Exact Quantum Values and Permutation-Blind Maximizers in Cyclic Bell Inequalities

Sharp operator bounds, equality structure, and limits of Bell-value randomness certification

Alec Kriebel, Independent Researcher – Version 1.1, 9 August 2026

Source credit and result. Perito, D'Avino, Jung, Mironowicz, Acín, and Augusiak introduced the two cyclic Bell families, supplied canonical strategies and lower bounds, proved the first-family bound $d\sqrt{2}$, and conjectured their sharper attained value. Their NPA bounds through $d = 6$ agree with it to the reported precision; they also SOS-proved the second reduced value. These source results are credited, not reclaimed. For $\omega = e^{2\pi i/d}$ and $\mathcal{I}_d = \sum_y \mathrm{Re}[(A_0 + \omega^y A_1)B_y]$, version 1.1 proves

$$\mathcal{I}_d \le 2\csc\left(\frac{\pi}{2d}\right) I, \qquad \beta_q = \beta_{qa} = \beta_{qc} = 2\csc\left(\frac{\pi}{2d}\right).$$

The upper bound uses only cross-party commutation; the source's finite order-d strategy gives attainment. The first Hermitian augmentation has value one more.

Polar and scalar mechanism. For the canonical polar decomposition $C = V|C|$ and commuting unitary B,

$$\frac{|C| + |C^\dagger|}{2} - \mathrm{Re}(CB) = \frac{1}{2}P^\dagger P, \qquad P = |C^\dagger|^{1/2} - V|C|^{1/2}B.$$

An explicit strong-limit argument puts V in Alice's generated von Neumann algebra, hence in Bob's commutant. Support projections handle kernels; no inverse or unitary polar extension is used. Functional calculus for $U = A_0^\dagger A_1$ reduces the remaining gap to

$$\sum_y |1 + \omega^y z| \le 2\csc\frac{\pi}{2d}, \qquad \text{equality iff } z^d = (-1)^{d-1}.$$

Finite-dimensional support rigidity. Every attained finite-dimensional *tensor-product* exact maximizer of the first *augmented* family has the following property: if $K = \operatorname{supp}\rho_A$, then K reduces U, every one of the d equality roots occurs in $U|_K$ with equal multiplicity, and $d \mid \dim K$. The proof uses the augmented stabilizer, Schmidt-support and kernel-safe polar cancellation, adjacent-phase reflections, and a rank bound. It says nothing about $K^\perp$, approximate or general commuting-operator maximizers, or a complete classification/self-test.

Conditional permutation orbit. Under the displayed product-one and paired polar-phase hypotheses, any permutation κ gives $A_0 = X$, $A_1 = X\,\mathrm{diag}(z_{\kappa_j})$ and matching Bob weighted shifts that remain full-spectrum order-d exact maximizers. All local first moments and complex first-harmonic correlators are permutation-independent. This is a sufficient family, not a classification.

Nonuniform maximizers and exact $d = 4$ entropy. With $q_0 = 1$, $q_{j+1} = z_{\kappa_j} q_j$,

$$p_\kappa(a, b \mid 1, d) = \frac{|\widehat{q}_{-(a+b)}|^2}{d^3}.$$

Both marginals are uniform. For every $d \geq 4$, swapping the final two roots gives $|R_2| = 4\sin(\pi/d)\sin(3\pi/d) > 0$, hence a nonuniform exact maximizer of both augmented families and

$$G \geq \frac{1}{d^2} + \frac{2\sin(\pi/d)\sin(3\pi/d)}{d^2(d-1)} > \frac{1}{d^2}.$$

At $d = 4$, exact $\mathbb{Q}(\zeta_{16})$ arithmetic gives probabilities $1/32$ and $3/32$, so $G = 3/32$. The displayed trivial-Eve realization has $H_{\min} = 5 - \log_2 3 = 3.415037\ldots$ bits, not four. This is an upper bound on value-only worst-case min-entropy, not its exact optimized value.

Precise Conjecture 2 consequence. Using the displayed first augmented operator, whose maximum is $M_d + 1$, the construction disproves exactly

$$\langle \overline{\mathcal{I}}_d \rangle = M_d + 1 \implies G(AB \mid 1, d, E) = 1/d^2 \qquad (d \geq 4).$$

Thus scalar maximum alone does not certify the proposed $2\log_2 d$ bits or identify the canonical behavior. This does not challenge an SDP fixing the complete canonical behavior, show that the canonical strategy lacks maximal randomness, or determine the worst-case guessing probability.

Second-family convention. The SOS of Perito et al., Eqs. (22)–(23), with prefactor $1/(2d)$, supplies global optimality. The exact compression $\widehat{B}_\ell = d\lambda_\ell D_\ell$ and $A_\ell = \overline{D_\ell}$ annihilate every factor. The alternative source- appendix convention adjoints all Bob observables and changes the target table only by $b \mapsto -b$; value, nonuniformity, guessing, and entropy are unchanged.

Secondary benchmarks. The prior-art binary 2×2 score has $q = qa = qc = 3\sqrt{3}$ by a two-square SOS; every attaining finite-dimensional tensor-product strategy satisfies $\sigma_E^{ab} = \rho_E/4$ at the target pair. Separately, the private-MUB lemma is only a sufficient composition criterion: a private reference PVM, perfect Bob matching, and a supported MUB sandwich imply $\sigma_E^{a,\pi(b)} = \rho_E/d^2$. It is not a low-setting existence theorem.

Focused replay.

```
python3 verification/verify_rigidity.py
python3 verification/verify_exact_benchmarks.py
python3 verification/verify_private_mub_binary.py
./reproduce.sh
```

Canonical page. Open the manuscript and complete artifact index.