Issue 933: Normalization of Quantum States in the LDT Formalization

2026-04-30

1 At a glance

This note concerns the use of normalized quantum states in [JNV+20, Section *Preliminaries*], and the corresponding state object in the formal development. It is tracked by GitHub issue #933. The paper passages consulted here are in `references/ldt-paper/preliminaries.tex`: the definition of consistency is at lines 255–268, the definition of state-dependent distance is at lines 348–360, the Cauchy–Schwarz transport proposition is at lines 509–536, the switch-sandwich proposition is at lines 748–840, and the completion proposition is at lines 1101–1173.

- **Difficulty.** Low as a statement-alignment issue, but high blast radius if one tried to bundle normalization into the base state type.
- **Estimated weight.** A local audit or citation repair is small; a type-level migration would be a major project-wide refactor.
- **Mathlib/project split.** Mostly project-local. Mathlib supplies positive semidefinite matrices, trace, and Cauchy–Schwarz infrastructure, while the project supplies the normalized-trace convention and the LDT expectation notation.
- **Key mathematical inputs.** The paper's vector-state notation assumes $\langle\psi,\psi\rangle = 1$; the formal density-matrix object separates positivity from the predicate $\mathsf{QuantumState.IsNormalized}$.

2 Key theorem forms

The paper writes the shared state as a vector $|\psi\rangle$. In this notation a state is normalized: $\langle\psi,\psi\rangle = 1$. This convention is used not only in the definitions of consistency and state-dependent distance, but also in the proofs where a complete measurement gives a scalar equal to 1. For example, in the proof that the two definitions of consistency agree for measurements, the term

$$\mathbb{E}_x \sum_b \langle \psi | I \otimes B_b^x | \psi \rangle$$

is replaced by 1. The same normalization convention is used when a Cauchy–Schwarz factor of the form $\sum_a \langle \psi | C_a^2 | \psi \rangle$ is bounded by 1.

The formal state object is deliberately more general. A formal quantum state on $\mathcal{H}$ is a positive semidefinite matrix ρ. It does not, by itself, assert that ρ has unit trace. The separate predicate ρ is normalized means that the normalized trace $\tau(\rho) = \operatorname{Tr}(\rho)/\dim \mathcal{H}$ is equal to 1. The formal expectation is

$$\operatorname{ev}_\rho(X) = \operatorname{Re} \tau(\rho X).$$

For a pure unit vector v, the formal density matrix is $\dim(\mathcal{H})\,|v\rangle\langle v|$, so that $\tau(\dim(\mathcal{H})\,|v\rangle\langle v|) = 1$ and the formal expectation agrees with the paper's bra-ket scalar. [1] Thus the formal development separates positivity from normalization. It also allows mixed density matrices. This is a harmless generalization of the vector notation when the normalization predicate is supplied.

At the strategy level the paper convention is already bundled. A symmetric or projective LDT strategy carries both its bipartite density matrix and a proof that it is normalized. [2] The possible mismatch is therefore not that strategies in the main theorem lack normalization. The mismatch is that many auxiliary lemmas are stated for an arbitrary positive semidefinite density matrix, and a paper-scale estimate is valid for such a lemma only if normalization is added explicitly or the estimate is rescaled by $\operatorname{ev}_\rho(I)$.

[1] The corresponding declarations are `QuantumState`, `QuantumState.IsNormalized`, `pureDensity`, `PureState.toQuantumState_isNormalized`, and `ev`, in `MIPStarRE/LDT/Basic/QuantumState.lean`, lines 14–32, 46–108, and 159–163. The identity expectation lemma is `ev_one_of_isNormalized` in `MIPStarRE/LDT/Basic/OperatorExpectations.lean`, lines 88–93.

[2] See the fields `SymStrat.isNormalized` and `ProjStrat.isNormalized` in `MIPStarRE/LDT/Test/StrategyCore.lean`, lines 235–246 and 371–382.

3 Why the hypothesis is mathematically necessary

Let ρ be a positive semidefinite density matrix in the formal sense, and let $c > 0$. Then $c\rho$ is again a positive semidefinite density matrix, and

$$\mathrm{ev}_{c\rho}(X) = c\ \mathrm{ev}_\rho(X)$$

for every operator X. Consequently the formal consistency defects, state-dependent squared distances, strong self-consistency defects, and scalar transport averages all scale linearly with c.

For a normalized state and an operator $0 \le X \le I$, one has $\mathrm{ev}_\rho(X) \le 1$. Without normalization the bound is instead $\mathrm{ev}_\rho(X) \le \mathrm{ev}_\rho(I) = \tau(\rho)$. This distinction is invisible in the paper because $|\psi\rangle$ is a unit vector. It is not invisible in the formal statement. A typical Cauchy–Schwarz step gives

$$|\text{transport error}| \le \sqrt{\delta}\sqrt{\mathrm{ev}_\rho(I)}.$$

The paper bound $\sqrt{\delta}$ follows from this only when $\mathrm{ev}_\rho(I) = 1$. Thus a theorem over arbitrary positive semidefinite ρ, with the same state-independent right-hand side as in the paper, would be an overgeneralization. It would prove more than the paper's proof supplies, and it is generally false under scalar rescaling of ρ.

This is a genuine mathematical issue, but it is not a defect in the chosen state parameterization. The parameterization is sound when theorems with paper-scale constants include the normalization hypothesis. It also gives a useful language for intermediate weighted states and trace-scaled estimates, where forcing normalization into the base state structure would be inconvenient.

4 Representative formal sites where normalization is used

The following list is representative, not exhaustive. The common pattern is that the proof uses the identity $\mathrm{ev}_\rho(I) = 1$, or uses it after monotonicity to turn an operator bound $X \le I$ into a scalar bound $\mathrm{ev}_\rho(X) \le 1$.

First, the basic scalar bounds require normalization. The diagonal mass of a submeasurement is at most 1 on a normalized state, and bipartite consistency or strong self-consistency defects are at most 1 under the same hypothesis. [3] These lemmas are the formal versions of the paper's repeated use of the fact that a bounded measurement probability is at most 1.

Second, the Cauchy–Schwarz transport lemmas matching the paper's $\sqrt{\delta}$-scale estimates carry an explicit normalization hypothesis. This includes the transfer of consistency using state-dependent distance, the easy approximation lemma, and both left and right clauses of the inner-product transport proposition. [4] Not every distance manipulation needs normalization. For example, a purely operator-order contraction of a state-dependent squared distance may scale linearly on both sides. The normalization hypothesis is needed precisely when a bounded operator is converted into the numerical constant 1, or when that constant is hidden in a square-root loss.

Third, the switch-sandwich family of estimates uses normalization at exactly the places where the paper bounds a Cauchy–Schwarz factor by 1. The formal statement therefore assumes that the state is normalized, and the left and right transfer lemmas pass that assumption to the per-question gap bounds. [5] These lemmas feed directly into the commutativity transport estimates, such as the scalar marginalization bounds for full slices. [6]

Fourth, the completion and projectivization chain uses normalization when the paper writes that the total mass of a completed measurement is 1, or when a truncation error is stated at the paper's scale. The formal completion theorem, the missing-mass estimate, the spectral-truncation

[3] See `subMeas_diagMass_le_one` in `MIPStarRE/LDT/Preliminaries/SwitchSandwichPrep/Core.lean`, lines 103–114; `qBipartiteConsDefect_le_one_of_isNormalized`, `bipartiteConsError_le_one_of_isProbability`, and `bipartiteConsError_uniform_le_one` in `MIPStarRE/LDT/Test/Defs.lean`, lines 352–414; and `qBipartiteSSCDefect_le_one` and `bipartiteSSCError_uniform_le_one` in `MIPStarRE/LDT/Pasting/ComparisonLemmas/Common.lean`, lines 179–212.

[4] Representative declarations are `triangleSub` in `MIPStarRE/LDT/Preliminaries/Triangles/Core.lean`; `easyApproxFromApproxDelta`, `closenessOfIP`, and `closenessOfIPAdjoint` in `MIPStarRE/LDT/Preliminaries/CauchySchwarz.lean`, lines 148–218; and the underlying pointwise bounds `question_easyApproxFromApproxDelta`, `closenessOfInnerProduct_left`, and `closenessOfInnerProduct_right`.

[5] See `switchSandwich` in `MIPStarRE/LDT/Preliminaries/SwitchSandwichMain/Completeness.lean`, lines 22–45; `switchSandwich_leftTransfer` and `switchSandwich_rightTransfer` in `MIPStarRE/LDT/Preliminaries/SwitchSandwichMain/LeftTransfer.lean` and `MIPStarRE/LDT/Preliminaries/SwitchSandwichMain/RightTransfer.lean`; and `question_switchSandwich_left_gap` and `question_switchSandwich_middle_gap` in `MIPStarRE/LDT/Preliminaries/SwitchSandwichGapBounds/Left.lean` and `MIPStarRE/LDT/Preliminaries/SwitchSandwichGapBounds/Middle.lean`.

[6] For example, see `fullSlice_scalar_marginalize_x` and `fullSlice_scalar_marginalize_y` in `MIPStarRE/LDT/Commutativity/Main/Auxiliary.lean`, lines 350–413 and 460–500, and the evaluated-side transport bound `fullSlice_closenessOfIP_CAB_hEval_sqrt` in `MIPStarRE/LDT/Commutativity/Main/Auxiliary/HEvalTransport.lean`.

statement construction, and the orthonormalization theorem all expose the normalized-state hypothesis. [7] This is the correct formal counterpart of the paper's pure-state convention.

5 Relation with the blueprint and Lean

The blueprint preliminaries follow the paper's vector-state convention. They write $|\psi\rangle$ in the definitions of consistency and state-dependent distance, and in the statements of the transfer lemmas, without introducing a separate predicate for normalization. [8] This matches the paper as mathematical prose: a state vector is understood to be unit normalized. The blueprint therefore inherits the paper convention rather than displaying an additional normalization hypothesis.

The corresponding Lean declarations make the suppressed convention explicit. For example, the blueprint statements for the triangle transfer, inner-product transport, switch-sandwich lemma, completion theorem, spectral truncation statement, and orthonormalization theorem all link to declarations whose signatures include an explicit normalized-state hypothesis. Thus the blueprint and Lean are compatible, but they place the convention at different levels: the blueprint uses the standard mathematical meaning of $|\psi\rangle$ as a normalized state, while Lean records the hypothesis as data because its base state type contains all positive semidefinite matrices.

This note supplies the missing comparison. Future blueprint edits should either continue to use the paper convention for $|\psi\rangle$, or state explicitly that a density matrix ρ is normalized whenever the prose is no longer written in vector-state notation. In either case, a Lean declaration linked from such a blueprint item should carry `IsNormalized` whenever the proof uses the scalar identity $\mathrm{ev}_\rho(I) = 1$.

[7] See `completionMissingMassBound` in `MIPStarRE/LDT/Preliminaries/SelfConsistency/Core.lean`, lines 66–149; `completingToMeasurement` in `MIPStarRE/LDT/Preliminaries/CompletionTransfer.lean`, lines 234–254; `spectralTruncationStatement_of_sourceAlmostProjective` in `MIPStarRE/LDT/MakingMeasurementsProjective/SpectralTruncation/ProjectiveNonMeasurement.lean`; `leftLiftedProjectivizationRepair` in `MIPStarRE/LDT/MakingMeasurementsProjective/LocalityPreservingRepair.lean`; and `orthonormalization` in `MIPStarRE/LDT/MakingMeasurementsProjective/Orthonormalization.lean`, lines 587–610.

[8] See `blueprint/src/chapter/ch03_preliminaries.tex`, lines 209–265 for `def:simeq` and `def:approx_delta`, lines 274–280 for `prop:triangle-sub`, lines 316–376 for `prop:closeness-of-ip` and `prop:easy-approx-from-approx-delta`, lines 508–536 for `prop:switch-sandwich`, and lines 698–730 for `prop:completing-to-measurement`.

6 Verdict and migration policy

The verdict is that issue #933 records a real normalization gap between the paper's convention and the base formal state structure, but not a theorem-level defect in the current strategy statements. The paper works with normalized states. The base formal object intentionally records only positivity, and normalization is supplied by a separate predicate. This is a harmless parameterization when the predicate is required at every theorem whose conclusion has a paper-scale, state-independent constant.

The development should therefore not immediately migrate `QuantumState` to bundle normalization. Such a change would have high blast radius, and it would remove a useful type for unnormalized or weighted positive semidefinite states. The correct policy is local: keep `QuantumState` as the positive cone, keep normalization bundled in strategy structures, and require an explicit normalization hypothesis on any lemma that proves a bound using $\mathrm{ev}_\rho(I) = 1$, a scalar upper bound ≤ 1, or a paper-scale $\sqrt{\delta}$ transport estimate.

When an auxiliary theorem is intentionally meant to hold for unnormalized states, its statement should show the trace dependence explicitly. For example, a Cauchy–Schwarz transport estimate should contain the factor $\sqrt{\mathrm{ev}_\rho(I)}$, or an equivalent trace factor, rather than silently using 1. Conversely, if the theorem is meant to match a displayed statement in the paper, it should assume that the state is normalized. Future audits should look especially for occurrences of $\mathrm{ev}_\rho(I) = 1$, scalar conclusions bounded by 1, and square-root transport bounds stated over an arbitrary `QuantumState` without a normalization hypothesis.

7 Conclusion

The normalization hypothesis is a faithful formal boundary condition whenever a Lean theorem is meant to reproduce a paper-scale estimate for a state $|\psi\rangle$. The base QuantumState type should remain the positive cone; source-facing estimates should instead expose QuantumState.IsNormalized exactly at the points where the paper uses the unit-vector convention.

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.