No universal finite collapse of the upper PPT entanglement-cost hierarchy

Dev D. Goyal

30 October 2026; revised 1 November 2026

Abstract

Lami, Mele and Regula introduced converging upper and lower semidefinite hierarchies for the exact PPT entanglement cost, conjectured that the lower one collapses at a finite level and asked the analogous question for the upper one [1]; Brinster et al. proved arbitrary-depth noncollapse of the lower χ hierarchy and explicitly left the upper κ question open [2]. We answer the upper question: for every integer $q \geq 1$, there is a bipartite density matrix ρ_q for which $\kappa_{q+1}(\rho_q) < \kappa_q(\rho_q)$ in the original upper hierarchy. The witnesses can be chosen as full-rank product-Bell diagonal states on $\mathbb{C}^{d_q} \otimes \mathbb{C}^{d_q}$, where $d_q = 2^{(2q+4)4^q}$. Consequently, no fixed finite upper level is exact for all bipartite states.

The proof constructs arbitrarily long strictly positive reverse orbits of absolute partial transposition. A finite weighted-sign argument extends an orbit while controlling every coordinate. Strict positivity then turns a known telescoping dual identity into a quantitative obstruction for every shorter feasible upper chain. Bell expectations transfer the obstruction to unrestricted complex auxiliary matrices. The central all-depth theorem, including its physical state assumptions and the original SDP constraints, is formalized in Lean 4 with pinned dependencies; its operational consequences rest on cited published theorems.

The originality assessment concerns the positive inverse-orbit construction and the resulting all-depth upper separation. It is supported by targeted searches before and after the proof, but historical priority remains qualified. The dimension bound is deliberately generous; the result does not claim efficient small witnesses, noncollapse in one fixed dimension, or impossibility of other single-letter formulas.

Contents

1 Introduction and main theorem **3**

2 Bell coordinates and the unrestricted operator problem **4**

3 A positive reverse-step lemma **5**

4 Explicit parameters and arbitrarily long positive orbits **6**

5 Why full support forces an upper-level gap **7**

6 Normalization and proof of the main theorem **8**

7 Novelty assessment and significance **9**

7.1 What was already known 9

7.2 The contribution assessed here 9

7.3 Search coverage and qualified priority conclusion 9

8 Lean correspondence and reproducibility **10**

8.1 What is formalized 10

8.2 Pinned clean replay 11

9 Independent scrutiny and remaining limitations **11**

10 Research process and AI disclosure **12**

1 Introduction and main theorem

Lami, Mele and Regula expressed the exact PPT entanglement cost through two converging hierarchies of semidefinite programs, an upper hierarchy κ_q and a lower hierarchy χ_q. They conjectured finite collapse of the lower hierarchy at a specific level and asked the analogous finite-collapse question for the upper hierarchy [1]. Brinster et al. showed that the lower hierarchy does not collapse at any finite depth and explicitly left the upper question open [2]. Theorem 1 answers it: every adjacent pair of upper levels is separated by a full-rank state. The new ingredient is a positive inverse-orbit construction for absolute partial transposition (Sections 3–4); the Bell transform formalism and the telescoping dual identity are due to Brinster et al. and are credited where used (Section 7). The Lean development of Section 8 proves the all-depth theorem for the actual semidefinite programs, while the operational consequences below rest on published theorems of Lami, Mele and Regula that are cited, not formalized.

For partial transposition Γ on the second local system, define

$$\kappa_q(\rho) = \inf\left\{ \operatorname{Tr} S_{q-1} : \begin{array}{ll} S_{-1} = \rho, & S_i \succeq \pm S_{i-1}^{\Gamma} \quad (0 \le i < q), \\ S_{q-1}^{\Gamma} \succeq 0 & \end{array} \right\}. \tag{1}$$

This is the original indexing and terminal constraint in Lami–Mele–Regula [1]. The auxiliary matrices range over the entire complex matrix space. The PSD sandwich constraints force Hermiticity and positivity of the auxiliaries, so their traces are real. The formal definition uses the real part of the trace and proves this correspondence.

Theorem 1 (Strictness at every adjacent upper level)**.** *For every $q \ge 1$, put*

$$n_q = (2q+4)4^q, \qquad d_q = 2^{n_q}.$$

There is a positive-definite, trace-one matrix ρ_q on $\mathbb{C}^{d_q} \otimes \mathbb{C}^{d_q}$, diagonal in a product-Bell basis with strictly positive coefficients, such that

$$\kappa_{q+1}(\rho_q) < \kappa_q(\rho_q). \tag{2}$$

Thus there is no integer $q \ge 1$ for which $\kappa_q = \kappa_{q+1}$ holds on every finite-dimensional bipartite quantum state.

The construction is an existence proof over a finite collection of sign vectors, with an explicit dimension bound. It does not provide a short closed formula for the successful signs or a practical table of all state coefficients. This distinction does not affect the statement in Theorem 1: Lean proves that the required state exists, rather than accepting a missing numerical certificate as an assumption.

The operational consequence uses the published bound

$$E_{c,\mathrm{PPT}}^{\mathrm{exact}}(\rho) \le \log_2 \kappa_r(\rho) \qquad (r \ge 1). \tag{3}$$

For the witness at level q, this implies $E_{c,\mathrm{PPT}}^{\mathrm{exact}}(\rho_q) < \log_2 \kappa_q(\rho_q)$. Therefore no fixed upper level gives a universal exact PPT-cost formula. Equation (3) is established prior mathematics [1, Theorem S26, Eq. (S89)]; the operational theory is not introduced as a Lean axiom or claimed to be re-formalized in this project.

A further consequence uses the published regularization identity $E_{\kappa,q}^{\infty} = E_{c,\mathrm{PPT}}^{\mathrm{exact}}$ [1, Remark S27, Eq. (S98)]. For each finite q, the witness must satisfy $\kappa_q(\rho_q^{\otimes m}) < \kappa_q(\rho_q)^m$ for at least one integer $m \ge 2$; otherwise its regularized logarithmic value could not be smaller than its one-copy value. Thus every finite upper level fails tensor additivity somewhere. This is a consequence of the

formally proved separation and a cited prior theorem, not a separately formalized operational result.

The quantifiers matter. The dimension is allowed to grow with q. This does not prove that one fixed state or one fixed dimension has infinitely many strict upper improvements. The construction itself reaches a simple terminal point after $q+1$ steps. The theorem also does not rule out a different single-letter characterization of exact PPT cost.

2 Bell coordinates and the unrestricted operator problem

Use the one-pair Bell order $(\Phi^+, \Phi^-, \Psi^+, \Psi^-)$ and the matrix

$$T_1 = \frac{1}{2}\begin{pmatrix} 1 & 1 & 1 & -1 \\ 1 & 1 & -1 & 1 \\ 1 & -1 & 1 & 1 \\ -1 & 1 & 1 & 1 \end{pmatrix}, \qquad T = T_1^{\otimes n}. \tag{4}$$

Let $d = 2^n$ and $N = 4^n = d^2$. The matrix T is symmetric, orthogonal, involutive, fixes $\mathbf{1}$, and has entries of absolute value $1/d$. In particular,

$$T^2 = I, \quad T\mathbf{1} = \mathbf{1}, \quad |T_{ij}| = d^{-1}, \quad \sum_i (Tx)_i^2 = \sum_i x_i^2. \tag{5}$$

Let P_j be the rank-one product-Bell projectors and write

$$\mathcal{R}(x) = \sum_{j=1}^{N} x_j P_j.$$

The Bell bridge proves

$$\mathcal{R}(x)^\Gamma = \mathcal{R}(Tx), \quad \operatorname{Tr}\mathcal{R}(x) = \sum_j x_j, \quad x \geq 0 \Longrightarrow \mathcal{R}(x) \succeq 0. \tag{6}$$

Strictly positive coefficients give a positive-definite synthesized matrix.

For an arbitrary complex matrix X, define its Bell expectations by

$$\mathcal{E}(X)_j = \operatorname{Re}\operatorname{Tr}(P_j X).$$

The formal bridge establishes

$$X \succeq 0 \Longrightarrow \mathcal{E}(X) \geq 0, \qquad \mathcal{E}(X^\Gamma) = T\mathcal{E}(X), \qquad \sum_j \mathcal{E}(X)_j = \operatorname{Re}\operatorname{Tr} X. \tag{7}$$

Also $\mathcal{E}(\mathcal{R}(x)) = x$.

Consequently, every feasible complex-matrix chain for $\mathcal{R}(x)$ yields a coefficient chain

$$s_0 = x, \qquad s_{i+1} \geq |Ts_i| \ (0 \leq i < q), \qquad Ts_q \geq 0, \tag{8}$$

with the same real trace objective $\sum_j s_q(j)$. Conversely, every such coefficient chain synthesizes a feasible matrix chain. Lower bounds below therefore apply to all complex auxiliaries. No assertion that an optimizer is Bell diagonal is needed.

The manuscript uses the convenient chain indexing $s_0 = x$ and s_q terminal. It matches (1) by $s_i = \mathcal{E}(S_{i-1})$. The Lean definitions and compatibility lemmas make this shift explicit.

3 A positive reverse-step lemma

Define the absolute-transform map

$$F(x) = |Tx|,$$

with absolute value coordinatewise. The central new construction extends a reverse orbit $F(w) = v$ while preserving strict positivity of w.

Lemma 1 (Finite weighted-sign inverse)**.** *Suppose* $0 \le \varepsilon < 1$,

$$|v_i - 1| \le \varepsilon \quad \textit{for every } i, \qquad \sum_i v_i^2 = N.$$

Let $a, t > 0$ *satisfy*

$$\frac{2t}{d} \le 1, \qquad 2N \exp(2\varepsilon t^2) < \exp(ta). \tag{9}$$

There is a sign vector $s \in \{-1, 1\}^N$ *such that*

$$w = T(s \odot v), \qquad F(w) = v, \qquad |w_j - 1| < \varepsilon + a \quad \textit{for every } j, \qquad \sum_j w_j^2 = N.$$

If $\varepsilon + a < 1$*, then* w *is strictly positive.*

Proof. Put $c = 1 - \varepsilon$. Since $v_i \ge c > 0$, assign the two signs the weights

$$\Pr(s_i = 1) = \frac{v_i + c}{2v_i}, \qquad \Pr(s_i = -1) = \frac{v_i - c}{2v_i}. \tag{10}$$

These weights are nonnegative and sum to one. Use their finite product distribution. For each coordinate,

$$\mathbb{E}(s_i v_i) = c, \qquad \mathbb{E}(s_i v_i - c)^2 = v_i^2 - c^2.$$

All expectations here can be read simply as explicit finite weighted sums; no measure-theoretic existence result is required.

For a fixed row j, set

$$X_i = T_{ji}(s_i v_i - c), \qquad Z_j = \sum_i X_i = w_j - c.$$

The row-sum identity in (5) gives the last equality. The variables are centered, and flatness gives

$$\sum_i \mathbb{E} X_i^2 = \frac{1}{d^2} \sum_i (v_i^2 - c^2) = 1 - c^2 \le 2\varepsilon. \tag{11}$$

Moreover,

$$|X_i| \le \frac{v_i + c}{d} \le \frac{2}{d}.$$

Thus $|tX_i| \le 1$ by (9). The elementary inequality $e^z \le 1 + z + z^2$ for $|z| \le 1$ implies

$$\mathbb{E} e^{\pm tX_i} \le 1 + t^2 \mathbb{E} X_i^2 \le e^{t^2 \mathbb{E} X_i^2}.$$

Factoring the finite product distribution and using (11),

$$\mathbb{E} e^{tZ_j} \le e^{2\varepsilon t^2}, \qquad \mathbb{E} e^{-tZ_j} \le e^{2\varepsilon t^2}. \tag{12}$$

Consider the nonnegative potential

$$\Phi(s) = \sum_{j=1}^{N} \left(e^{tZ_j} + e^{-tZ_j} \right).$$

Its weighted average is at most $2Ne^{2\varepsilon t^2} < e^{ta}$. Hence at least one sign vector has $\Phi(s) < e^{ta}$. Every summand is nonnegative, so $e^{tZ_j} < e^{ta}$ and $e^{-tZ_j} < e^{ta}$ for all j. Since $t > 0$, this gives $|Z_j| < a$, and consequently $|w_j - 1| < a + \varepsilon$.

Finally, $Tw = s \odot v$ by involutivity. Since $v > 0$, $F(w) = |s \odot v| = v$. Orthogonality and $s_i^2 = 1$ preserve the squared Euclidean norm exactly. □

The small variance in (11) is essential to this argument. A bound that uses only the ranges of the signs would discard the gain obtained from v being close to $\mathbf{1}$. The Lean proof establishes the two-point moments, product factorization, exponential estimates, and existence of a successful sign vector directly. It does not invoke an unproved concentration inequality or a numerical random search.

4 Explicit parameters and arbitrarily long positive orbits

Fix $q \geq 1$ and put

$$p = 4^q, \qquad h = 16p = 2^{2q+4}, \qquad d = h^p = 2^{(2q+4)4^q}, \qquad N = d^2. \tag{13}$$

For $k \geq 0$, define

$$\varepsilon_k = \frac{2}{h^{4^k}}, \qquad a_k = \frac{1}{h^{4^k}}, \qquad t_k = \frac{(h^{4^k})^3}{8}. \tag{14}$$

The subscript k records the number of remaining error contractions, so inverse steps use $k = q-1, q-2, \ldots, 0$.

Lemma 2 (The concentration conditions hold)**.** *For $0 \leq k < q$, Lemma 1 applies with $\varepsilon = \varepsilon_{k+1}$, $a = a_k$, and $t = t_k$. Moreover,*

$$\varepsilon_{k+1} + a_k \leq \varepsilon_k \leq \frac{1}{8}.$$

Proof. Write $b = h^{4^k}$. Then $\varepsilon_{k+1} = 2/b^4$, $a_k = 1/b$, and $t_k = b^3/8$. Since $b \geq h \geq 16$, $2/b^4 + 1/b \leq 2/b \leq 1/8$. Also $3 \cdot 4^k \leq 4^q$, so $2t_k/d = b^3/(4d) \leq 1$.

The exponent identities are

$$2\varepsilon_{k+1}t_k^2 = \frac{b^2}{16}, \qquad t_k a_k = \frac{b^2}{8}.$$

It remains to prove $2d^2 < e^{b^2/16}$. Using $q \leq 4^q = p$, $p \geq 1$, and $\log 2 < 1$,

$$\begin{aligned}\log(2d^2) &= (1 + 4pq + 8p)\log 2 \\ &< 1 + 4pq + 8p \leq 13p^2 < 16p^2 = \frac{h^2}{16} \leq \frac{b^2}{16}.\end{aligned}$$

Multiplying the resulting strict inequality by $e^{b^2/16}$ proves $2d^2 e^{2\varepsilon_{k+1}t_k^2} < e^{t_k a_k}$. □

Choose any Bell label j and define

$$v_0 = \mathbf{1}, \qquad v_1 = T(\mathbf{1} - 2e_j) = \mathbf{1} - 2Te_j. \tag{15}$$

Then

$$F(v_1) = v_0, \quad Tv_1 = \mathbf{1} - 2e_j, \quad \|v_1\|_2^2 = N, \quad |v_1(i) - 1| = \frac{2}{d} = \varepsilon_q.$$

In particular $v_1 > 0$ and its partial-transpose coefficient at j is -1.

Starting from v_1, apply the inverse lemma q times, with the conditions provided by Lemma 2. This produces

$$v_0, v_1, \ldots, v_{q+1}, \qquad F(v_{i+1}) = v_i, \tag{16}$$

such that every coordinate of every v_i is at least $7/8$, and every $\|v_i\|_2^2 = N$. The formal orbit interface retains the weaker bound $1/2$, which is enough for the subsequent quantitative argument.

Every nonterminal reverse stage is also non-PPT. Indeed, if $Tv_i \geq 0$ and $F(v_i) = v_{i-1}$, then $Tv_i = v_{i-1}$ and $Tv_{i-1} = v_i \geq 0$. Induction would contradict the known negative coordinate of Tv_1. Only the defect at v_1 is needed by the central Lean theorem.

To obtain a definite finite specification, one may select at each step the first sign vector, in a fixed order, satisfying the rational coordinate bound for the next error level. Existence is proved by the preceding argument. This selection is finite but astronomically expensive at the displayed dimensions. The proof artifacts check the general existence argument without expanding these state tables.

5 Why full support forces an upper-level gap

Reverse the constructed orbit:

$$A_i = v_{q+1-i} \quad (0 \leq i \leq q+1).$$

Then

$$A_{i+1} = F(A_i), \quad A_{q+1} = \mathbf{1}, \quad A_i \geq \tfrac{1}{2}\mathbf{1}, \quad \|A_0\|_2^2 = N, \quad (TA_q)_j = -1. \tag{17}$$

The greedy chain $A_0, A_1, \ldots, A_{q+1}$ is feasible at upper level $q+1$ and has terminal trace N.

We prove a strict lower bound on *every* level-q chain. This avoids any need to infer a strict infimum gap merely from nonattainment of equality.

Lemma 3 (Quantitative obstruction)**.** *If $s_0 = A_0$, $s_{i+1} \geq |Ts_i|$ for $0 \leq i < q$, and $Ts_q \geq 0$, then*

$$\sum_j s_q(j) \geq N + \delta_q, \qquad \delta_q = \frac{1}{2(d+1)^{q+1}} > 0. \tag{18}$$

Proof. For each stage define $D_i = \mathrm{diag}(\mathrm{sign}(TA_i))$, choosing sign $+1$ at zero. Thus

$$A_{i+1} = D_i T A_i, \qquad T D_i A_{i+1} = A_i.$$

Extend the candidate chain by $s_{q+1} = Ts_q$. Since $Ts_q \geq 0$, this extension satisfies the final sandwich as well and preserves total trace. The active slack

$$r_i = s_{i+1} - D_i T s_i$$

is coordinatewise nonnegative. Orthogonality gives the exact identity

$$A_{i+1} \cdot s_{i+1} - A_i \cdot s_i = A_{i+1} \cdot r_i \geq 0. \tag{19}$$

This underlying telescoping identity is the mechanism already used in Brinster et al., Appendix C, Lemma 7 [2]. The present argument uses an orbit with a uniform positive coordinate margin.

Put $\Delta = \sum_j s_q(j) - N$. Telescoping (19) from 0 to q shows $\Delta \geq 0$. Each individual increment is at most Δ. Since $A_{i+1} \geq \frac{1}{2}\mathbf{1}$, for every coordinate and $i < q$,

$$0 \leq r_i(j) \leq 2\Delta.$$

Let $e_i = s_i - A_i$. Then $e_0 = 0$ and

$$e_{i+1} = D_i T e_i + r_i. \tag{20}$$

Flatness gives $\|Tz\|_\infty \le d\|z\|_\infty$. A direct induction in (20) yields

$$\|e_i\|_\infty \le 2\Delta(d+1)^i \qquad (0 \le i \le q).$$

At the selected terminal coordinate, $Ts_q \ge 0$ and $(TA_q)_j = -1$, hence

$$1 \le |(Te_q)_j| \le d\|e_q\|_\infty \le 2\Delta d(d+1)^q \le 2\Delta(d+1)^{q+1}.$$

This proves (18). □

The crucial distinction is between an arbitrarily deep absolute-transpose orbit and an optimized SDP gap. An orbit with zero coordinates does not control all active slacks. The positive margin in (17) provides exactly that control. The proof does not assume a greedy optimizer.

A Euclidean-norm refinement gives a stronger conventional estimate, but it is unnecessary for the central result. The theorem stated here and the verification discussion use the conservative bound (18), which is the one implemented in Lean.

6 Normalization and proof of the main theorem

Let $x = A_0 = v_{q+1}$ and $\tau = \sum_j x_j$. Since $x_j \ge 1/2$, we have $\tau > 0$. Define

$$\rho_q = \mathcal{R}(x/\tau). \tag{21}$$

All Bell coefficients are strictly positive and sum to one, so ρ_q is a full-rank density matrix on the promised bipartite space.

The greedy chain synthesizes a feasible upper-level $(q+1)$ witness with trace N/τ. Thus

$$\kappa_{q+1}(\rho_q) \le N/\tau.$$

For any complex-matrix level-q feasible chain, extract its Bell expectations using (7) and multiply the coefficients by τ. This is a feasible raw chain with input x. Lemma 3 gives

$$\operatorname{Re}\operatorname{Tr} S_{q-1} \ge \frac{N+\delta_q}{\tau}.$$

Taking the infimum gives

$$\kappa_q(\rho_q) \ge \frac{N+\delta_q}{\tau} > \frac{N}{\tau} \ge \kappa_{q+1}(\rho_q), \tag{22}$$

proving Theorem 1.

The formal proof establishes that the sets of feasible objective values are nonempty and bounded below before applying conditional-infimum lemmas. Nonemptiness follows, for example, by taking the first auxiliary coefficient vector to be a sufficiently large constant and keeping that constant for all subsequent stages. Its transform is the same nonnegative constant. The terminal PPT condition implies a nonnegative trace, giving a lower bound.

The formalization therefore does not hide an empty optimization problem, an assumed optimizer, an unproved duality theorem, or an artificial restriction on the auxiliary matrices. The sandwich constraints also prove that the real-trace formal objective agrees with the usual Hermitian SDP objective.

7 Novelty assessment and significance

7.1 What was already known

Lami, Mele, and Regula introduced converging upper and lower SDP hierarchies and proved their relation to exact PPT cost [1]. Their Supplemental Material, Conjecture S33, proposed finite collapse at a specific lower level, with a corresponding upper-level formula. The same subsection asks the analogous finite-collapse question for the upper hierarchy.

Brinster et al. construct arbitrary-depth separations for the lower χ hierarchy [2]. Section IV, immediately after Eqs. (47)–(48), explicitly observes that their states do not show upper κ noncollapse and leaves a finite upper formula open. Their Appendix C supplies the Bell transform formalism, multinegative examples, and the telescoping dual argument. These components receive prior credit; they are not claimed as discoveries of the present work.

Earlier unpublished computations in the same research process had verified only $\kappa_2 > \kappa_3$ and $\kappa_3 > \kappa_4$ on explicit states, with related tensor nonmultiplicativity, and found one further finite example without a complete Lean proof. None of those finite examples establishes the quantifier over every q proved here, and the present proof does not use them.

7.2 The contribution assessed here

The new mechanism is a positive inverse-orbit construction using biased signs and a small-variance finite exponential potential. It supplies exactly the full-support property that makes the upper optimization quantitatively rigid. Together, these steps yield a strict gap at every adjacent finite upper level for actual quantum states.

This resolves the specific remaining universal upper-hierarchy question in the located primary literature. It is a structural result rather than a small change in a constant or one additional numerical level. It explains why optimized upper bounds can require arbitrarily large depth and removes every fixed upper level as a universal exact-cost candidate.

The significance is bounded by the theorem's scope. The proof gives a very large dimension bound and a small guaranteed normalized gap. It does not improve the published approximation algorithm, optimize witness size, show unbounded depth at a fixed dimension, or establish a computational-hardness result. It leaves open other possible exact formulas. A specialist paper on this result would need normal external review; no venue acceptance or independent expert endorsement is asserted.

7.3 Search coverage and qualified priority conclusion

The audit refreshed the primary arXiv records and full texts, searched the upper/lower terminology and normalized-body formulations, and checked finite collapse, multinegativity, full-rank and dense examples, inverse absolute partial transposition, positive signed-Hadamard orbits, and related entanglement-cost additivity formulations. The normalized upper bodies $\widetilde{\mathrm{PPT}}_q = \{X \geq 0 : \kappa_q(X) \leq 1\}$ were distinguished from the lower bodies used in earlier contraction results. Computable lower bounds on entanglement cost, such as that of Wang, Jing and Zhu [3], are related background.

Repository checks included the index and priority disclosures of this collection, OpenAI/math, and the public trureturing problem/proof registry. The relevant quantum subtrees of the registry were compared with an earlier detailed audit, and searches of its issues and pull requests were refreshed. The located first-level 2×3 problem #17758 concerns a distinct κ_1 versus logarithmic-negativity question and does not supply the present arbitrary-depth upper theorem.

A related September 2026 doctoral-thesis abstract was located in the Scuola Normale Superiore repository; its full PDF returned an access error and was not audited. Targeted MathOverflow and mathematical-database searches did not locate a matching upper theorem. These access and indexing limitations are retained in the novelty record.

The fresh Brinster arXiv record exposed only version 1, submitted 17 September 2026, when checked. No inspected source contained the positive inverse-orbit result or all-depth upper separation proved here. The search records of that audit are not part of this release.

Priority assessment: the theorem has a defensible, specifically supported originality claim relative to the located literature. Historical priority is still qualified. Search cannot exclude unindexed theses, unpublished manuscripts, private ongoing work, or a proof filed under unexpected terminology. Absence of a search hit is not treated as proof of novelty. Lean verifies the mathematical implication, not historical priority.

8 Lean correspondence and reproducibility

8.1 What is formalized

The central statement ranges over every natural $q \geq 1$ and produces a normalized complex bipartite density matrix. The proof includes the finite sign existence lemma, all explicit parameter inequalities, the full finite induction, the concrete Bell partial-transpose instance, the quantitative coefficient obstruction, and its transfer to arbitrary complex-matrix SDP chains. It is not a proof merely of a finite table, a numerical residual, or an abstract involution without a physical instance.

The proof is organized as follows.

Module	Mathematical role
`FoldSystem`	Flat symmetric involution; norm, trace, and dual-pairing identities.
`SignConcentration`	Two-point weights and moments; finite product exponential bound.
`InverseSign`	Simultaneous coordinate control and an exact positive inverse step.
`ParameterBounds`	Explicit dimension, error, tilt, and strict exponential conditions.
`ReverseOrbit`	Initial defect and the complete finite reverse-orbit induction.
`BellFoldSystem`	Actual product-Bell partial transposition satisfies every abstract assumption.
`ChainGap`	Positive lower bound for every shorter feasible coefficient chain.
`ArbitraryHierarchy`	Original arbitrary-depth complex SDP, nonempty infima, normalization, and matrix gap.
`BellPositivity`	Strictly positive Bell weights imply positive definiteness.
`AllDepth`	Final all-depth physical theorem and no universal finite collapse.

Earlier small structural modules provide Bell projectors, synthesis, expectations, and PSD facts. The earlier research's large finite data tables and optimization scripts are not dependencies of the new theorem.

The use of `noncomputable` marks ordinary real-valued definitions and existential choices. It does not add axioms, bypass the kernel, or replace proofs with a solver result. Every sign-existence claim is proved. The project uses neither `sorry`/`admit`, custom axioms, nor a native decision oracle. Only the usual Lean foundations `propext`, `Classical.choice`, and `Quot.sound` are allowed by the audit.

8.2 Pinned clean replay

The release pins Lean 4.19.0, commit `6caaee842e94`, and mathlib [4] at

`c44e0c8ee63ca166450922a373c7409c5d26b00b`.

Every transitive dependency is pinned by full git commit. The verifier rejects dirty dependency checkouts, changed frozen sources, unlisted project imports, prohibited proof shortcuts, unexpected axioms, and incomplete audit output.

The replay freshly compiles 17 mathematical modules and the generated axiom audit, checking 256 public source declarations. For this release, the all-depth theorem and the full-rank state statement were replayed from source in this way, with the kernel at trust zero: all 17 modules and the complete 256-declaration axiom audit succeeded, and every audited declaration depends only on `propext`, `Classical.choice` and `Quot.sound`. Each compiler invocation checks the kernel at trust zero and disables automatic implicit parameters. The exact flags are

`-j1 --trust=0 -DwarningAsError=true -DautoImplicit=false`

The replay recorded with this release used a Mathlib checkout at the pinned revision whose modules had been built from source on the same machine; it took about seven minutes with two parallel jobs.

From the `lean/` directory of this package, after preparing the exact Lean release and dependencies, run:

```
lake exe cache get
python3 verify_release.py \
  --lean-home /absolute/path/to/lean-4.19.0-linux \
  --mathlib .lake/packages/mathlib \
  --packages-dir .lake/packages \
  --out /absolute/path/to/new-empty-verification-directory \
  --jobs 2
```

Do not run `lake update`, which changes the reviewed dependency manifest. A pristine standalone mathlib checkout at the same pin may also be supplied; `REPRODUCE.txt` gives that command.

Success requires a receipt with status `PASS`, zero exit status for every module, and an axiom inventory matching the complete reviewed list. The receipt records exact commands, source/object/log hashes, dependency revisions, and timing. The receipt and logs of the replay for this release are in the package's `verification/` directory, and the continuous integration workflow repeats the replay on every change.

9 Independent scrutiny and remaining limitations

In the research process, AI reviewers challenged the proof on four points: whether the inverse sign distribution is genuinely defined and normalized; whether its variance and simultaneous bounds have the correct dimension factors; whether full support controls an optimized upper chain; and whether the formal statement is the actual complex-matrix SDP. Each issue is addressed by explicit lemmas, and the final theorem discharges the conditional orbit hypotheses.

The development explored finite reverse-mask searches, dual exposed faces, and fixed-dimension nonpolyhedral possibilities before the general weighted sign argument. Those explorations did not establish the all-depth theorem and are not used as proof evidence. Earlier shortcut approaches, namely flag padding, face extraction and greedy depth, are not used as assumptions. The

successful proof proceeds through the complete inverse lemma and quantitative slack estimate given above.

The final artifacts establish mathematical correctness subject to the standard trust in Lean's kernel and the pinned library foundations. They do not replace specialist peer review of the theorem's presentation, priority, or practical interpretation. The main unresolved quantitative questions are much smaller witness dimensions, useful normalized gaps, and simple successful sign patterns. These are future questions rather than qualifications to the proved all-depth existence theorem.

10 Research process and AI disclosure

This is fully autonomous agentic research. Dev D. Goyal is responsible for publication. The work has not been externally peer reviewed, no expert validation is claimed, and the novelty assessment is provisional.

The theorem, its proof, the Lean formalization, the literature searches and an earlier version of this note were produced by autonomous AI research agents in an earlier session. The records of that session do not name the AI system used.

Opus 5.5 by Anthropic used with Claude Code assessed the package for this release. It checked the definition of κ_q in the Lean sources against the hierarchy stated in Section 1, including the partial transpose on the second system, the terminal PPT condition and the unrestricted complex auxiliaries; rederived by hand the inverse-sign lemma, the parameter inequalities, the initial defect and the quantitative chain gap; ran the fresh Lean replay described in Section 8; and revised this manuscript. On 1 November 2026 it also revised the presentation (front matter, abstract, introduction and this section) without changing the mathematics. The primary arXiv records could not be opened from its environment, so the comparison with the cited literature relies on the earlier session's audit. AI checks are not human peer review.

Lean verification establishes only the formal statements described in Section 8. The remaining limitations are stated in Sections 7 and 9.

References

[1] L. Lami, F. A. Mele, and B. Regula. *Computable entanglement cost under positive partial transpose operations.* Physical Review Letters **134**, 090202 (2025). arXiv:2405.09613v2, 7 March 2025.

[2] R. Brinster, T. V. Kondra, H. Kampermann, D. Bruß, and N. Wyderka. *Multinegativity and single-letter formulas for asymptotic entanglement.* arXiv:2609.20698v1, 17 September 2026.

[3] X. Wang, M. Jing, and C. Zhu. *Computable and Faithful Lower Bound on Entanglement Cost.* arXiv:2311.10649v4, 7 May 2025.

[4] The Lean and mathlib communities. *Lean 4.19.0 and mathlib*, pinned source and dependency manifest supplied with this release. Project proof sources, the replay receipt and the full axiom inventory are included with this manuscript's package.