Naimark dilation: projective-submeasurement form and the full bipartite assembly

2026-05-08

1 At a glance

- **Difficulty: low.** The mathematical issue is now resolved in Lean; the note remains to record the faithful projective-submeasurement reading of the paper statement.
- **Estimated weight:** no remaining proof construction is expected for Naimark itself. Downstream work should cite the checked tensor-product theorem or the questionwise adapter according to the required conclusion.
- **Mathlib/project split:** approximately **60% Mathlib** and **40% project-local**. The proof uses finite-dimensional matrix and trace algebra from Mathlib, with project-local packaging for measurements and bipartite tensor placement.
- **Status: formalized in the projective-submeasurement form.** The one-measurement Naimark step, the restriction to the original outcomes, and the two-sided tensor-product trace identity are now proved in Lean.
- **Remaining note:** the paper's word "measurement" must be read in the projective-submeasurement sense supplied by the helper lemma. The complete-measurement form on the original outcome type is false for arbitrary submeasurements because the residual mass is carried by the additional $\perp$ outcome.
- **Key Mathlib inputs:** finite-dimensional matrix trace identities, Naimark's one-measurement compression identity, tensor-product

states, and the register permutation used by the product-extension state.

2 Key theorem forms

The formalization isolates two related but distinct theorem forms.

1. **Source theorem form: Naimark dilation (`thm:naimark`, displayed at `references/ldt-paper/orthonormalization.tex`, lines 36–115).** For every collection of bipartite submeasurements $\{A^x\}_{x\in Q_A}$, $\{B^y\}_{y\in Q_B}$ on a finite Hilbert space $\mathcal{H}$, there is an enlarged Hilbert space $\widetilde{\mathcal{H}} \supseteq \mathcal{H}$, an isometry $V : \mathcal{H} \to \widetilde{\mathcal{H}}$, and projective submeasurements $\{\widetilde{A}^x\}$, $\{\widetilde{B}^y\}$ on $\widetilde{\mathcal{H}} \otimes \widetilde{\mathcal{H}}$ with
$$\langle\psi|A^x_a \otimes B^y_b|\psi\rangle = \langle V\psi|\widetilde{A}^x_a \otimes \widetilde{B}^y_b|V\psi\rangle. \tag{\texttt{thm:naimark}}$$
This is the displayed paper statement at `references/ldt-paper/orthonormalization.tex`, lines 36–115, read through the helper lemma at lines 121–159. The helper produces a projective submeasurement on the original outcomes; the residual mass is the additional $\perp$ outcome in the auxiliary register.

2. **Adapter theorem form: per-question one-measurement Naimark.** For each question x and submeasurement A^x there exists a one-measurement Naimark dilation on the enlarged space $\mathcal{H} \otimes (\mathbb{C} \oplus \mathbb{C}^{|\mathsf{Outcome}|})$ that preserves all single-outcome expectations:
$$\tau\big(\rho\, A^x_a\big) = \widetilde{\tau}\big(\widetilde{\rho}\, \widetilde{A}^x_{(a,\text{some})}\big).$$
This questionwise interface is also formalized; see the Lean structure NaimarkStatement and the Lean theorem questionwiseNaimark. The derived declaration OneMeasNaimarkData.toProjSubMeas turns the checked Option-outcome dilation into the projective submeasurement on the original outcomes.[1] The source-facing theorem naimarkTensorProductCorrelation constructs the auxiliary spaces, auxiliary product state, dilated state, and projective submeasurements from the checked one-measurement theorem. In Lean these auxiliary spaces are the standard one-register spaces Option($\mathsf{Outcome}_A$) and Option($\mathsf{Outcome}_B$), with a separate Naimark unitary chosen for each question; the theorem does not require different questions'

[1] Defined at `MIPStarRE/LDT/MakingMeasurementsProjective/NaimarkOneMeas.lean`.

dilated measurements to commute. The four-register trace identity OneMeasNaimarkData.twoSidedCorrelationPreservation is proved.

3. **Downstream theorem form.** The downstream uses inside the LDT pipeline never call the bipartite tensor product of two Naimark dilations; they call the per-question single-outcome marginal-preservation conclusion at Alice or Bob separately, with the unmodified state on the unaffected side. This is supplied directly by leftMarginalPreservation and rightMarginalPreservation.

3 The source assertion

We examine [JNV+20], the Naimark dilation theorem labelled `thm:naimark`.[2]

Notation. We write $X \approx_\delta Y$ on the state $|\psi\rangle$ to mean the state-dependent operator distance $\langle\psi|(X-Y)^\dagger(X-Y)|\psi\rangle \leq \delta$, defined as in [JNV+20] at `def:approx_delta` (`references/ldt-paper/preliminaries.tex`, line 348).

In the notation of the paper, the source assumes

$$\mathcal{H} = \text{a finite-dimensional Hilbert space}, \quad \{A^x\}_{x\in Q_A}, \{B^y\}_{y\in Q_B} \text{ submeasurements on } \mathcal{H},$$

and concludes the bipartite preservation identity above. The proof proceeds question by question via `lem:naimark-helper`. That helper explicitly produces a projective submeasurement indexed by the original outcomes. This is the mathematically necessary form for arbitrary submeasurements. If one asks instead for a complete projective measurement indexed by the original outcomes, the assertion is false: take a one-outcome submeasurement $A_* = 0$ on Alice's space and a one-outcome measurement $B_* = I$ on Bob's space. The original correlation is zero. But a one-outcome projective measurement on the enlarged Alice and Bob spaces must have its unique effect equal to the identity, so any normalized dilated state has correlation one.

The example `ex:easy-but-long` demonstrates that the dilation does *not* preserve the state-dependent distance $\approx_\delta$ defined just above, which is why downstream arguments in the LDT paper carry the dilation only as far as the single-outcome marginals.

[2] Tracked alongside ledger GitHub issue #1379. The paper passages are `references/ldt-paper/orthonormalization.tex`, lines 36–115 (the displayed theorem statement), 121–159 (the helper lemma `lem:naimark-helper` and its proof), 161–187 (the proof of `thm:naimark` that tensors the per-question helpers), and 189–272 (the genuine non-preservation example `ex:easy-but-long`).

4 Comparison with the formalization

The blueprint and formal development record both the per-question one-measurement dilations and the global tensor-product assembly into one bipartite statement. The source-facing Lean statement is stated in terms of projective submeasurements on the original outcome types. Equivalently, one could state a completed version using projective measurements indexed by Option(Outcome), but that is not the outcome type used in the source display. The checked assembly uses the standard one-register auxiliary space Option(Outcome) on each side, with different unitaries for different questions. This proves the same correlation statement as the paper's product-over-questions auxiliary construction because the theorem does not impose simultaneous commutation between dilated measurements for distinct questions.

The formal development also proves the adapter form

$$\forall x, a, \rho,\ \tau\big(\rho A^x_a\big) = \widetilde{\tau}\big(\widetilde{\rho}_x\,(\widetilde{A}^x)_{(a,\mathrm{some})}\big),$$

where $\widetilde{\rho}_x$ is the one-measurement-lifted density at question x. The corresponding statement for Bob is symmetric. The bipartite identity of the source theorem form is now proved by the four-register trace identity OneMeasNaimarkData.twoSidedCorrelationPreservation.

The questionwise interface is still sufficient for downstream LDT use: every consumer reduces to the per-side single-outcome marginal-preservation conclusion, applied either to Alice's submeasurements with B left untouched or vice versa. The source theorem is also available as an axiom-clean Lean theorem.

5 What remains documented

The former bipartite-assembly gap is closed. The note remains because it records a genuine correction to the source wording:

- **Distance non-preservation.** Example `ex:easy-but-long` shows that the dilation does not preserve the state-dependent distance $\approx_\delta$ recalled in Section 3, so any bipartite statement formulated in terms of state-dependent distances would need to track an explicit error inflation that is not used downstream.
- **Projective-submeasurement conclusion.** The helper lemma produces projectors on Option($\mathcal{A}$). Restricting to the original outcomes

gives a projective submeasurement. It does not give a complete projective measurement on the original outcome type unless the original submeasurement was already complete.

- **Blueprint interpretation.** The blueprint therefore states and links the theorem in the projective-submeasurement form, with this note recording why that is the faithful formal reading of the cited argument.

6 Conclusion

The Lean structure NaimarkStatement and theorem questionwiseNaimark record the per-question one-measurement Naimark data and the single-outcome marginal-preservation conclusions used by downstream LDT lemmas. The declaration OneMeasNaimarkData.toProjSubMeas proves the local passage from the completed Option-outcome projective measurement to the projective submeasurement on the original outcomes. The source-facing tensor theorem should be read in the projective-submeasurement form supplied by the paper's helper lemma; the complete-measurement form on the original outcome type would be false for submeasurements. The simultaneous bipartite tensor-product assembly is now proved by naimarkTensorProductCorrelation; no Naimark-specific 'sorryAx' remains.

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.