

Optimal Quantum de Finetti Theorems via Argmax Rounding

Fernando Granha Jeronimo*

Pei Wu[†]

Haochen Xu[‡]

Abstract

We prove optimal finite quantum de Finetti upper bounds. Given a bosonic state $\rho_N \in \mathcal{D}(\text{Sym}^N(\mathbb{C}^d))$, there is a probability measure ν on the unit sphere such that

$$\left\| \rho_N^{(2)} - \int_{\|u\|=1} |u\rangle\langle u|^{\otimes 2} d\nu(u) \right\|_1 \leq \frac{\sqrt{d-1}}{N-1}.$$

By purification, the bosonic theorem also gives the optimal $O(d/N)$ upper bound for arbitrary exchangeable states. These results settle the dimension dependence left open by Christandl, König, Mitchison, and Renner (CMP 2007). The proof casts de Finetti approximation as sum-of-squares rounding and applies the argmax method of Jeronimo, Wu, and Xu (manuscript 2026).

Among the consequences of the optimal bounds, we record three advances on long-standing problems in quantum complexity that have resisted progress for decades.

1. We refute Watrous’s disentangler conjecture. A disentangler is a quantum channel whose outputs are all close to separable states and whose image approximates every separable state. For every fixed $\varepsilon \in (0, 1)$, we construct such a channel with exact coverage, that is, an $(\varepsilon, 0)$ -disentangler $\mathcal{D}(\mathbb{C}^D) \rightarrow \mathcal{D}(\mathbb{C}^d \otimes \mathbb{C}^d)$ with

$$D = \exp(O_\varepsilon(\sqrt{d} \log d)) = \exp(o(d)).$$

The conjecture, recorded by Aaronson, Beigi, Drucker, Fefferman, and Shor (CCC 2008), predicted instead that every constant-error disentangler must have $D = \exp(\Omega(d))$.

2. We give the first algorithm for general explicit Best Separable State with arbitrary completeness and soundness parameters whose running time is subexponential in the local dimension: at every fixed accuracy, it runs in time $\exp(\tilde{O}(\sqrt{d}))$. This removes the perfect-completeness assumption from the earlier algorithm of Barak, Kothari, and Steurer (STOC 2017), which achieved such a running time only in that case.
3. We give the first algorithm for trace-norm separability testing of an explicitly given state whose running time is subexponential in the local dimension: at every fixed additive accuracy, it runs in time $\exp(\tilde{O}(\sqrt{d}))$. Computational separability has been studied since at least the work of Gurvits (STOC 2003), yet no such algorithm was previously known for trace-norm testing.

We also prove the first dimension-free bosonic quantum de Finetti theorem in Hilbert–Schmidt norm, with upper bound $O(N^{-1/2})$. This dimension-free result follows from our optimal trace-norm theorem by spectral truncation and recovers the dimension-free classical phenomenon of Diaconis and Freedman (*The Annals of Probability* 1980), which is impossible in quantum trace norm.

*University of Illinois Urbana–Champaign. Email: granha@illinois.edu

[†]The Pennsylvania State University. Email: pei.wu@psu.edu

[‡]The Pennsylvania State University. Email: hpx5065@psu.edu

Contents

1	Introduction and main results	3
1.1	Background	3
1.2	Our results	4
1.3	Technical overview	10
2	Preliminaries	16
2.1	Separable and extendible states	17
2.2	Hausdorff notation and convex duality	18
2.3	Linear-algebra estimates	20
2.4	Tensor norms and symmetric tensors	21
3	Two-site de Finetti upper bounds	23
3.1	Bosonic states: the argmax principle	23
3.2	Two-sided extendibility	28
4	t-site de Finetti upper bounds	31
4.1	Two-sided extendibility across the $1 : (t - 1)$ cut	32
4.2	From product vectors to tensor powers	32
4.3	Bounding the integrality gap	34
4.4	Finishing the proof	36
5	Applications of the optimal de Finetti theorem	37
5.1	Subexponential disentanglers	37
5.2	Best Separable State without perfect completeness	38
5.3	Trace-norm separability testing	39
6	A dimension-free bosonic Hilbert–Schmidt de Finetti theorem	41
6.1	Spectral truncation and the dimension-free Hilbert–Schmidt bound	41
6.2	Two-sided Bose-symmetric extendible states	48
7	Hilbert–Schmidt separability algorithms	51
7.1	Separable optimization with bounded Hilbert–Schmidt norm	51
7.2	Hilbert–Schmidt separability testing	52
A	Rectangular argmax rounding	57
B	Lower bounds for the de Finetti theorems	59
B.1	Optimality of the t -site bounds	59
B.2	Dimension-free bosonic Hilbert–Schmidt lower bounds	65

1 Introduction and main results

1.1 Background

The quantum de Finetti theorem is one of the fundamental structural results in quantum information. It gives a quantitative expression of the monogamy of entanglement: highly symmetric quantum systems cannot share much entanglement among many subsystems. The theorem captures this principle by showing that the marginal on any fixed number of sites approaches a mixture of tensor-power states as the total number of sites grows.

The standard finite bounds make this statement quantitative. Let $\rho_N \in \mathcal{D}((\mathbb{C}^d)^{\otimes N})$ be permutation invariant, and write $\rho_N^{(t)}$ for its t -site marginal. The quantum de Finetti theorem of Christandl, König, Mitchison, and Renner [CKMR07, Theorems II.7 and II.8] states that, for every $1 \leq t \leq N$, there is a probability measure μ on $\mathcal{D}(\mathbb{C}^d)$ such that

$$\frac{1}{2} \left\| \rho_N^{(t)} - \int_{\mathcal{D}(\mathbb{C}^d)} \sigma^{\otimes t} d\mu(\sigma) \right\|_1 \leq \frac{2d^2 t}{N}.$$

If ρ_N is supported on $\text{Sym}^N(\mathbb{C}^d)$, there is instead a probability measure ν on the unit sphere of $\mathbb{C}^d$ such that

$$\frac{1}{2} \left\| \rho_N^{(t)} - \int_{\|u\|=1} |u\rangle\langle u|^{\otimes t} d\nu(u) \right\|_1 \leq \frac{2dt}{N}.$$

The infinite quantum de Finetti theorem goes back to Størmer and to Hudson and Moody [Stø69, HM76]; Caves, Fuchs, and Schack later recast it in quantum-information language [CFS02]. Finite quantitative versions developed in the early 2000s, with the work of Christandl, König, Mitchison, and Renner forming a culmination of the theory of approximation by mixtures of exact tensor powers [KR05, CKMR07]. In a complementary direction, Renner’s exponential de Finetti theorem obtained exponentially small error by allowing almost-product states, with applications to quantum cryptography [Ren05, Ren07].

In quantum complexity, finite de Finetti bounds are a basic tool for analyzing SDP hierarchies based on symmetric extensions: they certify convergence to separable or product states and thereby yield algorithms for separability membership and product-state optimization [DPS04, NOP09, BCY11b, BH13, FF21]. They also enter the study of QMA with multiple unentangled proofs and of disentangler constructions [KMY09, ABD⁺08, BH13, LS15, HNW19, JLW26].

An important refinement changes the metric so that the approximation need only fool a restricted class of measurements. Brandão, Christandl, and Yard obtained an improved de Finetti-type bound with only logarithmic local-dimension dependence under one-way LOCC measurements, yielding quasipolynomial separability algorithms for the one-way LOCC and Hilbert–Schmidt norms [BCY11a, BCY11b]. Brandão and Christandl extended these quantitative symmetric-extension bounds to the multipartite setting [BC12]. Brandão and Harrow subsequently proved stronger de Finetti bounds under local measurements, with applications to quantum complexity and polynomial optimization [BH13]. Li and Smith extended this theory to fully adaptive one-way LOCC measurements; among other consequences, they show that polynomially many unentangled proofs add no power to QMA when the verifier is restricted to one-way LOCC [LS15]. Such results are closely related to distinguishability norms and quantum data hiding [MWW09, BCY11a, BH13, LS15].

The quantitative content of a finite de Finetti theorem is its convergence rate. Classically, the t -variable marginal of an exchangeable sequence of length N is within $O(t^2/N)$ in total variation of a mixture of i.i.d. laws, independently of the size of the sample space [DF80]. Quantum trace-norm bounds necessarily depend on the local Hilbert-space dimension, but the correct dependence was left

open. Christandl, König, Mitchison, and Renner devoted a section to optimality, asked whether their bounds could be improved, and concluded that the dimension dependence was unclear [CKMR07, Secs. II.C and V]. For bosonic states on $\text{Sym}^N(\mathbb{C}^d)$, the standard constructive theorem gives error $O(d/N)$ for two-site marginals [CKMR07, Chi11, LNR15]; Lewin, Nam, and Rougerie later singled out improving this d -dependence as an open problem [LNR15, Remark 2.2].

We settle this question by proving optimal trace-norm upper bounds. The main result improves the two-site bosonic error from $O(d/N)$ to $O(\sqrt{d}/N)$, and purification gives an $O(d/N)$ upper bound for arbitrary exchangeable states. The square-root improvement is exactly what moves several quantum-complexity problems from exponential to subexponential dependence on the local dimension.

The sharp bosonic theorem is the engine behind the paper, but it is not the end of the development. We next prove a de Finetti theorem for two-sided Bose-symmetric (n, m) -extendible states: every such bipartite state on $A \otimes B$ is within

$$O\left(\sqrt{\frac{\min\{d_A, d_B\}}{nm}}\right)$$

in trace distance of the separable states. For balanced extensions with equal local dimensions, this becomes $O(\sqrt{d}/n)$. The two-sided theorem provides a common interface for three complexity applications. Realizing the two-sided relaxation as the image of a quantum channel gives a subexponential disentangler. Optimizing an accepting operator over the same relaxation gives the subexponential Best Separable State algorithm, and the Harrow–Montanaro reduction then yields trace-norm separability testing [HM13].

The optimal bosonic trace-norm theorem is also the starting point for the first dimension-free bosonic de Finetti theorem in Hilbert–Schmidt norm, with upper bound $O(N^{-1/2})$. We also obtain trace-norm bounds for higher marginals, with error $O(t\sqrt{d}/N)$ for bosonic states and $O(td/N)$ for exchangeable states. We state these results precisely next.

1.2 Our results

1.2.1 Optimal trace-norm de Finetti

Let $\mathcal{H} = \mathbb{C}^d$. We use the following notations for trace distance

$$D_{\text{tr}}(\rho, \sigma) := \frac{1}{2} \|\rho - \sigma\|_1, \quad D_{\text{tr}}(\rho, \mathcal{K}) := \inf_{\tau \in \mathcal{K}} D_{\text{tr}}(\rho, \tau).$$

A state in $\text{D}(\text{Sym}^N \mathcal{H})$ is *bosonic*, while a state on $\mathcal{H}^{\otimes N}$ is *exchangeable* if it is permutation invariant. For either type, we write $\rho_N^{(t)}$ for the t -site marginal.

For $2 \leq t \leq N$, define the Hartree mixtures and tensor-power mixtures by

$$\text{SEP}_H^{(t)}(d) := \text{conv}\{|u\rangle\langle u|^{\otimes t} : u \in \mathbb{C}^d, \|u\| = 1\}, \quad (1.1)$$

$$\text{SEP}_P^{(t)}(d) := \text{conv}\{\sigma^{\otimes t} : \sigma \in \text{D}(\mathbb{C}^d)\}. \quad (1.2)$$

The subscripts H and P stand for Hartree and tensor power, respectively. The sets of t -site marginals of N -site bosonic and exchangeable states are

$$\text{BOS}_N^{(t)}(d) := \left\{ \rho_N^{(t)} : \rho_N \in \text{D}(\text{Sym}^N(\mathbb{C}^d)) \right\}, \quad (1.3)$$

$$\text{EXCH}_N^{(t)}(d) := \left\{ \sigma_N^{(t)} : \begin{array}{l} \sigma_N \in \text{D}((\mathbb{C}^d)^{\otimes N}), \\ U_\pi \sigma_N U_\pi^\dagger = \sigma_N \text{ for every } \pi \in \mathfrak{S}_N \end{array} \right\}. \quad (1.4)$$

Every Hartree mixture has an N -site bosonic extension, and every tensor-power mixture has an N -site exchangeable extension. Thus both marginal sets are outer relaxations:

$$\text{SEP}_H^{(t)}(d) \subseteq \text{BOS}_N^{(t)}(d), \quad \text{SEP}_P^{(t)}(d) \subseteq \text{EXCH}_N^{(t)}(d).$$

For the nested pair $\mathcal{L} \subseteq \mathcal{K}$ below, $D_H^{\text{tr}}(\mathcal{K}, \mathcal{L}) := \sup_{\rho \in \mathcal{K}} D_{\text{tr}}(\rho, \mathcal{L})$ denotes the worst-case trace distance from a state in $\mathcal{K}$ to $\mathcal{L}$. We call this the *worst-case de Finetti error*.

Theorem 1.1 (Bosonic two-site de Finetti upper bound; cf. Theorem 3.1). *For every $N, d \geq 2$,*

$$D_H^{\text{tr}}(\text{BOS}_N^{(2)}(d), \text{SEP}_H^{(2)}(d)) \leq \frac{\sqrt{d-1}}{2(N-1)}.$$

Theorem 1.2 (Exchangeable two-site de Finetti upper bound; cf. Corollary 3.7). *For every $N, d \geq 2$,*

$$D_H^{\text{tr}}(\text{EXCH}_N^{(2)}(d), \text{SEP}_P^{(2)}(d)) \leq \frac{\sqrt{d^2-1}}{2(N-1)}.$$

The exchangeable upper bound follows by applying the standard purification and then using the bosonic theorem with dimension d^2 ; the rectangular Werner examples of Christandl, König, Mitchison, and Renner give the matching $\Omega(d/N)$ lower bound, and their symmetric purifications give the matching $\Omega(\sqrt{d}/N)$ bosonic lower bound [CKMR07, Lemma II.5 and Corollary III.9].

Higher marginals. The two-site upper-bound dependence on N and d persists for arbitrary t , with only a linear loss in t . The case $t = 2$ is covered by the preceding two-site theorems; the higher-marginal argument treats $t \geq 3$.

Theorem 1.3 (t -site de Finetti upper bounds; cf. Theorem 4.1). *For every $d \geq 2$ and $2 \leq t \leq N$,*

$$\begin{aligned} D_H^{\text{tr}}(\text{BOS}_N^{(t)}(d), \text{SEP}_H^{(t)}(d)) &\leq \frac{4t\sqrt{d-1}}{N}, \\ D_H^{\text{tr}}(\text{EXCH}_N^{(t)}(d), \text{SEP}_P^{(t)}(d)) &\leq \frac{4t\sqrt{d^2-1}}{N}. \end{aligned}$$

Thus no higher tensor-power loss is necessary: for fixed t , the two-site dependence on N and d persists, and the two-site examples above show that this dependence is optimal. The linear factor in t is also unavoidable in a joint large-dimension regime; see Theorem B.5 in the appendix.

1.2.2 Two-sided Bose-symmetric extendible states and trace-norm applications

Let $\mathcal{H}_A = \mathbb{C}^{d_A}$ and $\mathcal{H}_B = \mathbb{C}^{d_B}$. Recall that a bipartite state is *separable* if it is a convex combination of product states, and write $\text{SEP}(d_A, d_B)$ for the set of such states. For $n, m \geq 1$, let $\text{EXT}_{n,m}(d_A, d_B)$ be the set of states ρ_{AB} admitting an extension on $\mathcal{H}_A^{\otimes n} \otimes \mathcal{H}_B^{\otimes m}$ that is invariant under permutations within the two blocks and whose $A_1 B_1$ marginal is ρ_{AB} [TDS03, JV13, JSZ22]. Its Bose-symmetric subfamily $\text{BEXT}_{n,m}(d_A, d_B)$ consists of the states admitting such an extension supported on $\text{Sym}^n(\mathcal{H}_A) \otimes \text{Sym}^m(\mathcal{H}_B)$. This is the natural two-sided analogue of the one-sided Bose-symmetric extension hierarchy [DPS04, NOP09]. Here “Bose-symmetric” denotes the same support condition called “bosonic” above. Thus $\text{BEXT}_{n,m}$ differs from $\text{EXT}_{n,m}$ only by this symmetric-support requirement. Since pure product states have tensor-power extensions,

$$\text{SEP}(d_A, d_B) \subseteq \text{BEXT}_{n,m}(d_A, d_B) \subseteq \text{EXT}_{n,m}(d_A, d_B).$$

The familiar one-sided k -extendibility hierarchy is the $(1, k)$ case, up to exchanging the two parties. Our bound applies to every pair of extension levels, depends on the smaller local dimension, and captures the multiplicative gain when both extension levels grow. At the one-sided endpoints, specialized de Finetti bounds can be stronger.

Theorem 1.4 (De Finetti upper bound for two-sided Bose-symmetric extendible states; cf. Theorem 3.9). *For every $n, m, d_A, d_B \geq 1$,*

$$D_H^{\text{tr}}(\text{BEXT}_{n,m}(d_A, d_B), \text{SEP}(d_A, d_B)) \leq \frac{1}{2} \sqrt{\frac{\min\{d_A, d_B\} - 1}{nm}}.$$

For balanced extension levels and equal local dimensions, the $\sqrt{d}/n$ dependence is optimal up to universal constants. The rectangular Werner construction of Christandl, König, Mitchison, and Renner gives the matching order after symmetric purification and grouping into two equal blocks [CKMR07].

Corollary 1.5 (De Finetti upper bound for two-sided extendible states; cf. Corollary 3.12). *For every $n, m, d_A, d_B \geq 1$,*

$$D_H^{\text{tr}}(\text{EXT}_{n,m}(d_A, d_B), \text{SEP}(d_A, d_B)) \leq \frac{1}{2} \sqrt{\frac{\min\{d_A^2, d_B^2\} - 1}{nm}}.$$

This follows by taking the canonical purification of the two-block extension and then applying the preceding theorem with local dimensions d_A^2 and d_B^2 . In the balanced case, the resulting d/n dependence for arbitrary two-sided extensions is also optimal in a joint high-dimensional regime, as witnessed by the rectangular Werner examples of [CKMR07].

The Bose-symmetric theorem above is the common interface for the three trace-norm applications. For the disentangler, $\text{BEXT}_{n,n}(d_A, d_B)$ is realized as the image of a channel: the theorem controls every possible output, while tensor-power inputs give exact coverage of the separable states. For BSS, we optimize the verifier's accepting operator over $\text{BEXT}_{n,n}(d_A, d_B)$. Because every feasible state is close to separable, the resulting relaxation approximates $h_{\text{SEP}}(M)$ with an additive error independent of the accepting operator M . Trace-norm weak membership then follows from this BSS approximation.

All three applications therefore have resource bounds at the same scale, subexponential in the local dimension. When $d_A = d_B = d$, the BSS SDP acts on $\text{Sym}^n(\mathbb{C}^d) \otimes \text{Sym}^n(\mathbb{C}^d)$, whose dimension is

$$\binom{d+n-1}{n}^2.$$

The disentangler instead takes $2n$ unrestricted d -dimensional input registers, grouped as $A_1, \dots, A_n$ and $B_1, \dots, B_n$, so its input dimension is $D = d^{2n}$. Both resources have logarithm $O(n \log d)$. The sharp error $O(\sqrt{d}/n)$ allows $n = \Theta(\sqrt{d}/\varepsilon)$, so both resource bounds are $\exp(\tilde{O}(\sqrt{d}/\varepsilon))$. An $O(d/n)$ convergence bound would instead require $n = \Theta(d/\varepsilon)$, returning to the exhaustive-search scale.

Disentanglers. A disentangler was proposed as a way to replace two unentangled quantum proofs by one ordinary quantum proof. It is a channel $\Lambda : D(\mathbb{C}^D) \rightarrow D(\mathbb{C}^d \otimes \mathbb{C}^d)$ with two guarantees. Every output of Λ is ε -close to some separable state, so an entangled input cannot create much entanglement at the output. Conversely, every separable state is δ -close to some output, so the image covers the entire separable set. Such a channel is called an (ε, δ) -disentangler. For sufficiently small errors, if $\log D$ were polylogarithmic in d and the channel could be implemented efficiently,

composing it with a QMA(2) verifier would simulate the two unentangled proofs using one ordinary QMA witness.

Watrous’s conjecture, recorded by Aaronson, Beigi, Drucker, Fefferman, and Shor in 2008 [ABD⁺08], proposed an information-theoretic obstruction: it predicted that $D = \exp(\Omega(d))$ whenever the fixed errors satisfy $\varepsilon + \delta < 1$, where d is the local output dimension. The conjecture had survived for almost two decades. The two-sided Bose-symmetric de Finetti theorem gives a channel with exact coverage, $\delta = 0$, whose input dimension is subexponential in the local output dimension d .

Theorem 1.6 (Subexponential disentanglers). *For every fixed $\varepsilon \in (0, 1)$ and every $d \geq 2$, there is an $(\varepsilon, 0)$ -disentangler*

$$\Lambda : D(\mathbb{C}^D) \longrightarrow D(\mathbb{C}^d \otimes \mathbb{C}^d)$$

with

$$\log D = O_\varepsilon(\sqrt{d} \log d).$$

In particular, there is a $(1/4, 0)$ -disentangler satisfying the same dimension bound.

The image of the channel is $\text{BEXT}_{n,n}(d, d)$. Theorem 1.4 gives soundness, while tensor-power inputs give exact coverage; the explicit construction appears in Subsection 5.1. For an m -qubit output register, where $d = 2^m$, our channel still needs $\Theta_\varepsilon(2^{m/2}m)$ input qubits, so the result does not imply $\text{QMA}(2) = \text{QMA}$. Its significance is instead structural: the conjectured exponential obstruction in the local dimension fails, leaving open the possibility of substantially smaller disentanglers.

Akibue, Kato, and Tani recently proved the conjectured exponential lower bound for *strong disentanglers* [AKT24]. For every reference system R and every joint input, their definition requires the output on ABR to be close to a state separable across $A : BR$; that is, one output must be approximately disentangled from both the other output and the reference. Our channel guarantees only that the reduced output on AB is close to separable across $A : B$ and need not satisfy this stronger requirement. Thus the two results separate the original and strong notions of disentangling rather than contradicting one another.

Best Separable State. The second application asks for the largest acceptance probability of a measurement on a separable state. Given its accepting operator M , satisfying $0 \preceq M \preceq I$ on $\mathbb{C}^d \otimes \mathbb{C}^d$, the *Best Separable State* (BSS) value is

$$h_{\text{SEP}}(M) := \max_{\sigma \in \text{SEP}(d, d)} \text{Tr}(M\sigma) = \max_{\substack{\|x\|=1 \\ \|y\|=1}} \langle x \otimes y | M | x \otimes y \rangle.$$

The equality holds because the linear objective $\text{Tr}(M\sigma)$ is maximized over the separable set by a pure product state. Operationally, M is the accepting operator of a verifier, and $h_{\text{SEP}}(M)$ is the largest acceptance probability achievable by two unentangled messages. This is exactly the optimization core of QMA(2). From a classical perspective, it is also a basic tensor-optimization problem, linked through standard reductions to the landscape surrounding the $2 \rightarrow 4$ norm and Unique Games [BBH⁺12].

For constant accuracy, the direct net-search algorithm takes $\exp(O(d))$ time. Before this work, no $\exp(o(d))$ algorithm was known for general explicit BSS—in other words, no nontrivial improvement over exhaustive search was known. Barak, Kothari, and Steurer came closest: under the perfect-completeness promise $h_{\text{SEP}}(M) = 1$, they obtained an algorithm running in time subexponential in the local dimension and explicitly asked whether that assumption could be removed [BKS17, Remark 1.4]. The two-sided Bose-symmetric extension SDP hierarchy removes it.

Theorem 1.7 (Best separable state without perfect completeness; cf. Theorem 5.2). *Let $0 \preceq M \preceq I$ be explicitly given on $\mathbb{C}^d \otimes \mathbb{C}^d$. For every $\varepsilon \in (0, 1]$, the value $h_{\text{SEP}}(M)$ can be approximated to additive error ε deterministically. Up to a factor polynomial in the input bit complexity, the running time is*

$$\exp\left(\tilde{O}\left(\frac{\sqrt{d}}{\varepsilon}\right)\right).$$

The guarantee in Theorem 1.7 is proved in Subsection 5.2 by evaluating $h_{\text{BEXT}_{n,n}}(M)$. The additive error is independent of the value of $h_{\text{SEP}}(M)$, so the guarantee applies to every choice of completeness and soundness parameters $c > s$ and removes perfect completeness for explicit BSS. The generic de Finetti argument does not exploit perfect completeness and therefore does not recover the ETH-optimal bound established by the direct SoS analysis [JWX26].

Trace-norm separability testing. The third application is separability testing: given an explicitly described bipartite state ρ , decide whether it is separable. Since exact membership is not numerically robust, we use the standard weak-membership formulation, with distance to $\text{SEP}(d, d)$ measured in trace distance. Given a tolerance $\varepsilon > 0$, the task is to distinguish

$$\rho \in \text{SEP}(d, d) \quad \text{from} \quad D_{\text{tr}}(\rho, \text{SEP}(d, d)) \geq \varepsilon,$$

with no requirement when ρ lies in the intervening ε -neighborhood. Approximating the distance to the separable set is a stronger, quantitative version of this problem.

The direction we need—from BSS optimization to this weak-membership problem—is due to Harrow and Montanaro [HM13, Proposition 16]. They show that an additive approximation to $h_{\text{SEP}}(M)$ valid for every accepting operator M gives trace-norm weak membership for the separable set with only a constant-factor loss in accuracy. This is sharper than generic convex-optimization reductions, which lose polynomial factors in the dimension [GLS93]. In our setting, the reduction is visible from trace-distance duality:

$$D_{\text{tr}}(\rho, \text{SEP}(d, d)) = \max_{0 \preceq M \preceq I} \{\text{Tr}(M\rho) - h_{\text{SEP}}(M)\}.$$

Here M is the accepting operator of the two-outcome measurement that best separates ρ from every separable state, while $h_{\text{SEP}}(M)$ is the largest acceptance probability of that same measurement on a separable state. Our new input to the Harrow–Montanaro reduction is a subexponential-time approximation to $h_{\text{SEP}}(M)$ valid for every accepting operator $0 \preceq M \preceq I$. Optimizing the resulting estimate over M gives the distance to separability.

Weak membership for the separable set with promise gap $1/\text{poly}(d)$ is NP-hard [Gur03, Gha10]. Much less is known in the constant-gap regime. In particular, before this work no trace-norm algorithm running in time subexponential in the local dimension was known nor ruled out. The standard $O(d/N)$ de Finetti analysis reaches constant error only at level $N = \Omega(d)$, again at exponential cost; the quasipolynomial algorithms in weaker LOCC and Hilbert–Schmidt geometries do not imply a trace-norm algorithm [BCY11b]. Our subexponential algorithm for $h_{\text{SEP}}(M)$ gives the first such subexponential trace-norm weak-membership algorithm.

Theorem 1.8 (Trace-norm separability testing; cf. Theorem 5.4). *Let ρ be an explicitly given state on $\mathbb{C}^d \otimes \mathbb{C}^d$. For every $\varepsilon \in (0, 1]$, the trace distance from ρ to the separable states can be approximated to additive error ε deterministically. Up to a factor polynomial in the input bit complexity, the running time is*

$$\exp\left(\tilde{O}\left(\frac{\sqrt{d}}{\varepsilon}\right)\right).$$

Harrow and Montanaro connected BSS through reductions with more than a dozen problems, including injective tensor norms, channel output norms, mean-field ground-state energy, minimum output Rényi-2 and min-entropies, and planted clique [HM13, Sec. 4.2]. Our algorithm therefore transfers to every problem in this network that reduces to BSS, with the resulting running time determined by the dimension and accuracy losses of the reduction.

1.2.3 From optimal trace norm to dimension-free Hilbert–Schmidt norm

The classical finite de Finetti theorem is dimension free, while dimension-free quantum convergence cannot hold in trace norm. More practically, the familiar $O(\sqrt{d}/N)$ trace-norm bounds require $N = \Omega(\sqrt{d})$ even for constant accuracy, making de Finetti-based algorithms prohibitively expensive in many applications. These two obstacles motivated de Finetti theorems under restricted measurements [BC12, BH13, LS15]. Rougerie further asked whether a dimension-independent quantum analogue of the classical Diaconis–Freedman theorem could be found [Rou15, Remark 4.2(4)]. The optimal trace-norm theorem motivates another look at this question. We show that Hilbert–Schmidt geometry restores dimension-free convergence, with the optimal dimension-independent rate.

We use Hilbert–Schmidt distance

$$D_{\text{HS}}(\rho, \sigma) := \|\rho - \sigma\|_2, \quad D_{\text{HS}}(\rho, \mathcal{K}) := \inf_{\tau \in \mathcal{K}} D_{\text{HS}}(\rho, \tau),$$

and use $D_{\text{H}}^{\text{HS}}(\mathcal{K}, \mathcal{L})$ for the analogous worst-case de Finetti error in Hilbert–Schmidt distance.

Theorem 1.9 (Dimension-free bosonic Hilbert–Schmidt de Finetti upper bound; cf. Theorem 6.1). *For every $N, d \geq 2$,*

$$D_{\text{H}}^{\text{HS}}\left(\text{BOS}_N^{(2)}(d), \text{SEP}_{\text{H}}^{(2)}(d)\right) = O(N^{-1/2}),$$

with a universal implied constant.

This gives a qualitatively new, dimension-free de Finetti theorem for bosonic states. Its $N^{-1/2}$ dependence is optimal when the dimension is allowed to grow: the lower bound in Theorem B.6 adapts the bosonic pair construction of Jiang, Tacla, and Caves [JTC17]. Unlike in trace norm, the purification argument does not give the same conclusion for arbitrary exchangeable states: tracing out the purifying registers need not contract Hilbert–Schmidt norm. Whether an analogous dimension-free theorem holds in the exchangeable setting remains open.

Theorem 1.10 (Dimension-free Hilbert–Schmidt upper bound for two-sided Bose-symmetric extendible states; cf. Corollary 6.8). *For every $n, d_A, d_B \geq 1$,*

$$D_{\text{H}}^{\text{HS}}\left(\text{BEXT}_{n,n}(d_A, d_B), \text{SEP}(d_A, d_B)\right) = O(n^{-1/2}),$$

with a universal implied constant.

This rate is likewise optimal when the local dimensions may grow. Theorem B.7 groups the same adapted pair construction into two equal blocks.

Because the Hilbert–Schmidt norm is its own dual, the state approximation also controls $h_{\text{SEP}}(M)$ for every Hermitian objective matrix M of bounded Hilbert–Schmidt norm. Adding a multiple of the identity to M merely shifts every feasible value by the same known constant, so only its traceless part matters. For positive-semidefinite objectives, Shi and Wu previously obtained a polynomial-time approximation at fixed accuracy and bounded Hilbert–Schmidt norm by an ε -net method [SW15]. Here the conclusion for arbitrary Hermitian objectives follows instead from a

constant-level symmetric-extension SDP: at every fixed accuracy, it suffices to take $n = O(\varepsilon^{-2})$, so the resulting semidefinite program is polynomial in the local dimensions. The same de Finetti estimate also improves the quasipolynomial weak-membership guarantee of Brandão, Christandl, and Yard [BCY11b]. Concurrent work of Malavolta [Mal26] gives a randomized polynomial-time algorithm for constant-gap Euclidean separability.

Theorem 1.11 (Polynomial-time h_{SEP} approximation; cf. Theorem 7.1). *For every fixed $\varepsilon > 0$, $h_{\text{SEP}}(M)$ admits a deterministic additive- ε approximation in time polynomial in the local dimensions whenever M is explicitly given and its traceless part has Hilbert–Schmidt norm $O(1)$.*

Theorem 1.12 (Polynomial-time Hilbert–Schmidt weak membership; cf. Theorem 7.2). *For every fixed $\varepsilon > 0$, given an explicitly described state $\rho \in D(\mathbb{C}^{d_A} \otimes \mathbb{C}^{d_B})$, its Hilbert–Schmidt distance $D_{\text{HS}}(\rho, \text{SEP}(d_A, d_B))$ can be approximated to additive error ε in deterministic time polynomial in the local dimensions. In the same running time one can solve Hilbert–Schmidt weak membership, distinguishing*

$$\rho \in \text{SEP}(d_A, d_B) \quad \text{from} \quad D_{\text{HS}}(\rho, \text{SEP}(d_A, d_B)) \geq \varepsilon,$$

with no requirement in the intervening regime.




Figure 1: Dependency structure of the main results. Solid arrows show the proof organization used in the main body. Blue nodes are de Finetti statements; tan nodes are applications. The dashed gray branch records further consequences available through the reductions of Harrow–Montanaro [HM13, Sec. 4.2]; their parameters depend on the losses in those reductions.

1.3 Technical overview

The sharp two-site bosonic trace-norm theorem is the engine of the paper. We first explain its argmax proof, obtain the exchangeable theorem by purification, and develop the rectangular version

of argmax rounding. The rectangular estimate directly gives the two-sided de Finetti theorem and its trace-norm applications. The higher-marginal proof then combines the same estimate with a symmetry test. Finally, a separate spectral-truncation argument builds on the sharp trace-norm bound to prove the dimension-free bosonic Hilbert–Schmidt theorem.

1.3.1 Optimal two-site trace-norm de Finetti

We begin with the optimal bosonic and exchangeable two-site bounds stated in Theorems 1.1 and 1.2.

From de Finetti to a symmetric-tensor SDP integrality gap. The route to the proof began with a different problem. In our recent work on sum-of-squares algorithms [JWX26], we developed argmax rounding in pursuit of an ETH-optimal analysis of the SoS algorithm for Best Separable State in the perfect-completeness case. There the idea is quite natural: build a high-moment polynomial from the SoS solution, choose the direction where it is largest, and use the local optimality conditions at this direction for rounding.

The connection between symmetric extensions and SoS/SDP relaxations reaches back at least to the Doherty–Parrilo–Spedalieri (DPS) hierarchy, which casts symmetric extendibility as an SDP test for separability [DPS04]. The duality between SoS relaxations on the sphere and the DPS hierarchy is explicitly recorded by Fang and Fawzi [FF21]. By comparison, the associated rounding viewpoint seems to have been explored less extensively; see, however, Rao [Rao24] for related recent work on Hermitian SoS rounding and quantum de Finetti theorems.

In our setting, trace-norm duality makes this interface concrete: the Hartree optimum of a distinguishing witness is a polynomial optimization over the sphere, while the level- N bosonic optimum is the corresponding symmetric-tensor SDP relaxation. The argmax viewpoint is especially natural here because the de Finetti theorem is information-theoretic: unlike in the algorithmic setting, there is no need to find the maximizing direction efficiently. We may choose it exactly and exploit its local optimality conditions without regard to running time. This shared geometry leads to the optimal de Finetti rate, and the parallel with the earlier analyses goes further.

The BKS argument obtains its rounding direction by Gaussian rounding, averaging over random directions and then conditioning [BKS17]. Standard constructive de Finetti proofs likewise first choose a universal rounding map, usually by averaging tensor powers over Haar-random directions. Given a bosonic state, this fixed map produces one mixture that must approximate its marginal against every measurement at once. The resulting argument gives the familiar $O(d/N)$ bosonic trace-norm bound [CKMR07, Chi11, Har13, LNR15].

Our proof replaces this averaging by an objective-dependent extremal choice: the rounding is tailored to the distinguishing witness. To make that choice, we reverse the order. Trace-norm duality first exposes a bounded Hermitian witness M , with $\|M\|_\infty \leq 1$, for the approximation error; only then do we round an optimizer of the symmetric-tensor SDP relaxation for M . That such dependence can reveal extra information is already suggested by measurement-restricted de Finetti theorems [BH13, LS15] and the value de Finetti theorem for stoquastic verification [GGJ26]. We do not need a single map that succeeds simultaneously for all witnesses. This is a weaker demand, but the dual formulation below says that it is exactly enough.

To make this precise, let $M \in \text{Herm}(\text{Sym}^2 \mathcal{H})$ and define

$$h(M) := \max_{\|u\|=1} \langle u^{\otimes 2} | M | u^{\otimes 2} \rangle,$$

$$M^{[N]} := \binom{N}{2}^{-1} \sum_{1 \leq i < j \leq N} M_{ij} \quad \text{on } \text{Sym}^N \mathcal{H},$$

where M_{ij} acts as M on sites i, j and as the identity on every other site. Thus $h(M)$ is the original quartic optimization, while

$$\lambda_{\max}(M^{[N]}) = \max_{\rho_N \in \mathcal{D}(\text{Sym}^N \mathcal{H})} \text{Tr}(M^{[N]} \rho_N)$$

is the value of its level- N symmetric-tensor SDP relaxation, the bosonic form of the corresponding SoS relaxation. After dualizing trace distance and using permutation symmetry, the worst-case de Finetti error is exactly

$$D_{\text{H}}^{\text{tr}}(\text{BOS}_N^{(2)}(d), \text{SEP}_H^{(2)}(d)) = \frac{1}{2} \sup_{\|M\|_{\infty} \leq 1} \left\{ \lambda_{\max}(M^{[N]}) - h(M) \right\}. \quad (1.5)$$

So the de Finetti problem has become a familiar TCS question: bound the integrality gap of an SDP relaxation. We no longer need to construct the approximating mixture. For each bounded Hermitian witness M , it is enough to round an optimizer of the corresponding relaxation with loss $O(\sqrt{d}/N)$.

The argmax rounding object. Fix M and let $\psi \in \text{Sym}^N(\mathcal{H})$ be a unit top eigenvector of $M^{[N]}$. The argmax principle gives a natural rounding object: choose the tensor power $u^{\otimes N}$ having maximum overlap with ψ . After fixing its phase, write

$$c := \langle u^{\otimes N}, \psi \rangle > 0.$$

The argmax lemma provides the basic tools for analyzing this rounding object. Take the partial inner product of ψ with u on all but two sites:

$$\eta_u := (\langle u |^{\otimes (N-2)} \otimes I_{\mathcal{H}^{\otimes 2}}) \psi \in \text{Sym}^2(\mathcal{H}).$$

Concretely, the lemma says

$$\eta_u = c u^{\otimes 2} + \zeta_u, \quad \zeta_u \in \text{Sym}^2(u^{\perp}), \quad \sup_{\substack{v \perp u \\ \|v\|=1}} |\langle v^{\otimes 2}, \zeta_u \rangle| \leq \frac{c}{N-1}. \quad (1.6)$$

Here is why. Perturb u toward any orthogonal direction v and differentiate the overlap with ψ . The first derivative says that the component with one copy of v vanishes. The second derivative says that the component with two copies of v is only $O(c/N)$. Doing this for every $v \perp u$ gives (1.6). In words, the partial inner product is the tensor square $u^{\otimes 2}$, scaled by c , plus an error living entirely in directions orthogonal to u and having small overlap with every orthogonal square $v^{\otimes 2}$. These are precisely the two tools supplied by the argmax lemma: the mixed component vanishes, and the fully orthogonal component is small against every square. This is the bosonic counterpart of the pseudo-covariance lemma underlying our SoS argmax rounding [JWX26].

Bounding the integrality gap. We now use the information supplied by the argmax lemma to compare $\lambda_{\max}(M^{[N]})$ with $h(M)$. By symmetry, the lift $M^{[N]}$ can be evaluated on sites 1, 2, and hence

$$\lambda_{\max}(M^{[N]})c = \langle M u^{\otimes 2}, \eta_u \rangle. \quad (1.7)$$

Let

$$B_u := P_{\text{Sym}^2(u^{\perp})} M u^{\otimes 2}.$$

The component of $Mu^{\otimes 2}$ along $u^{\otimes 2}$ contributes the feasible tensor-square value $\langle u^{\otimes 2} | M | u^{\otimes 2} \rangle \leq h(M)$; the component in $u \vee u^\perp$ is orthogonal to η_u . Equations (1.6) and (1.7) therefore leave the exact identity

$$(\lambda_{\max}(M^{[N]}) - \langle u^{\otimes 2} | M | u^{\otimes 2} \rangle) c = \langle B_u, \zeta_u \rangle.$$

Choose an orthonormal basis of $u^\perp$ and identify symmetric two-tensors with complex symmetric matrices. Write C_u for the matrix representing B_u and Z_u for the matrix representing ζ_u . The argmax bound (1.6) says $\|Z_u\|_\infty \leq c/(N-1)$. Moreover, C_u is a $(d-1) \times (d-1)$ matrix and $\|C_u\|_2 = \|B_u\| \leq \|M\|_\infty$. Its trace norm is therefore at most $\sqrt{d-1} \|M\|_\infty$. Schatten duality now gives

$$|\langle B_u, \zeta_u \rangle| = \left| \text{Tr}(C_u^\dagger Z_u) \right| \leq \frac{c\sqrt{d-1}}{N-1} \|M\|_\infty.$$

Consequently,

$$\lambda_{\max}(M^{[N]}) \leq h(M) + \frac{\sqrt{d-1}}{N-1} \|M\|_\infty.$$

Together with (1.5), this proves the sharp $O(\sqrt{d}/N)$ theorem. The important point is that argmax rounding leaves the single matrix pairing $\text{Tr}(C_u^\dagger Z_u)$. Passing from the Hilbert–Schmidt norm of the $(d-1) \times (d-1)$ matrix C_u to its trace norm costs only $\sqrt{d-1}$, which is exactly the improvement we need.

The exchangeable theorem follows by purifying each site into dimension d^2 , applying the bosonic theorem, and tracing out the purifying registers.

Two-sided extendibility. The same proof extends directly to two symmetric blocks. For a Hermitian operator M on $\mathcal{H}_A \otimes \mathcal{H}_B$, average M over the nm cross-block pairs to form the rectangular lift $M^{[n,m]}$. Maximizing the overlap of a top eigenvector with $x^{\otimes n} \otimes y^{\otimes m}$, the same first- and second-order argument gives

$$\lambda_{\max}(M^{[n,m]}) - h_{\text{SEP}}(M) \leq \sqrt{\frac{\min\{d_A, d_B\} - 1}{nm}} \|M\|_\infty.$$

Trace-distance duality yields Theorem 1.4, and a canonical purification of the two blocks yields Corollary 1.5.

The disentangler from the two-sided theorem. The disentangler in Theorem 1.6 is obtained by realizing $\text{BEXT}_{n,n}(d, d)$ as the image of a channel: project the $A_1, \dots, A_n$ registers and the $B_1, \dots, B_n$ registers onto $\text{Sym}^n(\mathbb{C}^d)$, sending rejected mass to a fixed product state, and then retain only $A_1 B_1$. Every output is Bose-symmetric (n, n) -extendible, while every separable state has an exact preimage obtained from tensor powers of its product-state decomposition. Choosing $n = O(\sqrt{d}/\varepsilon)$ gives an $(\varepsilon, 0)$ -disentangler with input dimension $D = d^{2n}$, and hence $\log D = O_\varepsilon(\sqrt{d} \log d)$.

BSS and weak membership from the same theorem. For an accepting operator $0 \preceq M \preceq I$ on AB , optimize over $\text{BEXT}_{n,n}(d_A, d_B)$ and write

$$h_{\text{BEXT}_{n,n}}(M) := \max_{\rho \in \text{BEXT}_{n,n}(d_A, d_B)} \text{Tr}(M\rho).$$

Every separable state is feasible, while the two-sided theorem places every feasible state within $O(\sqrt{\min\{d_A, d_B\}}/n)$ of the separable states. Consequently,

$$h_{\text{SEP}}(M) \leq h_{\text{BEXT}_{n,n}}(M) \leq h_{\text{SEP}}(M) + O\left(\frac{\sqrt{\min\{d_A, d_B\}}}{n}\right).$$

This bound, valid for every M , gives the BSS approximation in Theorem 1.7. The Harrow–Montanaro reduction [HM13, Proposition 16] then gives trace-norm weak membership, and the equivalent trace-distance dual formulation gives the additive distance approximation in Theorem 1.8.

1.3.2 t -site trace-norm de Finetti

The $1 : (t - 1)$ cut. For the t -site bounds in Theorem 1.3, trace-distance duality asks us to compare the value of a witness M with its Hartree optimum. For $0 \preceq M \preceq I$ on $\text{Sym}^t(\mathcal{H})$, this optimum is

$$h_t(M) := \max_{\|u\|=1} \langle u^{\otimes t} | M | u^{\otimes t} \rangle.$$

The two-site argument suggests first enforcing unentanglement across a bipartition. The natural choice is the $1 : (t - 1)$ cut

$$\mathcal{H}_A := \mathcal{H}, \quad \mathcal{H}_B := \text{Sym}^{t-1}(\mathcal{H}), \quad \text{Sym}^t(\mathcal{H}) \subseteq \mathcal{H}_A \otimes \mathcal{H}_B.$$

This cut exposes one factor of dimension d , so the loss in rectangular argmax rounding depends on $\sqrt{d-1}$ rather than on the much larger dimension of $\text{Sym}^{t-1}(\mathcal{H})$. It also fits the bosonic extension naturally: organizing the N sites into n individual sites and m blocks of $t - 1$ sites makes $\rho_N^{(t)}$ Bose-symmetric (n, m) -extendible across $\mathcal{H}_A : \mathcal{H}_B$. The associated rectangular lift is therefore the two-sided symmetric-extension relaxation of product optimization across this cut.

The argmax principle supplies the main rounding step. The value of the relaxation is the largest eigenvalue of the rectangular lift. Let ψ be a unit top eigenvector of the rectangular lift. Argmax rounding chooses

$$(x, y) \in \arg \max_{\substack{\|v\|=1 \\ \|w\|=1}} |\langle v^{\otimes n} \otimes w^{\otimes m}, \psi \rangle|,$$

where $v, x \in \mathcal{H}_A$ and $w, y \in \mathcal{H}_B$, and returns the product vector $x \otimes y$. This enforces unentanglement across the $1 : (t - 1)$ cut, but it does not yet imply that $y = x^{\otimes(t-1)}$, as required for a Hartree vector.

Symmetry propagates unentanglement. To propagate this unentanglement to the remaining cuts, we use the symmetry test in the spirit of the disentangler from unentanglement of Jeronimo and Wu [JW24]. Let Π_t be the orthogonal projector from $\mathcal{H}_A \otimes \mathcal{H}_B$ onto $\text{Sym}^t(\mathcal{H})$. A product vector in this symmetric subspace must be a tensor power. Define the symmetry-test rejection probability by

$$\delta_{\text{sym}} := 1 - \langle x \otimes y | \Pi_t | x \otimes y \rangle.$$

Permutation averaging gives

$$1 - \left| \langle x^{\otimes(t-1)}, y \rangle \right|^2 \leq t \delta_{\text{sym}}.$$

Thus, when the rounded product passes the symmetry test with high probability, its unentanglement across the $1 : (t - 1)$ cut propagates to the remaining cuts, and $x \otimes y$ is close to the tensor power $x^{\otimes t}$. This is the content of Lemma 4.4; its linear dependence on t sharpens the corresponding one-versus-many estimate in [JW24, Theorem 23].

Bounding the integrality gap. We incorporate the symmetry test before rectangular argmax rounding so that the witness value and the rejection probability are controlled through the same rounded product. Extend M by zero outside $\text{Sym}^t(\mathcal{H})$ and introduce the mixed test

$$M_\theta := (1 - \theta)\Pi_t + \theta M, \quad 0 < \theta < 1.$$

It performs the symmetry test with weight $1 - \theta$ and evaluates M with weight θ . Every bosonic t -site marginal passes the symmetry test with certainty. If $\tilde{\rho}_{n,m}$ is its grouped extension, then

$$1 - \theta + \theta \operatorname{Tr}(M \rho_N^{(t)}) = \operatorname{Tr}(M_\theta^{[n,m]} \tilde{\rho}_{n,m}) \leq \lambda_{\max}(M_\theta^{[n,m]}).$$

The right-hand side is the rectangular symmetric-extension relaxation of $h_{\text{SEP}}(M_\theta)$. Applying rectangular argmax rounding to this lift produces $x \otimes y$. On that product vector,

$$\langle x \otimes y | M_\theta | x \otimes y \rangle = (1 - \theta)(1 - \delta_{\text{sym}}) + \theta \langle x \otimes y | M | x \otimes y \rangle,$$

so rejection lowers the mixed-test value. Lemma 4.4 controls the replacement of $x \otimes y$ by $x^{\otimes t}$ in terms of the same quantity; the proof keeps these two contributions coupled rather than first establishing a separate bound on δ_{sym} .

Let β denote the error in rectangular argmax rounding. Combining the rectangular argmax estimate with the symmetry-test bounds gives

$$\theta \left(\operatorname{Tr}(M \rho_N^{(t)}) - h_t(M) \right) \leq -(1 - \theta)\delta_{\text{sym}} + (\theta\sqrt{t} + \beta(1 - \theta))\sqrt{\delta_{\text{sym}}} + \beta\theta. \quad (1.8)$$

Choosing the two extension levels and θ appropriately in this inequality gives the claimed $O(t\sqrt{d}/N)$ bound. Trace-distance duality proves the bosonic theorem, and symmetric purification gives the exchangeable theorem.

1.3.3 From optimal trace norm to dimension-free Hilbert–Schmidt norm

The dimension-free bosonic Hilbert–Schmidt de Finetti theorem stated in Theorem 1.9 uses the sharp trace-norm theorem in a different way. On a subspace of dimension less than N , the $\sqrt{d}/N$ trace-norm bound is already $O(N^{-1/2})$. The spectral subspace of the one-site marginal corresponding to eigenvalues above $1/N$ has dimension less than N . The proof applies the trace-norm theorem on this spectral head and controls every block meeting the complementary tail directly in Hilbert–Schmidt norm.

Spectral truncation. We first reduce to a pure bosonic state $\rho_N = |\psi\rangle\langle\psi|$ and look at its one-site marginal ρ_1 . Define

$$P := \mathbf{1}_{(1/N, \infty)}(\rho_1), \quad Q := I - P.$$

We call $P\mathcal{H}$ the spectral head and $Q\mathcal{H}$ the spectral tail. Since $\operatorname{Tr} \rho_1 = 1$,

$$\operatorname{rank} P < N, \quad Q\rho_1 Q \preceq \frac{1}{N}Q.$$

Let

$$\rho_2^{\text{head}} := (P \otimes P)\rho_2(P \otimes P)$$

be the compression of the two-site marginal to the head. The proof constructs a Hartree mixture τ_{head} , supported on $\operatorname{Sym}^2(P\mathcal{H})$ and having the same trace as ρ_2^{head} . A second Hartree mixture τ_{res} , supported on $\operatorname{Sym}^2(Q\mathcal{H})$, fills the missing trace. Thus

$$\tau := \tau_{\text{head}} + \tau_{\text{res}} \in \operatorname{SEP}_{\text{H}}^{(2)}(d).$$

The three estimates enter through the error decomposition

$$\|\rho_2 - \tau\|_2 \leq \underbrace{\left\| \rho_2^{\text{head}} - \tau_{\text{head}} \right\|_2}_{\text{round the high-eigenvalue compression}} + \underbrace{\left\| \rho_2 - \rho_2^{\text{head}} \right\|_2}_{\text{remove all blocks meeting the spectral tail}} + \underbrace{\left\| \tau_{\text{res}} \right\|_2}_{\text{restore the missing trace}}. \quad (1.9)$$

The head compression is not itself an N -site marginal, so we decompose ψ according to the number k of sites in the head. In the component with exactly k sites in $P\mathcal{H}$, the probability that a prescribed pair lies in the head is $\frac{k(k-1)}{N(N-1)}$, whereas the sharp trace-norm theorem costs only the square root of the head dimension divided by $k-1$. The factor $k-1$ cancels, and the fact that the head dimension is below N leaves an $O(N^{-1/2})$ Hilbert–Schmidt error. This cancellation is exactly where the optimal two-site theorem enters the dimension-free argument.

For the remaining terms, decompose $\text{Sym}^2 \mathcal{H}$ as

$$\text{Sym}^2(P\mathcal{H}) \oplus (P\mathcal{H} \vee Q\mathcal{H}) \oplus \text{Sym}^2(Q\mathcal{H}).$$

The cutoff $Q\rho_1 Q \preceq Q/N$ gives operator-norm control on the two diagonal compressions containing at least one Q direction. An ordinary Schmidt decomposition handles $P\mathcal{H} \vee Q\mathcal{H}$, while its symmetric form, or Takagi decomposition, handles $\text{Sym}^2(Q\mathcal{H})$. These decompositions reduce the estimates to partial inner products of ψ with orthogonal vectors in $Q\mathcal{H}$, each controlled by the cutoff. Positive-block estimates first combine the mixed and tail–tail diagonal bounds into an operator-norm bound on the full tail block and then control the off-diagonal head–tail blocks. Thus removing every block meeting $Q\mathcal{H}$ costs only $O(N^{-1/2})$ in Hilbert–Schmidt norm. Finally, restore the missing trace by adding the Haar-averaged Hartree mixture on $\text{Sym}^2(Q\mathcal{H})$; its Hilbert–Schmidt norm is $O(1/N)$. Substitution into (1.9) proves the dimension-free $O(N^{-1/2})$ theorem.

Two-sided extensions. The dimension-free Hilbert–Schmidt de Finetti theorem above extends to two-sided Bose-symmetric extendible states:

$$D_{\text{H}}^{\text{HS}}(\text{BEXT}_{n,n}(d_A, d_B), \text{SEP}(d_A, d_B)) = O(n^{-1/2}).$$

This is the content of Theorem 1.10. Since the Hilbert–Schmidt norm is self-dual, the two-sided Hilbert–Schmidt bound gives analogues of the BSS and weak-membership results above. Taking $n = O(\varepsilon^{-2})$ yields the polynomial-time h_{SEP} approximation and separability testing in Theorems 1.11 and 1.12.

2 Preliminaries

All Hilbert spaces in this note are finite-dimensional and complex. We denote them by $\mathcal{H}$, adding subscripts to indicate their roles, and take their inner products to be conjugate-linear in the first argument. For $z \in \mathbb{C}$, we write $\Re(z)$ and $\Im(z)$ for its real and imaginary parts, respectively. We write $\text{L}(\mathcal{H})$ for the space of linear operators on $\mathcal{H}$ and $\text{Herm}(\mathcal{H})$ for its real subspace of Hermitian operators; and $\text{D}(\mathcal{H}) \subseteq \text{Herm}(\mathcal{H})$ the density operators. For a vector v , the notation $\|v\|$ denotes its Hilbert-space norm. For $u \in \mathcal{H}$, we write $u^\perp := \{v \in \mathcal{H} : \langle u, v \rangle = 0\}$; more generally, $\mathcal{U}^\perp$ denotes the orthogonal complement of a subspace $\mathcal{U}$, with the ambient Hilbert space understood from context. For $X \in \text{L}(\mathcal{H})$, let $|X| := \sqrt{X^\dagger X}$ and define the Schatten norms by

$$\|X\|_p := (\text{Tr } |X|^p)^{\frac{1}{p}} \quad (1 \leq p < \infty), \quad \|X\|_\infty := \max_j s_j(X),$$

where $s_j(X)$ are the singular values of X . Thus $\|X\|_1$, $\|X\|_2$, and $\|X\|_\infty$ are the trace, Hilbert–Schmidt, and operator norms, respectively. We use normalized trace distance and unnormalized Hilbert–Schmidt distance:

$$D_{\text{tr}}(\rho, \sigma) := \frac{1}{2} \|\rho - \sigma\|_1, \tag{2.1}$$

and

$$D_{\text{HS}}(\rho, \sigma) := \|\rho - \sigma\|_2. \quad (2.2)$$

For probability distributions, D_{TV} denotes total-variation distance. Thus every displayed trace-distance constant is twice as large when the error is measured by $\|\rho - \sigma\|_1$ instead.

For $k \geq 1$, let $\mathfrak{S}_k$ act on $\mathcal{H}^{\otimes k}$ by permuting the tensor factors, and write U_π for the operator corresponding to $\pi \in \mathfrak{S}_k$. The symmetric tensor power is

$$\text{Sym}^k \mathcal{H} := \left\{ \Psi \in \mathcal{H}^{\otimes k} : U_\pi \Psi = \Psi \text{ for every } \pi \in \mathfrak{S}_k \right\},$$

and we set $\text{Sym}^0 \mathcal{H} := \mathbb{C}$.

2.1 Separable and extendible states

The de Finetti theorems compare t -site marginals of bosonic or exchangeable states, and bipartite states admitting two-sided extensions, with specific separable state families. We now define the states appearing in these comparisons.

Separable state families. Fix $d \geq 1$ and $t \geq 1$. The two separable families used below are

$$\begin{aligned} \text{SEP}_{\text{H}}^{(t)}(d) &:= \text{conv} \left\{ |u\rangle\langle u|^{\otimes t} : u \in \mathbb{C}^d, \|u\| = 1 \right\}, \\ \text{SEP}_{\text{P}}^{(t)}(d) &:= \text{conv} \left\{ \sigma^{\otimes t} : \sigma \in \text{D}(\mathbb{C}^d) \right\}. \end{aligned}$$

The subscripts H and P stand for Hartree and tensor power, respectively. Since pure states are density operators,

$$\text{SEP}_{\text{H}}^{(t)}(d) \subseteq \text{SEP}_{\text{P}}^{(t)}(d).$$

The families $\text{SEP}_{\text{H}}^{(t)}(d)$ and $\text{SEP}_{\text{P}}^{(t)}(d)$ are the approximants in the bosonic and exchangeable de Finetti theorems, respectively.

Bosonic states. Let $N \geq t$. A state ρ_N is *bosonic* if $\rho_N \in \text{D}(\text{Sym}^N(\mathbb{C}^d))$. Its marginal on the first t sites is

$$\rho_N^{(t)} := \text{Tr}_{t+1, \dots, N} \rho_N.$$

The set of these marginals is

$$\text{BOS}_N^{(t)}(d) := \left\{ \rho_N^{(t)} : \rho_N \in \text{D}(\text{Sym}^N(\mathbb{C}^d)) \right\}.$$

Extending Hartree states by tensor powers and tracing out extension sites give

$$\text{SEP}_{\text{H}}^{(t)}(d) \subseteq \text{BOS}_{N+1}^{(t)}(d) \subseteq \text{BOS}_N^{(t)}(d).$$

Exchangeable states. A state $\rho_N \in \text{D}((\mathbb{C}^d)^{\otimes N})$ is *exchangeable* if

$$U_\pi \rho_N U_\pi^\dagger = \rho_N \quad \text{for every } \pi \in \mathfrak{S}_N.$$

The set of all such t -site marginals is

$$\text{EXCH}_N^{(t)}(d) := \left\{ \rho_N^{(t)} : \begin{array}{l} \rho_N \in \text{D}((\mathbb{C}^d)^{\otimes N}), \\ U_\pi \rho_N U_\pi^\dagger = \rho_N \text{ for every } \pi \in \mathfrak{S}_N \end{array} \right\}.$$

Tensor powers extend exchangeably, and every bosonic state is exchangeable. Thus

$$\begin{aligned} \text{SEP}_{\text{P}}^{(t)}(d) &\subseteq \text{EXCH}_{N+1}^{(t)}(d) \subseteq \text{EXCH}_N^{(t)}(d), \\ \text{BOS}_N^{(t)}(d) &\subseteq \text{EXCH}_N^{(t)}(d). \end{aligned}$$

Two-sided extendible states. For $d_A, d_B \geq 1$, define the bipartite separable states by

$$\text{SEP}(d_A, d_B) := \text{conv} \left\{ |u\rangle\langle u| \otimes |v\rangle\langle v| : u \in \mathbb{C}^{d_A}, v \in \mathbb{C}^{d_B}, \|u\| = \|v\| = 1 \right\}.$$

For $n, m \geq 1$, the (n, m) -extendible states are the marginals

$$\text{EXT}_{n,m}(d_A, d_B) := \left\{ \tilde{\rho}_{A_1 B_1} : \begin{array}{l} \tilde{\rho} \in \mathcal{D}((\mathbb{C}^{d_A})^{\otimes n} \otimes (\mathbb{C}^{d_B})^{\otimes m}), \\ \tilde{\rho} \text{ is invariant under permutations within each block} \end{array} \right\}.$$

This is the standard notion of two-sided extendibility [TDS03, JV13, JSZ22]. Requiring the extension to be supported on the symmetric subspace of each block gives the Bose-symmetric extendible states

$$\text{BEXT}_{n,m}(d_A, d_B) := \left\{ \tilde{\rho}_{A_1 B_1} : \tilde{\rho} \in \mathcal{D}(\text{Sym}^n(\mathbb{C}^{d_A}) \otimes \text{Sym}^m(\mathbb{C}^{d_B})) \right\}. \quad (2.3)$$

This is the two-sided analogue of Bose-symmetric extendibility in [NOP09]. If $n' \geq n$ and $m' \geq m$, tracing out extension sites and extending every term in a pure-product decomposition by tensor powers give

$$\begin{aligned} \text{SEP}(d_A, d_B) &\subseteq \text{BEXT}_{n',m'}(d_A, d_B) \subseteq \text{BEXT}_{n,m}(d_A, d_B) \subseteq \text{EXT}_{n,m}(d_A, d_B), \\ \text{EXT}_{n',m'}(d_A, d_B) &\subseteq \text{EXT}_{n,m}(d_A, d_B). \end{aligned} \quad (2.4)$$




Figure 2: State-set containments underlying the de Finetti theorems. Nesting records containment. From left to right, the dotted markers represent the worst-case de Finetti errors for $(\text{EXCH}_N^{(t)}, \text{SEP}_P^{(t)})$, $(\text{BOS}_N^{(t)}, \text{SEP}_H^{(t)})$, $(\text{BEXT}_{n,m}, \text{SEP})$, and $(\text{EXT}_{n,m}, \text{SEP})$. Marker lengths are schematic.

2.2 Hausdorff notation and convex duality

For $\star \in \{\text{tr}, \text{HS}\}$ and a nonempty compact set $\mathcal{K}$ of states, write

$$D_\star(\rho, \mathcal{K}) := \inf_{\sigma \in \mathcal{K}} D_\star(\rho, \sigma). \quad (2.5)$$

For two nonempty compact sets $\mathcal{K}$ and $\mathcal{L}$ of states, their Hausdorff distance is

$$D_H^\star(\mathcal{K}, \mathcal{L}) := \max \left\{ \sup_{\rho \in \mathcal{K}} D_\star(\rho, \mathcal{L}), \sup_{\sigma \in \mathcal{L}} D_\star(\sigma, \mathcal{K}) \right\}. \quad (2.6)$$

In particular, if $\mathcal{L} \subseteq \mathcal{K}$, then

$$D_H^\star(\mathcal{K}, \mathcal{L}) = \sup_{\rho \in \mathcal{K}} D_\star(\rho, \mathcal{L}). \quad (2.7)$$

Figure 2 summarizes the relevant containments. Consequently, every de Finetti error in the paper has the form $D_H^*(\mathcal{K}, \mathcal{L})$, with the larger finite-extension set $\mathcal{K}$ first and its separable subset $\mathcal{L}$ second. The four instances are

$$\begin{aligned} &(\text{BOS}_N^{(t)}(d), \text{SEP}_H^{(t)}(d)), & (\text{EXCH}_N^{(t)}(d), \text{SEP}_P^{(t)}(d)), \\ &(\text{BEXT}_{n,m}(d_A, d_B), \text{SEP}(d_A, d_B)), & (\text{EXT}_{n,m}(d_A, d_B), \text{SEP}(d_A, d_B)). \end{aligned}$$

These quantities are exactly the corresponding worst-case de Finetti errors.

For a nonempty compact set $\mathcal{C} \subseteq \text{D}(\mathcal{H})$, write

$$h_{\mathcal{C}}(M) := \max_{\rho \in \mathcal{C}} \text{Tr}(M\rho) \quad (M \in \text{Herm}(\mathcal{H})) \quad (2.8)$$

for its support function. The following lemma collects the distance dualities used throughout the paper.

Lemma 2.1 (Distance duality for compact convex state sets). *Let $\mathcal{L} \subseteq \text{D}(\mathcal{H})$ be nonempty, compact, and convex. For every $\rho \in \text{D}(\mathcal{H})$,*

$$\begin{aligned} D_{\text{tr}}(\rho, \mathcal{L}) &= \max_{0 \preceq M \preceq I} \{\text{Tr}(M\rho) - h_{\mathcal{L}}(M)\} \\ &= \frac{1}{2} \max_{\substack{M \in \text{Herm}(\mathcal{H}) \\ \|M\|_{\infty} \leq 1}} \{\text{Tr}(M\rho) - h_{\mathcal{L}}(M)\}, \end{aligned} \quad (2.9)$$

$$\begin{aligned} D_{\text{HS}}(\rho, \mathcal{L}) &= \max_{\substack{M \in \text{Herm}(\mathcal{H}) \\ \|M\|_2 \leq 1}} \{\text{Tr}(M\rho) - h_{\mathcal{L}}(M)\} \\ &= \max_{\substack{M \in \text{Herm}(\mathcal{H}) \\ \text{Tr } M=0, \|M\|_2 \leq 1}} \{\text{Tr}(M\rho) - h_{\mathcal{L}}(M)\}. \end{aligned} \quad (2.10)$$

If $\mathcal{K} \subseteq \text{D}(\mathcal{H})$ is also nonempty, compact, and convex, and $\mathcal{L} \subseteq \mathcal{K}$, then

$$\begin{aligned} D_H^{\text{tr}}(\mathcal{K}, \mathcal{L}) &= \max_{0 \preceq M \preceq I} \{h_{\mathcal{K}}(M) - h_{\mathcal{L}}(M)\} \\ &= \frac{1}{2} \max_{\substack{M \in \text{Herm}(\mathcal{H}) \\ \|M\|_{\infty} \leq 1}} \{h_{\mathcal{K}}(M) - h_{\mathcal{L}}(M)\}, \end{aligned} \quad (2.11)$$

$$D_H^{\text{HS}}(\mathcal{K}, \mathcal{L}) = \max_{\substack{M \in \text{Herm}(\mathcal{H}) \\ \text{Tr } M=0, \|M\|_2 \leq 1}} \{h_{\mathcal{K}}(M) - h_{\mathcal{L}}(M)\}. \quad (2.12)$$

Proof. For states ρ and σ , trace-norm duality and $\text{Tr}(\rho - \sigma) = 0$ give

$$D_{\text{tr}}(\rho, \sigma) = \max_{0 \preceq M \preceq I} \text{Tr}(M(\rho - \sigma)) = \frac{1}{2} \max_{\substack{M \in \text{Herm}(\mathcal{H}) \\ \|M\|_{\infty} \leq 1}} \text{Tr}(M(\rho - \sigma)).$$

Hilbert–Schmidt self-duality similarly gives

$$D_{\text{HS}}(\rho, \sigma) = \max_{\substack{M \in \text{Herm}(\mathcal{H}) \\ \|M\|_2 \leq 1}} \text{Tr}(M(\rho - \sigma)).$$

In each case, Sion’s finite-dimensional minimax theorem [Sio58] exchanges the minimum over $\sigma \in \mathcal{L}$ with the maximum over the corresponding compact convex witness set. This proves both identities in (2.9) and the first identity in (2.10).

To obtain the formula with $\text{Tr } M = 0$, let $q = \dim \mathcal{H}$ and set

$$M_0 := M - \frac{\text{Tr } M}{q} I.$$

Replacing M by M_0 leaves $\text{Tr}(M\rho) - h_{\mathcal{L}}(M)$ unchanged, while orthogonality to the identity gives $\|M_0\|_2 \leq \|M\|_2$. Thus the maximum may be restricted to traceless witnesses.

Finally, take the maximum over $\rho \in \mathcal{K}$ in the state-to-set formulas. The two remaining maxima may be interchanged, and maximization over ρ replaces $\text{Tr}(M\rho)$ by $h_{\mathcal{K}}(M)$. Equation (2.7) then gives (2.11) and (2.12). $\square$

2.3 Linear-algebra estimates

Fact 2.2 (Positive block estimates). *Let $A \succeq 0$ act on $\mathcal{H}_A$, let $B \succeq 0$ act on $\mathcal{H}_B$, and let $X : \mathcal{H}_B \rightarrow \mathcal{H}_A$. Then*

$$\begin{pmatrix} A & X \\ X^\dagger & B \end{pmatrix} \succeq 0$$

if and only if

$$X = A^{1/2} K B^{1/2}$$

for a contraction $K : \mathcal{H}_B \rightarrow \mathcal{H}_A$. In this case,

$$\left\| \begin{pmatrix} A & X \\ X^\dagger & B \end{pmatrix} \right\|_\infty \leq \|A\|_\infty + \|B\|_\infty, \quad \|X\|_2^2 \leq \|B\|_\infty \text{Tr } A. \quad (2.13)$$

Proof. The factorization statement is [Bha07, Proposition 1.3.2]. It remains to verify the norm estimates. For $a := \|A\|_\infty$ and $b := \|B\|_\infty$,

$$\langle u \oplus v | \begin{pmatrix} A & X \\ X^\dagger & B \end{pmatrix} | u \oplus v \rangle \leq \left(\sqrt{a} \|u\| + \sqrt{b} \|v\| \right)^2 \leq (a+b) (\|u\|^2 + \|v\|^2),$$

which proves the operator-norm bound. The factorization also gives

$$X X^\dagger = A^{1/2} K B K^\dagger A^{1/2} \preceq \|B\|_\infty A,$$

and taking traces proves the bound on $\|X\|_2$. $\square$

Fact 2.3 (Discarding a positive block). *Under the hypotheses of Fact 2.2, suppose that*

$$M := \begin{pmatrix} A & X \\ X^\dagger & B \end{pmatrix} \succeq 0.$$

Then

$$\left\| M - \begin{pmatrix} A & 0 \\ 0 & 0 \end{pmatrix} \right\|_2^2 \leq \|B\|_\infty (2 \text{Tr } A + \text{Tr } B). \quad (2.14)$$

In particular, if $\text{Tr } M = 1$, then the right-hand side is at most $2 \|B\|_\infty$.

Proof. Block orthogonality in Hilbert–Schmidt inner product gives

$$\left\| \begin{pmatrix} 0 & X \\ X^\dagger & B \end{pmatrix} \right\|_2^2 = 2 \|X\|_2^2 + \|B\|_2^2.$$

Since $\|B\|_2^2 \leq \|B\|_\infty \text{Tr } B$, the bound on $\|X\|_2$ in (2.13) proves (2.14). $\square$

2.4 Tensor norms and symmetric tensors

Injective and projective tensor norms. Let $\mathcal{X}$ and $\mathcal{Y}$ be finite-dimensional Hilbert spaces. For $z \in \mathcal{X} \otimes \mathcal{Y}$, define the injective and projective tensor norms by

$$\|z\|_\varepsilon := \sup_{\substack{\|x\|=1 \\ \|y\|=1}} |\langle x \otimes y, z \rangle|, \quad (2.15)$$

$$\|z\|_\pi := \inf \left\{ \sum_j |a_j| : z = \sum_j a_j x_j \otimes y_j, \|x_j\| = \|y_j\| = 1 \right\}. \quad (2.16)$$

The unadorned norm $\|z\|$ remains the Hilbert-space norm. After choosing orthonormal bases of $\mathcal{X}$ and $\mathcal{Y}$, let Z be the coefficient matrix of z . The singular-value decomposition, equivalently the Schmidt decomposition of z , gives

$$\|z\|_\varepsilon = \|Z\|_\infty, \quad \|z\| = \|Z\|_2, \quad \|z\|_\pi = \|Z\|_1. \quad (2.17)$$

In particular, if $\text{Schmidt-rank}(z)$ denotes the number of nonzero Schmidt coefficients of z , Schatten-norm duality and Cauchy–Schwarz give

$$|\langle z, w \rangle| \leq \|z\|_\pi \|w\|_\varepsilon, \quad \|z\|_\pi \leq \sqrt{\text{Schmidt-rank}(z)} \|z\|. \quad (2.18)$$

The orthogonal projection onto $\text{Sym}^k \mathcal{H}$ is the symmetrization operator

$$\Pi_{\text{sym}}^{(k)} := \frac{1}{k!} \sum_{\pi \in \mathfrak{S}_k} U_\pi. \quad (2.19)$$

For a standard treatment of the symmetric subspace, see Harrow [Har13].

Orthogonal decompositions. An orthogonal decomposition of the one-site space induces a corresponding decomposition of every symmetric tensor power; see, for example, [FH91, Appendix B, Eq. (B.2)]. We record the precise form and its number-operator interpretation used in Section 6.

Fact 2.4 (Symmetric power of an orthogonal sum). *Let $\mathcal{H} = \mathcal{H}_0 \oplus \mathcal{H}_1$ be an orthogonal decomposition. For every integer $n \geq 1$, there is a natural orthogonal decomposition*

$$\text{Sym}^n \mathcal{H} \cong \bigoplus_{m=0}^n \text{Sym}^m(\mathcal{H}_0) \otimes \text{Sym}^{n-m}(\mathcal{H}_1). \quad (2.20)$$

If P is the projection onto $\mathcal{H}_0$, the summand indexed by m is the eigenvalue- m eigenspace of $\sum_{j=1}^n P_j$ on $\text{Sym}^n \mathcal{H}$, where

$$P_j = I^{\otimes(j-1)} \otimes P \otimes I^{\otimes(n-j)}.$$

For $n = 2$, the decomposition is

$$\text{Sym}^2 \mathcal{H} = \text{Sym}^2(\mathcal{H}_0) \oplus (\mathcal{H}_0 \vee \mathcal{H}_1) \oplus \text{Sym}^2(\mathcal{H}_1), \quad (2.21)$$

where $\mathcal{H}_0 \vee \mathcal{H}_1$ is spanned by the vectors

$$p \vee q := \frac{p \otimes q + q \otimes p}{\sqrt{2}}, \quad p \in \mathcal{H}_0, \quad q \in \mathcal{H}_1.$$

Partial inner products. For $k \geq 1$ and $u \in \mathcal{H}$, define the partial-inner-product map

$$C_u : \text{Sym}^k \mathcal{H} \longrightarrow \text{Sym}^{k-1} \mathcal{H}, \quad \Psi \longmapsto (\langle u| \otimes I) \Psi.$$

For $k \geq 2$ and $\phi \in \text{Sym}^2 \mathcal{H}$, similarly define

$$C_\phi : \text{Sym}^k \mathcal{H} \longrightarrow \text{Sym}^{k-2} \mathcal{H}, \quad \Psi \longmapsto (\langle \phi| \otimes I) \Psi.$$

Here I acts on the remaining tensor factors, and the symmetric degree will always be clear from context. When C_u acts on a tensor of symmetric degree zero, we set it equal to zero.

Equivalent calculations in the standard bosonic formalism appear in [LNR15, Sec. 4]; we use only partial-inner-product notation here.

Fact 2.5 (Partial inner products and marginals). *For $\Xi \in \text{Sym}^{k-1} \mathcal{H}$, the adjoint of C_u is*

$$C_u^*(\Xi) = \Pi_{\text{sym}}^{(k)}(u \otimes \Xi) = \frac{1}{k} \sum_{j=1}^k U_{1j}(u \otimes \Xi), \quad (2.22)$$

where U_{1j} swaps the first and j th tensor factors, with $U_{11} = I$. Moreover, if $k \geq 2$, $\Psi \in \text{Sym}^k \mathcal{H}$ is a unit vector, and ρ_1, ρ_2 are the one- and two-site marginals of $|\Psi\rangle\langle\Psi|$, then

$$\langle u|\rho_1|u\rangle = \|C_u(\Psi)\|^2, \quad \langle \phi|\rho_2|\phi\rangle = \|C_\phi(\Psi)\|^2. \quad (2.23)$$

Fact 2.6 (Bounds for symmetric partial inner products). *For every $u \in \mathcal{H}$ and $\Xi \in \text{Sym}^{k-1} \mathcal{H}$,*

$$\|C_u^*(\Xi)\|^2 = \frac{1}{k} \left(\|u\|^2 \|\Xi\|^2 + (k-1) \|C_u(\Xi)\|^2 \right). \quad (2.24)$$

For every orthonormal family (u_i) and every symmetric tensor Ψ ,

$$\sum_i \|C_{u_i}(\Psi)\|^2 \leq \|\Psi\|^2. \quad (2.25)$$

Consequently, if $u_1, \dots, u_\ell$ are orthonormal, then

$$\sum_{i=1}^\ell \|C_{u_i}^*(\Xi)\|^2 \leq \frac{\ell + k - 1}{k} \|\Xi\|^2. \quad (2.26)$$

In particular, the right-hand side is at most $2 \|\Xi\|^2$ when $\ell \leq k + 1$.

Proof. The case $k = 1$ of (2.24) is immediate. For $k \geq 2$, expanding the symmetrizer in (2.22) gives

$$\begin{aligned} \|C_u^*(\Xi)\|^2 &= \frac{1}{k} \sum_{j=1}^k \langle u \otimes \Xi, U_{1j}(u \otimes \Xi) \rangle \\ &= \frac{1}{k} \left(\|u\|^2 \|\Xi\|^2 + (k-1) \|C_u(\Xi)\|^2 \right). \end{aligned}$$

Indeed, symmetry of Ξ makes the terms with $j \geq 2$ equal, and their common value is $\langle \Xi, (|u\rangle\langle u| \otimes I) \Xi \rangle = \|C_u(\Xi)\|^2$. This proves (2.24).

For (2.25), let $P_U := \sum_i |u_i\rangle\langle u_i|$ be the orthogonal projection onto the span of the orthonormal family (u_i) . Then

$$\begin{aligned} \sum_i \|C_{u_i}(\Psi)\|^2 &= \sum_i \langle \Psi, (|u_i\rangle\langle u_i| \otimes I) \Psi \rangle \\ &= \langle \Psi, (P_U \otimes I) \Psi \rangle \\ &\leq \|\Psi\|^2, \end{aligned}$$

because $0 \preceq P_U \preceq I$. Summing (2.24) over $u_1, \dots, u_\ell$ and applying (2.25) to Ξ gives (2.26). $\square$

3 Two-site de Finetti upper bounds

We first prove the bosonic two-site bound by argmax rounding. Symmetric purification gives the exchangeable bound. Rectangular argmax rounding gives the two-sided Bose-symmetric bound, and a second symmetric purification removes the Bose-symmetry assumption.

3.1 Bosonic states: the argmax principle

Theorem 3.1 (Bosonic two-site de Finetti upper bound). *For every $N, d \geq 2$,*

$$D_{\text{H}}^{\text{tr}}\left(\text{BOS}_N^{(2)}(d), \text{SEP}_{\text{H}}^{(2)}(d)\right) \leq \frac{\sqrt{d-1}}{2(N-1)}. \quad (3.1)$$

To prove the theorem, associate with each $M \in \text{Herm}(\text{Sym}^2(\mathcal{H}))$ the quantities a_M , $h(M)$, and the normalized N -site lift $M^{[N]}$ defined by

$$a_M(u) := \langle u^{\otimes 2} | M | u^{\otimes 2} \rangle, \quad (3.2)$$

$$h(M) := \max_{\|u\|=1} a_M(u), \quad (3.3)$$

$$M^{[N]} := \binom{N}{2}^{-1} \sum_{1 \leq i < j \leq N} M_{ij} \quad \text{on } \text{Sym}^N(\mathcal{H}). \quad (3.4)$$

Here M_{ij} acts as M on sites i, j and as the identity elsewhere. Since every two-site marginal of a bosonic state is supported on $\text{Sym}^2(\mathcal{H})$, no extension of M outside that subspace is needed.

The proof has three steps. First, trace-norm duality identifies the de Finetti error with the largest gap $\lambda_{\max}(M^{[N]}) - h(M)$ over bounded Hermitian M . Next, given a top eigenvector ψ of $M^{[N]}$, the argmax principle chooses as the rounding object a unit vector u maximizing $|\langle u^{\otimes N}, \psi \rangle|$. The complex-sphere argmax lemma provides the basic tools for analyzing this rounding object. After taking the partial inner product with u on the other $N-2$ sites, first-order optimality eliminates the mixed component of the resulting two-site vector, while second-order optimality controls its component in $\text{Sym}^2(u^\perp)$. Finally, Takagi decomposition converts this structural information into the integrality-gap bound

$$\lambda_{\max}(M^{[N]}) - h(M) \leq \frac{\sqrt{d-1}}{N-1} \|M\|_\infty.$$

3.1.1 An exact dual problem

Trace-norm duality reduces the approximation problem to a gap detected by a bounded Hermitian witness M . Its maximum over Hartree states is $h(M)$, whereas its maximum over two-site marginals of bosonic N -site states is $\lambda_{\max}(M^{[N]})$. The following lemma makes this reduction exact.

Lemma 3.2 (Exact trace-norm duality). *For every $N, d \geq 2$,*

$$D_{\text{H}}^{\text{tr}}\left(\text{BOS}_N^{(2)}(d), \text{SEP}_{\text{H}}^{(2)}(d)\right) = \frac{1}{2} \sup_{\substack{M \in \text{Herm}(\text{Sym}^2 \mathcal{H}) \\ \|M\|_\infty \leq 1}} \left\{ \lambda_{\max}(M^{[N]}) - h(M) \right\}. \quad (3.5)$$

Proof. The support function of the Hartree mixtures is

$$h_{\text{SEP}_{\text{H}}^{(2)}(d)}(M) = h(M).$$

For the bosonic marginal set, permutation symmetry gives

$$h_{\text{BOS}_N^{(2)}(d)}(M) = \max_{\rho_N \in \text{D}(\text{Sym}^N \mathcal{H})} \text{Tr}(M^{[N]} \rho_N) = \lambda_{\max}(M^{[N]}).$$

By duality (Lemma 2.1), these two identities prove (3.5). $\square$

3.1.2 The argmax rounding object

The tensor-power vectors $\{u^{\otimes N} : u \in \mathcal{H}\}$ span $\text{Sym}^N(\mathcal{H})$; see also [Har13, Theorem 3]. Thus, for every nonzero $\psi \in \text{Sym}^N(\mathcal{H})$, the argmax principle gives a natural rounding object: choose

$$u \in \arg \max_{\|v\|=1} |\langle v^{\otimes N}, \psi \rangle|,$$

and use $u^{\otimes N}$ as the rounding object. Fix a unit vector $\psi \in \text{Sym}^N(\mathcal{H})$, and let $u^{\otimes N}$ be this argmax rounding object. After adjusting the phase of u , set

$$c := \langle u^{\otimes N}, \psi \rangle > 0. \quad (3.6)$$

Taking the partial inner product with u on $N - 2$ sites gives

$$\eta_u := (\langle u|^{\otimes(N-2)} \otimes I_{\mathcal{H}^{\otimes 2}}) \psi \in \text{Sym}^2(\mathcal{H}). \quad (3.7)$$

Write

$$u \vee u^\perp := \text{span}\{u \otimes v + v \otimes u : v \perp u\} \subseteq \text{Sym}^2(\mathcal{H}).$$

The next lemma provides the basic tools for analyzing this rounding object. It is the bosonic counterpart of the reweighted pseudo-covariance lemma underlying our SoS argmax rounding [JWX26]. Here η_u plays the role of the reweighted pseudo-covariance. The mechanism is concrete: first-order optimality eliminates its mixed $u \vee u^\perp$ component, while second-order optimality bounds the remaining $\text{Sym}^2(u^\perp)$ component against every square $v^{\otimes 2}$.

Lemma 3.3 (Complex-sphere argmax estimate). *With the notation above,*

$$\eta_u = c u^{\otimes 2} + \zeta_u, \quad \zeta_u \in \text{Sym}^2(u^\perp), \quad (3.8)$$

and

$$\sup_{\substack{v \perp u \\ \|v\|=1}} |\langle v^{\otimes 2}, \zeta_u \rangle| \leq \frac{c}{N-1}. \quad (3.9)$$

Proof. Fix a unit vector $v \perp u$ and, for $t \in \mathbb{C}$, set

$$u_t := \frac{u + tv}{\sqrt{1 + |t|^2}}, \quad g(t) := \langle u_t^{\otimes N}, \psi \rangle.$$

By the choice of u , the function $|g(t)|^2$ is maximized at $t = 0$. For $0 \leq k \leq N$, write

$$\gamma_k := \langle v^{\otimes k} \otimes u^{\otimes(N-k)}, \psi \rangle.$$

Symmetry of ψ gives the exact expansion

$$g(t) = (1 + |t|^2)^{-N/2} \sum_{k=0}^N \binom{N}{k} \bar{t}^k \gamma_k.$$

In particular, $\gamma_0 = c$. Set

$$\begin{aligned}\alpha &:= \gamma_1 = \langle v \otimes u^{\otimes(N-1)}, \psi \rangle, \\ \beta &:= \gamma_2 = \langle v^{\otimes 2} \otimes u^{\otimes(N-2)}, \psi \rangle.\end{aligned}$$

The exact expansion yields

$$g(t) = (1 + |t|^2)^{-N/2} \left(c + N\bar{t}\alpha + \binom{N}{2}\bar{t}^2\beta + O(|t|^3) \right). \quad (3.10)$$

The normalization factor has no linear term, and hence the first-order term in $|g(t)|^2$ is

$$2Nc\Re(\bar{t}\alpha).$$

To see why this forces $\alpha = 0$, fix any $\xi \in \mathbb{C}$ and restrict to the real one-parameter curve $t = r\xi$, $r \in \mathbb{R}$. Since $t = 0$ is a local maximum of $|g(t)|^2$, the differentiable function $r \mapsto |g(r\xi)|^2$ has derivative zero at $r = 0$. Thus

$$0 = \frac{d}{dr} |g(r\xi)|^2 \Big|_{r=0} = 2Nc\Re(\bar{\xi}\alpha).$$

Taking $\xi = 1$ gives $\Re(\alpha) = 0$, while taking $\xi = i$ gives $\Im(\alpha) = 0$. Therefore $\alpha = 0$. Keeping the normalization factor explicit in the first line and then Taylor-expanding it in the second, we obtain

$$\begin{aligned}|g(t)|^2 &= (1 + |t|^2)^{-N} \left(c^2 + 2\binom{N}{2}c\Re(\bar{t}^2\beta) + O(|t|^3) \right) \\ &= \left(1 - N|t|^2 + O(|t|^4) \right) \left(c^2 + 2\binom{N}{2}c\Re(\bar{t}^2\beta) + O(|t|^3) \right) \\ &= c^2 - Nc^2|t|^2 + 2\binom{N}{2}c\Re(\bar{t}^2\beta) + O(|t|^3).\end{aligned} \quad (3.11)$$

Here we used $(1 + |t|^2)^{-N} = 1 - N|t|^2 + O(|t|^4)$. Its product with the constant term c^2 contributes $-Nc^2|t|^2$. Its products with terms of order $|t|^2$ or higher are $O(|t|^4)$ and are therefore absorbed into the $O(|t|^3)$ remainder.

Now write $t = re^{i\theta}$ and choose θ so that $e^{-2i\theta}\beta = |\beta|$. Along this real one-parameter perturbation,

$$\left| g(re^{i\theta}) \right|^2 = c^2 + \left(-Nc^2 + 2\binom{N}{2}c|\beta| \right) r^2 + O(r^3).$$

Since $r = 0$ is a local maximum, the quadratic coefficient is nonpositive:

$$2\binom{N}{2}c|\beta| \leq Nc^2.$$

Therefore $|\beta| \leq c/(N-1)$.

It remains to translate these variational conditions into a statement about η_u . The orthogonal decomposition

$$\text{Sym}^2(\mathcal{H}) = \mathbb{C}u^{\otimes 2} \oplus (u \vee u^\perp) \oplus \text{Sym}^2(u^\perp)$$

is adapted to this partial inner product. By definition,

$$\langle u^{\otimes 2}, \eta_u \rangle = c, \quad \langle v \otimes u, \eta_u \rangle = \alpha = 0 \quad (v \perp u).$$

Since η_u is symmetric, the second identity eliminates its entire $u \vee u^\perp$ component. Thus $\eta_u = cu^{\otimes 2} + \zeta_u$ for some $\zeta_u \in \text{Sym}^2(u^\perp)$. Finally, for every unit $v \perp u$,

$$\langle v^{\otimes 2}, \zeta_u \rangle = \langle v^{\otimes 2}, \eta_u \rangle = \beta,$$

so the bound on β is exactly (3.9). □

3.1.3 Bounding the integrality gap

We now use the information supplied by the argmax lemma to bound the integrality gap. The remaining ingredient is the Takagi decomposition, the symmetric form of the singular-value, or Schmidt, decomposition. It says that a symmetric two-tensor can be written as

$$\zeta = \sum_i s_i u_i^{\otimes 2},$$

where the u_i are orthonormal and $s_i \geq 0$. This form identifies the relevant tensor norms with matrix Schatten norms.

To record the corresponding matrix statement, let $\mathcal{H}$ be an m -dimensional complex Hilbert space with an orthonormal basis $(e_i)_{i=1}^m$. Every $\zeta \in \text{Sym}^2(\mathcal{H})$ has a unique representation

$$\zeta = \sum_{i,j=1}^m Z_{ij} e_i \otimes e_j, \quad Z^\top = Z. \quad (3.12)$$

We write $Z = \text{Mat}(\zeta)$. Although Z need not be Hermitian, its complex symmetry implies a Takagi decomposition [HJ12, Cor. 4.4.4]

$$Z = U \Sigma U^\top, \quad (3.13)$$

where U is unitary and Σ is diagonal with nonnegative entries. The diagonal entries of Σ are the singular values of Z . A change of orthonormal basis replaces Z by a unitary congruence, so its Schatten norms are basis independent.

The injective and projective tensor norms from Section 2.4 have especially simple descriptions on $\text{Sym}^2(\mathcal{H})$. Namely, Takagi decomposition shows that they may be computed using only symmetric rank-one tensors.

Fact 3.4 (Takagi decomposition and symmetric tensor norms). *For every $\zeta \in \text{Sym}^2(\mathcal{H})$,*

$$\|\zeta\|_\varepsilon = \sup_{\|v\|=1} |\langle v^{\otimes 2}, \zeta \rangle| = \|\text{Mat}(\zeta)\|_\infty, \quad (3.14)$$

$$\|\zeta\|_\pi = \inf \left\{ \sum_k |a_k| : \zeta = \sum_k a_k v_k^{\otimes 2}, \|v_k\| = 1 \right\} = \|\text{Mat}(\zeta)\|_1. \quad (3.15)$$

Proof. Write $Z = U \Sigma U^\top$ as in (3.13). Our inner-product convention gives

$$\langle v^{\otimes 2}, \zeta \rangle = \bar{v}^\top Z v.$$

As v ranges over the unit sphere, so does $y = U^\top v$. Hence

$$|\langle v^{\otimes 2}, \zeta \rangle| = \left| \sum_i \Sigma_{ii} y_i^2 \right| \leq \sum_i \Sigma_{ii} |y_i|^2 \leq \max_i \Sigma_{ii}.$$

Taking y to be a coordinate vector corresponding to the largest diagonal entry shows that the supremum over squares equals $\|Z\|_\infty$. By (2.17), this is precisely $\|\zeta\|_\varepsilon$, proving (3.14).

If u_i are the columns of U , the Takagi decomposition gives

$$\zeta = \sum_i \Sigma_{ii} u_i^{\otimes 2},$$

so the infimum over square decompositions is at most $\sum_i \Sigma_{ii} = \|Z\|_1 = \|\zeta\|_\pi$. Conversely, every square decomposition is admissible in the definition of the ordinary projective norm (2.16), so the infimum is at least $\|\zeta\|_\pi$. This proves (3.15). $\square$

For $M \in \text{Herm}(\text{Sym}^2(\mathbb{C}^d))$, define

$$\tau(M) := \sup_{\|u\|=1} \left\| P_{\text{Sym}^2(u^\perp)} M u^{\otimes 2} \right\|_\pi, \quad (3.16)$$

where the projective norm is taken across the two copies of $u^\perp$.

Proposition 3.5 (Integrality gap from argmax rounding). *For every $N, d \geq 2$ and every $M \in \text{Herm}(\text{Sym}^2(\mathbb{C}^d))$,*

$$\lambda_{\max}(M^{[N]}) - h(M) \leq \frac{\tau(M)}{N-1}. \quad (3.17)$$

Consequently,

$$\lambda_{\max}(M^{[N]}) - h(M) \leq \frac{\sqrt{d-1}}{N-1} \|M\|_\infty. \quad (3.18)$$

Proof. Let $\psi \in \text{Sym}^N(\mathcal{H})$ be a unit top eigenvector of $M^{[N]}$, with eigenvalue λ . Apply the argmax principle to ψ , and let $u^{\otimes N}$ be the resulting rounding object, with c and η_u as in Lemma 3.3. Since both $u^{\otimes N}$ and ψ are symmetric,

$$\lambda c = \langle u^{\otimes N}, M^{[N]} \psi \rangle = \langle u^{\otimes N}, M_{12} \psi \rangle = \langle M u^{\otimes 2}, \eta_u \rangle. \quad (3.19)$$

Decompose orthogonally

$$M u^{\otimes 2} = a_M(u) u^{\otimes 2} + \xi_u + B_u, \quad \xi_u \in u \vee u^\perp, \quad B_u \in \text{Sym}^2(u^\perp). \quad (3.20)$$

Combining (3.19) and (3.8) with the orthogonal decomposition (3.20) gives

$$\begin{aligned} (\lambda - a_M(u))c &= \langle M u^{\otimes 2}, \eta_u \rangle - a_M(u)c \\ &= \langle a_M(u) u^{\otimes 2} + \xi_u + B_u, c u^{\otimes 2} + \zeta_u \rangle - a_M(u)c \\ &= \langle B_u, \zeta_u \rangle. \end{aligned} \quad (3.21)$$

Lemma 3.3 and (3.14) imply

$$\|\zeta_u\|_\varepsilon \leq \frac{c}{N-1}.$$

Taking absolute values in (3.21), using (2.18), and dividing by $c > 0$ gives

$$|\lambda - a_M(u)| \leq \frac{\|B_u\|_\pi}{N-1} \leq \frac{\tau(M)}{N-1}.$$

Because $a_M(u) \leq h(M)$,

$$\lambda - h(M) \leq \lambda - a_M(u) \leq |\lambda - a_M(u)|,$$

which proves (3.17).

Choose an orthonormal basis of $u^\perp$, and let $\text{Mat}_u(B_u)$ be the corresponding complex symmetric matrix from (3.12). Its Schatten norms are independent of this choice. By (3.15), and because $\text{Mat}_u(B_u)$ has size $(d-1) \times (d-1)$,

$$\begin{aligned} \|B_u\|_\pi &= \|\text{Mat}_u(B_u)\|_1 \\ &\leq \sqrt{d-1} \|\text{Mat}_u(B_u)\|_2 \\ &= \sqrt{d-1} \|B_u\| \\ &\leq \sqrt{d-1} \|M\|_\infty. \end{aligned}$$

Taking the supremum over u gives (3.18). $\square$

Proof of Theorem 3.1. Insert (3.18) into (3.5). This gives

$$D_{\mathcal{H}}^{\text{tr}}\left(\text{BOS}_N^{(2)}(d), \text{SEP}_{\mathcal{H}}^{(2)}(d)\right) \leq \frac{\sqrt{d-1}}{2(N-1)}. \quad \square$$

Remark 3.6. The dimension factor enters only through the inequality $\|\xi\|_{\pi} \leq \sqrt{d-1} \|\xi\|$ for $\xi \in \text{Sym}^2(u^{\perp})$. Thus the stronger estimate (3.17) is dimension independent for classes of Hermitian operators satisfying $\tau(M) = O(\|M\|_{\infty})$.

3.1.4 Exchangeable states via symmetric purification

Corollary 3.7 (Exchangeable two-site de Finetti upper bound). *For every $N, d \geq 2$,*

$$D_{\mathcal{H}}^{\text{tr}}\left(\text{EXCH}_N^{(2)}(d), \text{SEP}_{\mathcal{P}}^{(2)}(d)\right) \leq \frac{\sqrt{d^2-1}}{2(N-1)}. \quad (3.22)$$

The corollary follows from Theorem 3.1 through the standard symmetric-purification lemma of Christandl, König, Mitchison, and Renner [CKMR07, Lemma II.5].

Lemma 3.8 (Symmetric purification). *Let σ_N be exchangeable on $\mathcal{H}^{\otimes N}$, where $\dim \mathcal{H} = d$. There is a Hilbert space $\mathcal{H}_E \cong \mathcal{H}$ and a unit vector*

$$\psi_N \in \text{Sym}^N(\mathcal{H}_E \otimes \mathcal{H})$$

such that

$$\text{Tr}_{\mathcal{H}_E^{\otimes N}} |\psi_N\rangle\langle\psi_N| = \sigma_N.$$

Proof of Corollary 3.7. Apply Lemma 3.8 and set $\tilde{\rho}_N = |\psi_N\rangle\langle\psi_N|$. This is a bosonic state with one-site dimension d^2 . Theorem 3.1 gives a probability measure μ on unit vectors $\phi \in \mathcal{H}_E \otimes \mathcal{H}$ such that

$$D_{\text{tr}}\left(\tilde{\rho}_N^{(2)}, \int |\phi\rangle\langle\phi|^{\otimes 2} d\mu(\phi)\right) \leq \frac{\sqrt{d^2-1}}{2(N-1)}.$$

For each ϕ , define $\tau_{\phi} = \text{Tr}_{\mathcal{H}_E} |\phi\rangle\langle\phi| \in \text{D}(\mathcal{H})$. Sitewise partial trace maps $|\phi\rangle\langle\phi|^{\otimes 2}$ to $\tau_{\phi}^{\otimes 2}$, so contractivity gives

$$D_{\text{tr}}\left(\sigma_N^{(2)}, \int \tau_{\phi}^{\otimes 2} d\mu(\phi)\right) \leq \frac{\sqrt{d^2-1}}{2(N-1)}.$$

The approximating state belongs to $\text{SEP}_{\mathcal{P}}^{(2)}(d)$, which proves (3.22). $\square$

3.2 Two-sided extendibility

Theorem 3.9 (De Finetti upper bound for two-sided Bose-symmetric extendible states). *For every $n, m, d_A, d_B \geq 1$,*

$$D_{\mathcal{H}}^{\text{tr}}\left(\text{BEXT}_{n,m}(d_A, d_B), \text{SEP}(d_A, d_B)\right) \leq \frac{1}{2} \sqrt{\frac{\min\{d_A, d_B\} - 1}{nm}}. \quad (3.23)$$

The proof follows the same pattern as the bosonic argument above: rectangular argmax rounding bounds the integrality gap, and the same estimate will also be used for higher marginals in Section 4. Trace-distance duality then gives the theorem, while a two-block symmetric purification removes the Bose-symmetry assumption.

3.2.1 Rectangular argmax rounding

Let $\mathcal{H}_A$ and $\mathcal{H}_B$ be finite-dimensional Hilbert spaces, and write $d_A := \dim \mathcal{H}_A$ and $d_B := \dim \mathcal{H}_B$. For $M \in \text{Herm}(\mathcal{H}_A \otimes \mathcal{H}_B)$, define

$$a_M(x, y) := \langle x \otimes y | M | x \otimes y \rangle, \quad (3.24)$$

$$h_{\text{SEP}}(M) := \max_{\substack{\|x\|=1 \\ \|y\|=1}} a_M(x, y). \quad (3.25)$$

Since $\text{SEP}(d_A, d_B)$ is the convex hull of pure product states,

$$h_{\text{SEP}}(M) = \max_{\sigma \in \text{SEP}(d_A, d_B)} \text{Tr}(M\sigma). \quad (3.26)$$

For $n, m \geq 1$, its normalized rectangular lift is

$$M^{[n,m]} := \frac{1}{nm} \sum_{i=1}^n \sum_{j=1}^m M_{A_i B_j} \quad \text{on } \text{Sym}^n(\mathcal{H}_A) \otimes \text{Sym}^m(\mathcal{H}_B). \quad (3.27)$$

Here $M_{A_i B_j}$ acts as M on $A_i B_j$ and as the identity on all other registers. If $\omega_{n,m} \in \text{D}(\text{Sym}^n(\mathcal{H}_A) \otimes \text{Sym}^m(\mathcal{H}_B))$, permutation symmetry within the A - and B -blocks gives

$$\text{Tr}(M^{[n,m]} \omega_{n,m}) = \text{Tr}(M(\omega_{n,m})_{A_1 B_1}). \quad (3.28)$$

Given a unit top eigenvector of $M^{[n,m]}$, rectangular argmax rounding selects a product vector from the corresponding overlap maximization. The following estimate is the direct rectangular analogue of Proposition 3.5. Its proof is deferred to Appendix A.

Proposition 3.10 (Rectangular argmax estimate). *Let $M \in \text{Herm}(\mathcal{H}_A \otimes \mathcal{H}_B)$ and $n, m \geq 1$, and let ψ be a unit top eigenvector of $M^{[n,m]}$. Choose the argmax rounding object $x \otimes y$ by*

$$(x, y) \in \arg \max_{\substack{v \in \mathcal{H}_A, w \in \mathcal{H}_B \\ \|v\|=\|w\|=1}} |\langle v^{\otimes n} \otimes w^{\otimes m}, \psi \rangle|. \quad (3.29)$$

Then

$$\lambda_{\max}(M^{[n,m]}) - a_M(x, y) \leq \frac{1}{\sqrt{nm}} \|(P_{x^\perp} \otimes P_{y^\perp}) M(x \otimes y)\|_\pi. \quad (3.30)$$

As in the symmetric case, the pointwise estimate immediately gives an integrality-gap bound. Define

$$\tau(M) := \sup_{\substack{\|x\|=1 \\ \|y\|=1}} \|(P_{x^\perp} \otimes P_{y^\perp}) M(x \otimes y)\|_\pi. \quad (3.31)$$

Corollary 3.11 (Integrality gap from rectangular argmax rounding). *For every $M \in \text{Herm}(\mathcal{H}_A \otimes \mathcal{H}_B)$ and $n, m \geq 1$,*

$$\lambda_{\max}(M^{[n,m]}) - h_{\text{SEP}}(M) \leq \frac{\tau(M)}{\sqrt{nm}}, \quad (3.32)$$

and hence

$$\lambda_{\max}(M^{[n,m]}) - h_{\text{SEP}}(M) \leq \sqrt{\frac{\min\{\dim \mathcal{H}_A, \dim \mathcal{H}_B\} - 1}{nm}} \|M\|_\infty. \quad (3.33)$$

Proof. Apply Proposition 3.10 and use the definitions of $h_{\text{SEP}}(M)$ and $\tau(M)$. For the second bound, the projected vector $(P_{x^\perp} \otimes P_{y^\perp})M(x \otimes y)$ has Schmidt rank at most $\min\{\dim \mathcal{H}_A, \dim \mathcal{H}_B\} - 1$. Therefore, (2.18) gives

$$\begin{aligned} \|(P_{x^\perp} \otimes P_{y^\perp})M(x \otimes y)\|_\pi &\leq \sqrt{\min\{\dim \mathcal{H}_A, \dim \mathcal{H}_B\} - 1} \|(P_{x^\perp} \otimes P_{y^\perp})M(x \otimes y)\| \\ &\leq \sqrt{\min\{\dim \mathcal{H}_A, \dim \mathcal{H}_B\} - 1} \|M\|_\infty. \end{aligned}$$

□

3.2.2 Two-sided de Finetti bounds

The rectangular gap immediately gives a de Finetti bound for arbitrary extension levels on the two sides.

Proof of Theorem 3.9. Fix $\rho \in \text{BEXT}_{n,m}(d_A, d_B)$ and choose an extension $\tilde{\rho}$ from (2.3). By duality (Lemma 2.1),

$$D_{\text{tr}}(\rho, \text{SEP}(d_A, d_B)) = \frac{1}{2} \max_{\substack{M \in \text{Herm}(\mathcal{H}_A \otimes \mathcal{H}_B) \\ \|M\|_\infty \leq 1}} \{\text{Tr}(M\rho) - h_{\text{SEP}}(M)\}. \quad (3.34)$$

For every Hermitian M , the rectangular-lift identity (3.28) and Corollary 3.11 give

$$\begin{aligned} \text{Tr}(M\rho) - h_{\text{SEP}}(M) &= \text{Tr}(M^{[n,m]}\tilde{\rho}) - h_{\text{SEP}}(M) \\ &\leq \lambda_{\max}(M^{[n,m]}) - h_{\text{SEP}}(M) \\ &\leq \sqrt{\frac{\min\{d_A, d_B\} - 1}{nm}} \|M\|_\infty. \end{aligned} \quad (3.35)$$

Inserting this estimate into (3.34) and taking the supremum over ρ proves (3.23). □

We now remove the Bose-symmetry assumption by symmetric purification.

Corollary 3.12 (De Finetti upper bound for two-sided extendible states). *For every $n, m, d_A, d_B \geq 1$,*

$$D_{\text{H}}^{\text{tr}}(\text{EXT}_{n,m}(d_A, d_B), \text{SEP}(d_A, d_B)) \leq \frac{1}{2} \sqrt{\frac{\min\{d_A^2, d_B^2\} - 1}{nm}}. \quad (3.36)$$

Proof of Corollary 3.12. Let $\tilde{\rho}$ be an (n, m) -extension of $\rho \in \text{EXT}_{n,m}(d_A, d_B)$. Choose copies $\mathcal{H}_{E_A} \cong \mathcal{H}_A$ and $\mathcal{H}_{E_B} \cong \mathcal{H}_B$, with copied orthonormal bases, and let $|\Omega\rangle$ be the unnormalized maximally entangled vector between $\mathcal{H}_{E_A}^{\otimes n} \otimes \mathcal{H}_{E_B}^{\otimes m}$ and $\mathcal{H}_A^{\otimes n} \otimes \mathcal{H}_B^{\otimes m}$. The canonical purification

$$|\Psi\rangle := (I \otimes \sqrt{\tilde{\rho}})|\Omega\rangle$$

satisfies $\text{Tr}_{E_A^{\otimes n} E_B^{\otimes m}} |\Psi\rangle\langle\Psi| = \tilde{\rho}$.

Regroup the registers into the composite sites

$$\mathcal{X} := \mathcal{H}_{E_A} \otimes \mathcal{H}_A, \quad \mathcal{Y} := \mathcal{H}_{E_B} \otimes \mathcal{H}_B,$$

where the composite registers $X_i = E_{A,i}A_i$ and $Y_j = E_{B,j}B_j$ carry $\mathcal{X}$ and $\mathcal{Y}$, respectively. Fix $\pi \in \mathfrak{S}_n$ and regard the B - and E_B -registers as spectators. Invariance of $\tilde{\rho}$ under the A -permutation

U_π^A implies that U_π^A commutes with $\sqrt{\rho}$. The same permutation on the copied E_A -registers and on the original A -registers fixes $|\Omega\rangle$, so the canonical-purification identity gives

$$(U_\pi^{\mathcal{X}} \otimes I_{\mathcal{Y}^{\otimes m}})|\Psi\rangle = |\Psi\rangle,$$

where $U_\pi^{\mathcal{X}}$ permutes the n composite X -sites. The same argument with the two blocks exchanged gives

$$(I_{\mathcal{X}^{\otimes n}} \otimes U_\tau^{\mathcal{Y}})|\Psi\rangle = |\Psi\rangle \quad (\tau \in \mathfrak{S}_m).$$

Thus the symmetrizing projectors for the X - and Y -sites both fix $|\Psi\rangle$, and hence

$$|\Psi\rangle \in \text{Sym}^n(\mathcal{X}) \otimes \text{Sym}^m(\mathcal{Y}).$$

The $X_1 Y_1$ marginal therefore belongs to $\text{BEXT}_{n,m}(d_A^2, d_B^2)$ and reduces to ρ after tracing out $E_{A,1} E_{B,1}$. Applying Theorem 3.9 in local dimensions d_A^2 and d_B^2 , then using contractivity of trace distance and preservation of separability under partial trace, gives (3.36). $\square$

4 t -site de Finetti upper bounds

We now extend the two-site de Finetti bounds to t -site marginals.

Theorem 4.1 (t -site de Finetti upper bounds). *For every $d \geq 2$ and $3 \leq t \leq N$,*

$$D_{\text{H}}^{\text{tr}}\left(\text{BOS}_N^{(t)}(d), \text{SEP}_H^{(t)}(d)\right) \leq \frac{4t\sqrt{d-1}}{N}, \quad (4.1)$$

$$D_{\text{H}}^{\text{tr}}\left(\text{EXCH}_N^{(t)}(d), \text{SEP}_P^{(t)}(d)\right) \leq \frac{4t\sqrt{d^2-1}}{N}. \quad (4.2)$$

We first prove the bosonic bound; symmetric purification then gives the exchangeable bound. For $M \in \text{Herm}(\text{Sym}^t \mathcal{H})$, define

$$h_t(M) := \max_{\|u\|=1} \langle u^{\otimes t} | M | u^{\otimes t} \rangle.$$

This is the maximum expectation of M over $\text{SEP}_H^{(t)}(d)$.

Trace-distance duality reduces the problem to bounding $\text{Tr}(M \rho_N^{(t)}) - h_t(M)$. We distinguish one tensor factor from the other $t-1$ factors and view $\rho_N^{(t)}$ as a bipartite state on

$$\mathcal{H} \otimes \text{Sym}^{t-1}(\mathcal{H}).$$

Grouping the N sites into n individual sites and m blocks of $t-1$ sites makes this state Bose-symmetric (n, m) -extendible. This two-sided extendibility lets rectangular argmax rounding produce a product vector $x \otimes y$, enforcing unentanglement across the $1 : (t-1)$ cut. The challenge is to propagate this unentanglement to the remaining cuts: y need not equal $x^{\otimes(t-1)}$, so $x \otimes y$ need not be a Hartree vector.

The symmetry test is what propagates this unentanglement. We incorporate it into the objective before performing argmax rounding. The bosonic marginal passes the test with certainty, whereas the rounded product loses value according to its symmetry-test rejection probability. The same rejection probability controls the cost of replacing $x \otimes y$ by $x^{\otimes t}$. Choosing the weight of the symmetry test to balance this cost against the rectangular argmax rounding error gives the claimed rate. The rest of the section makes these steps precise.

We begin with the dual formulation of the bosonic bound.

Lemma 4.2 (Exact trace-distance duality). *For every $d \geq 2$, $t > 2$, and every $\gamma \in \mathcal{D}(\text{Sym}^t \mathcal{H})$,*

$$D_{\text{tr}}(\gamma, \text{SEP}_{\mathcal{H}}^{(t)}(d)) = \sup_{\substack{M \in \text{Herm}(\text{Sym}^t \mathcal{H}) \\ 0 \leq M \leq I}} \{\text{Tr}(M\gamma) - h_t(M)\}. \quad (4.3)$$

Proof. The support function of $\text{SEP}_{\mathcal{H}}^{(t)}(d)$ is $h_t(M)$. By duality (Lemma 2.1), this gives (4.3). Since all states involved are supported on $\text{Sym}^t \mathcal{H}$, the witness may be restricted to that subspace. $\square$

4.1 Two-sided extendibility across the $1 : (t-1)$ cut

Every Hartree vector factors across this cut as

$$u^{\otimes t} = u \otimes u^{\otimes (t-1)}.$$

The first step drops the requirement $y = x^{\otimes (t-1)}$ and allows arbitrary products $x \otimes y$. We choose the $1 : (t-1)$ cut because its smaller side has dimension d , which is the favorable dimension in the rectangular argmax estimate, and because Bose symmetry supplies exactly the corresponding two-sided extension.

Fix $t > 2$ and set

$$\mathcal{H}_A := \mathcal{H}, \quad \mathcal{H}_B := \text{Sym}^{t-1}(\mathcal{H}).$$

We use the natural inclusion $\text{Sym}^t(\mathcal{H}) \subseteq \mathcal{H}_A \otimes \mathcal{H}_B$ obtained by distinguishing the first tensor factor.

Fact 4.3 (Two-sided extendibility of the marginal). *Let $n, m \geq 1$ satisfy $n + (t-1)m \leq N$. For every $\rho_N \in \mathcal{D}(\text{Sym}^N \mathcal{H})$, the marginal $\rho_N^{(t)}$, viewed as a state on $\mathcal{H}_A \otimes \mathcal{H}_B$, is Bose-symmetric (n, m) -extendible.*

Proof. Trace out $N - n - (t-1)m$ sites and rebracket the remaining tensor factors as n individual sites and m blocks of $t-1$ sites:

$$\begin{array}{c} \bullet \quad \cdots \quad \bullet \\ A_1 \quad \quad A_n \end{array} \quad \begin{array}{c} t-1 \text{ sites} \\ [\bullet \quad \cdots \quad \bullet] \\ B_1 \end{array} \quad \cdots \quad \begin{array}{c} t-1 \text{ sites} \\ [\bullet \quad \cdots \quad \bullet] \\ B_m \end{array}.$$

Let $\tilde{\rho}_{n,m}$ be the resulting marginal. Because ρ_N is supported on $\text{Sym}^N(\mathcal{H})$, this marginal is supported on $\text{Sym}^{n+(t-1)m}(\mathcal{H})$. It is therefore supported on $\mathcal{H}_B$ within each B_j , and every permutation of the n A -sites or of the m B -blocks acts as the identity on its support. Consequently, $\tilde{\rho}_{n,m} \in \mathcal{D}(\text{Sym}^n(\mathcal{H}_A) \otimes \text{Sym}^m(\mathcal{H}_B))$. Moreover, every $A_i B_j$ consists of t distinct original sites, so $(\tilde{\rho}_{n,m})_{A_i B_j} = \rho_N^{(t)}$ for all i, j . $\square$

4.2 From product vectors to tensor powers

Rectangular argmax rounding produces a product vector $x \otimes y$ across the chosen $1 : (t-1)$ cut, but it need not satisfy $y = x^{\otimes (t-1)}$. The symmetry test propagates this unentanglement to the remaining cuts.

Let Π_t be the orthogonal projector from $\mathcal{H}_A \otimes \mathcal{H}_B = \mathcal{H} \otimes \text{Sym}^{t-1}(\mathcal{H})$ onto $\text{Sym}^t(\mathcal{H})$. It tests whether the product vector lies in the symmetric subspace. The following lemma quantifies the test: a high probability of passing forces y to be close to $x^{\otimes (t-1)}$. It also bounds the component of $\Pi_t(x \otimes y)$ and $M(x \otimes y)$ orthogonal to both rounded factors.

Lemma 4.4 (Approximate symmetry of a product vector). *Let $x \in \mathcal{H}$ and $y \in \text{Sym}^{t-1}(\mathcal{H})$ be unit vectors, and define the symmetry-test rejection probability*

$$\delta_{\text{sym}}(x, y) := 1 - \langle x \otimes y | \Pi_t | x \otimes y \rangle.$$

Set $P_{\perp} := P_{x^{\perp}} \otimes P_{y^{\perp}}$. Then

$$1 - \left| \langle x^{\otimes(t-1)}, y \rangle \right|^2 \leq t \delta_{\text{sym}}(x, y). \quad (4.4)$$

If $0 \preceq M \preceq I$ acts on $\text{Sym}^t \mathcal{H}$ and is extended by zero to $\mathcal{H}_A \otimes \mathcal{H}_B$, then

$$\langle x \otimes y | M | x \otimes y \rangle \leq h_t(M) + \sqrt{t \delta_{\text{sym}}(x, y)}. \quad (4.5)$$

Moreover,

$$\|P_{\perp} \Pi_t(x \otimes y)\|_{\pi} \leq \sqrt{d-1} \sqrt{\delta_{\text{sym}}(x, y)}, \quad (4.6)$$

$$\|P_{\perp} M(x \otimes y)\|_{\pi} \leq \sqrt{d-1}. \quad (4.7)$$

Proof. Because y is symmetric in its $t-1$ factors, averaging all permutations reduces to averaging over the t possible locations of the distinguished first factor:

$$\Pi_t = \frac{1}{t} \sum_{j=1}^t U_{1j}, \quad U_{11} = I.$$

Let γ_y be the one-site marginal of $|y\rangle\langle y|$. Since y is symmetric,

$$\langle x \otimes y | U_{1j} | x \otimes y \rangle = \langle x | \gamma_y | x \rangle \quad (2 \leq j \leq t).$$

It follows that

$$\langle x \otimes y | \Pi_t | x \otimes y \rangle = \frac{1 + (t-1) \langle x | \gamma_y | x \rangle}{t}. \quad (4.8)$$

Put $Q_x = I - |x\rangle\langle x|$. On $\text{Sym}^{t-1}(\mathcal{H})$, the operator

$$\mathcal{N}_x := \sum_{j=1}^{t-1} (Q_x)_j$$

counts how many of the $t-1$ factors lie in $x^{\perp}$. Its kernel is $\mathbb{C}x^{\otimes(t-1)}$, and all other eigenvalues are at least one. Therefore

$$I - |x^{\otimes(t-1)}\rangle\langle x^{\otimes(t-1)}| \preceq \mathcal{N}_x.$$

Taking the expectation in y and using (4.8),

$$\begin{aligned} 1 - \left| \langle x^{\otimes(t-1)}, y \rangle \right|^2 &\leq (t-1)(1 - \langle x | \gamma_y | x \rangle) \\ &= t \delta_{\text{sym}}(x, y), \end{aligned}$$

which proves (4.4).

The trace distance between the pure states $|x \otimes y\rangle\langle x \otimes y|$ and $|x\rangle\langle x|^{\otimes t}$ is

$$\sqrt{1 - \left| \langle x^{\otimes(t-1)}, y \rangle \right|^2} \leq \sqrt{t \delta_{\text{sym}}(x, y)}.$$

Since $0 \preceq M \preceq I$ and $\langle x^{\otimes t} | M | x^{\otimes t} \rangle \leq h_t(M)$, trace-distance duality gives (4.5).

For the bounds on the projected components, first note that orthogonality of Π_t gives

$$\|(I - \Pi_t)(x \otimes y)\|^2 = \delta_{\text{sym}}(x, y).$$

Since $P_\perp(x \otimes y) = 0$,

$$P_\perp \Pi_t(x \otimes y) = -P_\perp(I - \Pi_t)(x \otimes y).$$

Thus its Hilbert-space norm is at most $\sqrt{\delta_{\text{sym}}(x, y)}$. Every vector in $x^\perp \otimes y^\perp$ has Schmidt rank at most $d - 1$, so (2.18) proves (4.6). The same rank bound and

$$\|P_\perp M(x \otimes y)\| \leq \|M\|_\infty \leq 1$$

prove (4.7). $\square$

A related, quantitatively weaker estimate appears in the one-versus-many slicing de Finetti theorem of Jeronimo and Wu [JW24, Theorem 23]. After normalizing $\Pi_t(x \otimes y)$, that theorem gives squared overlap $1 - O(t^3 \delta_{\text{sym}}(x, y))$ with some tensor power $v^{\otimes t}$. Estimate (4.4) improves the cubic loss to a linear one and identifies the direction as x . This sharper dependence is essential for the linear dependence on t in Theorem 4.1.

4.3 Bounding the integrality gap

We now combine rectangular argmax rounding with the symmetry estimate. For $M_\theta = (1 - \theta)\Pi_t + \theta M$, the grouped bosonic marginal is feasible for the rectangular lift and retains the score $1 - \theta + \theta \text{Tr}(M \rho_N^{(t)})$. Rectangular argmax rounding produces $x \otimes y$, and Lemma 4.4 controls its replacement by $x^{\otimes t}$. Optimizing over θ balances the two estimates and gives the integrality-gap bound required by Lemma 4.2.

Proposition 4.5 (Rounding across the $1 : (t - 1)$ cut). *Let $d \geq 2$, $3 \leq t \leq N$, and let $n, m \geq 2$ satisfy $n + (t - 1)m \leq N$. For every $\rho_N \in \text{D}(\text{Sym}^N \mathcal{H})$ and every $0 \preceq M \preceq I_{\text{Sym}^t \mathcal{H}}$,*

$$\text{Tr}(M \rho_N^{(t)}) - h_t(M) \leq (\sqrt{t} + 1) \sqrt{\frac{d - 1}{nm}}.$$

Proof. Let $\tilde{\rho}_{n,m}$ be the grouped state from Fact 4.3. For $0 < \theta < 1$, extend M by zero outside $\text{Sym}^t \mathcal{H}$ and consider the mixed test

$$M_\theta := (1 - \theta)\Pi_t + \theta M \quad \text{on } \mathcal{H}_A \otimes \mathcal{H}_B.$$

Because $\rho_N^{(t)}$ is supported on $\text{Sym}^t \mathcal{H}$, the projector Π_t acts as the identity on its support. Therefore

$$\begin{aligned} \text{Tr}(M_\theta \rho_N^{(t)}) &= (1 - \theta) \text{Tr}(\Pi_t \rho_N^{(t)}) + \theta \text{Tr}(M \rho_N^{(t)}) \\ &= 1 - \theta + \theta \text{Tr}(M \rho_N^{(t)}). \end{aligned}$$

By Fact 4.3, $(\tilde{\rho}_{n,m})_{A_1 B_1} = \rho_N^{(t)}$. Hence the rectangular-lift identity (3.28) gives

$$\begin{aligned} 1 - \theta + \theta \text{Tr}(M \rho_N^{(t)}) &= \text{Tr}(M_\theta \rho_N^{(t)}) = \text{Tr}(M_\theta^{[n,m]} \tilde{\rho}_{n,m}) \\ &\leq \lambda_{\max}(M_\theta^{[n,m]}). \end{aligned} \tag{4.9}$$

Rectangular argmax rounding. Proposition 3.10, applied to M_θ , gives the first-stage product vector $x \otimes y$. Set $P_\perp := P_{x^\perp} \otimes P_{y^\perp}$. Moving its product-value term to the right-hand side gives

$$\lambda_{\max}(M_\theta^{[n,m]}) \leq \langle x \otimes y | M_\theta | x \otimes y \rangle + \frac{1}{\sqrt{nm}} \|P_\perp M_\theta(x \otimes y)\|_\pi. \quad (4.10)$$

We use the symmetry test to round this product vector to $x^{\otimes t}$ and estimate the two terms on the right using the same rejection probability. Write $\delta = \delta_{\text{sym}}(x, y)$ for this probability and

$$\beta := \sqrt{\frac{d-1}{nm}}.$$

The product value of M_θ is

$$\begin{aligned} \langle x \otimes y | M_\theta | x \otimes y \rangle &= (1-\theta)(1-\delta) + \theta \langle x \otimes y | M | x \otimes y \rangle \\ &\leq (1-\theta)(1-\delta) + \theta(h_t(M) + \sqrt{t\delta}) \\ &= 1 - \theta + \theta h_t(M) - (1-\theta)\delta + \theta\sqrt{t\delta}, \end{aligned} \quad (4.11)$$

where the inequality is (4.5). The triangle inequality and the two bounds on projected components in Lemma 4.4 give

$$\begin{aligned} \frac{1}{\sqrt{nm}} \|P_\perp M_\theta(x \otimes y)\|_\pi &\leq \frac{1}{\sqrt{nm}} ((1-\theta) \|P_\perp \Pi_t(x \otimes y)\|_\pi + \theta \|P_\perp M(x \otimes y)\|_\pi) \\ &\leq \sqrt{\frac{d-1}{nm}} ((1-\theta)\sqrt{\delta} + \theta) \\ &= \beta((1-\theta)\sqrt{\delta} + \theta). \end{aligned} \quad (4.12)$$

Both estimates concern the pair selected in (4.10), and hence the same quantity δ . Combining (4.9), (4.10), (4.11), and (4.12), we obtain

$$1 - \theta + \theta \text{Tr}(M\rho_N^{(t)}) \leq 1 - \theta + \theta h_t(M) - (1-\theta)\delta + \theta\sqrt{t\delta} + \beta(1-\theta)\sqrt{\delta} + \beta\theta.$$

Subtracting $1 - \theta + \theta h_t(M)$ from both sides gives

$$\theta \left(\text{Tr}(M\rho_N^{(t)}) - h_t(M) \right) \leq -(1-\theta)\delta + (\theta\sqrt{t} + \beta(1-\theta))\sqrt{\delta} + \beta\theta. \quad (4.13)$$

Keeping the negative symmetry-test term and the two error terms coupled through the same δ is what allows the optimization to retain linear dependence on t .

Optimizing the error bound. Since $0 \leq \delta \leq 1$, the variable $s := \sqrt{\delta}$ lies in $[0, 1]$. In terms of s , (4.13) reads

$$\theta \left(\text{Tr}(M\rho_N^{(t)}) - h_t(M) \right) \leq -(1-\theta)s^2 + (\theta\sqrt{t} + \beta(1-\theta))s + \beta\theta.$$

For every real s , completing the square gives

$$-(1-\theta)s^2 + (\theta\sqrt{t} + \beta(1-\theta))s \leq \frac{(\theta\sqrt{t} + \beta(1-\theta))^2}{4(1-\theta)}.$$

After dividing by $\theta > 0$ and expanding the square, we obtain

$$\text{Tr}(M\rho_N^{(t)}) - h_t(M) \leq \frac{(\theta\sqrt{t} + \beta(1-\theta))^2}{4\theta(1-\theta)} + \beta$$

$$= \frac{\theta t}{4(1-\theta)} + \frac{\beta\sqrt{t}}{2} + \frac{\beta^2(1-\theta)}{4\theta} + \beta. \quad (4.14)$$

Set

$$\lambda := \frac{\theta}{1-\theta}.$$

As θ ranges over $(0, 1)$, the parameter λ ranges over $(0, \infty)$. Since $\beta > 0$, we may choose

$$\lambda = \frac{\beta}{\sqrt{t}}, \quad \text{equivalently} \quad \theta = \frac{\beta}{\beta + \sqrt{t}}.$$

With this choice, (4.14) becomes

$$\begin{aligned} \text{Tr}(M\rho_N^{(t)}) - h_t(M) &\leq \frac{t\lambda}{4} + \frac{\beta\sqrt{t}}{2} + \frac{\beta^2}{4\lambda} + \beta \\ &= \beta(\sqrt{t} + 1). \end{aligned}$$

Recalling the definition of β , we conclude that

$$\text{Tr}(M\rho_N^{(t)}) - h_t(M) \leq (\sqrt{t} + 1)\sqrt{\frac{d-1}{nm}}.$$

This proves the proposition. $\square$

4.4 Finishing the proof

It remains to choose the two extension levels so that nm is large.

Proof of Theorem 4.1. Both trace distances are at most one. If $N < 4(t-1)$, then both right-hand sides in (4.1) and (4.2) exceed one, so the result is immediate. We may therefore suppose that $N \geq 4(t-1)$.

For the bosonic bound, fix $\rho_N \in \text{D}(\text{Sym}^N \mathcal{H})$ and an operator M on $\text{Sym}^t \mathcal{H}$ satisfying $0 \preceq M \preceq I$. Define

$$m := \left\lfloor \frac{N}{2(t-1)} \right\rfloor, \quad n := N - (t-1)m.$$

Then $n, m \geq 2$ and

$$n \geq \frac{N}{2}, \quad m \geq \frac{N}{4(t-1)}, \quad nm \geq \frac{N^2}{8(t-1)}. \quad (4.15)$$

Proposition 4.5 and (4.15) give

$$\begin{aligned} \text{Tr}(M\rho_N^{(t)}) - h_t(M) &\leq (\sqrt{t} + 1)\sqrt{\frac{d-1}{nm}} \\ &\leq 2\sqrt{2} \frac{(\sqrt{t} + 1)\sqrt{(t-1)(d-1)}}{N} \\ &< \frac{4t\sqrt{d-1}}{N}. \end{aligned}$$

Here the last inequality follows from $(\sqrt{t} + 1)^2 \leq 2(t+1)$, which gives $\sqrt{t-1}(\sqrt{t} + 1) < \sqrt{2}t$. Taking the supremum over the witnesses $0 \preceq M \preceq I$ and using Lemma 4.2 gives the same bound for $D_{\text{tr}}(\rho_N^{(t)}, \text{SEP}_H^{(t)}(d))$. Taking the supremum over the bosonic states ρ_N therefore proves

$$D_H^{\text{tr}}(\text{BOS}_N^{(t)}(d), \text{SEP}_H^{(t)}(d)) \leq \frac{4t\sqrt{d-1}}{N}.$$

For the exchangeable bound, let σ_N be exchangeable on $(\mathbb{C}^d)^{\otimes N}$. By Lemma 3.8, it has a symmetric purification $\tilde{\rho}_N \in D(\text{Sym}^N(\mathbb{C}^d \otimes \mathbb{C}^d))$. Applying the bosonic bound just proved, now in one-site dimension d^2 , gives a probability measure μ on unit vectors $\phi \in \mathbb{C}^d \otimes \mathbb{C}^d$ such that

$$D_{\text{tr}}\left(\tilde{\rho}_N^{(t)}, \int |\phi\rangle\langle\phi|^{\otimes t} d\mu(\phi)\right) \leq \frac{4t\sqrt{d^2-1}}{N}.$$

For $\tau_\phi := \text{Tr}_E |\phi\rangle\langle\phi|$, sitewise partial trace maps $|\phi\rangle\langle\phi|^{\otimes t}$ to $\tau_\phi^{\otimes t}$ and $\tilde{\rho}_N^{(t)}$ to $\sigma_N^{(t)}$. Hence contractivity of trace distance gives

$$D_{\text{tr}}\left(\sigma_N^{(t)}, \int \tau_\phi^{\otimes t} d\mu(\phi)\right) \leq \frac{4t\sqrt{d^2-1}}{N}.$$

The integral belongs to $\text{SEP}_P^{(t)}(d)$. Taking the supremum over σ_N proves (4.2) and completes the proof. $\square$

5 Applications of the optimal de Finetti theorem

The optimal trace-norm theorem has three applications whose resource bounds are subexponential in the local dimension at fixed accuracy. The disentangler construction uses the statement that every two-sided Bose-symmetric extendible state is close to separable. The same two-sided extension relaxation approximates Best Separable State, and its error guarantee also gives an algorithm for trace-norm separability testing.

5.1 Subexponential disentanglers

Recall that a channel

$$\Lambda : D(\mathbb{C}^D) \longrightarrow D(\mathbb{C}^d \otimes \mathbb{C}^d)$$

is an (ε, δ) -disentangler if [ABD⁺08]

$$\sup_{\rho \in D(\mathbb{C}^D)} D_{\text{tr}}(\Lambda(\rho), \text{SEP}(d, d)) \leq \varepsilon, \quad (5.1)$$

$$\sup_{\sigma \in \text{SEP}(d, d)} \inf_{\rho \in D(\mathbb{C}^D)} D_{\text{tr}}(\Lambda(\rho), \sigma) \leq \delta. \quad (5.2)$$

The first condition is soundness: every output is close to separable. The second is coverage: every separable state is close to some output. A sufficiently small and efficiently implementable disentangler would simulate two unentangled quantum proofs using one ordinary QMA witness. The original Watrous disentangler conjecture predicts $\log D = \Omega(d)$ for fixed ε, δ satisfying $\varepsilon + \delta < 1$ [ABD⁺08]. The construction below therefore disproves the original Watrous disentangler conjecture.

Theorem 5.1 (Subexponential disentanglers). *Let $d \geq 2$ and $n \geq 1$, and set $\mathcal{H}_A = \mathcal{H}_B = \mathbb{C}^d$. There is a channel*

$$\Lambda_{d,n} : D(\mathcal{H}_A^{\otimes n} \otimes \mathcal{H}_B^{\otimes n}) \longrightarrow D(\mathcal{H}_A \otimes \mathcal{H}_B)$$

which is an $(\varepsilon_{d,n}, 0)$ -disentangler, where

$$\varepsilon_{d,n} := \min \left\{ 1, \frac{\sqrt{d-1}}{2n} \right\}. \quad (5.3)$$

Consequently, for every fixed $\varepsilon \in (0, 1)$, there is an $(\varepsilon, 0)$ -disentangler with input dimension D satisfying $\log D = O_\varepsilon(\sqrt{d} \log d)$. In particular, one may take $(\varepsilon, \delta) = (1/4, 0)$.

Proof. Let Π_n^A and Π_n^B be the orthogonal projections onto $\text{Sym}^n(\mathcal{H}_A)$ and $\text{Sym}^n(\mathcal{H}_B)$, respectively, and set

$$P_n := \Pi_n^A \otimes \Pi_n^B.$$

Choose unit vectors $a_0 \in \mathcal{H}_A$ and $b_0 \in \mathcal{H}_B$, and let

$$\omega_n := |a_0\rangle\langle a_0|^{\otimes n} \otimes |b_0\rangle\langle b_0|^{\otimes n}.$$

The trace-preserving projection onto $\text{Sym}^n(\mathcal{H}_A) \otimes \text{Sym}^n(\mathcal{H}_B)$ is the channel

$$\mathcal{P}_n(\rho) := P_n \rho P_n + \text{Tr}((I - P_n)\rho) \omega_n. \quad (5.4)$$

Indeed, the first term retains the symmetric component, while the second replaces the rejected component by the fixed pure product state ω_n in the same subspace. Define the disentangler by retaining the $A_1 B_1$ registers:

$$\Lambda_{d,n}(\rho) := \text{Tr}_{A_{2:n} B_{2:n}} \mathcal{P}_n(\rho) = (\mathcal{P}_n(\rho))_{A_1 B_1}. \quad (5.5)$$

The state $\mathcal{P}_n(\rho)$ is supported on $\text{Sym}^n(\mathcal{H}_A) \otimes \text{Sym}^n(\mathcal{H}_B)$. Hence every output of $\Lambda_{d,n}$ belongs to $\text{BEXT}_{n,n}(d, d)$, and Theorem 3.9 gives

$$D_{\text{tr}}(\Lambda_{d,n}(\rho), \text{SEP}(d, d)) \leq \varepsilon_{d,n}. \quad (5.6)$$

Coverage is exact: if $\sigma = \sum_j p_j |a_j\rangle\langle a_j| \otimes |b_j\rangle\langle b_j|$ is separable, then

$$\tilde{\sigma}_n := \sum_j p_j |a_j\rangle\langle a_j|^{\otimes n} \otimes |b_j\rangle\langle b_j|^{\otimes n}$$

is supported on the range of P_n . Therefore $\mathcal{P}_n(\tilde{\sigma}_n) = \tilde{\sigma}_n$ and $\Lambda_{d,n}(\tilde{\sigma}_n) = \sigma$.

For fixed $\varepsilon \in (0, 1)$, choose

$$n := \left\lceil \frac{\sqrt{d-1}}{2\varepsilon} \right\rceil.$$

The input space has dimension $D_{d,n} = d^{2n}$, so

$$\log D_{d,n} = 2n \log d = \Theta_\varepsilon(\sqrt{d} \log d).$$

This proves Theorem 5.1. Taking $n = \lceil 2\sqrt{d-1} \rceil$ gives the stated $(1/4, 0)$ parameters. $\square$

5.2 Best Separable State without perfect completeness

Let $\mathcal{H}_A \cong \mathbb{C}^{d_A}$ and $\mathcal{H}_B \cong \mathbb{C}^{d_B}$. Best Separable State (BSS) asks for the largest expectation of an operator $0 \preceq M \preceq I$ on a separable state. When M is the accepting operator of a QMA(2) verifier, this is its largest acceptance probability over two unentangled proofs. For $M \in \text{Herm}(\mathcal{H}_A \otimes \mathcal{H}_B)$, write

$$h_{\text{SEP}}(M) := \max_{\sigma \in \text{SEP}(d_A, d_B)} \text{Tr}(M\sigma) = \max_{\substack{\|x\|=1 \\ \|y\|=1}} \langle x \otimes y | M | x \otimes y \rangle. \quad (5.7)$$

For an accepting operator $0 \preceq M \preceq I$ and $n \geq 1$, set

$$M^{[n,n]} := \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n M_{A_i B_j}, \quad (5.8)$$

$$h_{\text{BEXT}_{n,n}}(M) := \max_{\rho \in \text{BEXT}_{n,n}(d_A, d_B)} \text{Tr}(M\rho) = \lambda_{\max}(M^{[n,n]}),$$

on $\text{Sym}^n(\mathcal{H}_A) \otimes \text{Sym}^n(\mathcal{H}_B)$, with each summand acting as the identity on all unlisted tensor factors. The last equality follows from permutation symmetry within each block.

Theorem 5.2 (Subexponential BSS). *Let $0 \preceq M \preceq I$ act on $\mathcal{H}_A \otimes \mathcal{H}_B$. For every $n \geq 1$,*

$$h_{\text{SEP}}(M) \leq h_{\text{BEXT}_{n,n}}(M) \leq h_{\text{SEP}}(M) + \frac{\sqrt{\min\{d_A, d_B\} - 1}}{2n}. \quad (5.9)$$

Consequently, for an explicitly given accepting operator M on $\mathbb{C}^d \otimes \mathbb{C}^d$, $h_{\text{SEP}}(M)$ can be approximated to additive error $\varepsilon \in (0, 1]$ deterministically. Up to a factor polynomial in the bit complexity of M , the running time is

$$\exp \left(O \left(\frac{\sqrt{d}}{\varepsilon} \log(2 + \varepsilon\sqrt{d}) \right) \right). \quad (5.10)$$

The promise $h_{\text{SEP}}(M) \geq c$ versus $h_{\text{SEP}}(M) \leq s$ is decidable within the same bound with $\varepsilon = \Theta(c - s)$.

Proof. Write

$$\delta_n := \frac{\sqrt{\min\{d_A, d_B\} - 1}}{2n}.$$

The inclusion (2.4) gives the lower bound in (5.9). For the upper bound, every $\rho \in \text{BEXT}_{n,n}(d_A, d_B)$ is within δ_n of a separable state by Theorem 3.9; since $0 \preceq M \preceq I$, $\text{Tr}(M\rho) \leq h_{\text{SEP}}(M) + \delta_n$.

The two-sided lift acts on a space of dimension

$$L_{d_A, d_B, n} = \binom{d_A + n - 1}{n} \binom{d_B + n - 1}{n}.$$

Its entries are computable in time polynomial in $L_{d_A, d_B, n}$ and the bit complexity of M . For $d_A = d_B = d$, take $n = \lceil \sqrt{d-1}/\varepsilon \rceil$ and compute its largest eigenvalue to precision $\varepsilon/2$. Then $\delta_n \leq \varepsilon/2$, and the standard binomial estimate gives

$$\log L_{d,d,n} = O \left(\frac{\sqrt{d}}{\varepsilon} \log(2 + \varepsilon\sqrt{d}) \right).$$

This proves (5.10); taking $\varepsilon = \Theta(c - s)$ proves the promise statement. $\square$

The algorithm of Barak, Kothari, and Steurer for explicitly given BSS instances assumes perfect completeness, meaning $h_{\text{SEP}}(M) = 1$ in the yes case. Theorem 5.2 removes this assumption while retaining a running time subexponential in the local dimension [BKS17, Remark 1.4].

5.3 Trace-norm separability testing

Given a bipartite state ρ , trace-norm separability testing asks whether ρ is separable or is at trace distance at least ε from every separable state. This is the weak-membership problem for the separable set in trace distance. The BSS algorithm from the preceding subsection gives a semidefinite relaxation that approximates $h_{\text{SEP}}(M)$ with an additive error independent of M . General reductions between weak optimization and weak membership can lose polynomial factors in the dimension [GLS93]. Harrow and Montanaro proved the direction needed here with only a constant-factor loss in accuracy [HM13, Proposition 16]. Applying their reduction to our BSS approximation gives the claimed algorithm. By duality (Lemma 2.1),

$$D_{\text{tr}}(\rho, \text{SEP}(d_A, d_B)) = \max_{0 \preceq M \preceq I} \{\text{Tr}(M\rho) - h_{\text{SEP}}(M)\}. \quad (5.11)$$

For every operator $0 \preceq M \preceq I$, Theorem 5.2 gives

$$h_{\text{SEP}}(M) \leq h_{\text{BEXT}_{n,n}}(M) \leq h_{\text{SEP}}(M) + \delta_n, \quad \delta_n := \frac{\sqrt{\min\{d_A, d_B\} - 1}}{2n}. \quad (5.12)$$

Replacing h_{SEP} by this relaxation in (5.11), define

$$\widehat{D}_n(\rho) := \max_{0 \preceq M \preceq I} \{ \text{Tr}(M\rho) - h_{\text{BEXT}_{n,n}}(M) \}. \quad (5.13)$$

Proposition 5.3 (From weak optimization to weak membership). *For every state $\rho \in \mathcal{D}(\mathcal{H}_A \otimes \mathcal{H}_B)$ and every $n \geq 1$,*

$$D_{\text{tr}}(\rho, \text{SEP}(d_A, d_B)) - \delta_n \leq \widehat{D}_n(\rho) \leq D_{\text{tr}}(\rho, \text{SEP}(d_A, d_B)). \quad (5.14)$$

In particular, $\widehat{D}_n(\rho) = 0$ whenever ρ is separable. Moreover, (5.13) is the semidefinite program

$$\begin{aligned} & \text{maximize} && \text{Tr}(M\rho) - t \\ & \text{subject to} && 0 \preceq M \preceq I, \\ & && M^{[n,n]} \preceq tI_{\text{Sym}^n(\mathcal{H}_A) \otimes \text{Sym}^n(\mathcal{H}_B)}, \end{aligned} \quad (5.15)$$

with variables $M \in \text{Herm}(\mathcal{H}_A \otimes \mathcal{H}_B)$ and $t \in \mathbb{R}$.

Proof. For every operator $0 \preceq M \preceq I$, (5.12) gives

$$\text{Tr}(M\rho) - h_{\text{SEP}}(M) - \delta_n \leq \text{Tr}(M\rho) - h_{\text{BEXT}_{n,n}}(M) \leq \text{Tr}(M\rho) - h_{\text{SEP}}(M).$$

Maximizing over M and applying (5.11) proves (5.14). If ρ is separable, the upper bound gives $\widehat{D}_n(\rho) \leq 0$, while $M = 0$ gives the reverse inequality. Hence $\widehat{D}_n(\rho) = 0$.

By (5.8), $h_{\text{BEXT}_{n,n}}(M) = \lambda_{\max}(M^{[n,n]})$. For any Hermitian operator B , the inequality $B \preceq tI$ holds exactly when $\lambda_{\max}(B) \leq t$. Therefore,

$$M^{[n,n]} \preceq tI_{\text{Sym}^n(\mathcal{H}_A) \otimes \text{Sym}^n(\mathcal{H}_B)} \iff t \geq h_{\text{BEXT}_{n,n}}(M).$$

For fixed M , the objective $\text{Tr}(M\rho) - t$ decreases as t increases, so its maximum is attained at the smallest feasible value, $t = h_{\text{BEXT}_{n,n}}(M)$. Thus eliminating t from (5.15) recovers exactly (5.13). Finally, $M \mapsto M^{[n,n]}$ is linear, so the matrix in the displayed semidefinite constraint depends linearly on the optimization variables. $\square$

Theorem 5.4 (Subexponential trace-norm separability algorithm). *Let ρ be an explicitly given state on $\mathbb{C}^d \otimes \mathbb{C}^d$. For every $\varepsilon \in (0, 1]$, the distance $D_{\text{tr}}(\rho, \text{SEP}(d, d))$ can be approximated to additive error ε deterministically. Up to a factor polynomial in the bit complexity of ρ , the running time is*

$$\exp \left(O \left(\frac{\sqrt{d}}{\varepsilon} \log(2 + \varepsilon\sqrt{d}) \right) \right). \quad (5.16)$$

In the same running time one can distinguish

$$\rho \in \text{SEP}(d, d) \quad \text{from} \quad D_{\text{tr}}(\rho, \text{SEP}(d, d)) \geq \varepsilon. \quad (5.17)$$

Proof. Choose

$$n := \left\lceil \frac{\sqrt{d-1}}{\varepsilon} \right\rceil.$$

Then $\delta_n \leq \varepsilon/2$, and Proposition 5.3 gives

$$D_{\text{tr}}(\rho, \text{SEP}(d, d)) - \delta_n \leq \widehat{D}_n(\rho) \leq D_{\text{tr}}(\rho, \text{SEP}(d, d)). \quad (5.18)$$

Solve (5.15) to additive precision $\varepsilon/2$, and denote the resulting value by $\widetilde{D}_n(\rho)$. By (5.18),

$$\left| \widetilde{D}_n(\rho) - D_{\text{tr}}(\rho, \text{SEP}(d, d)) \right| \leq \frac{\varepsilon}{2} + \delta_n \leq \varepsilon.$$

Thus $\widetilde{D}_n(\rho)$ is the required approximation.

The semidefinite constraint $M^{[n,n]} \preceq t I_{\text{Sym}^n(\mathbb{C}^d) \otimes \text{Sym}^n(\mathbb{C}^d)}$ has matrix dimension

$$L_{d,d,n} = \binom{d+n-1}{n}^2.$$

The remaining constraints $0 \preceq M \preceq I$ have dimension d^2 . The construction and dimension estimate in the proof of Theorem 5.2, together with standard semidefinite optimization, therefore give (5.16).

For the promise problem, run the additive approximation just proved with accuracy $\varepsilon/3$. Its output is at most $\varepsilon/3$ when ρ is separable and at least $2\varepsilon/3$ when $D_{\text{tr}}(\rho, \text{SEP}(d, d)) \geq \varepsilon$. Comparing the output with $\varepsilon/2$ distinguishes the two cases. Replacing ε by $\varepsilon/3$ changes the running time only by constant factors in the exponent. $\square$

6 A dimension-free bosonic Hilbert–Schmidt de Finetti theorem

Brandão and Christandl identified local-dimension dependence as a crucial bottleneck for quantum de Finetti theorems and showed that the linear barrier can be broken under one-way LOCC measurements, obtaining only logarithmic dependence on the local dimension [BC12]. In Hilbert–Schmidt distance, the theorem below removes the local-dimension dependence entirely for bosonic approximation.

Theorem 6.1 (Dimension-free bosonic Hilbert–Schmidt de Finetti upper bound). *For every $N, d \geq 2$,*

$$D_{\text{H}}^{\text{HS}}\left(\text{BOS}_N^{(2)}(d), \text{SEP}_H^{(2)}(d)\right) \leq \min \left\{ \frac{\sqrt{d-1}}{N-1}, \frac{8}{\sqrt{N-1}} \right\}. \quad (6.1)$$

For each fixed d , the first term in (6.1) retains the faster rate $O(\sqrt{d}/N)$.

6.1 Spectral truncation and the dimension-free Hilbert–Schmidt bound

The dimension-dependent bound in (6.1) requires no new argument. Theorem 3.1 gives a state $\tau \in \text{SEP}_H^{(2)}(d)$ such that

$$\left\| \rho_N^{(2)} - \tau \right\|_2 \leq \left\| \rho_N^{(2)} - \tau \right\|_1 \leq \frac{\sqrt{d-1}}{N-1}.$$

It remains to prove the dimension-free bound.

Setup. By convexity of $D_{\text{HS}}(\cdot, \text{SEP}_H^{(2)}(d))$, it is enough to consider a pure bosonic state $\rho_N = |\psi\rangle\langle\psi|$ with $\psi \in \text{Sym}^N \mathcal{H}$ a unit vector. Set

$$\rho_1 := \rho_N^{(1)}, \quad \rho_2 := \rho_N^{(2)}.$$

Split the spectrum of ρ_1 at the scale $1/N$:

$$P := \mathbf{1}_{(1/N, \infty)}(\rho_1), \quad Q := I - P, \quad r := \text{rank } P. \quad (6.2)$$

We call $P\mathcal{H}$ the spectral head and $Q\mathcal{H}$ the spectral tail. The cutoff has two immediate consequences: since $\text{Tr } \rho_1 = 1$, the head has dimension $r < N$, and the tail is spectrally flat,

$$Q\rho_1 Q \preceq \frac{1}{N}Q. \quad (6.3)$$

By Fact 2.4, the induced two-site decomposition is

$$\text{Sym}^2 \mathcal{H} = \text{Sym}^2(P\mathcal{H}) \oplus (P\mathcal{H} \vee Q\mathcal{H}) \oplus \text{Sym}^2(Q\mathcal{H}). \quad (6.4)$$

Let Π_{PP} , Π_{PQ} , and Π_{QQ} be the projections onto the three summands. Group the two summands containing at least one tail direction by setting

$$\Pi_{\text{tail}} := \Pi_{PQ} + \Pi_{QQ} = I - \Pi_{PP}.$$

The corresponding diagonal blocks of ρ_2 are

$$\rho_2^{\text{head}} := \Pi_{PP}\rho_2\Pi_{PP}, \quad \rho_2^{\text{tail}} := \Pi_{\text{tail}}\rho_2\Pi_{\text{tail}}. \quad (6.5)$$

Let X denote the off-diagonal head–tail block,

$$X := \Pi_{PP}\rho_2\Pi_{\text{tail}}.$$

Relative to $\text{range } \Pi_{PP} \oplus \text{range } \Pi_{\text{tail}}$, the state has the positive block form

$$\rho_2 = \begin{pmatrix} \rho_2^{\text{head}} & X \\ X^\dagger & \rho_2^{\text{tail}} \end{pmatrix} \succeq 0. \quad (6.6)$$

Since $\text{Tr } \rho_2 = 1$, Fact 2.3 gives

$$\left\| \rho_2 - \rho_2^{\text{head}} \right\|_2^2 \leq 2 \left\| \rho_2^{\text{tail}} \right\|_\infty.$$

Thus the tail need not have small trace, and the off-diagonal block X need not be estimated separately: an operator-norm bound on the diagonal tail block controls the entire discarded part. The rank bound $r < N$ makes the head accessible to the optimal trace-norm de Finetti theorem, while (6.3) will force ρ_2^{tail} to be small in operator norm.

Proof architecture. The proof consists of two substantive estimates followed by a trace-restoration step. First, we construct a Hartree mixture τ_{head} that approximates ρ_2^{head} and has the same trace. Second, we transfer the one-site flatness (6.3) to the two-site estimate $\left\| \rho_2^{\text{tail}} \right\|_\infty = O(N^{-1})$. Fact 2.3 then removes the tail at Hilbert–Schmidt cost $O(N^{-1/2})$. The missing trace is restored by a Hartree mixture τ_{res} supported on $\text{Sym}^2(Q\mathcal{H})$. Set

$$\tau := \tau_{\text{head}} + \tau_{\text{res}}. \quad (6.7)$$

Their traces will add to one, so $\tau \in \text{SEP}_H^{(2)}(d)$. The proof is organized around the following error decomposition:

$$\begin{aligned} \left\| \rho_2 - \tau \right\|_2 &= \left\| (\rho_2^{\text{head}} - \tau_{\text{head}}) + (\rho_2 - \rho_2^{\text{head}}) - \tau_{\text{res}} \right\|_2 \\ &\leq \left\| \rho_2^{\text{head}} - \tau_{\text{head}} \right\|_2 + \left\| \rho_2 - \rho_2^{\text{head}} \right\|_2 + \left\| \tau_{\text{res}} \right\|_2 \\ &\leq \frac{1}{\sqrt{N-1}} + \sqrt{\frac{24}{N-1}} + \frac{2\sqrt{2}}{N} < \frac{8}{\sqrt{N-1}}. \end{aligned} \quad (6.8)$$

The three terms in the triangle inequality are bounded, in order, by Lemmas 6.2, 6.4, and 6.5.

6.1.1 Hartree approximation of the spectral head

We begin with the low-dimensional head, where the optimal trace-norm de Finetti theorem enters.

Lemma 6.2 (Rounding the spectral head). *There is a Hartree mixture τ_{head} , supported on $\text{Sym}^2(P\mathcal{H})$, such that*

$$\text{Tr } \tau_{\text{head}} = \text{Tr } \rho_2^{\text{head}}, \quad \left\| \rho_2^{\text{head}} - \tau_{\text{head}} \right\|_2 \leq \frac{1}{\sqrt{N-1}}.$$

Proof. Fact 2.4, applied to $\mathcal{H} = P\mathcal{H} \oplus Q\mathcal{H}$, gives

$$\text{Sym}^N \mathcal{H} \cong \bigoplus_{k=0}^N \text{Sym}^k(P\mathcal{H}) \otimes \text{Sym}^{N-k}(Q\mathcal{H}).$$

Let Π_k be the projection onto the k th summand and set

$$p_k := \text{Tr}(\Pi_k \rho_N).$$

The PP -compression is block diagonal with respect to this decomposition and vanishes on the summands with $k < 2$. For $k \geq 2$, after projecting two sites into $P\mathcal{H}$, the residual component lies in $\text{Sym}^{k-2}(P\mathcal{H}) \otimes \text{Sym}^{N-k}(Q\mathcal{H})$. These residual spaces are orthogonal for distinct values of k , so only the diagonal blocks $\Pi_k \rho_N \Pi_k$ contribute to ρ_2^{head} .

For $p_k > 0$, define

$$\sigma_k := \frac{\Pi_k \rho_N \Pi_k}{p_k} \in \text{D}(\text{range } \Pi_k), \quad \rho_{P,k} := \text{Tr}_{\text{Sym}^{N-k}(Q\mathcal{H})}(\sigma_k) \in \text{D}(\text{Sym}^k(P\mathcal{H})).$$

Write $\rho_{P,k}^{(2)}$ for its two-site marginal. Within the k th slice, exactly k of the N sites lie in $P\mathcal{H}$. Thus the probability that two sites sampled without replacement both lie in $P\mathcal{H}$ is

$$\frac{\binom{k}{2}}{\binom{N}{2}} = \frac{k(k-1)}{N(N-1)}.$$

Here P_i denotes P acting on the i th site. Equivalently, permutation symmetry and the identity $\sum_{i \neq j} P_i P_j = k(k-1)I$ on $\text{range } \Pi_k$ give

$$\text{Tr}(P_1 P_2 \sigma_k) = \frac{1}{N(N-1)} \text{Tr} \left(\left(\sum_{i \neq j} P_i P_j \right) \sigma_k \right) = \frac{k(k-1)}{N(N-1)}.$$

Under the canonical identification of the k th slice, its normalized PP -block is $\rho_{P,k}^{(2)}$. Therefore

$$\Pi_{PP} \text{Tr}_{3,\dots,N}(\sigma_k) \Pi_{PP} = \frac{k(k-1)}{N(N-1)} \rho_{P,k}^{(2)}.$$

Using $\Pi_k \rho_N \Pi_k = p_k \sigma_k$ and summing the diagonal slice contributions gives

$$\rho_2^{\text{head}} = \sum_{\substack{2 \leq k \leq N \\ p_k > 0}} p_k \frac{k(k-1)}{N(N-1)} \rho_{P,k}^{(2)}. \quad (6.9)$$

If $r \leq 1$, every term in (6.9) is Hartree, so we may take $\tau_{\text{head}} := \rho_2^{\text{head}}$. Assume now that $r \geq 2$. Applying Theorem 3.1 to $\rho_{P,k}$ gives a Hartree mixture τ_k supported on $\text{Sym}^2(P\mathcal{H})$ and satisfying

$$\left\| \rho_{P,k}^{(2)} - \tau_k \right\|_1 \leq \frac{\sqrt{r-1}}{k-1}. \quad (6.10)$$

Define the Hartree mixture

$$\tau_{\text{head}} := \sum_{\substack{2 \leq k \leq N \\ p_k > 0}} \frac{p_k k(k-1)}{N(N-1)} \tau_k. \quad (6.11)$$

This operator has the same trace as ρ_2^{head} . The key point is that the factor $k-1$ in the de Finetti error cancels the same factor in the pair weight:

$$\begin{aligned} \left\| \rho_2^{\text{head}} - \tau_{\text{head}} \right\|_2 &\leq \sum_{\substack{2 \leq k \leq N \\ p_k > 0}} \frac{p_k k(k-1)}{N(N-1)} \left\| \rho_{P,k}^{(2)} - \tau_k \right\|_1 \\ &\leq \frac{\sqrt{r-1}}{N(N-1)} \sum_{k=2}^N p_k k \\ &\leq \frac{\sqrt{r-1}}{N-1} \leq \frac{1}{\sqrt{N-1}}. \end{aligned} \quad (6.12)$$

For $r \leq 1$, the same estimate holds with zero left-hand side. $\square$

6.1.2 Operator-norm control of the spectral tail

The mixed and tail–tail diagonal blocks are bounded directly, after which positive block estimates assemble them into an operator-norm bound on ρ_2^{tail} .

Lemma 6.3 (Mixed and tail–tail block estimates). *With the notation above,*

$$\|\Pi_{PQ}\rho_2\Pi_{PQ}\|_\infty \leq \frac{4}{N}, \quad (6.13)$$

$$\|\Pi_{QQ}\rho_2\Pi_{QQ}\|_\infty \leq \frac{8}{N-1}. \quad (6.14)$$

Proof. Recall that $\rho_N = |\psi\rangle\langle\psi|$, where $\psi \in \text{Sym}^N \mathcal{H}$ is a unit vector.

The mixed block. Let $\phi \in P\mathcal{H} \vee Q\mathcal{H}$ be a unit vector. Its Schmidt decomposition is

$$\phi = \sum_{i=1}^{r_\phi} \sqrt{\mu_i} p_i \vee q_i, \quad \mu_i > 0, \quad \sum_i \mu_i = 1, \quad (6.15)$$

where (p_i) and (q_i) are orthonormal families in $P\mathcal{H}$ and $Q\mathcal{H}$, respectively, and $r_\phi \leq r < N$. Since ψ is symmetric,

$$C_{p_i \vee q_i}(\psi) = \frac{1}{\sqrt{2}} ((\langle p_i | \otimes \langle q_i | \otimes I) + (\langle q_i | \otimes \langle p_i | \otimes I)) \psi = \sqrt{2} C_{p_i}(C_{q_i}(\psi)).$$

Although the Schmidt directions q_i need not be eigenvectors of ρ_1 , each is a unit vector in $Q\mathcal{H}$, so (6.3) gives

$$\|C_{q_i}(\psi)\|^2 = \langle q_i | \rho_1 | q_i \rangle = \langle q_i | Q \rho_1 Q | q_i \rangle \leq \frac{1}{N}.$$

Since the Schmidt weights sum to one, it follows that

$$\sum_i \mu_i \|C_{q_i}(\psi)\|^2 \leq \frac{1}{N} \sum_i \mu_i = \frac{1}{N}. \quad (6.16)$$

For every unit vector $\Xi \in \text{Sym}^{N-2} \mathcal{H}$, use the adjoint relation

$$\langle \Xi, C_{p_i}(C_{q_i}(\psi)) \rangle = \langle C_{p_i}^*(\Xi), C_{q_i}(\psi) \rangle.$$

Together with Cauchy–Schwarz, Fact 2.6, specifically (2.26), with $k = N - 1$ and $\ell = r_\phi$, and (6.16), this gives

$$\begin{aligned} |\langle \Xi, C_\phi(\psi) \rangle| &= \left| \sum_i \sqrt{\mu_i} \langle \Xi, C_{p_i \vee q_i}(\psi) \rangle \right| \\ &= \sqrt{2} \left| \sum_i \sqrt{\mu_i} \langle \Xi, C_{p_i}(C_{q_i}(\psi)) \rangle \right| \\ &= \sqrt{2} \left| \sum_i \sqrt{\mu_i} \langle C_{p_i}^*(\Xi), C_{q_i}(\psi) \rangle \right| \\ &\leq \sqrt{2} \left(\sum_i \|C_{p_i}^*(\Xi)\|^2 \right)^{1/2} \left(\sum_i \mu_i \|C_{q_i}(\psi)\|^2 \right)^{1/2} \\ &\leq \sqrt{2} \sqrt{\frac{r_\phi + N - 2}{N - 1}} \frac{1}{\sqrt{N}} \\ &\leq \frac{2}{\sqrt{N}}. \end{aligned}$$

The last inequality uses $r_\phi < N$. Taking the supremum over unit vectors $\Xi \in \text{Sym}^{N-2} \mathcal{H}$ therefore gives

$$\|C_\phi(\psi)\| \leq \frac{2}{\sqrt{N}}.$$

Since $\phi \in P\mathcal{H} \vee Q\mathcal{H}$, we have $\Pi_{PQ}\phi = \phi$. Fact 2.5, specifically the partial-inner-product identity (2.23), now gives

$$\langle \phi | \Pi_{PQ} \rho_2 \Pi_{PQ} | \phi \rangle = \langle \phi | \rho_2 | \phi \rangle = \|C_\phi(\psi)\|^2 \leq \frac{4}{N}.$$

The compression $\Pi_{PQ} \rho_2 \Pi_{PQ}$ is positive semidefinite, so its operator norm is the supremum of the left-hand side over unit vectors $\phi \in P\mathcal{H} \vee Q\mathcal{H}$. This proves (6.13).

The tail–tail block. Let $\phi \in \text{Sym}^2(Q\mathcal{H})$ be a unit vector. A Takagi decomposition gives

$$\phi = \sum_{i=1}^m \sqrt{\mu_i} q_i^{\otimes 2}, \quad \mu_1 \geq \mu_2 \geq \cdots \geq 0, \quad \sum_i \mu_i = 1, \quad (6.17)$$

for an orthonormal family $(q_i) \subseteq Q\mathcal{H}$. Split

$$\phi = \phi_{\leq N} + \phi_{> N}, \quad \phi_{\leq N} := \sum_{i \leq N} \sqrt{\mu_i} q_i^{\otimes 2}, \quad \phi_{> N} := \sum_{i > N} \sqrt{\mu_i} q_i^{\otimes 2},$$

with empty terms omitted. The leading N Takagi terms are controlled by applying Cauchy–Schwarz with the weights μ_i multiplying $\|C_{q_i}(\psi)\|^2$, exactly as in the mixed block. The same pointwise tail estimate gives

$$\sum_{i \leq N} \mu_i \|C_{q_i}(\psi)\|^2 \leq \frac{1}{N} \sum_{i \leq N} \mu_i \leq \frac{1}{N}.$$

Therefore, for every unit $\Xi \in \text{Sym}^{N-2} \mathcal{H}$,

$$\begin{aligned} |\langle \Xi, C_{\phi_{\leq N}}(\psi) \rangle| &\leq \left(\sum_{i \leq N} \|C_{q_i}^*(\Xi)\|^2 \right)^{1/2} \left(\sum_{i \leq N} \mu_i \|C_{q_i}(\psi)\|^2 \right)^{1/2} \\ &\leq \sqrt{\frac{2}{N}}. \end{aligned}$$

Hence

$$\|C_{\phi_{\leq N}}(\psi)\|^2 \leq \frac{2}{N}. \quad (6.18)$$

For the remaining Takagi terms, monotonicity of (μ_i) and $\sum_i \mu_i = 1$ imply $\mu_{N+1} \leq 1/(N+1)$. There may now be arbitrarily many terms, so we instead attach the weights μ_i to $\|C_{q_i}^*(\Xi)\|^2$ and leave $\|C_{q_i}(\psi)\|^2$ unweighted. Thus

$$|\langle \Xi, C_{\phi_{>N}}(\psi) \rangle| \leq \left(\sum_{i > N} \mu_i \|C_{q_i}^*(\Xi)\|^2 \right)^{1/2} \left(\sum_{i > N} \|C_{q_i}(\psi)\|^2 \right)^{1/2}.$$

To estimate the first factor, apply Fact 2.6, specifically (2.24), with $k = N - 1$:

$$\begin{aligned} \sum_{i > N} \mu_i \|C_{q_i}^*(\Xi)\|^2 &= \frac{1}{N-1} \left(\sum_{i > N} \mu_i + (N-2) \sum_{i > N} \mu_i \|C_{q_i}(\Xi)\|^2 \right) \\ &\leq \frac{1}{N-1} \left(1 + \frac{N-2}{N+1} \right) < \frac{2}{N-1}. \end{aligned}$$

Here $\sum_{i > N} \mu_i \leq 1$, while (2.25) and $\mu_i \leq \mu_{N+1} \leq 1/(N+1)$ for $i > N$ give

$$\sum_{i > N} \mu_i \|C_{q_i}(\Xi)\|^2 \leq \mu_{N+1} \sum_{i > N} \|C_{q_i}(\Xi)\|^2 \leq \frac{1}{N+1}.$$

The same Bessel bound, applied to ψ , gives

$$\sum_{i > N} \|C_{q_i}(\psi)\|^2 \leq 1.$$

Therefore

$$\|C_{\phi_{>N}}(\psi)\|^2 \leq \frac{2}{N-1}. \quad (6.19)$$

By linearity, $C_\phi = C_{\phi_{\leq N}} + C_{\phi_{>N}}$. Hence the triangle inequality, together with (6.18) and (6.19), gives

$$\|C_\phi(\psi)\| \leq \|C_{\phi_{\leq N}}(\psi)\| + \|C_{\phi_{>N}}(\psi)\| \leq \sqrt{\frac{2}{N}} + \sqrt{\frac{2}{N-1}}.$$

Since $\phi \in \text{Sym}^2(Q\mathcal{H})$, we have $\Pi_{QQ}\phi = \phi$. Fact 2.5, specifically (2.23), therefore implies

$$\langle \phi | \Pi_{QQ} \rho_2 \Pi_{QQ} | \phi \rangle = \langle \phi | \rho_2 | \phi \rangle = \|C_\phi(\psi)\|^2 \leq \left(\sqrt{\frac{2}{N}} + \sqrt{\frac{2}{N-1}} \right)^2 \leq \frac{8}{N-1}.$$

Finally, $\Pi_{QQ} \rho_2 \Pi_{QQ}$ is positive semidefinite, so its operator norm is the supremum of the left-hand side over unit vectors $\phi \in \text{Sym}^2(Q\mathcal{H})$. This proves (6.14). $\square$

Lemma 6.4 (Removing the spectral tail). *The head compression in (6.5) satisfies*

$$\|\rho_2 - \rho_2^{\text{head}}\|_2 \leq \sqrt{\frac{24}{N-1}}.$$

Proof. The range of Π_{tail} is the direct sum of the mixed and tail–tail spaces in (6.4). Relative to this decomposition,

$$\rho_2^{\text{tail}} = \begin{pmatrix} \Pi_{PQ} \rho_2 \Pi_{PQ} & \Pi_{PQ} \rho_2 \Pi_{QQ} \\ \Pi_{QQ} \rho_2 \Pi_{PQ} & \Pi_{QQ} \rho_2 \Pi_{QQ} \end{pmatrix} \succeq 0.$$

Fact 2.2 and Lemma 6.3 therefore give

$$\begin{aligned} \|\rho_2^{\text{tail}}\|_\infty &\leq \|\Pi_{PQ} \rho_2 \Pi_{PQ}\|_\infty + \|\Pi_{QQ} \rho_2 \Pi_{QQ}\|_\infty \\ &\leq \frac{4}{N} + \frac{8}{N-1} \leq \frac{12}{N-1}. \end{aligned} \tag{6.20}$$

Apply the trace-one case of Fact 2.3 to the block decomposition (6.6). This controls the diagonal tail and the off-diagonal head–tail blocks simultaneously:

$$\|\rho_2 - \rho_2^{\text{head}}\|_2^2 \leq 2 \|\rho_2^{\text{tail}}\|_\infty \leq \frac{24}{N-1}. \tag{6.21}$$

Thus positivity converts an $O(N^{-1})$ operator-norm estimate on the discarded block into an $O(N^{-1/2})$ Hilbert–Schmidt estimate for the entire discarded part, including the off-diagonal blocks. $\square$

6.1.3 Completion of the dimension-free bound

At this point both substantive estimates are complete: the head has a Hartree approximation of the correct trace, and the entire discarded part is small in Hilbert–Schmidt norm. It remains only to restore the missing trace. Spreading this mass uniformly over the symmetric tail costs only $O(N^{-1})$.

Lemma 6.5 (Restoring the missing trace). *There is a Hartree mixture τ_{res} , supported on $\text{Sym}^2(Q\mathcal{H})$, such that*

$$\text{Tr } \tau_{\text{res}} = 1 - \text{Tr } \rho_2^{\text{head}}, \quad \|\tau_{\text{res}}\|_2 \leq \frac{2\sqrt{2}}{N}.$$

Proof. Set

$$\delta := 1 - \text{Tr } \rho_2^{\text{head}}, \quad s := \dim(Q\mathcal{H}).$$

Since $\Pi_{\text{tail}} \preceq Q \otimes I + I \otimes Q$ on $\text{Sym}^2 \mathcal{H}$, the spectral-tail bound (6.3) gives

$$\delta = \text{Tr}(\Pi_{\text{tail}} \rho_2) \leq \text{Tr}((Q \otimes I + I \otimes Q) \rho_2) = 2 \text{Tr}(Q \rho_1) \leq \frac{2s}{N}.$$

If $s = 0$, then $\delta = 0$, and we take $\tau_{\text{res}} = 0$. Suppose that $s \geq 1$, and let μ_Q denote Haar probability measure on the unit sphere of $Q\mathcal{H}$. Define

$$\tau_{\text{res}} := \delta \frac{\Pi_{QQ}}{\dim \text{Sym}^2(Q\mathcal{H})} = \delta \int |u\rangle\langle u|^{\otimes 2} d\mu_Q(u).$$

Thus τ_{res} is a Hartree mixture supported on $\text{Sym}^2(Q\mathcal{H})$, and its trace is δ . Since $\dim \text{Sym}^2(Q\mathcal{H}) = s(s+1)/2$,

$$\|\tau_{\text{res}}\|_2 = \delta \sqrt{\frac{2}{s(s+1)}} \leq \frac{2\sqrt{2}}{N} \sqrt{\frac{s}{s+1}} \leq \frac{2\sqrt{2}}{N}. \quad \square$$

Completion of the proof of the dimension-free bound. Let τ_{head} and τ_{res} be supplied by Lemmas 6.2 and 6.5, and let τ be their sum as in (6.7). The two summands are Hartree mixtures whose traces add to one, so $\tau \in \text{SEP}_{\text{H}}^{(2)}(d)$. The three estimates used in (6.8) are precisely the conclusions of Lemmas 6.2, 6.4, and 6.5. Hence (6.8) gives $\|\rho_2 - \tau\|_2 < 8/\sqrt{N-1}$. This is the dimension-free term in (6.1). $\square$

6.2 Two-sided Bose-symmetric extendible states

Since Theorem 6.1 is dimension free, one might first try a tensor-product construction: pair $A_i B_i$ into n bosonic sites with local space $\mathcal{H}_A \otimes \mathcal{H}_B$ and apply the theorem. The desired cross marginal is then obtained from the resulting two-site marginal by tracing out $B_1 A_2$, and the image of each Hartree state under this partial trace is separable. However, partial trace destroys dimension independence.

To avoid this loss, we instead use a label-bundling, or the direct-sum construction: the two types of registers are placed in orthogonally labeled summands of a single bosonic system, so that the desired state is a compression of a two-site marginal rather than a partial trace. The label-bundling construction is inspired by the symmetrization idea of Liu and Wu [LW26, Sec. 4.2.1].

This label-bundling construction is a convenient technical workaround for the applications in the next section. A more systematic approach would repeat the spectral-truncation argument of the preceding subsection directly for two-sided extensions, using Theorem 3.9 on the low-dimensional spectral head; we leave such a direct proof for future work.

Let $\mathcal{H}_A = \mathbb{C}^{d_A}$ and $\mathcal{H}_B = \mathbb{C}^{d_B}$, and set

$$\mathcal{H}_{\oplus} = (|0\rangle \otimes \mathcal{H}_A) \oplus (|1\rangle \otimes \mathcal{H}_B).$$

Under the standard decomposition of a symmetric power of an orthogonal sum, $\text{Sym}^n(\mathcal{H}_A) \otimes \text{Sym}^n(\mathcal{H}_B)$ identifies with the summand of $\text{Sym}^{2n}(\mathcal{H}_{\oplus})$ containing n sites of each label.

The subspace of $\text{Sym}^2(\mathcal{H}_{\oplus})$ containing one site of each label is identified with $\mathcal{H}_A \otimes \mathcal{H}_B$ by the isometry

$$V : \mathcal{H}_A \otimes \mathcal{H}_B \longrightarrow \text{Sym}^2(\mathcal{H}_{\oplus}), \quad V(x \otimes y) = \frac{|0, x\rangle \otimes |1, y\rangle + |1, y\rangle \otimes |0, x\rangle}{\sqrt{2}}. \quad (6.22)$$

Lemma 6.6 (Label-bundling embedding). *Let $\rho \in \text{BEXT}_{n,n}(d_A, d_B)$. There is a bosonic state $\hat{\rho}_{2n} \in \text{D}(\text{Sym}^{2n}(\mathcal{H}_{\oplus}))$ whose two-site marginal satisfies*

$$V^\dagger \hat{\rho}_{2n}^{(2)} V = q_n \rho, \quad q_n := \frac{n}{2n-1}. \quad (6.23)$$

Thus the cross-label block on the left has trace q_n . Equivalently, measuring the two labels gives one label of each type with probability q_n , and conditioned on this event the state on $\mathcal{H}_A \otimes \mathcal{H}_B$ is exactly ρ .

Proof. Choose an extension $\tilde{\rho}$ from (2.3). By Fact 2.4,

$$\text{Sym}^{2n}(\mathcal{H}_{\oplus}) \cong \bigoplus_{k=0}^{2n} \text{Sym}^k(\mathcal{H}_A) \otimes \text{Sym}^{2n-k}(\mathcal{H}_B).$$

We may therefore regard $\tilde{\rho}$ as a bosonic state $\hat{\rho}_{2n}$ supported on the $k = n$ summand.

It remains to compute the cross-label block of the two-site marginal. Under this fixed-label-count identification, the block is the uniform average of the n^2 cross marginals:

$$V^\dagger \hat{\rho}_{2n}^{(2)} V = \binom{2n}{2}^{-1} \sum_{i=1}^n \sum_{j=1}^n \tilde{\rho}_{A_i B_j}.$$

Bose symmetry within each block makes every $\tilde{\rho}_{A_i B_j}$ equal to ρ . Hence

$$V^\dagger \hat{\rho}_{2n}^{(2)} V = \frac{n^2}{\binom{2n}{2}} \rho = \frac{n}{2n-1} \rho.$$

The coefficient $n^2/\binom{2n}{2} = n/(2n-1)$ is precisely the probability that a uniformly chosen pair contains one site of each label. This proves (6.23). $\square$

We also record how the same isometry acts on Hermitian operators.

Fact 6.7. *For $M \in \text{Herm}(\mathcal{H}_A \otimes \mathcal{H}_B)$, define*

$$A_M := V M V^\dagger \tag{6.24}$$

on $\text{Sym}^2(\mathcal{H}_{\oplus})$, acting as zero on $(\text{range } V)^\perp$. Then

$$h(A_M) = \frac{1}{2} \max\{h_{\text{SEP}}(M), 0\}, \quad \|A_M\|_2 = \|M\|_2. \tag{6.25}$$

If either $M \succeq 0$ or $\text{Tr } M = 0$, then $h_{\text{SEP}}(M) \geq 0$ and hence $h(A_M) = \frac{1}{2} h_{\text{SEP}}(M)$.

Proof. Write a unit vector in $\mathcal{H}_{\oplus}$ as

$$z = \sqrt{p} |0, x\rangle + \sqrt{q} |1, y\rangle,$$

where x, y are unit vectors and $p, q \geq 0$ satisfy $p + q = 1$. Since $V^\dagger z^{\otimes 2} = \sqrt{2pq} x \otimes y$,

$$\langle z^{\otimes 2} | A_M | z^{\otimes 2} \rangle = 2pq \langle x \otimes y | M | x \otimes y \rangle.$$

Maximizing over p, q, x, y gives the first identity in (6.25). Since V is an isometry, A_M is unitarily equivalent to $M \oplus 0$, which proves the Hilbert–Schmidt norm identity. Finally, $h_{\text{SEP}}(M) \geq 0$ is immediate when $M \succeq 0$. If $\text{Tr } M = 0$, then the average of $\langle x \otimes y | M | x \otimes y \rangle$ over independent Haar-random unit vectors x and y is $\text{Tr}(M)/(d_A d_B) = 0$, so its maximum is nonnegative. $\square$

Theorem 3.9 gives the dimension-dependent estimate below. Combining the label-bundling embedding in Lemma 6.6 with Theorem 6.1 gives the dimension-free estimate.

Corollary 6.8 (Dimension-free Hilbert–Schmidt de Finetti bound for two-sided Bose-symmetric extendible states). *For every $n, d_A, d_B \geq 1$,*

$$D_{\text{H}}^{\text{HS}}\left(\text{BEXT}_{n,n}(d_A, d_B), \text{SEP}(d_A, d_B)\right) \leq \min \left\{ \frac{\sqrt{\min\{d_A, d_B\} - 1}}{n}, \frac{8\sqrt{2n-1}}{n} \right\}. \tag{6.26}$$

Proof of Corollary 6.8. Fix $\rho \in \text{BEXT}_{n,n}(d_A, d_B)$. By Theorem 3.9, there is a state $\sigma \in \text{SEP}(d_A, d_B)$ such that

$$D_{\text{tr}}(\rho, \sigma) \leq \frac{\sqrt{\min\{d_A, d_B\} - 1}}{2n}.$$

Therefore,

$$D_{\text{HS}}(\rho, \sigma) \leq \|\rho - \sigma\|_1 = 2D_{\text{tr}}(\rho, \sigma) \leq \frac{\sqrt{\min\{d_A, d_B\} - 1}}{n}.$$

This proves the dimension-dependent estimate.

For the dimension-free estimate, by duality (Lemma 2.1),

$$D_{\text{HS}}(\rho, \text{SEP}(d_A, d_B)) = \max_{\substack{M=M^\dagger \\ \|M\|_2 \leq 1}} \{\text{Tr}(M\rho) - h_{\text{SEP}}(M)\}. \quad (6.27)$$

The maximum may be restricted to traceless M : replacing M by $M - \frac{\text{Tr } M}{d_A d_B} I$ leaves the objective unchanged and can only decrease its Hilbert–Schmidt norm. For $A_M := VMV^\dagger$, Fact 6.7 and the tracelessness of M give

$$h(A_M) = \frac{1}{2} h_{\text{SEP}}(M), \quad \|A_M\|_2 = \|M\|_2.$$

Let $\hat{\rho}_{2n}$ be the bosonic state supplied by Lemma 6.6, and set

$$\varepsilon_{2n} := D_{\text{H}}^{\text{HS}}\left(\text{BOS}_{2n}^{(2)}(d_A + d_B), \text{SEP}_{\text{H}}^{(2)}(d_A + d_B)\right).$$

There is a Hartree mixture $\tau \in \text{SEP}_{\text{H}}^{(2)}(d_A + d_B)$ such that $\|\hat{\rho}_{2n}^{(2)} - \tau\|_2 \leq \varepsilon_{2n}$. Since $\text{Tr}(A_M \tau) \leq h(A_M)$, Hilbert–Schmidt Cauchy–Schwarz and Fact 6.7 give

$$\text{Tr}(A_M \hat{\rho}_{2n}^{(2)}) \leq \frac{1}{2} h_{\text{SEP}}(M) + \varepsilon_{2n} \|M\|_2.$$

On the other hand, the label-bundling identity (6.23) gives $\text{Tr}(A_M \hat{\rho}_{2n}^{(2)}) = \frac{n}{2n-1} \text{Tr}(M\rho)$. Thus, for every traceless M with $\|M\|_2 \leq 1$,

$$\begin{aligned} \text{Tr}(M\rho) - h_{\text{SEP}}(M) &\leq \frac{2n-1}{n} \varepsilon_{2n} \|M\|_2 - \frac{1}{2n} h_{\text{SEP}}(M) \\ &\leq \frac{2n-1}{n} \varepsilon_{2n}, \end{aligned}$$

where the last inequality uses $h_{\text{SEP}}(M) \geq 0$. Taking the maximum in (6.27) and then the supremum over ρ gives, by (2.7),

$$D_{\text{H}}^{\text{HS}}\left(\text{BEXT}_{n,n}(d_A, d_B), \text{SEP}(d_A, d_B)\right) \leq \frac{2n-1}{n} \varepsilon_{2n}.$$

Theorem 6.1, applied with $N = 2n$, gives $\varepsilon_{2n} \leq 8/\sqrt{2n-1}$ and therefore

$$D_{\text{H}}^{\text{HS}}\left(\text{BEXT}_{n,n}(d_A, d_B), \text{SEP}(d_A, d_B)\right) \leq \frac{8\sqrt{2n-1}}{n}.$$

Together with the dimension-dependent estimate above, this proves (6.26). $\square$

7 Hilbert–Schmidt separability algorithms

The dimension-free estimate for two-sided Bose-symmetric extendible states has two algorithmic consequences. The two-sided extension relaxation approximates $h_{\text{SEP}}(M)$ for Hermitian objectives whose traceless part has bounded Hilbert–Schmidt norm. Through Hilbert–Schmidt duality, the same approximation also gives Hilbert–Schmidt separability testing. At fixed accuracy, the required extension level is independent of the local dimensions, so both programs have polynomial size in those dimensions.

Let $\mathcal{H}_A \cong \mathbb{C}^{d_A}$ and $\mathcal{H}_B \cong \mathbb{C}^{d_B}$, for $n \geq 1$, set

$$\delta_n^{\text{HS}}(d_A, d_B) := \min \left\{ \frac{\sqrt{\min\{d_A, d_B\} - 1}}{n}, \frac{8\sqrt{2n-1}}{n} \right\}. \quad (7.1)$$

7.1 Separable optimization with bounded Hilbert–Schmidt norm

For $M \in \text{Herm}(\mathcal{H}_A \otimes \mathcal{H}_B)$, set

$$m := \frac{\text{Tr } M}{d_A d_B}, \quad M_0 := M - mI. \quad (7.2)$$

Recall from Subsection 5.2 that

$$h_{\text{BEXT}_{n,n}}(M) := \max_{\rho \in \text{BEXT}_{n,n}(d_A, d_B)} \text{Tr}(M\rho) = \lambda_{\max}(M^{[n,n]}).$$

Adding a scalar multiple of the identity shifts the value of M by the same amount on every state. Hence

$$\begin{aligned} h_{\text{SEP}}(M) &= m + h_{\text{SEP}}(M_0), \\ h_{\text{BEXT}_{n,n}}(M) &= m + h_{\text{BEXT}_{n,n}}(M_0), \\ \|M_0\|_2 &\leq \|M\|_2. \end{aligned} \quad (7.3)$$

The last inequality holds because $M \mapsto M_0$ is the orthogonal projection onto the traceless subspace in the Hilbert–Schmidt inner product.

Theorem 7.1 (Separable optimization with bounded Hilbert–Schmidt norm). *For every Hermitian M and every $n \geq 1$,*

$$h_{\text{SEP}}(M) \leq h_{\text{BEXT}_{n,n}}(M) \leq h_{\text{SEP}}(M) + \delta_n^{\text{HS}}(d_A, d_B) \|M_0\|_2. \quad (7.4)$$

Consequently, given an explicit rational description of M and $\varepsilon \in (0, 1]$, one can approximate $h_{\text{SEP}}(M)$ to additive error ε deterministically in time

$$(d_A + d_B)^{O(\|M_0\|_2^2/\varepsilon^2)}. \quad (7.5)$$

In particular, for bounded $\|M_0\|_2$ and fixed ε , the running time is polynomial in the local dimensions.

Proof. The lower bound in (7.4) follows from $\text{SEP}(d_A, d_B) \subseteq \text{BEXT}_{n,n}(d_A, d_B)$. For the upper bound, fix $\rho \in \text{BEXT}_{n,n}(d_A, d_B)$. Corollary 6.8 gives $\sigma \in \text{SEP}(d_A, d_B)$ such that

$$\|\rho - \sigma\|_2 \leq \delta_n^{\text{HS}}(d_A, d_B).$$

Since M and M_0 differ by a scalar multiple of the identity, Hilbert–Schmidt Cauchy–Schwarz gives

$$\text{Tr}(M\rho) - h_{\text{SEP}}(M) = \text{Tr}(M_0\rho) - h_{\text{SEP}}(M_0)$$

$$\begin{aligned} &\leq \text{Tr}(M_0(\rho - \sigma)) \\ &\leq \delta_n^{\text{HS}}(d_A, d_B) \|M_0\|_2. \end{aligned}$$

Maximizing over ρ proves the upper bound.

For the algorithm, take

$$n := 1 + \left\lceil \frac{512 \|M_0\|_2^2}{\varepsilon^2} \right\rceil.$$

Then the error term in (7.4) is at most $\varepsilon/2$. Compute the largest eigenvalue of $M^{[n,n]}$ to additive precision $\varepsilon/2$. The lift acts on a space of dimension

$$L_{d_A, d_B, n} = \binom{d_A + n - 1}{n} \binom{d_B + n - 1}{n}.$$

The standard binomial estimate and $n = O(\|M_0\|_2^2 / \varepsilon^2)$ give (7.5). $\square$

7.2 Hilbert–Schmidt separability testing

Given a bipartite state ρ , Hilbert–Schmidt separability testing asks whether ρ is separable or is at Hilbert–Schmidt distance at least ε from every separable state. This is weak membership for the separable set in Hilbert–Schmidt distance. By duality (Lemma 2.1),

$$D_{\text{HS}}(\rho, \text{SEP}(d_A, d_B)) = \max_{\substack{M \in \text{Herm}(\mathcal{H}_A \otimes \mathcal{H}_B) \\ \text{Tr } M = 0, \|M\|_2 \leq 1}} \{\text{Tr}(M\rho) - h_{\text{SEP}}(M)\}. \quad (7.6)$$

We replace h_{SEP} in (7.6) by the relaxation from the preceding subsection.

Theorem 7.2 (Polynomial-time Hilbert–Schmidt separability testing). *Given an explicit rational state $\rho \in \text{D}(\mathbb{C}^{d_A} \otimes \mathbb{C}^{d_B})$ and $\varepsilon \in (0, 1]$, one can approximate $D_{\text{HS}}(\rho, \text{SEP}(d_A, d_B))$ to additive error ε deterministically in time*

$$(d_A + d_B)^{O(1/\varepsilon^2)}. \quad (7.7)$$

In the same running time one can distinguish

$$\rho \in \text{SEP}(d_A, d_B) \quad \text{from} \quad D_{\text{HS}}(\rho, \text{SEP}(d_A, d_B)) \geq \varepsilon. \quad (7.8)$$

Thus, for every fixed $\varepsilon > 0$, Hilbert–Schmidt separability testing is polynomial-time solvable in the explicit input dimensions.

Proof. For $n \geq 1$, define

$$\widehat{D}_n^{\text{HS}}(\rho) := \max_{\substack{M \in \text{Herm}(\mathcal{H}_A \otimes \mathcal{H}_B) \\ \text{Tr } M = 0, \|M\|_2 \leq 1}} \{\text{Tr}(M\rho) - h_{\text{BEXT}_{n,n}}(M)\}. \quad (7.9)$$

For every feasible M , Theorem 7.1 gives

$$h_{\text{SEP}}(M) \leq h_{\text{BEXT}_{n,n}}(M) \leq h_{\text{SEP}}(M) + \delta_n^{\text{HS}}(d_A, d_B).$$

Subtracting these inequalities from $\text{Tr}(M\rho)$, maximizing over M , and applying (7.6) gives

$$D_{\text{HS}}(\rho, \text{SEP}(d_A, d_B)) - \delta_n^{\text{HS}}(d_A, d_B) \leq \widehat{D}_n^{\text{HS}}(\rho) \leq D_{\text{HS}}(\rho, \text{SEP}(d_A, d_B)). \quad (7.10)$$

Since $h_{\text{BEXT}_{n,n}}(M) = \lambda_{\max}(M^{[n,n]})$, the relaxation (7.9) is the semidefinite program

$$\begin{aligned} & \text{maximize} && \text{Tr}(M\rho) - t \\ & \text{subject to} && \text{Tr } M = 0, \\ & && \begin{pmatrix} I_{(d_A d_B)^2} & \text{vec}(M) \\ \text{vec}(M)^\dagger & 1 \end{pmatrix} \succeq 0, \\ & && M^{[n,n]} \preceq t I_{\text{Sym}^n(\mathcal{H}_A) \otimes \text{Sym}^n(\mathcal{H}_B)}, \end{aligned} \tag{7.11}$$

with variables $M \in \text{Herm}(\mathcal{H}_A \otimes \mathcal{H}_B)$ and $t \in \mathbb{R}$, where $\text{vec}(M)$ denotes the vectorization of M . The first matrix inequality is equivalent to $\|M\|_2 \leq 1$ by the Schur complement, while the second is equivalent to $t \geq h_{\text{BEXT}_{n,n}}(M)$. Hence the SDP has value $\hat{D}_n^{\text{HS}}(\rho)$.

For the approximation algorithm, take

$$n := 1 + \left\lceil \frac{512}{\varepsilon^2} \right\rceil.$$

Then $\delta_n^{\text{HS}}(d_A, d_B) \leq \varepsilon/2$. Solving (7.11) to additive precision $\varepsilon/2$ therefore gives an additive- ε approximation by (7.10).

The two semidefinite constraints have dimensions $(d_A d_B)^2 + 1$ and

$$L_{d_A, d_B, n} = \binom{d_A + n - 1}{n} \binom{d_B + n - 1}{n},$$

respectively. Standard algorithms for semidefinite programs, together with the binomial estimate and $n = O(1/\varepsilon^2)$, give (7.7).

For the promise problem, increase n by a constant factor so that $\delta_n^{\text{HS}}(d_A, d_B) \leq \varepsilon/4$ and solve to precision $\varepsilon/8$. The relaxation value is zero when ρ is separable and is at least $3\varepsilon/4$ when $D_{\text{HS}}(\rho, \text{SEP}(d_A, d_B)) \geq \varepsilon$. Thus the numerical output is at most $\varepsilon/8$ in the separable case and at least $5\varepsilon/8$ in the far case. Comparing it with $\varepsilon/2$ distinguishes the two cases. $\square$

For fixed additive accuracy, this improves the quasipolynomial Hilbert–Schmidt weak-membership guarantee of Brandão, Christandl, and Yard [BCY11b] to polynomial time in the local dimensions. These algorithms take explicit matrices as input and do not give a copy-efficient test for an unknown quantum state. Since Hilbert–Schmidt distance is not contractive under general quantum channels, they do not replace trace- or LOCC-norm separability testing.

Acknowledgment

The authors used GPT-5.6 Pro for exploratory analysis, idea testing, editorial polishing, and assistance with Lean formalization. The authors reviewed and finalized all mathematical claims and proofs, and take full responsibility for the correctness, exposition, and attribution in the final manuscript.

References

- [ABD⁺08] Scott Aaronson, Salman Beigi, Andrew Drucker, Bill Fefferman, and Peter Shor. The power of unentanglement. In *Proceedings of the 23rd IEEE Conference on Computational Complexity (CCC)*, pages 223–236, 2008.

- [AKT24] Seiseki Akibue, Go Kato, and Seiichiro Tani. On the hardness of conversion from entangled proof into separable one, 2024.
- [BBH⁺12] Boaz Barak, Fernando G. S. L. Brandão, Aram W. Harrow, Jonathan A. Kelner, David Steurer, and Yuan Zhou. Hypercontractivity, sum-of-squares proofs, and their applications. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC)*, pages 307–326, New York, NY, USA, 2012. Association for Computing Machinery.
- [BC12] Fernando G. S. L. Brandão and Matthias Christandl. Detection of multiparticle entanglement: Quantifying the search for symmetric extensions. *Physical Review Letters*, 109(16):160502, 2012.
- [BCY11a] Fernando G. S. L. Brandão, Matthias Christandl, and Jon Yard. Faithful squashed entanglement. *Communications in Mathematical Physics*, 2011.
- [BCY11b] Fernando G. S. L. Brandão, Matthias Christandl, and Jon Yard. A quasipolynomial-time algorithm for the quantum separability problem. In *Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC)*, pages 343–351. ACM, 2011.
- [BH13] Fernando G.S.L. Brandão and Aram W. Harrow. Quantum de Finetti theorems under local measurements with applications. In *Proceedings of the 45th ACM Symposium on Theory of Computing (STOC)*, 2013.
- [Bha07] Rajendra Bhatia. *Positive Definite Matrices*. Princeton Series in Applied Mathematics. Princeton University Press, 2007.
- [BKS17] Boaz Barak, Pravesh K. Kothari, and David Steurer. Quantum entanglement, sum of squares, and the log rank conjecture. In *Proceedings of the 49th Annual ACM Symposium on Theory of Computing (STOC)*, pages 975–988, New York, NY, USA, 2017. Association for Computing Machinery.
- [CFS02] Carlton M. Caves, Christopher A. Fuchs, and Rüdiger Schack. Unknown quantum states: The quantum de Finetti representation. *Journal of Mathematical Physics*, 43(9):4537–4559, 2002.
- [Chi11] Giulio Chiribella. On quantum estimation, quantum cloning and finite quantum de Finetti theorems. In *Theory of Quantum Computation, Communication, and Cryptography*, volume 6519 of *Lecture Notes in Computer Science*, pages 9–25. Springer, 2011.
- [CKMR07] Matthias Christandl, Robert König, Graeme Mitchison, and Renato Renner. One-and-a-half quantum de Finetti theorems. *Communications in Mathematical Physics*, 273(2):473–498, 2007.
- [DF80] Persi Diaconis and David Freedman. Finite exchangeable sequences. *The Annals of Probability*, 8(4):745–764, 1980.
- [DPS04] Andrew C. Doherty, Pablo A. Parrilo, and Federico M. Spedalieri. Complete family of separability criteria. *Physical Review A*, 69(2):022308, 2004.

- [FF21] Kun Fang and Hamza Fawzi. The sum-of-squares hierarchy on the sphere and applications in quantum information theory. *Mathematical Programming*, 190(1–2):331–360, 2021.
- [FH91] William Fulton and Joe Harris. *Representation Theory: A First Course*, volume 129 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1991.
- [GGJ26] William Gay and Fernando Granha Jeronimo. The collapse of unentangled stoquastic Merlin–Arthur proof systems. *arXiv:2605.16249*, 2026.
- [Gha10] Sevag Gharibian. Strong NP-hardness of the quantum separability problem. *Quantum Information and Computation*, 10(3–4):343–360, 2010.
- [GLS93] Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric Algorithms and Combinatorial Optimization*, volume 2 of *Algorithms and Combinatorics*. Springer, Berlin, second edition, 1993.
- [Gur03] Leonid Gurvits. Classical deterministic complexity of Edmonds’ problem and quantum entanglement. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC)*, pages 10–19, New York, NY, USA, 2003. Association for Computing Machinery.
- [Har13] Aram W. Harrow. The church of the symmetric subspace. *arXiv preprint arXiv:1308.6595*, 2013.
- [HJ12] Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, second edition, 2012.
- [HM76] R. L. Hudson and G. R. Moody. Locally normal symmetric states and an analogue of de Finetti’s theorem. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 33:343–351, 1976.
- [HM13] Aram W. Harrow and Ashley Montanaro. Testing product states, quantum Merlin–Arthur games and tensor optimization. *Journal of the ACM*, 60(1):1–43, 2013.
- [HNW19] Aram W Harrow, Anand Natarajan, and Xiaodi Wu. Limitations of semidefinite programs for separable states and entangled games. *Communications in Mathematical Physics*, 366(2):423–468, 2019.
- [JLW26] Fernando Granha Jeronimo, Itai Leigh, and Pei Wu. The QMA(2) universe complexity, entanglement, and optimization. *SIGACT News*, 57(1):64–99, March 2026.
- [JSZ22] Dávid Jakab, Adrian Solymos, and Zoltán Zimborás. Extendibility of Werner states. *arXiv:2208.13743*, 2022.
- [JTC17] Zhang Jiang, Alexandre B. Tacla, and Carlton M. Caves. Bosonic particle-correlated states: A nonperturbative treatment beyond mean field. *Physical Review A*, 96(2):023621, 2017.
- [JV13] Peter D. Johnson and Lorenza Viola. Compatible quantum correlations: On extension problems for Werner and isotropic states. *Physical Review A*, 88(3):032323, 2013.

- [JW24] Fernando Granha Jeronimo and Pei Wu. Dimension independent disentanglers from unentanglement and applications. In *39th Computational Complexity Conference (CCC 2024)*, volume 300 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 26:1–26:28. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2024.
- [JWX26] Fernando Granha Jeronimo, Pei Wu, and Haonan Xu. An argmax principle for sum-of-squares relaxations on the sphere. Manuscript, 2026.
- [KMY09] Hirotada Kobayashi, Keiji Matsumoto, and Tomoyuki Yamakami. Quantum Merlin–Arthur proof systems: Are multiple merlins more helpful to arthur? *Chicago Journal of Theoretical Computer Science*, 2009(3), 2009.
- [KR05] Robert König and Renato Renner. A de Finetti representation for finite symmetric quantum states. *Journal of Mathematical Physics*, 46(12):122108, 2005.
- [LNR15] Mathieu Lewin, Phan Thành Nam, and Nicolas Rougerie. Remarks on the quantum de Finetti theorem for bosonic systems. *Applied Mathematics Research eXpress*, 2015(1):48–63, 2015.
- [LS15] Ke Li and Graeme Smith. Quantum de Finetti theorem under fully-one-way adaptive measurements. *Physical Review Letters*, 114(16):160503, 2015.
- [LW26] Yupan Liu and Pei Wu. Unentangled stoquastic Merlin–Arthur proof systems: The power of unentanglement without destructive interference. *arXiv preprint arXiv:2604.27886*, 2026.
- [Mal26] Giulio Malavolta. Quantum separability in polynomial time. *arXiv:2607.23773*, 2026.
- [MWW09] William Matthews, Stephanie Wehner, and Andreas Winter. Distinguishability of quantum states under restricted families of measurements with an application to quantum data hiding. *Communications in Mathematical Physics*, 291(3):813–843, 2009.
- [NOP09] Miguel Navascués, Masaki Owari, and Martin B. Plenio. Power of symmetric extensions for entanglement detection. *Physical Review A*, 80(5):052306, 2009.
- [OO97] Andrei Okounkov and Grigori Olshanski. Shifted Schur functions. *Algebra i Analiz*, 9(2):73–146, 1997. English translation: *St. Petersburg Mathematical Journal* 9 (1998), no. 2, 239–300.
- [OW16] Ryan O’Donnell and John Wright. Efficient quantum tomography. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, STOC ’16, pages 899–912. Association for Computing Machinery, 2016.
- [Rao24] Sujit Rao. A unified approach to quantum de Finetti theorems and SoS rounding via geometric quantization, 2024.
- [Ren05] Renato Renner. *Security of Quantum Key Distribution*. PhD thesis, ETH Zurich, 2005.
- [Ren07] Renato Renner. Symmetry of large physical systems implies independence of subsystems. *Nature Physics*, 3(9):645–649, 2007.
- [Rou15] Nicolas Rougerie. De Finetti theorems, mean-field limits and Bose–Einstein condensation, 2015. Lecture notes; revised 2020.

- [Sio58] Maurice Sion. On general minimax theorems. *Pacific Journal of Mathematics*, 8(1):171–176, 1958.
- [Stø69] Erling Størmer. Symmetric states of infinite tensor products of C^* -algebras. *Journal of Functional Analysis*, 3(1):48–68, 1969.
- [SW15] Yaoyun Shi and Xiaodi Wu. Epsilon-net method for optimizations over separable states. *Theoretical Computer Science*, 598:51–63, 2015.
- [TDS03] Barbara M. Terhal, Andrew C. Doherty, and David Schwab. Symmetric extensions of quantum states and local hidden variable theories. *Physical Review Letters*, 90(15):157903, 2003.

A Rectangular argmax rounding

This appendix proves Proposition 3.10. The eigenvector-contraction step is the same as in Proposition 3.5; the remaining ingredient is an argmax estimate on a product of two complex spheres. This product-sphere estimate is the complex pure-state form of the argmax argument underlying our argmax-SoS analysis of Best Separable State [JWX26, Sec. 3].

Let $\mathcal{H}_A$ and $\mathcal{H}_B$ be finite-dimensional Hilbert spaces, and use the rectangular lift $M^{[n,m]}$ from Subsubsection 3.2.1. We use throughout the injective and projective tensor norms defined in Section 2.4.

The vectors $x^{\otimes n} \otimes y^{\otimes m}$ span $\text{Sym}^n(\mathcal{H}_A) \otimes \text{Sym}^m(\mathcal{H}_B)$. Fix integers $n, m \geq 1$ and a unit vector ψ in this space. Choose unit vectors x, y maximizing $|\langle x^{\otimes n} \otimes y^{\otimes m}, \psi \rangle|$. After adjusting the phase of x , set

$$c := \langle x^{\otimes n} \otimes y^{\otimes m}, \psi \rangle > 0, \quad (\text{A.1})$$

and contract all but one tensor factor in each block:

$$\eta_{x,y} := \left[(\langle x|^{\otimes(n-1)} \otimes I_{\mathcal{H}_A}) \otimes (\langle y|^{\otimes(m-1)} \otimes I_{\mathcal{H}_B}) \right] \psi \in \mathcal{H}_A \otimes \mathcal{H}_B. \quad (\text{A.2})$$

Lemma A.1 (Product-sphere argmax estimate). *With the notation above,*

$$\eta_{x,y} = c x \otimes y + \zeta_{x,y}, \quad \zeta_{x,y} \in x^\perp \otimes y^\perp, \quad (\text{A.3})$$

and

$$\|\zeta_{x,y}\|_\varepsilon \leq \frac{c}{\sqrt{nm}}. \quad (\text{A.4})$$

Proof. For any unit vector $p \perp x$ and $z \in \mathbb{C}$, set

$$x_z := \frac{x + zp}{\sqrt{1 + |z|^2}}, \quad g_p(z) := \langle x_z^{\otimes n} \otimes y^{\otimes m}, \psi \rangle.$$

By the choice of x, y , the function $|g_p(z)|^2$ is maximized at $z = 0$. Its first-order expansion is

$$g_p(z) = c + n\bar{z} \langle p \otimes x^{\otimes(n-1)} \otimes y^{\otimes m}, \psi \rangle + O(|z|^2).$$

The normalization factor has no linear term. Varying the phase of z therefore gives

$$\langle p \otimes x^{\otimes(n-1)} \otimes y^{\otimes m}, \psi \rangle = 0.$$

The analogous variation in the y -block gives, for every $q \perp y$,

$$\langle x^{\otimes n} \otimes q \otimes y^{\otimes(m-1)}, \psi \rangle = 0.$$

Since these identities hold for every $p \perp x$ and $q \perp y$, and $\langle x \otimes y, \eta_{x,y} \rangle = c$, they prove (A.3).

It remains to bound the component orthogonal to both x and y . If either $x^\perp$ or $y^\perp$ is zero-dimensional, then $\zeta_{x,y} = 0$ and the claim is immediate. Otherwise, fix unit vectors $p \perp x$ and $q \perp y$. For $z, w \in \mathbb{C}$, set

$$x_z := \frac{x + zp}{\sqrt{1 + |z|^2}}, \quad y_w := \frac{y + wq}{\sqrt{1 + |w|^2}}, \quad g(z, w) := \langle x_z^{\otimes n} \otimes y_w^{\otimes m}, \psi \rangle.$$

Again, $|g(z, w)|^2$ is maximized at $(z, w) = (0, 0)$. Write

$$\beta := \langle p \otimes x^{\otimes(n-1)} \otimes q \otimes y^{\otimes(m-1)}, \psi \rangle = \langle p \otimes q, \zeta_{x,y} \rangle.$$

Expanding this overlap to second order gives

$$\begin{aligned} |g(z, w)|^2 &= c^2 - c^2(n|z|^2 + m|w|^2) \\ &\quad + 2c \Re \left(\binom{n}{2} \bar{z}^2 \alpha + \binom{m}{2} \bar{w}^2 \gamma + nm \bar{z} \bar{w} \beta \right) + O((|z| + |w|)^3), \end{aligned} \quad (\text{A.5})$$

When $n \geq 2$, here $\alpha := \langle p^{\otimes 2} \otimes x^{\otimes(n-2)} \otimes y^{\otimes m}, \psi \rangle$; when $n = 1$, the corresponding term is absent. Define γ analogously for the y -block.

Fix $a, b > 0$ and choose θ so that $e^{-2i\theta} \beta = |\beta|$. For each $\phi \in [0, 2\pi]$, restrict (A.5) to

$$z = rae^{i(\theta+\phi)}, \quad w = rbe^{i(\theta-\phi)}, \quad r \in \mathbb{R}.$$

Since $r = 0$ is a local maximum, the coefficient of r^2 is nonpositive. Averaging this inequality over ϕ eliminates the terms containing α and γ , while the term containing β is unchanged. Hence

$$2cnmab|\beta| \leq c^2(na^2 + mb^2).$$

Taking $a = \sqrt{m}$ and $b = \sqrt{n}$ yields

$$|\beta| \leq \frac{c}{\sqrt{nm}}.$$

Taking the supremum over unit $p \perp x$ and $q \perp y$ proves (A.4). $\square$

Proof of Proposition 3.10. Let $\psi \in \text{Sym}^n(\mathcal{H}_A) \otimes \text{Sym}^m(\mathcal{H}_B)$ be a unit top eigenvector of $M^{[n,m]}$, with eigenvalue λ . Choose $x, y, c, \eta_{x,y}$ as in Lemma A.1. Because ψ is invariant under permutations of the $\mathcal{H}_A$ factors and, independently, of the $\mathcal{H}_B$ factors,

$$\lambda c = \langle x^{\otimes n} \otimes y^{\otimes m}, M^{[n,m]} \psi \rangle = \langle M(x \otimes y), \eta_{x,y} \rangle. \quad (\text{A.6})$$

Decompose orthogonally

$$\begin{aligned} M(x \otimes y) &= a_M(x, y)x \otimes y + \xi_{x,y} + B_{x,y}, \\ \xi_{x,y} &\in (x^\perp \otimes \mathbb{C}y) \oplus (\mathbb{C}x \otimes y^\perp), \\ B_{x,y} &\in x^\perp \otimes y^\perp. \end{aligned} \quad (\text{A.7})$$

Combining (A.6), (A.3), and (A.7) gives

$$(\lambda - a_M(x, y))c = \langle M(x \otimes y), \eta_{x,y} \rangle - a_M(x, y)c$$

$$\begin{aligned}
&= \langle a_M(x, y)x \otimes y + \xi_{x,y} + B_{x,y}, cx \otimes y + \zeta_{x,y} \rangle - a_M(x, y)c \\
&= \langle B_{x,y}, \zeta_{x,y} \rangle.
\end{aligned} \tag{A.8}$$

By Lemma A.1 and (2.18),

$$|\lambda - a_M(x, y)| \leq \frac{1}{\sqrt{nm}} \|B_{x,y}\|_\pi.$$

Since $B_{x,y} = (P_{x^\perp} \otimes P_{y^\perp})M(x \otimes y)$, this proves (3.30). $\square$

B Lower bounds for the de Finetti theorems

B.1 Optimality of the t -site bounds

The two-site rectangular Werner examples of Christandl, König, Mitchison, and Renner, together with their symmetric purifications, already show that the dimension dependence of our bounds is optimal [CKMR07, Lemma II.5 and Corollary III.9]. By contractivity under partial trace, the same conclusion holds for every fixed marginal order t . We use their rectangular families below only to prove the new point: the linear dependence on t is also necessary.

Fix $N = ar$. Let $\omega_{ar,(ar)}$ be the rectangular Werner state: the normalized projector onto the (a^r) Schur–Weyl block of $(\mathbb{C}^r)^{\otimes ar}$. Let $\psi_{a,r}$ be its canonical symmetric purification in $\text{Sym}^{ar}(\mathbb{C}^r \otimes \mathbb{C}^r)$, equivalently the unit vector in the one-dimensional Cauchy summand

$$S_{(ar)}(\mathbb{C}^r) \otimes S_{(ar)}(\mathbb{C}^r).$$

We write $\gamma_{a,r}^{(t)}$ and $\omega_{ar,(ar)}^{(t)}$ for their respective t -site marginals. These are the rectangular families analyzed in [CKMR07]; the calculation below extracts their dependence on t .

B.1.1 Exact Schur-label distances

The same rectangular families also detect the dependence on the marginal order. The cleanest formulation uses the Schur label of the t -site marginal. For a partition $\mu \vdash t$, let f^μ be the dimension of the Specht module V_μ , let μ' be the conjugate partition, and write

$$c(\square) = j - i, \quad h_\mu(\square) = \mu_i - j + \mu'_j - i + 1$$

for the content and hook length of a box $\square = (i, j) \in \mu$. We use

$$(x)_t = x(x-1) \cdots (x-t+1)$$

for the falling factorial.

For $\mu \vdash t$ with $\ell(\mu) \leq r$, define

$$p_{a,r,t}(\mu) := \frac{(f^\mu)^2}{t!} \frac{\prod_{\square \in \mu} (a - c(\square))(r + c(\square))}{(ar)_t}, \tag{B.1}$$

$$q_{r,t}(\mu) := \frac{(f^\mu)^2}{t!} \frac{\prod_{\square \in \mu} (r + c(\square))}{r^t}, \tag{B.2}$$

$$R_{a,r,t}(\mu) := \frac{p_{a,r,t}(\mu)}{q_{r,t}(\mu)} = \frac{\prod_{\square \in \mu} \left(1 - \frac{c(\square)}{a}\right)}{\prod_{j=0}^{t-1} \left(1 - \frac{j}{ar}\right)}. \quad (\text{B.3})$$

A partition with $\mu_1 > a$ has zero weight in (B.1), since its first row contains a box of content a .

Lemma B.1 (Rectangular branching weights). *Under the Cauchy decomposition*

$$\text{Sym}^t(\mathbb{C}^r \otimes \mathbb{C}^r) = \bigoplus_{\substack{\mu \vdash t \\ \ell(\mu) \leq r}} S_\mu(\mathbb{C}^r) \otimes S_\mu(\mathbb{C}^r),$$

the block weights of $\gamma_{a,r}^{(t)}$ are $p_{a,r,t}(\mu)$. Under the Schur–Weyl decomposition of $(\mathbb{C}^r)^{\otimes t}$, the block weights of $\omega_{ar,(a^r)}^{(t)}$ are the same numbers. Moreover,

$$p_{a,r,t}(\mu) = \frac{\dim S_\mu(\mathbb{C}^r) \dim S_{\mu'}(\mathbb{C}^a)}{\binom{ar}{t}}. \quad (\text{B.4})$$

In particular, the numbers in (B.1) form a probability distribution.

Proof. Both restrictions are governed by the coherent down-transition measure from the rectangular Specht module $V_{(a^r)}$:

$$p_{a,r,t}(\mu) = \frac{f^\mu f^{(a^r)}/\mu}{f^{(a^r)}}. \quad (\text{B.5})$$

For the Werner state, this follows by restricting the normalized identity on $V_{(a^r)}$.

We give the corresponding argument for $\psi_{a,r}$. Set $\lambda = (a^r)$ and choose a real orthonormal model of the Specht module V_λ . In the Schur–Weyl realization of the Cauchy summand, the Specht factor of $\psi_{a,r}$ is the normalized diagonal invariant

$$\Omega_\lambda := \frac{1}{\sqrt{f^\lambda}} \sum_{j=1}^{f^\lambda} e_j \otimes e_j \in V_\lambda \otimes V_\lambda.$$

Upon restriction to S_t ,

$$V_\lambda \downarrow_{S_t} \cong \bigoplus_{\substack{\mu \vdash t \\ \mu \subseteq \lambda}} V_\mu \otimes M_{\lambda/\mu}, \quad \dim M_{\lambda/\mu} = f^{\lambda/\mu}.$$

Choose the basis (e_j) adapted to this orthogonal decomposition and let Q_μ be the projection onto $V_\mu \otimes M_{\lambda/\mu}$. Then

$$\|(Q_\mu \otimes Q_\mu)\Omega_\lambda\|^2 = \frac{\text{rank } Q_\mu}{f^\lambda} = \frac{f^\mu f^{\lambda/\mu}}{f^\lambda}.$$

Because Ω_λ is invariant under the diagonal action $g \otimes g$ of $\mathfrak{S}_t$ and Q_μ commutes with that action, $(Q_\mu \otimes Q_\mu)\Omega_\lambda$ is itself diagonal invariant. On the $\mu \otimes \mu$ isotypic component, the Cauchy-block projector for the first t paired factors is the projection onto $(V_\mu \otimes V_\mu)^{\text{diag}(\mathfrak{S}_t)}$, tensored with the identity on the multiplicity spaces. It therefore acts as the identity on $(Q_\mu \otimes Q_\mu)\Omega_\lambda$, while the projectors with other labels annihilate this component. Consequently, its expectation in the global

vector, or equivalently in the t -site marginal, is the displayed squared norm. This proves (B.5) for $\gamma_{a,r}^{(t)}$ as well.

The shifted-Schur skew-dimension formula, specialized to a rectangle, gives [OO97, Theorem 8.1]

$$\frac{f^{(a^r)/\mu}}{f^{(a^r)}} = \frac{1}{(ar)_t} \prod_{\square \in \mu} \frac{(a - c(\square))(r + c(\square))}{h_\mu(\square)}.$$

Using the hook-length identity $f^\mu = t! / \prod_{\square \in \mu} h_\mu(\square)$ gives (B.1). Applying the hook-content formula to $S_\mu(\mathbb{C}^r)$ and $S_{\mu'}(\mathbb{C}^a)$ gives (B.4). Summing the latter identity over μ is the dual Cauchy decomposition of $\bigwedge^t(\mathbb{C}^a \otimes \mathbb{C}^r)$, whose dimension is $\binom{ar}{t}$. $\square$

For a probability vector $x = (x_1, \dots, x_r)$, write

$$\text{SW}_{t,x}(\mu) := f^\mu s_\mu(x). \quad (\text{B.6})$$

This is the Schur–Weyl distribution of a length- t random word with letter law x . If $u \in \mathbb{C}^r \otimes \mathbb{C}^r$ has squared Schmidt coefficients x , then twirling $|u\rangle\langle u|^{\otimes t}$ under $U(r) \times U(r)$ gives Cauchy-block weights $\text{SW}_{t,x}$. Likewise, if a density operator σ has spectrum x , then twirling $\sigma^{\otimes t}$ under $U(r)$ gives Schur–Weyl block weights $\text{SW}_{t,x}$. For the uniform vector $\bar{x} = (1/r, \dots, 1/r)$, the hook-content formula gives

$$\text{SW}_{t,\bar{x}}(\mu) = q_{r,t}(\mu).$$

Proposition B.2 (Exact higher-marginal distance). *For every $a, r \in \mathbb{N}$, $r \geq 2$, and $2 \leq t \leq ar$,*

$$D_{\text{tr}}(\gamma_{a,r}^{(t)}, \text{SEP}_{\text{H}}^{(t)}(r^2)) = \frac{1}{2} \sum_{\substack{\mu \vdash t \\ \ell(\mu) \leq r}} q_{r,t}(\mu) |R_{a,r,t}(\mu) - 1|, \quad (\text{B.7})$$

$$D_{\text{tr}}(\omega_{ar,(ar)}^{(t)}, \text{SEP}_{\text{P}}^{(t)}(r)) = \frac{1}{2} \sum_{\substack{\mu \vdash t \\ \ell(\mu) \leq r}} q_{r,t}(\mu) |R_{a,r,t}(\mu) - 1|. \quad (\text{B.8})$$

A closest invariant Hartree mixture in the first line is the twirl of $|\Phi_r\rangle\langle\Phi_r|^{\otimes t}$, where $\Phi_r = r^{-1/2} \sum_i e_i \otimes e_i$; a closest invariant state of the form $\tau^{\otimes t}$ in the second line is $(I/r)^{\otimes t}$. An optimal separating operator $0 \preceq A \preceq I$ is the sum of the Schur projectors indexed by

$$\mathcal{L}_{a,r,t} := \{\mu : R_{a,r,t}(\mu) \geq 1\}. \quad (\text{B.9})$$

Proof. Twirling is trace-norm contractive, fixes the two target states, and preserves the respective approximation sets. On the Cauchy side, the $U(r) \times U(r)$ decomposition is multiplicity free. On the Werner side, the target and every twirled approximant commute with both the $U(r)$ and S_t actions and are therefore scalar on each joint Schur–Weyl block $S_\mu(\mathbb{C}^r) \otimes V_\mu$. Consequently, trace distance after twirling is the total-variation distance between the normalized block weights. It remains to show that $q_{r,t}$ is closest among convex combinations of $\text{SW}_{t,x}$.

Order partitions by dominance, with $\nu \supseteq \mu$ meaning that ν is more dominant. The likelihood ratio in (B.3) is decreasing in this order. Indeed, a covering move transfers one box from a lower row to a higher row and strictly increases its content; the corresponding factor $a - c(\square)$ decreases, while all other factors are unchanged. If the move leaves the $a \times r$ rectangle, the new likelihood ratio is zero. Consequently, $\mathcal{L}_{a,r,t}$ is a lower ideal in dominance order.

Every probability vector x majorizes $\bar{x}$. The Schur–Weyl majorization coupling states that if x majorizes y , one may couple $\mu_x \sim \text{SW}_{t,x}$ and $\mu_y \sim \text{SW}_{t,y}$ so that $\mu_x \succeq \mu_y$ almost surely [OW16, Thm. 1.11]. Therefore every lower ideal $\mathcal{L}$ satisfies

$$\text{SW}_{t,x}(\mathcal{L}) \leq q_{r,t}(\mathcal{L}). \quad (\text{B.10})$$

The same inequality holds for every convex combination of Schur–Weyl distributions.

Let $p = p_{a,r,t}$ and $q = q_{r,t}$. The set $\mathcal{L}_{a,r,t} = \{p \geq q\}$ is the Neyman–Pearson set for p versus q . For every feasible invariant approximant with block distribution s ,

$$\begin{aligned} D_{\text{TV}}(p, s) &\geq p(\mathcal{L}_{a,r,t}) - s(\mathcal{L}_{a,r,t}) \\ &\geq p(\mathcal{L}_{a,r,t}) - q(\mathcal{L}_{a,r,t}) = D_{\text{TV}}(p, q). \end{aligned}$$

Since q itself is feasible in both problems, equality follows, as does the optimal separating operator. $\square$

The exact formula reduces optimality to a scalar Schur-label calculation. We next extract a lower bound showing that the linear dependence on t in Theorem 4.1 is necessary. Let

$$\text{ct}(\mu) := \sum_{\square \in \mu} c(\square), \quad M_t := \binom{t}{2}. \quad (\text{B.11})$$

The transposition class sum $\mathsf{T}_t = \sum_{1 \leq i < j \leq t} (ij)$ acts on V_μ as the scalar $\text{ct}(\mu)$.

B.1.2 Optimal dependence on the marginal order

Lemma B.3 (Plancherel content fluctuations). *Let $t \geq 2$, and let μ have Plancherel law $\text{Pl}_t(\mu) = (f^\mu)^2/t!$. There is a universal constant $c_0 > 0$ such that*

$$c_0 t \leq \mathbb{E}_{\text{Pl}_t} |\text{ct}(\mu)| \leq \sqrt{M_t}. \quad (\text{B.12})$$

Proof. Use the normalized trace τ_{reg} on the regular representation of $\mathfrak{S}_t$, and write

$$\|X\|_{p,\text{reg}} := \tau_{\text{reg}}(|X|^p)^{1/p}.$$

Distinct group elements are orthonormal in the associated L_2 inner product, so

$$\mathbb{E} \text{ct}(\mu) = 0, \quad \mathbb{E} \text{ct}(\mu)^2 = M_t.$$

Moreover, $\mathbb{E} \text{ct}(\mu)^4$ is the number of ordered quadruples of transpositions whose product is the identity. Associate to such a quadruple the four-edge multigraph formed by its transpositions. A vertex of degree one cannot occur in a word whose product is the identity. Hence every nonisolated vertex has degree at least two, so the graph uses at most four vertices. There are only $O(t^4)$ such ordered quadruples, and therefore

$$\mathbb{E} \text{ct}(\mu)^4 \leq C_0 t^4$$

for a universal C_0 . Hölder’s inequality gives

$$\mathbb{E} |\text{ct}(\mu)| \geq \frac{(\mathbb{E} \text{ct}(\mu)^2)^{3/2}}{(\mathbb{E} \text{ct}(\mu)^4)^{1/2}} \geq c_0 t.$$

The upper bound is Cauchy–Schwarz. $\square$

Lemma B.4 (Perturbative separation of the rectangular law). *There are universal constants $c_1, C_1 > 0$ such that, for every integer $t \geq 2$, whenever*

$$a \geq C_1 t, \quad r \geq C_1 t^2, \quad (\text{B.13})$$

one has

$$\frac{1}{2} \sum_{\mu} q_{r,t}(\mu) |R_{a,r,t}(\mu) - 1| \geq c_1 \frac{t}{a}. \quad (\text{B.14})$$

Proof. We take $C_1 \geq 2$ throughout. In particular, $a \geq t$ and $r \geq t$, so every partition of t fits inside the $a \times r$ rectangle and all Schur labels appearing under Plancherel measure lie in the support of $q_{r,t}$. The distribution $q_{r,t}$ is the RSK shape of a uniform random word in $[r]^t$. With probability

$$\alpha_{r,t} := \frac{(r)_t}{r^t} \geq 1 - \frac{M_t}{r},$$

all letters are distinct. Conditional on this event, their relative order is a uniform permutation, so the shape has Plancherel law. Thus

$$q_{r,t} = \alpha_{r,t} \text{Pl}_t + (1 - \alpha_{r,t}) \nu_{r,t} \quad (\text{B.15})$$

for some probability distribution $\nu_{r,t}$. Under (B.13), $\alpha_{r,t} \geq 1/2$.

Let $J_k = \sum_{i < k} (ik)$ be the Jucys–Murphy elements and $\mathbb{T}_t = \sum_k J_k$. The Jucys identity gives

$$\prod_{k=1}^t \left(1 - \frac{J_k}{a}\right) = \sum_{\pi \in \mathfrak{S}_t} \text{sgn}(\pi) a^{-|\pi|} \pi, \quad |\pi| := t - \#\text{cycles}(\pi). \quad (\text{B.16})$$

On V_{μ} , the left-hand side has eigenvalue $\prod_{\square \in \mu} (1 - c(\square)/a)$. Put

$$B_{a,r,t} := \prod_{j=0}^{t-1} \left(1 - \frac{j}{ar}\right).$$

Thus the central element

$$\mathbb{R}_{a,r,t} := B_{a,r,t}^{-1} \prod_{k=1}^t \left(1 - \frac{J_k}{a}\right)$$

has eigenvalue $R_{a,r,t}(\mu)$ on V_{μ} .

Write

$$\prod_{k=1}^t \left(1 - \frac{J_k}{a}\right) = I - \frac{\mathbb{T}_t}{a} + \mathbb{H}_a, \quad \mathbb{H}_a := \sum_{|\pi| \geq 2} \text{sgn}(\pi) a^{-|\pi|} \pi.$$

If n_k is the number of permutations of transposition length k , then $n_k \leq M_t^k$, because every such permutation has a factorization into k transpositions. Orthogonality of the group basis in regular L_2 therefore gives, provided $a^2 \geq 2M_t$,

$$\|\mathbb{H}_a\|_{2,\text{reg}}^2 = \sum_{k \geq 2} \frac{n_k}{a^{2k}} \leq 2 \frac{M_t^2}{a^4}. \quad (\text{B.17})$$

Also,

$$1 - B_{a,r,t} \leq \frac{M_t}{ar}, \quad B_{a,r,t}^{-1} \leq 2$$

when the constants in (B.13) are sufficiently large. Since

$$\|\mathsf{T}_t\|_{1,\text{reg}} \leq \|\mathsf{T}_t\|_{2,\text{reg}} = \sqrt{M_t},$$

(B.17) implies

$$\|\mathsf{R}_{a,r,t} - I + \mathsf{T}_t/a\|_{1,\text{reg}} \leq \frac{2M_t}{ar} \left(1 + \frac{\sqrt{M_t}}{a}\right) + \frac{2\sqrt{2}M_t}{a^2}. \quad (\text{B.18})$$

Choose C_1 so that the right-hand side is at most $(c_0/2)t/a$, where c_0 is from Lemma B.3. Then

$$\begin{aligned} \mathbb{E}_{\text{Pl}_t} |R_{a,r,t}(\mu) - 1| &= \|\mathsf{R}_{a,r,t} - I\|_{1,\text{reg}} \\ &\geq \frac{1}{a} \mathbb{E}_{\text{Pl}_t} |\text{ct}(\mu)| - \|\mathsf{R}_{a,r,t} - I + \mathsf{T}_t/a\|_{1,\text{reg}} \\ &\geq \frac{c_0 t}{2a}. \end{aligned}$$

Combining this with (B.15) and $\alpha_{r,t} \geq 1/2$ proves (B.14). $\square$

Theorem B.5 (Optimal t -dependence of the rectangular examples). *There is a universal constant $C > 0$ such that, for every $a, r \in \mathbb{N}$ and $2 \leq t \leq ar$ satisfying $r \geq Ct^2$,*

$$D_{\text{tr}}(\gamma_{a,r}^{(t)}, \text{SEP}_{\text{H}}^{(t)}(r^2)) = \Omega\left(\min\left\{1, \frac{t}{a}\right\}\right) = \Omega\left(\min\left\{1, \frac{t\sqrt{d}}{N}\right\}\right), \quad d = r^2, \quad (\text{B.19})$$

$$D_{\text{tr}}(\omega_{ar,(ar)}^{(t)}, \text{SEP}_{\text{P}}^{(t)}(r)) = \Omega\left(\min\left\{1, \frac{t}{a}\right\}\right) = \Omega\left(\min\left\{1, \frac{td}{N}\right\}\right), \quad d = r. \quad (\text{B.20})$$

Together with Theorem 4.1, these estimates match up to universal constants. In particular, the linear dependence on the marginal order in (4.1) and (4.2) cannot be improved in general.

Proof. By Proposition B.2, the two distances are the same scalar quantity. If $a \geq C_1 t$, apply Lemma B.4.

Suppose $a < C_1 t$. Distance to the relevant de Finetti set cannot increase under partial trace, so the t -site distance is at least the s -site distance for every $2 \leq s \leq t$. If a is larger than a sufficiently large universal constant, take

$$s = \left\lfloor \frac{a}{2C_1} \right\rfloor.$$

Then $2 \leq s \leq t$, $a \geq C_1 s$, and $r \geq Ct^2 \geq C_1 s^2$ after increasing C . The perturbative lemma gives a universal constant lower bound because $s/a \geq 1/(4C_1)$. If a is bounded, use the two-site marginal instead. Proposition B.2, specialized to two sites, gives for both rectangular families

$$D_t \geq D_2 = \frac{r^2 - 1}{2r(ar - 1)} \geq \frac{3}{8a},$$

which is again bounded below by a universal constant. This proves the lower bounds in the regime $t/a \gtrsim 1$ and hence the displayed capped estimates. The matching upper bounds follow from Theorem 4.1, together with the trace-distance diameter bound, after substituting $N = ar$ and, respectively, $d = r^2$ or $d = r$. $\square$

B.2 Dimension-free bosonic Hilbert–Schmidt lower bounds

Theorem B.6 (Dimension-free bosonic Hilbert–Schmidt lower bound). *As $N \rightarrow \infty$,*

$$\sup_{d \geq 2} D_{\text{H}}^{\text{HS}}\left(\text{BOS}_N^{(2)}(d), \text{SEP}_H^{(2)}(d)\right) = \Omega(N^{-1/2}).$$

The lower bound is included for completeness. It adapts the Jiang–Tacla–Caves bosonic pair construction, based on symmetrized powers of a two-site vector [JTC17]. We give the short calculation in symmetric-tensor notation.

Proof of Theorem B.6. Fix $m \geq 1$, let $\mathcal{H}_m = \mathbb{C}^{m+1}$ with orthonormal basis $e_0, e_1, \dots, e_m$, and define

$$b_m := \frac{1}{\sqrt{m}} \sum_{i=1}^m e_i^{\otimes 2}, \quad A_m := |e_0^{\otimes 2}\rangle \langle b_m| + |b_m\rangle \langle e_0^{\otimes 2}|. \quad (\text{B.21})$$

The vectors $e_0^{\otimes 2}$ and b_m are orthonormal, so

$$\|A_m\|_2 = \sqrt{2}. \quad (\text{B.22})$$

If $u = \alpha e_0 + v$ with $v \perp e_0$, then

$$|\langle u^{\otimes 2} | A_m | u^{\otimes 2} \rangle| \leq \frac{2}{\sqrt{m}} |\alpha|^2 \|v\|^2 \leq \frac{1}{2\sqrt{m}}.$$

Equality is attained at $u = (e_0 + e_1)/\sqrt{2}$, and therefore

$$h(A_m) = \frac{1}{2\sqrt{m}}. \quad (\text{B.23})$$

Thus the Hartree value vanishes as $m \rightarrow \infty$.

For $q \geq 0$, set

$$x_q^{(m)} := \Pi_{\text{sym}}^{(2q)} b_m^{\otimes q}, \quad \chi_q^{(m)} := \frac{x_q^{(m)}}{\|x_q^{(m)}\|}, \quad (\text{B.24})$$

with $x_0^{(m)} = \chi_0^{(m)} = 1$. Expanding in the orthonormal monomial basis of $\text{Sym}^{2q} \mathbb{C}^m$ gives

$$\|x_q^{(m)}\|^2 = \prod_{j=0}^{q-1} \frac{m+2j}{m(2j+1)}. \quad (\text{B.25})$$

Let $C_{b_m} \xi := (\langle b_m | \otimes I) \xi$ denote the partial inner product with b_m on two tensor factors. Both b_m and $\chi_{q+1}^{(m)}$ are invariant under the natural $O(m)$ action, and C_{b_m} is $O(m)$ -equivariant. Hence $C_{b_m} \chi_{q+1}^{(m)}$ lies in $(\text{Sym}^{2q} \mathbb{C}^m)^{O(m)}$. This invariant subspace is one-dimensional, spanned by $\chi_q^{(m)}$; equivalently, every $O(m)$ -invariant homogeneous polynomial of degree $2q$ is a scalar multiple of $(z_1^2 + \dots + z_m^2)^q$. Thus

$$C_{b_m} \chi_{q+1}^{(m)} = \kappa_q \chi_q^{(m)}.$$

On symmetric tensors, the adjoint of C_{b_m} is $\xi \mapsto \Pi_{\text{sym}}^{(2q+2)}(b_m \otimes \xi)$. Taking the inner product with $\chi_{q+1}^{(m)}$ and using (B.25) gives

$$\kappa_q = \frac{\|x_{q+1}^{(m)}\|}{\|x_q^{(m)}\|} = \sqrt{\frac{m+2q}{m(2q+1)}}, \quad C_{b_m} \chi_{q+1}^{(m)} = \sqrt{\frac{m+2q}{m(2q+1)}} \chi_q^{(m)}. \quad (\text{B.26})$$

For $0 \leq q \leq \lfloor N/2 \rfloor$, define the normalized symmetric tensor

$$\psi_q^{(m,N)} := \sqrt{\frac{N!}{(N-2q)!(2q)!}} \Pi_{\text{sym}}^{(N)} \left(e_0^{\otimes(N-2q)} \otimes \chi_q^{(m)} \right) \in \text{Sym}^N \mathcal{H}_m. \quad (\text{B.27})$$

These vectors are mutually orthogonal. Fix $0 \leq q \leq \lfloor (N-2)/2 \rfloor$ and put $r = N - 2q$. After contracting either $e_0^{\otimes 2}$ from $\psi_q^{(m,N)}$ or b_m from $\psi_{q+1}^{(m,N)}$, the normalized remaining tensor is the same; denote it by $\theta_q^{(m,N)}$. Directly from (B.27) and (B.26),

$$(\langle e_0^{\otimes 2} | \otimes I) \psi_q^{(m,N)} = \sqrt{\frac{r(r-1)}{N(N-1)}} \theta_q^{(m,N)}, \quad (\text{B.28})$$

$$(\langle b_m | \otimes I) \psi_{q+1}^{(m,N)} = \frac{2\sqrt{(q+1)(q+m/2)}}{\sqrt{mN(N-1)}} \theta_q^{(m,N)}. \quad (\text{B.29})$$

Since the vectors in (B.27) are symmetric, every summand in the normalized lift has the same matrix element. Equations (B.28) and (B.29) therefore give

$$\langle \psi_{q+1}^{(m,N)} | A_m^{[N]} | \psi_q^{(m,N)} \rangle = t_{q,m,N}, \quad (\text{B.30})$$

where

$$t_{q,m,N} := \frac{2}{N(N-1)\sqrt{m}} \sqrt{(q+1)(q+m/2)} \sqrt{(N-2q)(N-2q-1)}. \quad (\text{B.31})$$

The diagonal matrix elements vanish. Hence the bosonic state

$$\rho_{q,m,N} := \frac{1}{2} |\psi_q^{(m,N)} + \psi_{q+1}^{(m,N)}\rangle \langle \psi_q^{(m,N)} + \psi_{q+1}^{(m,N)}|$$

satisfies $\text{Tr}(A_m \rho_{q,m,N}^{(2)}) = t_{q,m,N}$.

Now let $\tau \in \text{SEP}_H^{(2)}(m+1)$. Hilbert–Schmidt Cauchy–Schwarz, (B.22), and $\text{Tr}(A_m \tau) \leq h(A_m)$ give

$$\left\| \rho_{q,m,N}^{(2)} - \tau \right\|_2 \geq \frac{\text{Tr}(A_m \rho_{q,m,N}^{(2)}) - \text{Tr}(A_m \tau)}{\sqrt{2}} \geq \frac{t_{q,m,N}}{\sqrt{2}} - \frac{1}{2\sqrt{2m}}.$$

Taking the infimum over τ proves

$$D_H^{\text{HS}} \left(\text{BOS}_N^{(2)}(m+1), \text{SEP}_H^{(2)}(m+1) \right) \geq \frac{t_{q,m,N}}{\sqrt{2}} - \frac{1}{2\sqrt{2m}}. \quad (\text{B.32})$$

For fixed N and q ,

$$\lim_{m \rightarrow \infty} \frac{t_{q,m,N}}{\sqrt{2}} = \frac{\sqrt{(q+1)(N-2q)(N-2q-1)}}{N(N-1)}.$$

Since the second term in (B.32) tends to zero, taking the supremum over $d = m+1$ gives, for every admissible q ,

$$\sup_{d \geq 2} D_H^{\text{HS}} \left(\text{BOS}_N^{(2)}(d), \text{SEP}_H^{(2)}(d) \right) \geq \frac{\sqrt{(q+1)(N-2q)(N-2q-1)}}{N(N-1)}.$$

For $N \geq 6$, choose $q = \lfloor N/6 \rfloor$. Then

$$q+1 \geq \frac{N}{6}, \quad N-2q \geq \frac{2N}{3}, \quad N-2q-1 \geq \frac{N}{2}.$$

Substitution into the preceding display gives the claimed $\Omega(N^{-1/2})$ lower bound. $\square$

Theorem B.7 (Hilbert–Schmidt lower bound for two-sided Bose-symmetric extendible states). *As $n \rightarrow \infty$,*

$$\sup_{d \geq 2} D_{\text{H}}^{\text{HS}}\left(\text{BEXT}_{n,n}(d, d), \text{SEP}(d, d)\right) = \Omega(n^{-1/2}).$$

Proof. Use the witness A_m from (B.21). If $x = \alpha e_0 + v$ and $y = \beta e_0 + w$, then

$$|\langle x \otimes y | A_m | x \otimes y \rangle| \leq \frac{2}{\sqrt{m}} |\alpha| \|v\| |\beta| \|w\| \leq \frac{1}{2\sqrt{m}}.$$

Equality holds for $x = y = (e_0 + e_1)/\sqrt{2}$, and hence

$$h_{\text{SEP}}(A_m) = \frac{1}{2\sqrt{m}} = h(A_m). \quad (\text{B.33})$$

Set $q = \lfloor n/3 \rfloor$ and take the bosonic $2n$ -site state $\rho_{q,m,2n}$ from the proof of Theorem B.6. Grouping its first and last n sites gives a two-sided Bose-symmetric extension whose cross-block marginal is $\rho_{q,m,2n}^{(2)}$. The Hilbert–Schmidt-normalized witness $A_m/\sqrt{2}$ and (B.33) give the finite- m estimate (B.32), now with $N = 2n$ and the separable set in place of $\text{SEP}_{\text{H}}^{(2)}(m+1)$. With $q = \lfloor n/3 \rfloor$, the right-hand side tends, as $m \rightarrow \infty$, to at least $1/(6\sqrt{n})$ by the calculation in the proof of Theorem B.6. For each $n \geq 3$, choose a sufficiently large finite m and set $d_n = m+1$. The resulting distance is at least $1/(12\sqrt{n})$, and taking the supremum over finite dimensions proves the claim. $\square$