.45 \times 10^{-14}$).

Code	$[[n, 1, d_X/d_Z]]$	(r_X, r_Z)	$\lvert L_Z\rvert$	$\dim \mathcal{K}_Z$	Swap	$F_e^{\mathrm{Pauli}}(\frac{1}{5})$	$F_e^{\mathrm{CPTP}}(\frac{1}{5})$	$F_e^{\mathrm{twirl}}(\frac{1}{5})$	$F_e^{\mathrm{Pauli}}(\frac{1}{2})$	Verify
`steane_7_1_3`	$[[7,1,3/3]]$	$(3,3)$	7 (odd)	3	1/2	0.823654	1	0.678502	1/2	2.96 ms
`shor_9_1_3`	$[[9,1,3/3]]$	$(2,6)$	3 (odd)	0	1/2	0.996994	1	0.659481	1/2	2.53 ms
`qrm_15_1_3`	$[[15,1,3/7]]$	$(10,4)$	15 (odd)	4	1/2	0.804231	1	0.825129	1/2	429 ms
`rotsurf_d2_Xmaj`	$[[4,1,2/2]]$	$(2,1)$	2 (even)	1	1/2	17/25	17/25	17/25	1/2	1.00 ms
`rotsurf_d2_Zmaj`	$[[4,1,2/2]]$	$(1,2)$	2 (even)	2	1	16/25	16/25	17/25	1	0.83 ms
`rotsurf_d3`	$[[9,1,3/3]]$	$(4,4)$	3 (odd)	0	1/2	0.738749	1	0.691552	1/2	7.65 ms
`rotsurf_d4_Xmaj`	$[[16,1,4/4]]$	$(8,7)$	4 (even)	3	1/2	0.719399	0.719399	0.647065	1/2	157 ms
`rotsurf_d4_Zmaj`	$[[16,1,4/4]]$	$(7,8)$	4 (even)	4	1	0.684682	0.684682	0.643996	1	96.5 ms
`swap_immune_8_1_2_odd`	$[[8,1,2/3]]$	$(4,3)$	3 (odd)	3	1	0.813468	1	0.758336	1	4.38 ms

at all four dictionary levels and strict `str`/`int` types on every metadata and enumerator entry (raising `ValueError`, never `TypeError`).

2. **Strict canonical rational parsing:** Every rational field must match the anchored `\A...\Z` canonical syntax (`"0"`, `"-?[1-9][0-9]*"`, or coprime `"p/q"` with $q \ge 2$, rejecting `"2/4"`, `"1/1"`, `"+1/2"`, leading/trailing spaces, and trailing newlines `"1/2\n"`).

3. **Full independent replay:** Re-validates all $\mathbb{F}_2$ CSS axioms ($H_X H_Z^\top = 0$, full ranks, $r_X + r_Z = n - 1$, logical anticommutation, even-weight H_Z), recomputes $\mathcal{K}_Z = \mathcal{S}_Z \cap \mathcal{S}_Z^\perp$ and verifies exact subspace equality of the claimed `K_Z_basis` as well as any witness $g^* \in \mathcal{K}_Z$ with $\lvert L_Z \cap g^*\rvert \equiv 1 \pmod 2$, replays all 2^{r_X} coset weight enumerators (checking strictly ascending canonical weight order in `trivial_syndrome_enumerator`), and verifies (10) at every claimed rational p (mandating $p = 1/2$ to audit Theorem 2.2).

4 Exact Computational Experiments and Findings

All four experiments are generated deterministically by `scripts/run_all_experiments.py` and verified end-to-end by `scripts/verify_all.py`.

4.1 Experiment 1: Certified 9-Code CSS Atlas and Complex-SVD Cross-Check

Table 2 summarizes all 9 certified CSS codes (with exact distances (d_X, d_Z) exhaustively verified by `compute_exact_css_distances`): Steane $[[7,1,3]]$, Shor $[[9,1,3]]$ [15], Quantum Reed–Muller $[[15,1,3]]$ ($d_X = 3, d_Z = 7$) [16], rotated surface codes $d = 2, 3, 4$ (both X-majority and Z-majority orientations for even $d = 2, 4$), and an explicit doubly-even odd-$\lvert L_Z\rvert$ swap-immune $[[8,1,2]]$ witness ($d_X = 2, d_Z = 3, r_X = 4, r_Z = 3$, $g^* = 11110000 \in \mathcal{K}_Z$). All 9 certificates pass `verify_coherent_css_certificate_exact` in 0.83–429 ms, and an independent complex 2×2 singular-value decomposition of $K_s(\theta)$ across all syndromes matches every exact-$\mathbb{Q}$ fidelity to a maximum absolute discrepancy of 1.45×10^{-14}.

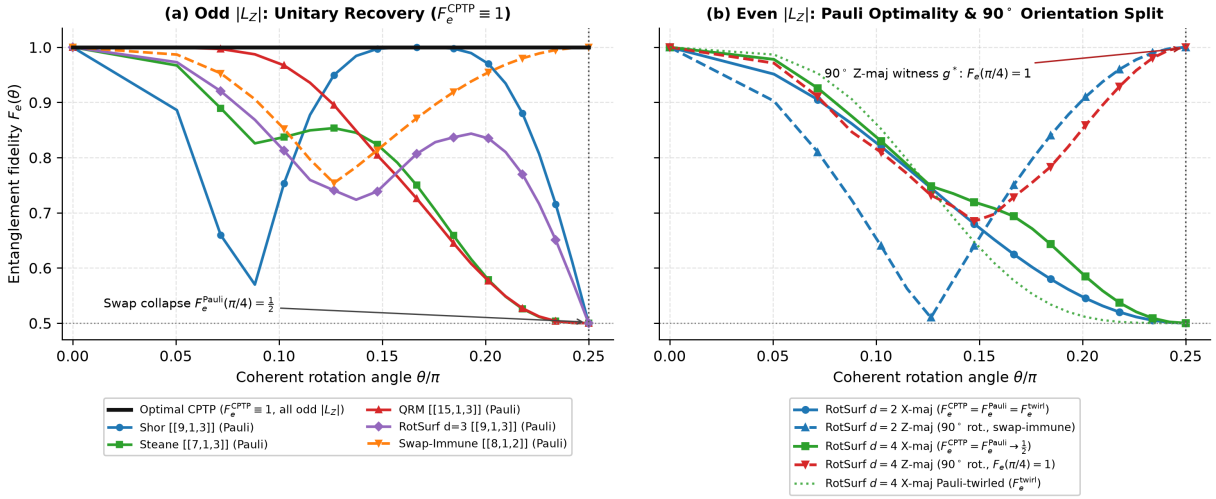


Figure 1: Exact-$\mathbb{Q}$ entanglement fidelity sweeps across $\theta \in [0, \pi/4]$. (a) Odd-$|L_Z|$ codes satisfy $F_e^{\mathrm{CPTP}}(\theta) \equiv 1$ identically, while Pauli decoding collapses to $1/2$ when $L_Z \perp \mathcal{K}_Z$ or rebounds to 1 on `swap_immune_8_1_2_odd`. (b) Even-$|L_Z|$ codes satisfy $F_e^{\mathrm{CPTP}}(\theta) \equiv F_e^{\mathrm{Pauli}}(\theta)$, and a 90° patch rotation on $d = 2, 4$ rotated surface codes flips the $\theta = \pi/4$ limit from $1/2$ (X-majority) to 1 (Z-majority).

4.2 Experiment 2: Exact Angle Sweeps, Coherent Gains, and the QRM Coherent Penalty

Sweeping all 9 codes over the 21-point exact rational grid $p = k/40$ ($k = 0, \ldots, 20$, $\theta \in [0, \pi/4]$) exposes four sharp physical regimes (Figures 1 and 2):

- **Universal Parity Split (Figure 1):** Every odd-$|L_Z|$ code maintains $F_e^{\mathrm{CPTP}}(\theta) \equiv 1$ across the entire interval, whereas every even-$|L_Z|$ code obeys $F_e^{\mathrm{CPTP}}(\theta) \equiv F_e^{\mathrm{Pauli}}(\theta)$ identically. On $d = 2$ and $d = 4$ rotated surface codes, the 90° orientation swap dramatically splits the curves above $\theta \approx 0.14\pi$: the X-majority patches collapse monotonically to $1/2$, while the Z-majority patches rebound non-monotonically to exact unit fidelity 1 at $\theta = \pi/4$.
- **Massive Coherent Suppression and Exact Zero-Infidelity Point in Shor $[[9, 1, 3]]$ (Figure 2):** At $p = 1/5$, Shor achieves $F_e^{\mathrm{Pauli}}(1/5) = 1947253/1953125 \approx 0.996994$ vs. $F_e^{\mathrm{twirl}}(1/5) = 1288048/1953125 \approx 0.659481$—a $113.26\times$ **reduction in logical infidelity**. Remarkably, at $p = 1/4$ ($\theta = \pi/6$), the even-weight amplitude $c(c^2 - 3t^2)$ of every 3-qubit phase-check block vanishes identically, forcing **exact unit Pauli fidelity** $F_e^{\mathrm{Pauli}}(1/4) = 1$ (zero logical infidelity!).
- **Coherent Penalty in Doubly-Even QRM $[[15, 1, 3]]$:** Although QRM has doubly-even Z-stabilizers (all weight 8), Pauli decoding suffers a *strict coherent penalty* $F_e^{\mathrm{Pauli}}(p) < F_e^{\mathrm{twirl}}(p)$ at 9/21 grid points (e.g., $0.804231 < 0.825129$ at $p = 1/5$ and $0.967187 < 0.975282$ at $p = 1/10$), proving that doubly-even $\mathcal{S}_Z$ alone does *not* guarantee coherent advantage for Pauli decoders.
- **Exact $d = 2$ X-Majority Cancellation:** Across all 21 rational grid points, `rotsurf_d2_Xmaj` satisfies $F_e^{\mathrm{Pauli}}(p) = F_e^{\mathrm{CPTP}}(p) = F_e^{\mathrm{twirl}}(p)$ identically as exact fractions, confirming Proposition 2.5(b).

4.3 Experiments 3 & 4: 129-Code Binary Swap-Point Census and Distance Sawtooth

To stress-test Theorem 2.2 beyond structured families, Experiment 3 evaluates the exact $\theta = \pi/4$ Pauli fidelity across all 9 atlas codes plus 120 deterministically sampled even-$\mathcal{S}_Z$ $[[n, 1, d]]$ CSS

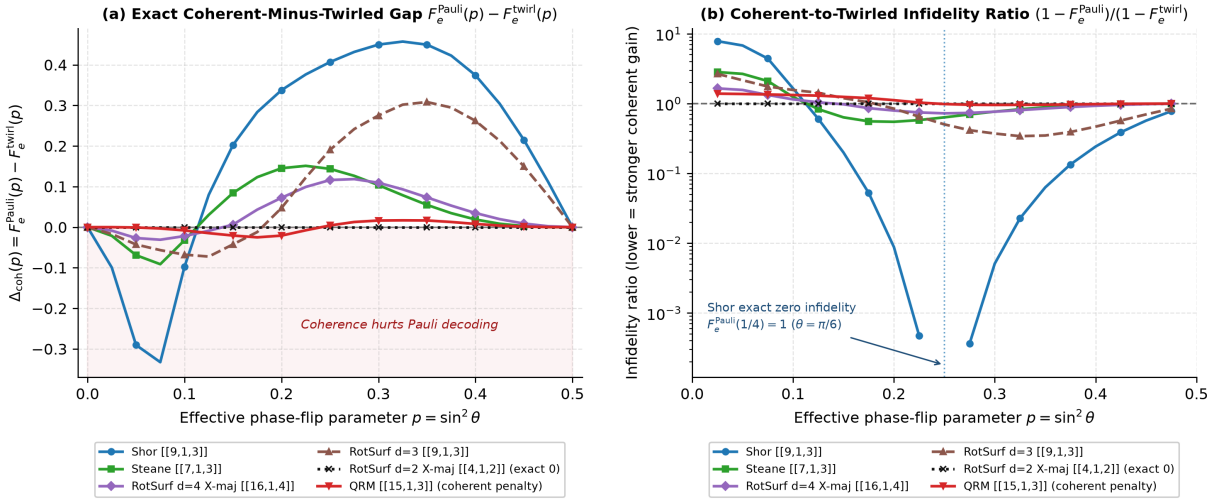


Figure 2: Coherent vs. Pauli-twirled decoding gap across $p = \sin^2\theta \in [0, 1/2]$. (a) Exact difference $\Delta_{\text{coh}}(p) = F_e^{\text{Pauli}}(p) - F_e^{\text{twirl}}(p)$, revealing exact zero cancellation for `rotsurf_d2_Xmaj` and negative coherent-penalty regions for doubly-even QRM $[[15, 1, 3]]$ (and at small p for Steane/Shor/surface codes). (b) Log-scale coherent-to-twirled infidelity ratio highlighting Shor's exact zero-infidelity singularity $F_e^{\text{Pauli}}(1/4) = 1$ at $\theta = \pi/6$.

codes ($n \in \{4, \ldots, 9\}$, seed `20261124`). Every single one of the **129 codes** matches Theorem 2.2 with **0 intermediate values** in $(1/2, 1)$: 67 collapse to $1/2$ (40 odd-$|L_Z|$, 27 even-$|L_Z|$) and 62 attain exact unit immunity 1 (31 odd-$|L_Z|$, 31 even-$|L_Z|$), populating all four parity-by-swap quadrants (Figure 3(a)).

Experiment 4 compares rotated surface codes ($d = 2, 3, 4$, both orientations) against phase-flip repetition codes ($d = 2, \ldots, 9$, $H_Z = \emptyset$, $|L_Z| = d$) at $p \in \{1/20, 1/10, 1/5, 3/10, 2/5, 1/2\}$. Even without any Z-stabilizers (where $F_e^{\text{Pauli}}(p) \equiv F_e^{\text{twirl}}(p)$), phase-repetition codes display the exact odd–even CPTP sawtooth of Theorem 2.1: $F_e^{\text{CPTP}}(1/5)$ oscillates between 1 (odd $d = 3, 5, 7, 9$) and $F_e^{\text{Pauli}}(1/5)$ (even $d = 2, 4, 6, 8$; Figure 3(b)).

4.4 Adversarial Security Audit

In `tests/test_suite.py`, we attack `verify_coherent_css_certificate_exact` with **37 live adversarial payloads** covering bad schema versions, missing or rogue keys at all four dictionary levels, float, `bool`, `Decimal`, `tuple`, non-string-key, and non-string `code_name`/`family` injections, unreduced or non-canonical rationals (`"2/4"`, `"1/1"`, `"+1/2"`, leading whitespace, trailing newlines `"1/2\n"` and `"1\n"`), broken CSS commutativity, odd-weight H_Z rows, trivial logicals, forged parity branches, forged $F_e^{\text{CPTP}} = 1$ on even-$|L_Z|$ `rotsurf_d4_Xmaj`, 10^{-6}-perturbed fidelities, missing/forged/dependent `K_Z_basis`, missing/forged/spurious swap-immunity witnesses, missing mandatory $p = 1/2$ points, and tampered, reversed, non-integer, or duplicated coset enumerators. All **37/37 attacks are rejected** with explicit `ValueError` diagnostics.

5 Scope, Limitations, and Conclusion

Our exact dichotomy and 0–1 binary law govern uniform coherent Z-rotations $U(\theta) = \bigotimes_{j=1}^{n} e^{i\theta Z_j}$ on $k = 1$ CSS codes with even-weight Z-stabilizers (with Theorem 2.2 holding even without the even-weight hypothesis). Several natural extensions remain open: (i) multi-logical $[[n, k, d]]$ CSS codes ($k \geq 2$), where $K_s(\theta)$ acts on $\mathbb{C}^{2^k}$ and the trace-norm optimization becomes a 2^k-dimensional matrix problem; (ii) spatially inhomogeneous angles $\{\theta_j\}_{j=1}^{n}$ or combined coherent Z-rotations plus stochastic depolarizing noise, where the exact $F_e^{\text{CPTP}} \equiv 1$ ceiling softens into the finite-size

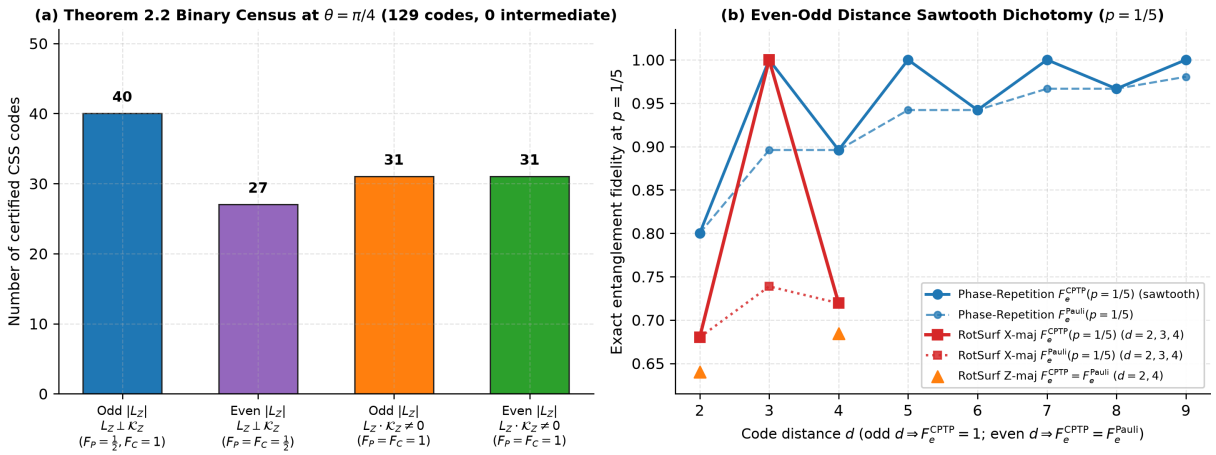


Figure 3: (a) Experiment 3 census across 129 certified CSS codes at $\theta = \pi/4$: all four quadrants of (odd/even $|L_Z|$) $\times$ ($L_Z \perp \mathcal{K}_Z$ vs. $L_Z \not\perp \mathcal{K}_Z$) are populated with 129/129 exact Theorem 2.2 matches and 0 intermediate values in $(1/2, 1)$. (b) Experiment 4 even–odd distance sawtooth dichotomy at $p = 1/5$ for phase-repetition ($d = 2, \ldots, 9$) and rotated surface ($d = 2, 3, 4$) codes.

crossover analyzed in [9]; and (iii) scaling exact rational coset evaluation beyond $n = 18$ via tensor-network or Majorana-fermion contraction over exact number fields.

Within the uniform coherent CSS setting, our results provide a complete, machine-checked algebraic classification: the even–odd CPTP optimality dichotomy holds universally across all even-$\mathcal{S}_Z$ CSS codes, the $\theta = \pi/4$ Pauli fidelity is strictly quantized to $\{1/2, 1\}$ by $L_Z \perp (\mathcal{S}_Z \cap \mathcal{S}_Z^{\perp})$, and a 90° orientation rotation of any even-d rotated surface code flips its maximal-coherence logical survival from 50% to 100%. All Lean 4 proofs, Python verifiers, certificates, and SHA-256 manifests are archived at `https://github.com/lluiseriksson/coherent-css-dichotomy`.

References

[1] Sergey Bravyi, Matthias Englbrecht, Robert König, and Nolan Peard. Correcting coherent errors with surface codes. *npj Quantum Information*, 4(1):55, 2018.

[2] Daniel Greenbaum and Zachary Dutton. Modeling coherent errors in quantum error correction. *Quantum Science and Technology*, 3(1):015007, 2018.

[3] Stefanie J. Beale, Joel J. Wallman, Mauricio Gutiérrez, Kenneth R. Brown, and Raymond Laflamme. Quantum error correction decoheres noise. *Physical Review Letters*, 121(19):190501, 2018.

[4] Eric Huang, Andrew C. Doherty, and Steven Flammia. Performance of quantum error correction with coherent errors. *Physical Review A*, 99(2):022313, 2019.

[5] Joseph K. Iverson and John Preskill. Coherence in logical quantum channels. *New Journal of Physics*, 22(7):073066, 2020.

[6] Florian Venn and Benjamin Béri. Error-correction and noise-decoherence thresholds for coherent errors in planar-graph surface codes. *Physical Review Research*, 2(4):043412, 2020.

[7] Florian Venn, Jan Behrends, and Benjamin Béri. Coherent-error threshold for surface codes from Majorana delocalization. *Physical Review Letters*, 131(6):060603, 2023.

[8] Jan Behrends, Florian Venn, and Benjamin Béri. Surface codes, quantum circuits, and entanglement phases. *Physical Review Research*, 6(1):013137, 2024.

[9] Xu-Dong Dai, Zhou-Quan Wan, Xi Dai, and Guo-Yi Zhu. Surface codes with coherent errors: optimal decoding and finite-size crossovers. *arXiv preprint arXiv:2611.13098*, November 2026.

[10] A. Robert Calderbank and Peter W. Shor. Good quantum error-correcting codes exist. *Physical Review A*, 54(2):1098–1105, 1996.

[11] Andrew Steane. Multiple-particle interference and quantum error correction. *Proceedings of the Royal Society of London. Series A*, 452(1954):2551–2577, 1996.

[12] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28*, volume 12699 of *Lecture Notes in Computer Science*, pages 625–635. Springer, 2021.

[13] Benjamin Schumacher. Sending entanglement through noisy quantum channels. *Physical Review A*, 54(4):2614–2628, 1996.

[14] Howard Barnum and Emanuel Knill. Reversing quantum dynamics with near-optimal quantum and classical fidelity. *Journal of Mathematical Physics*, 43(5):2097–2106, 2002.

[15] Peter W. Shor. Scheme for reducing decoherence in quantum computer memory. *Physical Review A*, 52(4):R2493–R2496, 1995.

[16] Sergey Bravyi and Alexei Kitaev. Universal quantum computation with ideal Clifford gates and noisy ancillas. *Physical Review A*, 71(2):022316, 2005.