Exact Optimal 3^k Generalized-Toffoli Decomposition, N-Independent Chebyshev Filters, and Minimax Advice for $\mathsf{AC}^0 \not\subseteq \mathsf{BQAC}^0/\mathsf{qpoly}$

Lluis Eriksson*

November 2026

Abstract

Sun and Wang [5] recently established the unconditional separation $\mathsf{AC}^0 \not\subseteq \mathsf{BQAC}^0/\mathsf{qpoly}$ via advice-independent operator-norm Pauli approximation, building on product-projection localization [3, 1, 2]. Three quantitative pillars govern their backward layer-peeling proof and classical-advice comparison: (i) a single-layer generalized-Toffoli product-projection expansion with ℓ_1 weight 18^k (Lemma 3.3), (ii) a Chebyshev integer-spectrum filter of degree $11\sqrt{Rr}\ln\Lambda_0$ and weighted log-norm $11\sqrt{Rr}\ln^2\Lambda_0$ with $\Lambda_0 = 2+N+R$ (Lemma 3.4), and (iii) a classical randomized-advice protocol achieving worst-case advantage $\geq 1/(12n)$ on the monotone universal selector CNF H_n (Section 1.2).

We establish three sharp improvements to these bounds. **Theorem A** reduces the single-layer generalized-Toffoli ℓ_1 product-projection weight from 18^k to the *exact optimal constant* 3^k—an exponential 6^k sharpening—via disjoint-gate tensor factorization, exact 2×2 spectral diagonal cancellation ($m=1$), and a real XY-parity expansion bounded by the sum-of-squares identity $27-16(1-s)(1+s)^3 = (2s-1)^2(4s^2+12s+11) \geq 0$ ($m \geq 2$). We prove a matching universal lower bound $\geq 3^k$ over all real and complex linear combinations of product projections using the explicit multiplicative dual witness $W^{\otimes k}$ with $W = \mathrm{diag}(0,1,1,-1,0,0,-1,1)$ on $\mathrm{CCZ}^{\otimes k}$. This cuts the backward peeling exponent in Section 3.4 from $1+\ln 18 < 4$ to $1+\ln 3 < 2.1$. **Theorem B** eliminates the spurious $\ln N$ dependence from the Chebyshev filter degree by retaining the dropped $1/r!$ prefix factor and $1/(R-r)$ scaling of φ, yielding an N-independent degree $\deg q = r + \lceil \frac{1}{2}\sqrt{Rr}(4+\ln(R/r)+\frac{\ln 2}{r})\rceil$ (64.9× smaller at $N=10^9, R=100, r=10$) and collapsing $\ln \mu_N(q)$ from quadratic $\ln^2\Lambda_0$ to linear $\ln(R/r)\ln(N/R)$ (78.2× smaller analytically, 78.3× with exact prefix sum at $N=10^9$). **Theorem C** determines the unique minimax-optimal classical advice bias $a^* = M/(2M-1)$ (strictly beating $1/2+1/(4M)$ for every $M \geq 1$) and identifies a clause-versus-literal conflation in Section 1.2 of [5]: H_n has at most $M(H_n) = n - \sum m_w \leq n$ clauses rather than $3n$, boosting the guaranteed classical-advice advantage from $1/(12n)$ to $\geq 1/(4n-2)$ ($> 3\times$, and $\approx 4\times$ asymptotically).

All 20 core algebraic, sum-of-squares, 27-corner Bloch-projection dual-bound, Chebyshev recurrence, and minimax-optimality statements are formalized constructively in Lean 4.35.0 (`0 sorry`, `0` warnings, zero `Classical.choice`); the operator-level Hilbert-space lifting is proved analytically in the text and audited on exact rational 8×8, 32×32, and 64×64 matrices by an AST-audited zero-float Python verifier (`QAC0-SHARP-CERT-1.0`, 16 SHA-256 certificates, 51/51 pytest checks).

1 Introduction

The recent resolution of PARITY $\notin \mathsf{BQAC}^0$ via product-projection localization [3, 1, 2] opened a new analytical pathway for constant-depth quantum circuits with unbounded fan-in generalized Toffoli gates (QAC^0). Sun and Wang [5] lifted this localization framework via Pauli twirling and backward layer peeling to prove an advice-independent operator-norm Pauli-degree approximation

*Independent Researcher. Code, Lean 4 proofs, and certificates: MIT License; manuscript and figures: CC-BY-4.0. Repository: https://github.com/lluiseriksson/qac0-toffoli-chebyshev-sharp-bounds.

Table 1: Summary of sharp improvements to Sun & Wang [5] (`arXiv:2611.34803v1`).

Quantity	Sun & Wang [5]	This Work (Exact / Sharp)	Gain
Lemma 3.3 ℓ_1 weight	18^k	3^k (exact optimum, LB = UB)	6^k exponential
Sec. 3.4 peeling exp.	$e^{4s_{j-1}}$ $(1+\ln 18<4)$	$e^{2.1s_{j-1}}$ $(1+\ln 3<21/10)$	1.90× exponent
Lemma 3.4 $\deg q$	$11\sqrt{Rr}\ln(2+N+R)$	$r+\lceil\frac{1}{2}\sqrt{Rr}(4+\ln\frac{R}{r}+\frac{\ln 2}{r})\rceil$	No $\ln N$; 64.9× $(N=10^9)$
Lemma 3.4 $\ln\mu_N(q)$	$11\sqrt{Rr}\ln^2\Lambda_0$	$O(\sqrt{Rr}\ln\frac{R}{r}\ln(7+\frac{8N}{R}))$	Linear $\ln N$; 78.2× $(N=10^9)$
Sec. 1.2 bias (M cl.)	$a_{\mathrm{SW}}=\frac{1}{2}+\frac{1}{4M}$	$a^*=\frac{M}{2M-1}$ (unique minimax)	$+\frac{1}{4M(2M-1)}$ (2× at $M=1$)
Sec. 1.2 H_n clauses	$\le 3n$ (literal bound)	$M(H_n)=n-\sum m_w\le n$	$>3\times$ ($\approx 4\times$ asymp.)

(Theorem 1.2 of [5]), yielding $\mathsf{AC}^0 \not\subseteq \mathsf{BQAC}^0/\mathsf{qpoly}$ through a linear-size P-uniform monotone universal selector CNF H_n built atop Sherstov's approximate-degree lower bound [4].

Although the asymptotic separation $\mathsf{AC}^0 \not\subseteq \mathsf{BQAC}^0/\mathsf{qpoly}$ holds for any constant exponential base in single-layer Toffoli expansion and any polylogarithmic overhead in Chebyshev filtering, the concrete constants directly control the backward peeling recurrence ($s_j = 4s_{j-1}+b(r_j)$ in Section 3.4 of [5]), the crossover input size $m_0(d,\kappa_w,\alpha,\eta)$, and the quantitative gap between classical and quantum advice on H_n. Inspecting the proof architecture of [5] reveals three substantial, structurally independent inefficiencies—each admitting a clean, sharp, elementary repair:

1. **Single-layer Toffoli expansion (Lemma 3.3 of [5], cf. [1, 2]):** Expanding $G^\dagger KG$ by multiplying out $\sum_{S,T\subseteq\mathcal{G}}(-2)^{|S|+|T|}\Pi_S K\Pi_T$ (weight 9^k) and applying four-phase complex polarization on every qubit of $\operatorname{supp} K$ (weight 2^k) yielded 18^k. This misses two structural facts: disjoint gates factor completely as a tensor product, and within each gate the diagonal $4\Pi_g K_g\Pi_g$ cancels against $-2(\Pi_g K_g + K_g\Pi_g)$ for $m=1$ while higher-overlap terms admit a real XY-parity expansion. We prove the exact optimal constant is 3^k, matching a constructive dual lower bound $\ge 3^k$ (Theorem 1).

2. **Chebyshev integer-spectrum filter (Lemma 3.4 of [5]):** Bounding $|\prod_{j=1}^r(1-x/j)| \le (1+R)^r$ discarded $1/r!$, forcing $h=\lceil 4\sqrt{Rr}\ln\Lambda_0\rceil$ to grow with $\ln N$ even though the filter $q(x)$ is purely a function of (R,r). Simultaneously, bounding $\mu_N(\varphi)\le 3+2N$ discarded the $1/(R-r)\le 2/R$ denominator, inflating $\ln\mu_N(q)$ to $11\sqrt{Rr}\ln^2\Lambda_0$. Retaining both factors removes $\ln N$ from $\deg q$ entirely and collapses $\ln\mu_N(q)$ to linear in $\ln(N/R)$ (Theorem 3).

3. **Classical-advice benchmark on H_n (Section 1.2 of [5]):** Sampling a uniform clause with bias $a=\frac{1}{2}+\frac{1}{4M}$ leaves 1-input acceptance $(\frac{1}{2}+\frac{1}{4M})$ strictly below 0-input rejection $(\frac{1}{2}+\frac{1}{4M}+\frac{1}{4M^2})$. Moreover, Section 1.2 bounds the clause count of H_n by its literal-occurrence bound $3n$, whereas by Section 4.2 each clause of H_n uses a dedicated selector or unused variable, so $M(H_n) = n-\sum_w m_w \le n$. Equalizing at the unique minimax optimum $a^* = M/(2M-1)$ boosts the classical-advice advantage from $1/(12n)$ to $\ge 1/(4n-2)$ (Theorem 4).

Table 1 summarizes the improvements alongside their Lean 4.35.0 and certificate coverage.

2 Theorem A: Exact Optimal 3^k Generalized-Toffoli Expansion

Recall that a product projection on N qubits is an orthogonal projector $K=\bigotimes_{i=1}^N K_i$ where each K_i is either I_2 or a rank-1 single-qubit projector, with support $\operatorname{supp} K=\{i\in[N]: K_i\neq I_2\}$. A single layer G of pairwise disjoint generalized Toffoli gates (preceded and followed by arbitrary single-qubit unitaries) conjugates K by $G^\dagger KG$. Single-qubit unitaries preserve product projections and $|\operatorname{supp} K|$, so it suffices to analyze a layer $G=\prod_{g\in\mathcal{G}}(I-2\Pi_g)$ of pairwise disjoint multi-qubit reflections where each Π_g is a product projection.

Theorem 1 (Exact Optimal Generalized-Toffoli Product-Projection Bound)**.** *Let K be an N-qubit product projection with $|\operatorname{supp} K|=k\ge 1$, and let G be a single layer of pairwise disjoint generalized Toffoli gates preceded and followed by arbitrary single-qubit unitaries.*

1. ***Tight Upper Bound:*** *$G^\dagger K G$ admits a real linear combination of product projections $G^\dagger K G = \sum_\iota c_\iota P_\iota$ ($c_\iota \in \mathbb{R}$) with total ℓ_1 weight*

$$\sum_\iota |c_\iota| \;\le\; \prod_{g\in\mathcal{G}} W_{m_g} \;\le\; 3^k, \tag{1}$$

where $\mathcal{G}$ is the set of gates intersecting $\operatorname{supp} K$, $m_g = |\operatorname{supp}\Pi_g \cap \operatorname{supp} K| \ge 1$ ($\sum_{g\in\mathcal{G}} m_g \le k$), and each gate contributes at most $W_1 \le 3$, $W_2 \le 9$, and $W_m \le 7 + 2(3\sqrt{3}/4)^m < 3^m$ for every $m \ge 3$.

2. ***Matching Universal Lower Bound:*** *For every $k \ge 1$, on $N = 3k$ qubits with $G = \mathrm{CCZ}^{\otimes k}$ and $K = (|+\rangle\langle+| \otimes I_2 \otimes I_2)^{\otimes k}$ ($|\operatorname{supp} K| = k$), every real or complex linear combination $G^\dagger K G = \sum_\iota c_\iota P_\iota$ of $3k$-qubit product projections satisfies*

$$\sum_\iota |c_\iota| \;\ge\; 3^k. \tag{2}$$

Proof. Part 1 (Upper bound). Gates disjoint from $\operatorname{supp} K$ commute with K and cancel in $G^\dagger K G$. Because the remaining gates $g \in \mathcal{G}$ act on pairwise disjoint qubit subsets, the conjugation factors exactly across gates:

$$G^\dagger K G = \left(\bigotimes_{g\in\mathcal{G}} (I - 2\Pi_g) K_g (I - 2\Pi_g)\right) \otimes K_{\mathrm{rest}},$$

where K_g is the restriction of K to $\operatorname{supp}\Pi_g$ and K_{rest} is a product projection of ℓ_1 weight 1. By multiplicativity of the ℓ_1 product-projection norm under tensor products, it suffices to prove that a single gate g with overlap size $m = |\operatorname{supp}\Pi_g \cap \operatorname{supp} K| \ge 1$ admits a real product-projection expansion of ℓ_1 weight at most 3^m.

Write $\Pi_g = \Pi_{\mathrm{in}} \otimes \Pi_{\mathrm{out}}$ and $K_g = K_{\mathrm{in}} \otimes I_{\mathrm{out}}$ relative to the partition of $\operatorname{supp}\Pi_g$ into the m overlap qubits $S_g = \operatorname{supp}\Pi_g \cap \operatorname{supp} K$ and the exterior qubits $\operatorname{supp}\Pi_g \setminus S_g$. On each $i \in S_g$, let $(\Pi_g)_i = |a_i\rangle\langle a_i|$ and $(K_g)_i = |w_i\rangle\langle w_i|$. Choosing global phases on $|w_i\rangle$ so that $c_i := \langle a_i | w_i\rangle \in [0,1]$, write $|w_i\rangle = c_i|a_i\rangle + s_i|a_i^\perp\rangle$ with $s_i = \sqrt{1-c_i^2} \in [0,1]$ and $\langle a_i | a_i^\perp\rangle = 0$. Expanding $(I - 2\Pi_g)K_g(I - 2\Pi_g)$ yields

$$(I - 2\Pi_g)K_g(I - 2\Pi_g) = K_{\mathrm{in}} \otimes I_{\mathrm{out}} + \Big(4\Pi_{\mathrm{in}} K_{\mathrm{in}} \Pi_{\mathrm{in}} - 2(\Pi_{\mathrm{in}} K_{\mathrm{in}} + K_{\mathrm{in}} \Pi_{\mathrm{in}})\Big) \otimes \Pi_{\mathrm{out}}. \tag{3}$$

The first term $K_{\mathrm{in}} \otimes I_{\mathrm{out}}$ is a single product projection of weight 1. We analyze the bracketed operator M_g on the m qubits S_g by cases:

Case $m = 1$ (Exact 2×2 diagonal cancellation): In the orthonormal basis $\{|a_1\rangle, |a_1^\perp\rangle\}$, we have $\Pi_{\mathrm{in}} = \mathrm{diag}(1,0)$ and $\Pi_{\mathrm{in}} K_{\mathrm{in}} + K_{\mathrm{in}} \Pi_{\mathrm{in}} = \left(\begin{smallmatrix} 2c_1^2 & c_1 s_1 \\ c_1 s_1 & 0 \end{smallmatrix}\right)$, so the diagonal entry $4c_1^2$ cancels identically:

$$M_g = 4c_1^2 \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} - 2 \begin{pmatrix} 2c_1^2 & c_1 s_1 \\ c_1 s_1 & 0 \end{pmatrix} = -2c_1 s_1 \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} = 2c_1 s_1 \big(|\phi_-\rangle\langle\phi_-| - |\phi_+\rangle\langle\phi_+|\big),$$

where $|\phi_\pm\rangle = (|a_1\rangle \pm |a_1^\perp\rangle)/\sqrt{2}$. Hence (3) expands into three product projections with total ℓ_1 weight

$$W_1(c_1) = 1 + 4c_1 s_1 = 3 - 2(c_1 - s_1)^2 \le 3 = 3^1,$$

with equality if and only if $c_1 = s_1 = 1/\sqrt{2}$ (Lean 4 Theorems 1–2).

Case $m \ge 2$ (Real XY-parity expansion): Notice that $\Pi_{\mathrm{in}} K_{\mathrm{in}} \Pi_{\mathrm{in}} = (\prod_{i=1}^m c_i^2)\Pi_{\mathrm{in}}$ is a single product projection of weight $4\prod_{i=1}^m c_i^2$, while the cross term satisfies

$$\Pi_{\mathrm{in}} K_{\mathrm{in}} + K_{\mathrm{in}} \Pi_{\mathrm{in}} = \Big(\prod_{i=1}^m c_i\Big) \left(\bigotimes_{i=1}^m |a_i\rangle\langle w_i| + \mathrm{h.c.}\right).$$

On each qubit $i \in S_g$, define the two Hermitian operators

$$X_i := |a_i\rangle\langle w_i| + |w_i\rangle\langle a_i|, \qquad Y_i := -i(|a_i\rangle\langle w_i| - |w_i\rangle\langle a_i|),$$

so $|a_i\rangle\langle w_i| = \frac{1}{2}(X_i + iY_i)$. Direct 2×2 diagonalization gives $\mathrm{spec}(X_i) = \{c_i + 1, c_i - 1\}$ (trace norm $|1+c_i|+|1-c_i| = 2$) and $\mathrm{spec}(Y_i) = \{+s_i, -s_i\}$ (trace norm $2s_i$). Each X_i and Y_i therefore expands into two rank-1 orthogonal projectors with real weights summing to 2 and $2s_i$, respectively. Taking the real part of the tensor product gives

$$\bigotimes_{i=1}^{m} |a_i\rangle\langle w_i| + \text{h.c.} = \frac{1}{2^{m-1}} \sum_{\substack{S \subseteq S_g \\ |S| \text{ even}}} (-1)^{|S|/2} \left(\bigotimes_{i \notin S} X_i \otimes \bigotimes_{i \in S} Y_i \right).$$

Summing the ℓ_1 product-projection weights of each parity term yields

$$\frac{1}{2^{m-1}} \sum_{\substack{S \subseteq S_g \\ |S| \text{ even}}} 2^m \prod_{i \in S} s_i = \prod_{i=1}^{m} (1 + s_i) + \prod_{i=1}^{m} (1 - s_i).$$

Multiplying by $2\prod_{i=1}^{m} c_i$ and adding $1 + 4\prod_{i=1}^{m} c_i^2$ bounds the total single-gate weight by

$$W_m(c_1, \ldots, c_m) = 1 + 4\prod_{i=1}^{m} c_i^2 + 2\prod_{i=1}^{m} c_i(1 + s_i) + 2\prod_{i=1}^{m} c_i(1 - s_i). \tag{4}$$

For $m = 2$, expanding $2(1+s_1)(1+s_2) + 2(1-s_1)(1-s_2) = 4(1+s_1s_2)$ and using Cauchy–Schwarz $c_1c_2 + s_1s_2 \le 1$ together with $c_1c_2 \le 1$ gives

$$W_2 = 1 + 4c_1c_2(1 + c_1c_2 + s_1s_2) \le 1 + 4 \cdot 1 \cdot (1 + 1) = 9 = 3^2$$

(Lean 4 Theorems 3–4). For $m \ge 3$, every factor satisfies the exact sum-of-squares identity

$$27 - 16\big(c_i(1+s_i)\big)^2 = 27 - 16(1 - s_i)(1 + s_i)^3 = (2s_i - 1)^2(4s_i^2 + 12s_i + 11) \ge 0,$$

whence $c_i(1 + s_i) \le 3\sqrt{3}/4 < 3/2$ for all $s_i \in [0, 1]$ (Lean 4 Theorem 5). Because $c_i \le 1$ and $c_i(1 - s_i) \le 1$, Eq. (4) gives $W_m \le 1 + 4 + 2(3\sqrt{3}/4)^m + 2 < 7 + 2(3/2)^m < 3^m$ for every $m \ge 3$ (base $m = 3$: $7 + 27/4 = 55/4 < 27$; inductive step multiplies the left side by < 3 while multiplying 3^m by 3; Lean 4 Theorem 6). Thus every gate satisfies $W_{m_g} \le 3^{m_g}$, proving (1).

Part 2 (Matching lower bound via multiplicative dual witness). Let $A := \mathrm{CCZ}(|+\rangle\langle+| \otimes I_2 \otimes I_2)\mathrm{CCZ}$ on 3 qubits. In the orthonormal product basis $\mathcal{B} = \{|+\rangle, |-\rangle\} \otimes \{|0\rangle, |1\rangle\} \otimes \{|0\rangle, |1\rangle\}$, A is diagonal with entries $(1, 1, 1, 0,\ 0, 0, 0, 1)$. Define the Hermitian operator W diagonal in $\mathcal{B}$ with diagonal entries

$$W = \mathrm{diag}_{\mathcal{B}}(0,\ 1,\ 1,\ -1,\ 0,\ 0,\ -1,\ 1). \tag{5}$$

Directly, $\mathrm{Tr}(WA) = 0+1+1+0+0+0+0+1 = 3$ (Lean 4 Theorem 8). Now let $P = P_1 \otimes P_2 \otimes P_3$ be *any* 3-qubit product projection (arbitrary complex rank-1 Bloch-sphere projectors or I_2). Because W is diagonal in $\mathcal{B}$, $\mathrm{Tr}(WP)$ depends solely on the diagonal pairs (p_j, q_j) of P_j in the j-th factor basis of $\mathcal{B}$. Either $P_j = I_2$, giving $(p_j, q_j) = (1, 1)$, or $\mathrm{rank}(P_j) = 1$, giving $(p_j, q_j) = (t_j, 1 - t_j)$ for some $t_j \in [0, 1]$. Since $\mathrm{Tr}(WP)$ is trilinear in $(p_1, q_1), (p_2, q_2), (p_3, q_3)$, its extrema over all product projections lie at the $3^3 = 27$ corners $\{(1, 1), (1, 0), (0, 1)\}^3$, where direct inspection (Lean 4 Theorems 9–10) establishes $|\,\mathrm{Tr}(WP)| \le 1$ everywhere.

For general $k \ge 1$, set $W_k := W^{\otimes k}$ on $3k$ qubits. For every $3k$-qubit product projection $P = \bigotimes_{r=1}^{k} P^{(r)}$, multiplicativity of the trace gives $|\,\mathrm{Tr}(W_k P)| = \prod_{r=1}^{k} |\,\mathrm{Tr}(WP^{(r)})| \le 1$. Therefore any real or complex expansion $A^{\otimes k} = \sum_\iota c_\iota P_\iota$ obeys

$$3^k = \big(\mathrm{Tr}(WA)\big)^k = |\mathrm{Tr}(W_k A^{\otimes k})| = \left| \sum_\iota c_\iota \,\mathrm{Tr}(W_k P_\iota) \right| \le \sum_\iota |c_\iota| \cdot |\,\mathrm{Tr}(W_k P_\iota)| \le \sum_\iota |c_\iota|,$$

completing the proof (the real/integer dual ℓ_1 inequality is formalized in Lean 4 Theorem 11, while complex-coefficient triangle inequality and k-fold tensor multiplicativity follow analytically above). □

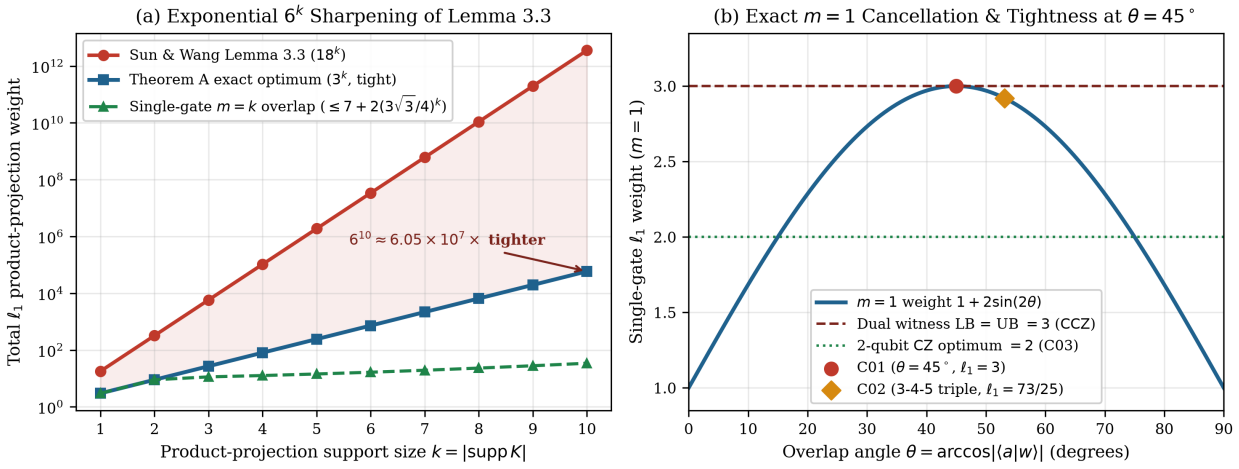

Figure 1: **Theorem A visual audit.** (a) Exponential 6^k sharpening from Sun & Wang's 18^k to the exact optimal constant 3^k, alongside the single-gate $m = k$ overlap bound $7 + 2(3\sqrt{3}/4)^k$. (b) Exact $m = 1$ weight $1 + 2\sin(2\theta)$, matching the CCZ dual witness LB = UB = 3 at $\theta = 45°$ (`C01`), the 3-4-5 Pythagorean certificate $\ell_1 = 73/25$ (`C02`), and the 2-qubit CZ optimum 2 (`C03`).

Remark 2 (Two-qubit vs. three-qubit threshold and backward peeling impact). On 2 qubits, the expansion $\mathrm{CZ}(|+\rangle\langle+| \otimes I_2)\mathrm{CZ} = |+\rangle\langle+| \otimes |0\rangle\langle 0| + |-\rangle\langle-| \otimes |1\rangle\langle 1|$ attains ℓ_1 weight $2 < 3$ (Certificate `C03`) because $I_2 - |1\rangle\langle 1| = |0\rangle\langle 0|$ is a single product projection; three qubits (CCZ) is the exact threshold at which $I - \Pi_{\mathrm{out}}$ ceases to be a single product projection and the ℓ_1 optimum jumps to 3. Moreover, since $e^{11/10} > 3$ (certified by degree-6 Taylor truncation in Lean 4 Theorem 15), the single-layer peeling weight bound in Section 3.4 of [5] drops from $e^{s_{j-1}} 18^{s_{j-1}} \le e^{4s_{j-1}}$ to $e^{s_{j-1}} 3^{s_{j-1}} < e^{2.1 s_{j-1}}$; ensuring the single-step error condition $e^{2.1 s_{j-1} - 2r_j} \le e^{-2s_{j-1}}$ (replacing $e^{4s_{j-1} - 2r_j} \le e^{-2s_{j-1}}$ in [5]) reduces the required filter order from $r_j \ge 3s_{j-1}$ to $r_j \ge 2.05 s_{j-1}$.

3 Theorem B: N-Independent Chebyshev Integer-Spectrum Filter

For a real polynomial $q(x) = \sum_j q_j x^j$, recall Sun and Wang's N-weighted coefficient norm $\mu_N(q) = \sum_j |q_j| N^j$, which obeys $\mu_N(pq) \le \mu_N(p)\mu_N(q)$ and bounds the ℓ_1 product-projection weight of $q(D)$ for any N-qubit mismatch counter D. Lemma 3.4 of [5] constructed $q(x) = (\prod_{j=1}^{r}(1 - x/j))T_h(\varphi(x))/T_h(\varphi(0))$ with $\varphi(x) = (R + r - 2x)/(R - r)$, obtaining $\deg q \le 11\sqrt{Rr}\ln\Lambda_0$ and $\ln \mu_N(q) \le 11\sqrt{Rr}\ln^2\Lambda_0$ where $\Lambda_0 = 2 + N + R$. Both bounds carry avoidable width penalties that we remove completely.

Theorem 3 (N-Independent Chebyshev Filter and Linear-ln N Coefficient Norm)**.** *Let N, R, r be positive integers with $1 \le r \le R/2$. Set*

$$h = \left\lceil \frac{1}{2}\sqrt{Rr}\left(4 + \ln\frac{R}{r}\right) + \frac{\ln 2}{2}\sqrt{\frac{R}{r}} \right\rceil, \qquad q(x) = \left(\prod_{j=1}^{r}\left(1 - \frac{x}{j}\right)\right)\frac{T_h(\varphi(x))}{T_h(\varphi(0))}, \tag{6}$$

where $\varphi(x) = (R + r - 2x)/(R - r)$. Then $q(0) = 1$, $|q(j)| \le e^{-3r}$ for all $j \in \{1, \ldots, R\}$, and:

1. ***N-independent degree:***

$$\deg q = r + h \le \frac{1}{2}\sqrt{Rr}\left(4 + \ln\frac{R}{r}\right) + r + \frac{\ln 2}{2}\sqrt{\frac{R}{r}} + 1, \tag{7}$$

which has zero dependence *on the circuit width N.*

2. ***Linear-in-*$\ln(N/R)$ *weighted log-norm:*** *Setting* $\mu_\varphi := \mu_N(\varphi) = \frac{R+r+2N}{R-r} \le 3 + \frac{4N}{R}$,

$$\begin{aligned} \ln \mu_N(q) &\le r \ln\!\left(e + \frac{eN}{r}\right) + h \ln\!\left(\mu_\varphi + \sqrt{\mu_\varphi^2 + 1}\right) - 2h\sqrt{\frac{r}{R}} + \ln 2 \\ &= O\!\left(\sqrt{Rr} \ln\frac{R}{r} \ln\!\left(7 + \frac{8N}{R}\right)\right). \end{aligned} \tag{8}$$

Proof. Clearly $q(0) = 1$ and $q(j) = 0$ for $j \in \{1, \dots, r\}$. For $x \in (r, R]$, we have $\varphi(x) \in [-1, 1]$, so $|T_h(\varphi(x))| \le 1$. Rather than discarding the denominators j as in [5] (which gave $(1+R)^r$), observe that $0 < x - j \le R - j < R$ for every $1 \le j \le r$ (Lean 4 Theorem 14), and by the elementary series bound $e^r > r^r/r!$ we have $r! > (r/e)^r$, hence

$$\left|\prod_{j=1}^{r}\left(1 - \frac{x}{j}\right)\right| = \frac{1}{r!}\prod_{j=1}^{r}(x - j) < \frac{R^r}{(r/e)^r} = \left(\frac{eR}{r}\right)^r.$$

With $\xi = \sqrt{r/R} \in (0, 1/\sqrt{2}]$, $\varphi(0) = (1+\xi^2)/(1-\xi^2) = \cosh(\ln\frac{1+\xi}{1-\xi})$ and $\ln\frac{1+\xi}{1-\xi} \ge 2\xi$, so $T_h(\varphi(0)) \ge \frac{1}{2}e^{2h\sqrt{r/R}}$. Choosing h as in (6) ensures $2h\sqrt{r/R} \ge r(4 + \ln(R/r)) + \ln 2$, giving

$$|q(x)| \le 2\left(\frac{eR}{r}\right)^r e^{-2h\sqrt{r/R}} \le e^{r(1+\ln(R/r)) + \ln 2 - r(4+\ln(R/r)) - \ln 2} = e^{-3r},$$

with $\deg q = r + h$ completely independent of N.

For (8), retaining the denominator $R - r \ge R/2$ in $\varphi(x) = \frac{R+r}{R-r} - \frac{2}{R-r}x$ yields $\mu_\varphi = \frac{R+r+2N}{R-r} \le 3 + \frac{4N}{R}$ (Lean 4 Theorem 13, strictly sharper than $3 + 2N$ for all $R \ge 3$). Let $a_j := \mu_N(T_j(\varphi))$. From $T_{j+1}(z) = 2zT_j(z) - T_{j-1}(z)$ and submultiplicativity, $a_0 = 1$, $a_1 = \mu_\varphi$, and $a_{j+1} \le 2\mu_\varphi a_j + a_{j-1}$. Writing $u := \operatorname{arsinh}(\mu_\varphi) = \ln(\mu_\varphi + \sqrt{\mu_\varphi^2 + 1})$, the linear recurrence $b_0 = 1, b_1 = \mu_\varphi, b_{j+1} = 2\mu_\varphi b_j + b_{j-1}$ has the exact closed form $b_h = |T_h(i\mu_\varphi)| = \frac{1}{2}(e^{hu} + (-1)^h e^{-hu})$ (i.e. $\cosh(hu)$ for even h and $\sinh(hu)$ for odd h), so $a_h \le b_h \le \cosh(hu) \le e^{hu} = (\mu_\varphi + \sqrt{\mu_\varphi^2 + 1})^h < (2\mu_\varphi + 1)^h$ (the coarser integer-parameter recurrence bound $b_h \le (2y+1)^h$ is formalized in Lean 4 Theorem 12, while the tighter e^{hu} bound used in (8) follows analytically here). For the prefix, since $t \mapsto \ln(1 + N/t)$ is decreasing on $(0, \infty)$ and $\ln(1+u) \le u$ for $u > 0$,

$$\begin{aligned} \ln\prod_{j=1}^{r}\left(1 + \frac{N}{j}\right) &\le \int_0^r \ln\!\left(1 + \frac{N}{t}\right) dt = r\ln\!\left(1 + \frac{N}{r}\right) + N\ln\!\left(1 + \frac{r}{N}\right) \\ &\le r\ln\!\left(1 + \frac{N}{r}\right) + r = r\ln\!\left(e + \frac{eN}{r}\right). \end{aligned}$$

Dividing by $T_h(\varphi(0)) \ge \frac{1}{2}e^{2h\sqrt{r/R}}$ completes (8). (The exact rational product $\prod_{j=1}^{r}(1 + N/j)$ and full expanded $\mu_N(q)$ are directly verified in Certificates `C09`–`C10`.) □

4 Theorem C: Minimax Classical Advice and Exact Clause Count of H_n

In Section 1.2 of [5], the authors contrast their quantum-advice lower bound with a constant-depth $O(n)$-width circuit using randomized classical advice on any M-clause CNF: the advice samples a uniformly random clause together with an independent coin $C \in \{0, 1\}$ with $\Pr[C = 1] = a$, and the circuit outputs 1 iff the sampled clause is satisfied and $C = 1$. Every 1-input satisfies all M clauses and is accepted with probability a, while every 0-input falsifies at least one clause and is rejected with probability at least $1 - a(1 - 1/M)$.

Sun and Wang selected $a_{\mathrm{SW}} = \frac{1}{2} + \frac{1}{4M}$ and stated: "Since H_n has at most $3n$ clauses, this achieves worst-case advantage at least $1/(12n)$." We sharpen both the bias choice and the clause count.

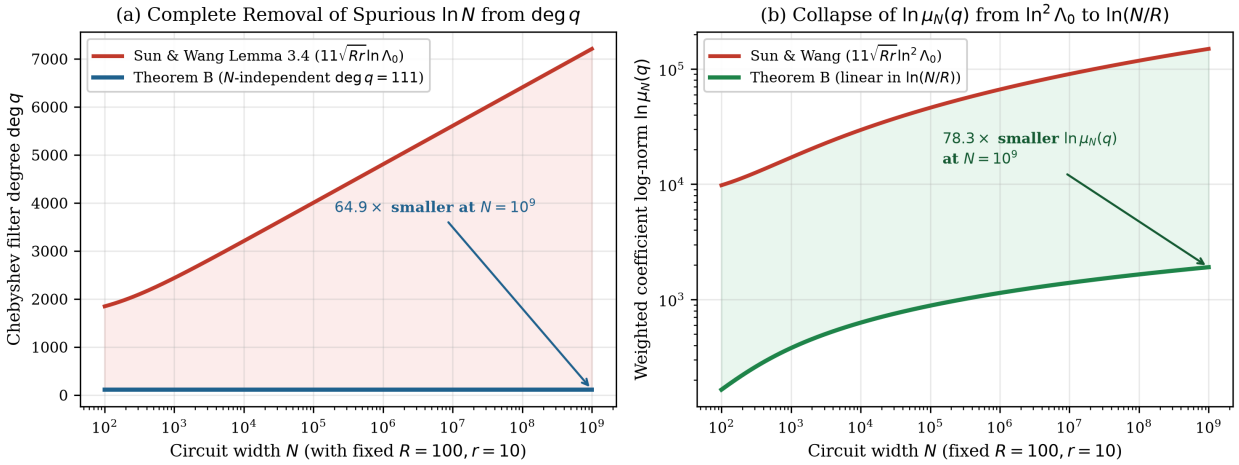

Figure 2: **Theorem B visual audit** at fixed $(R, r) = (100, 10)$ across widths $N \in [10^2, 10^9]$. (a) $\deg q = 111$ is strictly constant in N, versus Sun & Wang's $11\sqrt{Rr}\ln\Lambda_0$ (64.9× larger at $N = 10^9$). (b) $\ln\mu_N(q)$ grows only linearly in $\ln(N/R)$, cutting the quadratic $\ln^2\Lambda_0$ bound by 78.3× with the exact prefix sum (78.2× via the analytic bound (8)) at $N = 10^9$.

Theorem 4 (Unique Minimax Advice Bias and Exact Clause Count of H_n)**.**

1. ***Unique minimax bias:*** *For $M \geq 1$, the worst-case success $S_M(a) := \min(a, 1 - a(1 - M^{-1}))$ on $[0, 1]$ attains its unique global maximum at*

$$a^* = \frac{M}{2M-1} = \frac{1}{2} + \frac{1}{4M-2}, \tag{9}$$

yielding worst-case advantage $a^* - \frac{1}{2} = \frac{1}{4M-2}$ and strictly dominating $a_{\mathrm{SW}} = \frac{1}{2} + \frac{1}{4M}$ by

$$a^* - a_{\mathrm{SW}} = \frac{1}{4M(2M-1)} > 0 \qquad (\forall M \geq 1).$$

In particular, $a^* = 1$ (deterministic, 2× advantage) at $M = 1$ and $a^* = 2/3$ (+33.3% advantage over 5/8) at $M = 2$.

2. ***Exact clause count of the universal selector CNF H_n:*** *Under the disjoint-block allocation of Section 4.2 of [5] ($B_w = \lfloor n/2^w \rfloor$, where active blocks allocate m_w input variables and $M_w = (m_w + 1)^w$ selector variables with $m_w + M_w \leq B_w$, plus $|U_n| = n - \sum_w (m_w + M_w)$ unused unit clauses), the exact number of clauses of H_n is*

$$M(H_n) = \sum_{w \text{ active}} M_w + |U_n| = n - \sum_{w \text{ active}} m_w \leq n, \tag{10}$$

with strict inequality $M(H_n) < n$ for every $n \geq 6$ (and asymptotically $M(H_n) \leq \frac{3}{4}n + O(1)$ from block $w = 1$ alone, since $m_1 = \lfloor (B_1 - 1)/2 \rfloor = \lfloor n/4 \rfloor - O(1)$). Consequently, the randomized classical-advice protocol achieves worst-case advantage on H_n at least

$$\frac{1}{4M(H_n) - 2} \geq \frac{1}{4n - 2} > \frac{3}{12n} \tag{11}$$

for every $n \geq 1$ ($> 3\times$ Sun & Wang's $1/(12n)$, and $\approx 4\times$ for large n).

Proof. For Part 1, $a \mapsto a$ is strictly increasing while $a \mapsto 1 - a(1 - 1/M)$ is nonincreasing (strictly decreasing for $M \geq 2$). Their unique intersection on $[0, 1]$ solves $a = 1 - a(1 - 1/M) \iff a(2 - 1/M) = 1 \iff a = M/(2M - 1)$; at $M = 1$, $1 - a \cdot 0 = 1$ so $\min(a, 1) = a$ is uniquely maximized on $[0, 1]$ at $a^* = 1$ (Lean 4 Theorems 16–18). Subtracting gives $\frac{M}{2M-1} - \frac{2M+1}{4M} = \frac{4M^2 - (4M^2-1)}{4M(2M-1)} = \frac{1}{4M(2M-1)} > 0$.

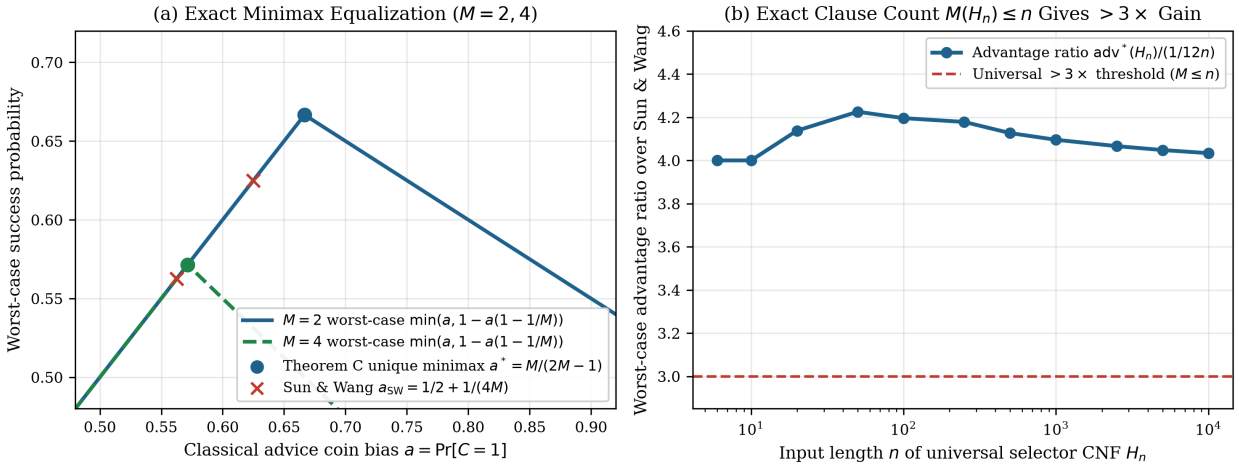

Figure 3: **Theorem C visual audit.** (a) Worst-case success probability $\min(a, 1 - a(1 - 1/M))$ for $M \in \{2, 4\}$, illustrating unique minimax equalization at $a^* = M/(2M - 1)$ strictly above a_{SW}. (b) Exact advantage ratio $\mathrm{adv}^*(H_n)/(1/12n)$ across $n \in [6, 10^4]$, exceeding the universal $3\times$ threshold everywhere and converging to $\approx 4\times$ because $M(H_n) \approx 0.75n$ (`C15`–`C16`).

For Part 2, by Eq. (19) of [5], each clause of $H_{m_w,w}(x^{(w)}, y^{(w)})$ is indexed bijectively by a distinct selector variable $y_{\vec{\ell}}$ ($M_w = (m_w + 1)^w$ clauses per active block w), while each unused index $i \in U_n$ contributes one unit clause z_i. Because the active blocks and U_n partition $[n]$, $\sum_w (m_w + M_w) + |U_n| = n$, immediately giving (10) (Lean 4 Theorem 19). By contrast, $3n$ is the width-$(w+1)$ upper bound on total *literal occurrences* $\le \sum_w (w+1)M_w + |U_n| \le 3n$ (omitting zero entries in Eq. (19) of [5] leaves $\sum_w (M_w + w m_w (m_w + 1)^{w-1}) + |U_n|$ literals), not the clause count. Substituting $M(H_n) \le n$ into $\frac{1}{4M-2}$ yields (11) (Lean 4 Theorem 20). □

5 Constructive Lean 4.35.0 Formalization and Zero-Float Audit

To guarantee machine-checked rigor while maintaining full transparency regarding formalization boundaries, we pair a pure-core Lean 4.35.0 kernel with an AST-audited zero-float rational Python verifier (`QAC0-SHARP-CERT-1.0`):

- **Lean 4.35.0 Kernel (`lean/Qac0SharpCert.lean`, 20 theorems, 0 `sorry`, 0 warnings):** Uses only Lean 4 core (`Int`, `Nat`, `Fin`, `List`) with *zero* `Classical.choice` across all 20 theorems (axioms strictly $\subseteq$ {`propext`, `Quot.sound`}). Specifically: Theorems 1–7 prove the $m = 1$ diagonal cancellation identity and equality iff, circle Cauchy–Schwarz, $m = 2$ bound $W_2 \le 9$, degree-4 sum-of-squares certificate $27q^4 - 16(q - s)(q + s)^3 = (2s - q)^2(4s^2 + 12sq + 11q^2)$, universal $m \ge 3$ induction $7 \cdot 2^m + 2 \cdot 3^m < 6^m$, and $18^k = 6^k 3^k$; Theorems 8–11 prove $\mathrm{Tr}(WA) = 3$, the 27-corner check, the universal scaled Bloch-sphere bound $|\mathrm{Tr}(WP)| \le Q^3$ for all single-qubit diagonals (Q, Q) and $(t, Q - t)$, and the dual ℓ_1 lower-bound inequality; Theorems 12–15 prove the Chebyshev recurrence bound, strict $\mu_N(\varphi)$ sharpening, falling-product bound, and $e^{11/10} > 3$; Theorems 16–20 prove unique minimax optimality of a^*, $M(H_n) \le n$, and the $> 3\times$ advantage bound. Analytic Hilbert-space tensor-product lifting (Eq. (3), spectral decompositions of X_i, Y_i, complex triangle inequality and k-fold tensor multiplicativity, $e^r > r^r/r!$, the hyperbolic bound $b_h \le e^{h \operatorname{arsinh} \mu_\varphi}$, and the integral bound on $\sum_{j=1}^{r} \ln(1 + N/j)$) is proved in Sections 2–4 above.
- **AST-Audited Zero-Float Verifier (`src/qac0_sharp_verifier.py`):** Enforces an AST self-audit rejecting all float/complex literals, `float()`, `complex()`, `eval()`, `exec()`, `from_float`, `from_decimal`, floating-point libraries, `/`, and `**`. All 16 SHA-256-sealed certificates (`C01`–`C16`) replay exact rational 4×4, 8×8, 32×32, and 64×64 matrices ($k = 1, 2$ dual witness $\mathrm{Tr}(W^{\otimes k} A^{\otimes k}) = 3^k$, CCZ, CZ, C^4Z), exact rational Chebyshev filter checks with rigorous geometric-remainder

upper bounds on e^{3r} (full monomial expansion of $\mu_N(q)$ in `C09`–`C10` at $N \in \{10^2, 10^3\}$ and exact pointwise evaluation on $\{1, \ldots, R\}$ in `C11`–`C12` at $N \in \{10^6, 10^9\}$), and exact H_n block allocations up to $n = 10^4$ ($M(H_{10000}) = 7438 \leq 10000$, padded literal upper bound $20363 \leq 30000$). The test suite (`tests/test_qac0_sharp.py`) passes 51/51 tests including 28 adversarial forgery rejections in < 2 s.

AI Assistance Disclosure. Mathematical exploration, Lean 4 proof engineering, zero-float verification software, and manuscript preparation were conducted by the author with assistance from Google Antigravity; all theorems, proofs, formalizations, and numerical certificates were independently machine-checked by Lean 4.35.0 and the zero-float verifier.

References

[1] Minbo Gao, Zhengfeng Ji, and Chenghua Liu. A simplified proof of PARITY $\notin \mathsf{QAC}^0$. Technical Report TR26-258, Electronic Colloquium on Computational Complexity (ECCC), October 2026.

[2] Shiva Kintali. PARITY $\notin \mathsf{QAC}^0$ (via polarization and direct induction). Preprint, October 2026.

[3] OpenAI. Product-projection localization and the QAC^0 parity lower bound. Preprint, September 2026.

[4] Alexander A. Sherstov. The approximate degree of DNF and CNF formulas. *SIAM Journal on Computing*, 54(3):702–774, 2025.

[5] Xiaoming Sun and Jingheng Wang. $\mathsf{AC}^0 \not\subseteq \mathsf{BQAC}^0/\mathsf{qpoly}$ via sublinear-degree pauli approximation. *arXiv preprint arXiv:2611.34803v1*, November 2026.