

Why the full-slice $ABAB$ transport carries an extra $2\sqrt{\zeta}$

August 24, 2026

1 The assertion

We examine the transport from the full-slice $ABAB$ term to the evaluated-slice $ABAB$ term in the proof of commutativity of G in [JNV⁺20, Section *Commutativity of G*].¹ The context is the main commutativity theorem for G ,² where the commutativity difference is expanded as

$$\mathbf{E}_{\mathbf{x}, \mathbf{y}} \sum_{g, h} \langle \psi | G_g^{\mathbf{x}} G_h^{\mathbf{y}} G_g^{\mathbf{x}} \otimes I | \psi \rangle - \mathbf{E}_{\mathbf{x}, \mathbf{y}} \sum_{g, h} \langle \psi | G_g^{\mathbf{x}} G_h^{\mathbf{y}} G_g^{\mathbf{x}} G_h^{\mathbf{y}} \otimes I | \psi \rangle. \quad (1)$$

The second term is the $ABAB$ term. The proof compares this term in the full-slice and evaluated-slice regimes.

The paper’s `closenessOfIP` step asserts

$$\mathbf{E}_{\mathbf{x}, \mathbf{y}} \sum_{g, h} \langle \psi | G_g^{\mathbf{x}} G_h^{\mathbf{y}} G_g^{\mathbf{x}} G_h^{\mathbf{y}} \otimes I | \psi \rangle \approx_{\sqrt{\zeta}} \mathbf{E}_{\mathbf{x}, \mathbf{y}} \sum_{g, h} \langle \psi | G_h^{\mathbf{y}} G_g^{\mathbf{x}} G_h^{\mathbf{y}} \otimes G_g^{\mathbf{x}} | \psi \rangle. \quad (2)$$

The right-hand side has the form $BAB \otimes G$ and is manifestly positive semidefinite as a tensor product of operators. The `closenessOfIP` proposition is applied with $A = G_g^{\mathbf{x}} \otimes I$, $B = I \otimes G_g^{\mathbf{x}}$, and $C = G_h^{\mathbf{y}} G_g^{\mathbf{x}} G_h^{\mathbf{y}}$; the closeness between A and B follows from self-consistency at cost ζ (hence $\sqrt{\zeta}$ under the `closenessOfIP` square root).

2 The formal scalar endpoint

In the formalization the top-level $ABAB$ averages are defined as *scalar* expectations on a single register.³ Concretely, the full-slice average is

$$\mathbf{E}_{\mathbf{x}, \mathbf{y}} \sum_{g, h} \text{ev}_{\psi}(\text{leftTensor}(G_g^{\mathbf{x}} G_h^{\mathbf{y}} G_g^{\mathbf{x}} G_h^{\mathbf{y}})),$$

¹The discrepancy is tracked in [GitHub issue #713](#). In the paper T_EX source consulted here, the transport appears in `references/ldt-paper/commutativity-G.tex`, lines 332–401, and the corresponding Lean file is `MIPStarRE/LDT/Commutativity/Main/Auxiliary.lean`.

²Theorem label `thm:com-main` in the cited paper.

³The two declarations are `fullSliceABABAvg` and `evaluatedSliceABABAvg`, in `MIPStarRE/LDT/Commutativity/Transport/FullSlice.lean`.

and the evaluated-slice average is

$$\mathbf{E}_{\mathbf{u}, \mathbf{v}, \mathbf{x}, \mathbf{y}} \sum_{a, b} \text{ev}_{\psi}(\text{leftTensor}(G_{[g(\mathbf{u})=a]}^{\mathbf{x}} G_{[h(\mathbf{v})=b]}^{\mathbf{y}} G_{[g(\mathbf{u})=a]}^{\mathbf{x}} G_{[h(\mathbf{v})=b]}^{\mathbf{y}})).$$

In the evaluation ev_{ψ} , a `leftTensor` is just an abbreviation for $X \otimes I$. Thus both endpoints are scalar: they average $\text{tr}(\rho_{\psi}(X \otimes I))$.

3 The scalar/tensor bridge

The paper’s `closenessOfIP` call in (2) produces a *tensor* endpoint $BAB \otimes G$, not a scalar $BAB \otimes I$. To connect the formal scalar endpoints to this tensor endpoint, the formalization inserts a scalar-to-tensor bridge.⁴ This bridge replaces I on the right register with $G_g^{\mathbf{x}}$ (or $G_{[h(\mathbf{v})=b]}^{\mathbf{y}}$) and is itself an application of `closenessOfIP`, costing an additional $\sqrt{\zeta}$.

Write A_{full}^{ABAB} and A_{eval}^{ABAB} for the two scalar averages displayed above. The full transport lemma⁵ then assembles the chain

$$\begin{aligned} |A_{\text{full}}^{ABAB} - A_{\text{eval}}^{ABAB}| &\leq \underbrace{\frac{md}{q} + \sqrt{\zeta}}_{\text{paper prefix}} + \underbrace{2\sqrt{\zeta}}_{\text{scalar/tensor bridge}} + \underbrace{\frac{md}{q} + \sqrt{\zeta}}_{\text{paper tail}} \\ &= 2 \cdot \frac{md}{q} + 4\sqrt{\zeta}. \end{aligned}$$

Without the scalar/tensor bridge, the bound would be $2md/q + 2\sqrt{\zeta}$, matching the paper’s implicit transport cost before the commutativity-bridge step. The extra $2\sqrt{\zeta}$ is the price of applying `closenessOfIP` twice more to shuttle between scalar and tensor endpoints.

4 Impact on the final estimate

The factor $2\sqrt{\zeta}$ is absorbed into the existing error budget. The paper’s total transport error without scalar/tensor overhead is $2md/q + 2\sqrt{\zeta}$ for the $ABAB$ term (one-way, before the factor 2 from (1)). The augmented error $2md/q + 4\sqrt{\zeta}$ is still bounded by $2md/q + 2\sqrt{\zeta}$ when ζ is replaced by a 4-times larger parameter, which does not change the asymptotic exponent in the final soundness cascade.

In the public interface of the formalization, all four scalar averages are exposed with the augmented bounds, while the tensor-form intermediates are kept private, consistent with the decision to keep the public interface purely scalar.⁶

⁴The bridge lemma carries the name `xEvaluatedFullSliceABABAvg_to_xEvaluatedFullSliceABABtensorAvg` and is located in `MIPStarRE/LDT/Commutativity/Main/Auxiliary.lean`.

⁵The corresponding declaration is `fullSlice_scalar_marginalize.y` in the same file.

⁶The public scalar averages are `fullSliceABAAvg`, `evaluatedSliceABAAvg`, `fullSliceABABAvg`, and `evaluatedSliceABABAvg`; the tensor intermediates are `xEvaluatedSliceBABAtensorAvg` and `xEvaluatedFullSliceABABtensorAvg`.

5 Conclusion

The overhead is genuine but does not represent a mathematical discrepancy. The formalization pays an extra $2\sqrt{\zeta}$ to cross between the scalar public interface and the tensor form required by `closenessOfIP`. This overhead is accounted for in the stated error bounds and does not obstruct the final soundness estimate. The divergence is present in every lemma that transports a scalar *ABAB* average through a tensor `closenessOfIP` step, and it is a consequence of the deliberate design choice to keep the public API in scalar form.

References

- [JNV⁺20] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. Quantum soundness of the classical low individual degree test, 2020. arXiv:2009.12982.