# Machine-checking the PPT discrimination value of twisted Bell states

Ali Asaria
ali@lab.cloud

**Abstract**

We formalize in Lean, against the Mathlib library, the value of optimally discriminating a twisted Bell basis by measurements with positive partial transpose, using the companion project's semidefinite-program certificates. Specific results that we formalize include: the exact formula $\frac{1}{2}\left(1+\sqrt{1-v_+^2}\right)$ with $v = 2\alpha\beta\lambda_0 - \lambda_1$ and $v_+ = \max\{v, 0\}$, established as a two-sided bound; the certificate lemmas on which it rests, namely the partial-transpose spectrum, a $4 \times 4$ matrix criterion, the template, and the dual witness; and a correction to the companion paper's template hypothesis, certified by an explicit counterexample, together with the clean build and the axiom audit under which every declaration depends only on the Mathlib baseline.

## 1 Introduction

A measurement performed by two parties on a bipartite system is said to be a *PPT measurement* when each of its elements has a positive semidefinite partial transpose with respect to the bipartition between the parties. The class of PPT measurements contains the class of separable measurements, and therefore every measurement that can be carried out by local operations and classical communication (LOCC) [BCJ+14]; unlike those two classes it is defined by linear and positive semidefinite constraints alone, so that the best PPT measurement for a given task is the value of a semidefinite program. This paper reports a machine check, in the Lean proof assistant against the Mathlib library, of one such value: the optimal probability of discriminating a twisted Bell basis with a partially entangled resource.

### State discrimination with an entangled resource

In the discrimination problems considered here, two parties, Alice and Bob, share a system in one of four known orthogonal two-qubit states, each with probability $1/4$, and must identify which one they hold. A measurement is a collection $\{P_0, P_1, P_2, P_3\}$ of positive semidefinite operators summing to the identity, and its success probability is

$$\frac{1}{4}\sum_{k=0}^{3}\langle P_k, \rho_k\rangle, \tag{1}$$

where $\langle A, B\rangle = \mathrm{Tr}(A^* B)$ denotes the Hilbert–Schmidt inner product. When the states are entangled, the parties may fail to discriminate them perfectly by any PPT measurement even though a global measurement succeeds with certainty, and the standard way to quantify this failure is to give the parties an additional entangled pair as a resource and to ask how the optimal success probability grows with the entanglement of that resource.

The four signal states studied here are the twisted Bell basis, the orbit of

$$|\psi_0\rangle = \alpha|00\rangle + \beta|11\rangle \tag{2}$$

under local Pauli unitaries, so that

$$|\psi_0\rangle = \alpha|00\rangle + \beta|11\rangle, \quad |\psi_1\rangle = \beta|00\rangle - \alpha|11\rangle, \quad |\psi_2\rangle = \alpha|01\rangle + \beta|10\rangle, \quad |\psi_3\rangle = \beta|01\rangle - \alpha|10\rangle, \tag{3}$$

with real Schmidt coefficients $\alpha \geq \beta \geq 0$ satisfying $\alpha^2 + \beta^2 = 1$. The resource is the partially entangled pair

$$|\tau_\varepsilon\rangle = \sqrt{\tfrac{1+\varepsilon}{2}}\,|00\rangle + \sqrt{\tfrac{1-\varepsilon}{2}}\,|11\rangle, \tag{4}$$

with $\varepsilon \in [0,1]$, so that the state to be identified is $|\psi_k\rangle \otimes |\tau_\varepsilon\rangle$. The relevant parameters are the Schmidt weights $\lambda_0 = (1+\varepsilon)/2$ and $\lambda_1 = (1-\varepsilon)/2$ of the resource, the concurrence $c = 2\alpha\beta$ of each signal state, and the excess

$$v = c\,\lambda_0 - \lambda_1 = 2\alpha\beta\lambda_0 - \lambda_1, \tag{5}$$

which is nonnegative exactly when the resource outweighs the imbalance of the signal state.

The companion project [Asa26] proved, on the page, that the optimal PPT discrimination probability is

$$\frac{1}{2}\big(1 + \sqrt{1 - v_+^2}\big), \qquad v_+ = \max\{v, 0\}, \tag{6}$$

by exhibiting a matching pair of primal and dual solutions to the associated semidefinite program, each in closed form. On the Bell line $\alpha = \beta = 1/\sqrt{2}$, where $c = 1$ and $v = (1+\varepsilon)/2 - (1-\varepsilon)/2 = \varepsilon$, this reduces to the known value $\frac{1}{2}(1 + \sqrt{1-\varepsilon^2})$ due to Bandyopadhyay, Cosentino, Johnston, Russo, Watrous and Yu [BCJ+14]; off the Bell line the companion paper [Asa26] presents the value as new. Other general results on this and related problems may be found in [DFXY07, BGK11, YDY11, Cos12, YDY12, Bae12, CR13, BCJ+14, YO15, BHN15, BHN17, GKR+01, WH02, EMV02, HDSH03, Nat04, SBJ+98].

### The class of measurements studied

A measurement $\{P_0, P_1, P_2, P_3\}$ on $\mathcal{X} \otimes \mathcal{Y}$ with $\mathcal{X} = \mathcal{X}_1 \otimes \mathcal{X}_2$ and $\mathcal{Y} = \mathcal{Y}_1 \otimes \mathcal{Y}_2$ is PPT when every element has a positive semidefinite partial transpose across the $(X_1X_2) : (Y_1Y_2)$ cut. The partial transpose $T_Y$ transposes Bob's two systems in the computational basis, after the systems are permuted to the order $X_1X_2Y_1Y_2$, so that a measurement is PPT when

$$T_Y(P_k) \succeq 0 \qquad \text{for every } k. \tag{7}$$

(Here, and elsewhere in the paper, $\mathrm{Pos}(\mathcal{X})$ denotes the cone of positive semidefinite operators on $\mathcal{X}$, and $A \succeq 0$ means $A \in \mathrm{Pos}(\mathcal{X})$.) Because LOCC and separable measurements are both contained in the PPT measurements, an upper bound proved for PPT measurements applies to all three classes, while a value attained by an explicit PPT measurement is a lower bound that an LOCC protocol may or may not also reach.

**The tool's previous use**

The tractability of the PPT class is exactly what permits a closed-form answer: the optimal PPT value is the solution of a semidefinite program, whose primal and dual exhibit the matching certificates, and the value of a discriminating measurement is read off from the inner products of its elements with the signal states. The method is standard [Cos12, CR13, BCJ+14]; what the companion paper adds is the closed-form solution for the twisted Bell basis off the Bell line, and what the present paper adds is the machine check of those certificates: every algebraic identity, every positivity assertion, and the certified attainment and bound steps verified in Lean, while weak duality, the value-optimality identification, the completeness condition, and the dual-witness lift remain on the page, as recorded in Section 7.

**Overview of results**

Our results are a formalization, in the Lean proof assistant with the Mathlib library, of the companion paper's theorem and its certificates, using the techniques of the latter. The following specific results are among those we prove:

- A machine-checked proof, for $0 < v < 1$, of the two-sided bound: the template attains $\frac{1}{2}(1+\sqrt{1-v^2})$ and the dual witness bounds every PPT measurement by that value. This answers, in machine-checked form, the question answered on the page in [Asa26]. The upper bound is closely connected to the PPT value computed by convex optimization in [BCJ+14].
- Machine-checked proofs of the certificate lemmas the theorem rests on: the partial-transpose spectrum of a pure two-qubit state, the $4 \times 4$ matrix criterion, the template's PPT blocks and value, the optimizer's attainment, and the dual witness. Each reduces a positivity assertion to a finite real matrix whose spectrum is computed in Lean, which answers the requirement, implicit in the certificates, that every PSD step be explicit.
- A correction to the companion paper's template hypothesis, certified by an explicit counterexample in the formalization: the template is a valid family of PPT operators only under the additional constraint $(x_0 - x_1)^2 \leq 1$, not $(x_0 + x_1)^2 \leq 1$ alone (Section 5). This correction answers the companion paper's own question of which hypothesis the template actually needs.
- A complete verification record: a clean build of the whole development, and an axiom audit showing that every declared theorem depends only on the baseline axioms `propext`, `Classical.choice` and `Quot.sound` that Mathlib itself assumes, together with a statement of exactly where the informal theorem exceeds the formal statements (Section 7).

## 2 A cone program for optimizing over PPT measurements

The value (6) is read off a semidefinite program, and the formalization follows the structure of that program without formalizing the full PPT cone, which Mathlib does not provide at the pinned revision.

**The general framework**

A cone program over real Euclidean spaces $\mathcal{X}$ and $\mathcal{W}$, with a cone $\mathcal{K} \subseteq \mathcal{X}$, a linear map $\Phi\colon \mathcal{X} \to \mathrm{Herm}(\mathcal{W})$, and Hermitian operators $A$ and $B$, has the pair

| Primal problem | | Dual problem | |
|---|---|---|---|
| maximize: | $\langle A, X\rangle$ | minimize: | $\langle B, Y\rangle$ |
| subject to: | $\Phi(X) = B,$ | subject to: | $\Phi^*(Y) - A \in \mathcal{K}^*,$ |
| | $X \in \mathcal{K}.$ | | $Y \in \mathrm{Herm}(\mathcal{W}).$ |

where $\mathcal{K}^* = \{\, Z \in \mathrm{Herm}(\mathcal{X}) : \langle Z, X\rangle \geq 0 \text{ for all } X \in \mathcal{K} \,\}$ is the dual cone and $\Phi^*$ the adjoint map.

**Proposition 1** (Weak duality for cone programs)**.** *For any primal feasible $X$ and dual feasible $Y$ it holds that $\langle A, X\rangle \leq \langle B, Y\rangle$.*

*Proof.* It will first be proved that the two inner products differ by a known nonnegative quantity. From the constraint $\Phi(X) = B$ and the self-adjointness of the Hilbert–Schmidt inner product one has $\langle B, Y\rangle = \langle \Phi(X), Y\rangle = \langle X, \Phi^*(Y)\rangle$, so that

$$\langle B, Y\rangle - \langle A, X\rangle = \langle X, \Phi^*(Y) - A\rangle \geq 0, \tag{8}$$

the final inequality because $X \in \mathcal{K}$ and $\Phi^*(Y) - A \in \mathcal{K}^*$. This establishes that $\langle A, X\rangle \leq \langle B, Y\rangle$. □

Every dual feasible operator $Y$ therefore provides an upper bound $\langle B, Y\rangle$ on the value achievable over all choices of a primal feasible $X$. Although the companion paper [Asa26] notes that the cone programs considered here do possess strong duality (the primal is strictly feasible and bounded), that fact is not needed for any of our results and is not itself machine-checked: a matching pair of feasible primal and dual solutions with equal objective value certifies optimality without it.

**The specialisation to the problem**

Let us now return to the state discrimination problem. Hereafter let the signal states and the resource be as in (3) and (4), let $\rho_k = |\psi_k\rangle\langle\psi_k| \otimes |\tau_\varepsilon\rangle\langle\tau_\varepsilon|$, and let the cone $\mathcal{K}$ be the cone of positive partial transpose, the set of positive semidefinite operators on $\mathcal{X} \otimes \mathcal{Y}$ whose partial transpose is also positive semidefinite. The optimal PPT value is the value of the program

$$\begin{aligned} \text{maximize:} \quad & \tfrac{1}{4}\sum_{k=0}^{3} \langle P_k, \rho_k\rangle \\ \text{subject to:} \quad & P_k \succeq 0, \quad \sum_k P_k = \mathbb{1}, \quad T_{\mathcal{Y}}(P_k) \succeq 0. \end{aligned} \tag{9}$$

and its dual is

$$\begin{aligned} \text{minimize:} \quad & \operatorname{Tr} Y \\ \text{subject to:} \quad & Q_k \succeq 0, \quad Y - \tfrac{1}{4}\rho_k - T_{\mathcal{Y}}(Q_k) \succeq 0 \ \text{ for every } k. \end{aligned} \tag{10}$$

For any Hermitian operator $H$ such that $H - \frac{1}{4}\rho_k$ has positive partial transpose for every $k$, the probability of a correct discrimination is upper-bounded by $\operatorname{Tr} H$; this is the sentence every later upper bound invokes.

## 3 The partial transpose and a matrix criterion

Two self-contained finite-matrix facts underpin every later certificate. Both are proved in Lean directly, without appealing to the PPT cone.

### The partial-transpose spectrum

The partial transpose of the pure state $|\psi_0\rangle\langle\psi_0|$ with diagonal coefficient matrix is computed explicitly, and its eigenvalues govern the PSD assertions that follow.

**Theorem 2** (`Proj.PartialTranspose`). *Let $\rho = |\psi\rangle\langle\psi| = \mathrm{diag}(m_0, m_1) \otimes \mathrm{diag}(m_0, m_1)^\intercal$ be the pure two-qubit state with diagonal coefficient matrix $\mathrm{diag}(m_0, m_1)$ in the span of $|00\rangle$ and $|11\rangle$. The partial transpose $T_Y(\rho)$ has eigenvalues*

$$m_0^2,\ m_1^2,\ \ m_0 m_1,\ -m_0 m_1. \tag{11}$$

*Consequently the partial transpose of $|\psi_0\rangle\langle\psi_0|$ has eigenvalues $\alpha^2$, $\beta^2$, $\pm\alpha\beta$, and the partial transpose of $|\tau_\varepsilon\rangle\langle\tau_\varepsilon|$ has eigenvalues $\lambda_0$, $\lambda_1$, $\pm\sqrt{\lambda_0\lambda_1}$.*

*Proof.* It is first shown that $T_Y(\rho)$ has the four stated eigenvectors with eigenvalues as in (11), verified componentwise in the computational basis; the factored characteristic determinant is computed at the same time, which establishes that no further eigenvalue is hidden. The two "consequently" cases follow by substituting $m_0 = \alpha$, $m_1 = \beta$ and $m_0 = \sqrt{\lambda_0}$, $m_1 = \sqrt{\lambda_1}$. □

**Remark 3.** The formal statement is specialised to the pure-Rho form that the certificates require; the eigenvalue formula (11) for an arbitrary diagonal coefficient matrix is the one-line generalisation, and this is recorded as an informal-formal gap in Section 7.

### A $4 \times 4$ matrix criterion

The positivity of the template and dual blocks reduces to the following criterion.

**Theorem 4** (`Proj.Omega`). *For a real symmetric $2 \times 2$ matrix $K$ and a real number $f$, let*

$$\Omega = \mathbb{1}_2 \otimes K - f\, \sigma_y \otimes \sigma_y, \tag{12}$$

*where $\sigma_y$ is the Pauli matrix. Then $\Omega \succeq 0$ if and only if $K \succeq 0$ and $\det K \geq f^2$.*

*Proof.* It will be proved by reducing the quadratic form of $\Omega$ to that of $K$. Over the two-parameter test family $z = (sp, sq, -q, p)$ the quadratic form of $\Omega$ evaluates to $(\det K - f^2)(p^2 + q^2)^2$ plus a manifestly nonnegative square, which forces $\det K \geq f^2$ in one direction; in the other direction the explicit SOS decomposition expresses $a \cdot z^\intercal \Omega z$ as a sum of two squares plus $(\det K - f^2)$ times a nonnegative quantity. The equivalence (12) follows. □

**Remark 5.** The criterion (12) is equivalent to the companion paper's eigenvalue form $\Omega$ has eigenvalues $\frac{k_1+k_2}{2} \pm \sqrt{\left(\frac{k_1-k_2}{2}\right)^2 + f^2}$, each twice; the formalization proves the positivity criterion, and the eigenvalue display is not a separate declaration (recorded in Section 7).

## 4 The twisted Bell states

The value (6) was proved on the page in the companion paper [Asa26] by a matching pair of SDP certificates; the present paper certifies the two halves of that pair in Lean, together with the algebra that assembles them.

### The optimizer

The lower bound is the template specialised to the optimizer. For $0 < v < 1$ the optimizer's success probability is computed exactly.

**Proposition 6** (`Proj.Optimizer`)**.** *For $0 < v < 1$ the template restricted to the optimizer attains the squared value*

$$F^2 = \left(\tfrac{\sqrt{1+v}+\sqrt{1-v}}{2}\right)^2 = \tfrac{1}{2}\left(1 + \sqrt{1-v^2}\right), \tag{13}$$

*and for $v \leq 0$ the optimizer achieves $F = 1$.*

*Proof.* The identity (13) is the bridge of Lemma 7 below applied to the optimizer's value; the feasibility of the optimizer under the template constraints is checked constraint by constraint, and the $v \leq 0$ case is the perfect-discrimination endpoint. The maximizing side of this statement, over the whole feasible family, is not part of the formalization and is recorded in Section 7; the two-sided theorem needs only the attainment. □

**Lemma 7** (`Proj.Assembly`)**.** *For $-1 \leq v \leq 1$ it holds that*

$$\left(\tfrac{\sqrt{1+v}+\sqrt{1-v}}{2}\right)^2 = \tfrac{1}{2}\left(1 + \sqrt{1-v^2}\right). \tag{14}$$

*Proof.* Squaring and collecting, the left-hand side is $\frac{1}{4}(2 + 2\sqrt{1-v^2})$, which equals the right-hand side. □

**Theorem 8** (`Proj.Assembly`)**.** *For $0 < v < 1$ the template measurement attains the success probability*

$$\tfrac{1}{2}\left(1 + \sqrt{1-v^2}\right), \tag{15}$$

*and the dual witness bounds every PPT measurement by this same value.*

*Proof.* It will first be proved that the value is attained, then that no PPT measurement exceeds it. The optimizer of Proposition 6 attains the value (13); this is the lower bound, and it is the algebraic identity that is machine-checked. The dual witness of Section 6 has trace $\frac{1}{2}(1+s)$ with $s^2 = 1 - v^2$; by Proposition 1, which is the general page-level fact of Section 2 and is not itself a Lean declaration, this trace bounds every PPT measurement from above. The two coincide, and the identification of this common value as the optimal PPT value rests on the companion paper's page-level weak-duality argument, not on the machine check, as recorded in Section 7. □

**Remark 9.** The upper bound is attained by the template operator family, so the value (15) is achieved within the PPT class; whether an LOCC protocol also reaches it off the Bell line is the open question of Section 8. On the Bell line the value matches the teleportation protocol of [BCJ+14].

**Corollary 10** (`Proj.Assembly`)**.** *The value* (6) *recovers the known special cases: for $\varepsilon = 0$ it is 1; for $\alpha = 1$ it is 1; for $\alpha = \beta$ it is $\frac{1}{2}(1 + \sqrt{1 - \varepsilon^2})$; for $\alpha = \beta$ and $\varepsilon = 1$ it is $\frac{1}{2}$; and for $\varepsilon = 1$ it is $\alpha^2$. Perfect discrimination is possible exactly when $c(1 + \varepsilon) \leq 1 - \varepsilon$, i.e. when $v \leq 0$.*

*Proof.* Each special case is a substitution into the definition of the value (6) with $v_+ = \max(v, 0)$, computed in Lean as a corollary of the main statement. The perfect-discrimination condition is the companion paper's [Asa26] boundary $c(1 + \varepsilon) \leq 1 - \varepsilon$, equivalent to $v \leq 0$; its formalized half is the optimizer's attainment of $F = 1$ at $v \leq 0$ together with the endpoint declaration, and the converse (that no PPT measurement exceeds the value) is the page-level bound that a success probability never exceeds 1. □

### Example

The interior point $(\alpha, \varepsilon) = (2/\sqrt{5}, 7/9)$ illustrates the formula. Here $\alpha = 2/\sqrt{5}$, $\beta = 1/\sqrt{5}$, $c = 4/5$, $\lambda_0 = 8/9$, $\lambda_1 = 1/9$, so that $v = c\lambda_0 - \lambda_1 = 3/5$ and

$$\tfrac{1}{2}\left(1 + \sqrt{1 - v^2}\right) = \tfrac{1}{2}\left(1 + \tfrac{4}{5}\right) = \tfrac{9}{10}. \tag{16}$$

This value, $9/10$, is the predicted success probability at this point, and it is recovered exactly by the formalized evaluation (the declaration `Proj.Counterexample.opt_value_at_work_example`).

## 5 The template measurement

The lower-bound half of the main theorem rests on a template: a covariant family of PPT operators whose value is $F^2$. The family is called a measurement here following the companion paper, though its completeness condition $\sum_i P_i = \mathbb{1}$ is not part of the formal statements and is recorded as a gap in Section 7; what is certified is the positivity of its PPT blocks and the value overlap. Its PPT blocks are the hardest, sign-fragile part of the development, and it is here that the formalization found a correction to the companion paper.

### The template and a correction

The template operator family $P_i(x, y)$ for vectors $x, y \in \mathbb{R}^2$ satisfying $\|x\|^2 + \|y\|^2 \leq 1$ has four blocks, whose positive semidefiniteness under partial transpose is equivalent, block by block, to scalar conditions on $x$ and $y$. One block is the $2 \times 2$ matrix

$$B_1 = \begin{pmatrix} L_{00} & x_0 x_1 \\ x_0 x_1 & L_{00} \end{pmatrix}, \qquad L_{00} = \tfrac{1-(x_0^2+x_1^2)}{2}, \tag{17}$$

together with a diagonal block $\mathrm{diag}(x_0^2, x_1^2)$.

**Theorem 11** (`Proj.Template`)**.** *The block $B_1$ of* (17) *is positive semidefinite under the corrected constraint*

$$(x_0 + x_1)^2 \leq 1 \quad \text{and} \quad (x_0 - x_1)^2 \leq 1, \tag{18}$$

*which is equivalent to $|x_0 x_1| \leq L_{00}$, and likewise for the block $B_2$ with $y$ and $L_{11} = \frac{1-(y_0^2+y_1^2)}{2}$.*

*Proof.* The claim reduces to the $2 \times 2$ criterion of Theorem 4: the block is PSD exactly when $|x_0 x_1| \leq L_{00}$, and the two inequalities (18) combine to give this bound by elementary algebra. □

**Remark 12.** The companion paper [Asa26] states the template hypothesis with $(x_0 + x_1)^2 \leq 1$ alone, which is insufficient: the off-diagonal $x_0 x_1$ is controlled by $L_{00}$ only when *both* of $(x_0 \pm x_1)^2 \leq 1$ hold, and with $x_0 x_1 < 0$ the single constraint does not imply the second. The formalization certifies this directly with the counterexample of Proposition 13. On the optimizer the corrected constraint (18) holds, and the companion paper's optimizers are said in [Asa26] to satisfy $x_0, x_1 \geq 0$, so the main theorem is unaffected; the formalization adopts the corrected hypothesis (18) throughout.

**Proposition 13** (`Proj.Counterexample`)**.** *At $x = (7/10, -7/10)$ and $y = 0$ the constraint $(x_0 + x_1)^2 \leq 1$ holds but the block $B_1$ is not positive semidefinite.*

*Proof.* The value $(x_0 + x_1)^2 = 0 \leq 1$ holds while $x_0 x_1 = -49/100 < 0$, so $|x_0 x_1| > L_{00} = \frac{1-98/100}{2} = 1/100$, and the block has a negative eigenvalue; this is the certified rational witness recorded in `Proj.Counterexample.parent_t3_hyp_insufficient`. □

## 6 The dual witness

The upper-bound half of the main theorem is a dual witness, an explicit Hermitian operator $Y$ together with positive operators $Q_k$ satisfying the dual constraints of (10), whose trace is $\frac{1}{2}(1+s)$ with $s^2 = 1 - v^2$. The witness is block-diagonal, and its positivity reduces to two blocks $C_-$ and $C_+$.

**Theorem 14** (`Proj.DualWitness`, `Proj.Cplus`)**.** *For $0 < v < 1$ the two blocks $C_-$ and $C_+$ of the dual operator $Y$ are positive semidefinite, and*

$$\operatorname{Tr} Y = \tfrac{1}{2}(1+s), \qquad s^2 = 1 - v^2. \tag{19}$$

*Combining this trace with Proposition 1 bounds every PPT measurement from above by $\frac{1}{2}(1 + \sqrt{1-v^2})$, the page-level duality step being what turns the certified trace identity into the universal bound.*

*Proof.* It will first be proved that the block $C_-$ is positive, then that the block $C_+$ is positive. The $C_-$ block is the operator $K'$ of Theorem 4 with parameters satisfying $\det K' = f^2$ exactly, so its positivity follows from the criterion; its trace is computed to be $(1 + s - v)/8$. The $C_+$ block is positive by its explicit kernel and a $2 \times 2$ restriction: two kernel vectors are shown to annihilate it, a change of basis reduces the quadratic form to that of a $2 \times 2$ matrix $Q$ whose determinant is $4A^2(A^2 - B^2)(B-D)^2/(A^2+D^2)^2 = 8A^2 v(B-D)^2/(A^2+D^2)^2 \geq 0$, and the criterion closes it. The two blocks positive gives the feasibility, and the trace identity (19) completes the reduction to $\frac{1}{2}(1 + \sqrt{1-v^2})$ via Lemma 7. The formal certificates are at the rationalized block level, in the abstract parameters $A, B, D$ below; the full $16 \times 16$ slack assembly of the dual constraint is the recorded gap of Section 7. □

**Remark 15.** The formal statement of the dual witness is at the rationalized block level: the parameters $A, B, D$ with $A^2 = 1 + v$, $B^2 = 1 - v$ and $D^2 = \varepsilon - v$ are treated as abstract radicals, so every proof is radical-free polynomial algebra. The final lift from the rational proxy to the literal $C_+$ through the diagonal congruence $\Lambda = \operatorname{diag}(\kappa, \sqrt{\lambda_0})$ is recorded as an informal-formal gap in Section 7.

| Declaration | Statement | File | Axioms |
|---|---|---|---|
| `Proj.PartialTranspose.ptB_pureRho` | Theorem 2 | `PartialTranspose` | baseline |
| `Proj.Omega.omegaMat_posSemidef_iff` | Theorem 4 | `Omega` | baseline |
| `Proj.Template.b1_posSemidef` | Theorem 11 | `Template` | baseline |
| `Proj.Optimizer.optimizer_value` | Proposition 6 | `Optimizer` | baseline |
| `Proj.DualWitness.cminus_posSemidef` | Theorem 14 | `DualWitness` | baseline |
| `Proj.Cplus.cplus_posSemidef` | Theorem 14 | `Cplus` | baseline |
| `Proj.Assembly.primal_dual_agree` | Theorem 8 | `Assembly` | baseline |
| `Proj.Assembly.optval_*` | Corollary 10 | `Assembly` | baseline |
| `Proj.Counterexample.parent_t3_hyp_insufficient` | Proposition 13 | `Counterexample` | baseline |

Table 1: Representative headline declarations and the informal statements they certify; a theorem may carry several declarations (for example, Theorem 11 covers both blocks $B_1$ and $B_2$, and Proposition 6 covers both the $0 < v < 1$ and the $v \leq 0$ optimizers). Every declaration audits to the baseline axioms only, in job `2102e28f`.

## 7 Verification

The development was checked with Lean $v$4.26.0 against Mathlib revision `2df2f0150c275ad53cb3c90f7c98ec15a` A clean `lake build` of the whole project, nine modules from source, completed in job `2102e28f` with no `sorry`. The headline declarations are as follows.

Every declared theorem depends only on `propext`, `Classical.choice` and `Quot.sound`, the same baseline axioms Mathlib itself assumes, and `sorryAx` appears nowhere in the audit of the forty-six headline declarations. The proof was certified module by module in a sequence of whole-project build jobs, concluding with the clean build and axiom audit of job `2102e28f`.

### The informal and formal statements

The paper's prose discusses several assertions that the formal statements do not cover literally, and each is recorded here once. The two-sided bound of Theorem 8 is not a proof over the PPT cone: Mathlib at the pinned revision provides no theory of the PPT cone or of SDP duality, so the statement is "a PPT measurement attains $\frac{1}{2}(1 + \sqrt{1 - v^2})$ and every PPT measurement is bounded by it", with the identification of this value as the optimal PPT value resting on the companion paper's page-level weak-duality argument. The main theorem is stated on $0 < v < 1$; the $v \leq 0$ branch, where the value is 1, is attested by the optimizer's attainment together with the page-level bound that a PPT success probability never exceeds 1. The dual witness is proved at the rationalized block level (Remark after Theorem 14), the lift to the literal $C_+$ through $\Lambda = \mathrm{diag}(\kappa, \sqrt{\lambda_0})$ being outside the formal statements. The completeness condition $\sum_k P_k = \mathbb{1}$ and the $16 \times 16$ slack-block assembly are not formalized; the template's positivity is checked block by block. The maximization (13) is formalized as attainment only, the max-over-the-family direction being unnecessary for the two-sided bound. The partial transpose of Theorem 2 is stated for the pure-Rho form, not an arbitrary diagonal matrix, and the $4 \times 4$ criterion of Theorem 4 is the positivity criterion, not the eigenvalue formula (the two are equivalent).

## 8 Conclusion and Future Work

In this paper we have used a machine-checked development in Lean, against the Mathlib library, to certify the companion paper's closed-form value for the PPT discrimination of the twisted Bell basis, together with the certificate lemmas on which it rests and a correction to the companion paper's template hypothesis.

Several interesting questions regarding the twisted Bell basis remain unsolved. Among them are the following questions.

- In Section 4 we proved the two-sided bound $\frac{1}{2}(1+\sqrt{1-v^2})$ on $0 < v < 1$. More generally, one could ask whether the full statement over the PPT cone, "the optimal PPT value equals the closed form (6) for all admissible $(\alpha, \varepsilon)$", can be formalized once a PPT-cone or SDP-duality theory is available in the library; this subsumes the two pieces left on the page here, the two-sided bound on the $v \leq 0$ branch and the maximization over the whole feasible family. The obstruction is the absence of such a theory at the pinned revision.
- Whether a separable or LOCC measurement reaches the PPT value off the Bell line is open even without the machine check, and it is [BCJ+14] that settles the Bell line by teleportation. A formal treatment of that question would require the separable and LOCC classes, which the present development does not model.
- The correctness of the template as a complete measurement, meaning the completeness condition $\sum_k P_k = \mathbb{1}$ together with the remaining blocks, is checked block by block rather than as the full $16 \times 16$ statement. A formalization of the completeness and the covariant assembly would make the lower bound fully self-contained.

The techniques presented in the paper are not intrinsically limited to the twisted Bell basis; applications of these techniques to other state-discrimination problems are topics for possible future work.

## Availability

Artifacts are available from the authors on request.

## Acknowledgements

This paper was made possible by the research tooling generated by Transformer Lab.

## References

[Asa26] Ali Asaria. PPT discrimination of twisted Bell states with a partially entangled resource. unpublished, 2026. Companion project twirled-bell-witness; the theorem this paper formalizes. Artifacts available from the authors on request.

[Bae12] Joonwoo Bae. Structure of minimum-error quantum state discrimination. arXiv:1210.2845 [quant-ph], 2012.

[BCJ+14] Somshubhro Bandyopadhyay, Alessandro Cosentino, Nathaniel Johnston, Vincent Russo, John Watrous, and Nengkun Yu. Limitations on separable measurements by convex optimization. arXiv:1408.6981 [quant-ph], 2014.

[BGK11] Somshubhro Bandyopadhyay, Sibasish Ghosh, and Guruprasad Kar. LOCC distinguishability of unilaterally transformable quantum states. arXiv:1102.0841 [quant-ph], 2011.

[BHN15] Somshubhro Bandyopadhyay, Saronath Halder, and Michael Nathanson. Entanglement as a resource for local state discrimination in multipartite systems. arXiv:1510.02443 [quant-ph], 2015.

[BHN17] Somshubhro Bandyopadhyay, Saronath Halder, and Michael Nathanson. Optimal resource states for local state discrimination. arXiv:1709.10302 [quant-ph], 2017.

[Cos12] Alessandro Cosentino. PPT-indistinguishable states via semidefinite programming. arXiv:1205.1031 [quant-ph], 2012.

[CR13] Alessandro Cosentino and Vincent Russo. Small sets of locally indistinguishable orthogonal maximally entangled states. arXiv:1307.3232 [quant-ph], 2013.

[DFXY07] Runyao Duan, Yuan Feng, Yu Xin, and Mingsheng Ying. Distinguishability of Quantum States by Separable Operations. arXiv:0705.0795 [quant-ph], 2007.

[EMV02] Yonina C. Eldar, Alexandre Megretski, and George C. Verghese. Designing Optimal Quantum Detectors Via Semidefinite Programming. arXiv:quant-ph/0205178 [quant-ph], 2002.

[GKR+01] Sibasish Ghosh, Guruprasad Kar, Anirban Roy, Aditi Sen De, and Ujjwal Sen. Distinguishability of the Bell states. arXiv:quant-ph/0106148 [quant-ph], 2001.

[HDSH03] Michal Horodecki, Aditi Sen De, Ujjwal Sen, and Karol Horodecki. Local indistinguishability: more nonlocality with less entanglement. arXiv:quant-ph/0301106 [quant-ph], 2003.

[Nat04] Michael Nathanson. Distinguishing Bipartite Orthogonal States using LOCC: Best and Worst Cases. arXiv:quant-ph/0411110 [quant-ph], 2004.

[SBJ+98] Masahide Sasaki, Stephen M. Barnett, Richard Jozsa, Masao Osaki, and Osamu Hirota. Accessible information and optimal strategies for real symmetrical quantum sources. arXiv:quant-ph/9812062 [quant-ph], 1998.

[WH02] Jonathan Walgate and Lucien Hardy. Nonlocality, Asymmetry, and Distinguishing Bipartite States. arXiv:quant-ph/0202034 [quant-ph], 2002.

[YDY11] Nengkun Yu, Runyao Duan, and Mingsheng Ying. Four Locally Indistinguishable Ququad-Ququad Orthogonal Maximally Entangled States. arXiv:1107.3224 [quant-ph], 2011.

[YDY12] Nengkun Yu, Runyao Duan, and Mingsheng Ying. Distinguishability of Quantum States by Positive Operator-Valued Measures with Positive Partial Transpose. arXiv:1209.4222 [quant-ph], 2012.

[YO15] Sixia Yu and C. H. Oh. Detecting the local indistinguishability of maximally entangled states. arXiv:1502.01274 [quant-ph], 2015.