Canonical-Form Weight Normalization and the Proportional MPV Comparison

2026-05-10

Verdict: unfaithful (open).

1 The source argument

Cirac–Pérez-García–Schuch–Verstraete 2016 builds the multi-block route to the Fundamental Theorem on top of the canonical form [CPGSV16, Section II]. A tensor $A = (A^i)_{i=1}^d$ is in canonical form when, after a basis change of the bond, it decomposes as

$$A^i = \bigoplus_{k=1}^{r} \mu_k \, A_k^i, \qquad \texttt{(eq:II\_CF1)}$$

where each block tensor A_k is normal and the scalars μ_k are nonzero. The text immediately fixes a normalization on the weights:

> "we can always choose $|\mu_k| \leq 1$ and at least one of them equals one, something which we will assume from now on." [CPGSV16, paragraph after `eq:II_CF1`]

This is a definitional convention applied throughout the rest of the paper; it is not derived later. Combined with the strict modulus ordering used by the BNT arrangement, it implies that the dominant block has $|\mu_1| = 1$ and every sub-dominant block has $|\mu_k| < 1$.

The same paper proves the proportional Fundamental Theorem and its equal-MPV corollary [CPGSV16, Section II]. Suppose two tensors A and B are presented in their basis-of-normal-tensor expansions, with representative families $(A_j)_{j=1}^g$ and $(B_j)_{j=1}^g$ and BNT coefficients μ_{j,q_a} and ν_{j,q_b} for $q_a = 1, \ldots, r_{a,j}$ and $q_b = 1, \ldots, r_{b,j}$. After the proportional Fundamental Theorem identifies the representatives, $B_j = e^{i\phi_j} Y_j A_j Y_j^{-1}$, the trace-pair argument in the proof of the equal-MPV corollary reduces to the unlabeled scalar identity

$$\sum_{q=1}^{r_{a,j}} \mu_{j,q}^N = \sum_{q=1}^{r_{b,j}} (\nu_{j,q} e^{i\phi_j})^N, \qquad (1)$$

The cited elementary power-sum lemma then concludes that $r_{a,j} = r_{b,j} =: r_j$ and that the multisets of weights coincide, $\mu_{j,q} = \nu_{j,q} e^{i\phi_j}$. The matching dimensions and the global gauge formula follow.

The mathematical structure of this proof is that the coefficient information is read off at the level of power sums, using all N simultaneously as input to a separation lemma for sums of Nth powers. No convergence of any N-indexed quantity is invoked.

2 The formal canonical form

The formal counterpart of the canonical form for the BNT route is the structure `IsNormalCanonicalFormBNT`.[1] Its mathematical content is: a strict-modulus ordering on the weight family $\mu : \mathrm{Fin}\, r \to \mathbb{C}$, the nonvanishing $\mu_k \neq 0$ for every k, and the per-block normality data inherited from the underlying canonical form. The dominant weight is unconstrained in modulus.

Comparing this with the source definition, the field $|\mu_1| = 1$ (in source enumeration; index 0 in the formal one) is missing. The blueprint BNT chapter records the strict-modulus ordering in the BNT-separated canonical-form definitions. Separately, the canonical-form chapter records the decomposition $\mu_k = \eta_k |\mu_k|$ with η_k on the unit circle. Neither place imposes $|\mu_1| = 1$.[2]

The omission is mathematically meaningful, because the source definition is used in two ways downstream: it fixes the dominant block on the unit circle, which is what the spectral and irreducibility arguments need, and it forces every sub-dominant block strictly inside the unit disk. The latter is the conventional content used by the proportional MPV comparison below.

3 The discarded formal proportional comparison

The discarded formal route to the proportional Fundamental Theorem replaced the source's unlabeled power-sum identity by a different mathematical object: two coefficient sequences indexed by chain length, together with the assumption that they converge to nonzero limits. In that route, each block index j had a function $N \mapsto a_{N,j} \in \mathbb{C}$ and a designated limit $a_j^\infty \in \mathbb{C}$ satisfying

$$a_{N,j} \to a_j^\infty, \qquad a_j^\infty \neq 0 \tag{2}$$

as $N \to \infty$, together with the analogous data on the B-side and a sequence of proportionality scalars c_N converging to a nonzero limit c. The identification of the scalars is then derived from these convergence hypotheses by a comparison of limits, rather than from the source power-sum identity.

[1] `TNLean/MPS/BNT/Construction.lean`, lines 181–189.

[2] The BNT-separated definitions are `def:cf_bnt` and `def:normal_cf_bnt` in `blueprint/src/chapter/ch10_bnt.tex`, lines 70–100 and 119–141, respectively. The phase–modulus decomposition is recorded in `blueprint/src/chapter/ch09_canonical.tex`, lines 125–139.

For the natural realization in which the jth source sector has one selected representative, or one copy in the coefficient multiplicity notation ($r_j = 1$), and the coefficient is the source weight itself, the formal choice of array[3] is

$$a_{N,j} = (\mu_j^A)^N. \tag{3}$$

With this choice, the convergence requirement $a_{N,j} \to a_j^\infty \neq 0$ becomes a condition on the powers $(\mu_j^A)^N$.

4 The design history of the limit hypothesis

The convergence-to-nonzero-limits hypothesis is not present in [CPGSV16]. It is also not present in the formalization's own planning record. A proof-plan memo committed on 2026-02-22 sets out the intended route for the proportional Fundamental Theorem and its equal-MPV corollary in [CPGSV16].[4] Its concrete recommendation is a span-equality bridge, together with a finite-length scalar power-sum identity for the μ-coefficient matching.

The memo records the source identity in the form

$$\sum_{k=1}^{x_a} \lambda_{a,k}^N = \sum_{k=1}^{x_b} \lambda_{b,k}^N \tag{4}$$

and proposes formalizing it as a fixed-length theorem requiring only finitely many values of N. The planned application to the BNT μ-coefficients reads, in the memo's own notation,

$$\sum_{q=1}^{r_{a,j}} \mu_{j,q}^N = \sum_{q=1}^{r_{b,j}} (\nu_{j,q} e^{i\phi_j})^N, \tag{5}$$

to be applied via this fixed-length scalar identity to obtain $r_{a,j} = r_{b,j}$ and the multiset equality. No convergence statement appears in the memo's plan, and no nonzero-limit hypothesis is contemplated.

The memo also identifies an apparent circularity in the span-equality route: to deduce span equality from proportionality, one wants linear independence at the witness length, while the surrounding BNT permutation step uses span equality to extract the linear independence it needs. The memo's own resolution is to weaken the global span-equality hypothesis to its eventual form,

$$\operatorname{span}\{\, V^{(N)}(A_j)\}_j = \operatorname{span}\{\, V^{(N)}(B_k)\}_k, \tag{6}$$

and to observe that the surrounding step already only consumes the eventual form. The memo records this as a backward-compatible weakening that discharges the circularity at the type level.

[3] The corresponding blueprint passage is the remark `rem:cf_coeff_arrays` following `lem:ft_proportional` in `blueprint/src/chapter/ch11_fundamental_theorem_proof.tex`, lines 172–202.

[4] `docs/audits/2026-02-22_thm4_4_proof_plan_ref1606.txt`, committed in `64c018ab` earlier the same day as the introducing implementation `dd54b79e`.

The implementation committed later the same day departs from both directions of the memo. It does not formalize the fixed-length scalar power-sum identity proposed for the μ-coefficient step. It also does not adopt the eventual span-equality form proposed for the bridge. Instead it introduced, on the BNT bridge theorems and later in a single formal data structure, a global convergence hypothesis on the coefficient sequences and a separate hypothesis that those limits are nonzero. The mathematical object packaged by these fields is not the power-sum identity from the source proof of the equal-MPV corollary, and it is not the eventual span equality from the memo's bridge plan. It is a third object, introduced without a corresponding written rationale. A later commit on 2026-05-09 describes the structure as "MPV-level Newton–Girard convergence data, not a §II 283–302 Wedderburn-style similarity," but this description records what the structure became, not why the change from the memo's plan was made.[5]

The recorded design history therefore shows that the limit hypothesis was not chosen to follow the source argument and was not the route the formalization itself planned for. It is a third route whose introduction is not explained in the project record.

5 The mathematical incompatibility

Throughout this section, the comparison is restricted to the strict-weight, one-representative-per-block specialization just described, namely $a_{N,j} = (\mu_j^A)^N$ and $b_{N,k} = (\mu_k^B)^N$. It is not a claim about the more general sector coefficient arrays $c_{S,j}^{(N)} = \sum_q \mu_{j,q}^N$ that retain repeated copies within a BNT sector.

Under the source normalization $|\mu_1^A| = 1$ and $|\mu_j^A| < 1$ for $j \geq 2$, the powers $(\mu_j^A)^N$ behave as follows. For each sub-dominant block,

$$|(\mu_j^A)^N| = |\mu_j^A|^N \to 0 \tag{7}$$

as $N \to \infty$, so the only candidate limit is $a_j^\infty = 0$, contradicting the formal hypothesis $a_j^\infty \neq 0$. For the dominant block, the modulus stays on the unit circle, and the sequence $(\mu_1^A)^N$ converges only in the special case $\mu_1^A = 1$; nontrivial roots of unity are periodic rather than convergent, and a generic phase has no limit.

The conclusion is a structural mismatch. With the source normalization in force, the discarded formal data cannot be instantiated for a multi-block ($r > 1$) family using $a_{N,j} = (\mu_j^A)^N$, because the limit hypothesis fails on every sub-dominant block. Without the source normalization, the formal hypothesis

[5] The introducing commit is `dd54b79e`, "Add ScalarPowerSumIdentity, BNTConstruction, BNTPermutationThm44." Despite the name, the `ScalarPowerSumIdentity` module added by that commit does not correspond to the memo's planned fixed-length scalar identity from the elementary power-sum lemma in [CPGSV16]. The structure refactor that bundled the limit fields is `9e096a14` (2026-03-22), and the post-hoc Newton–Girard description is in the message of `8ab1a5db` (2026-05-09).

is sometimes satisfiable, but the resulting object is no longer the canonical-form datum from [CPGSV16].

The paper does not face this difficulty, because it never asks for the $N \to \infty$ limit of $(\mu_j^A)^N$. The proof reads off the multiset of weights from a finite family of polynomial identities in N, rather than from an $N \to \infty$ limit. In the strict one-representative specialization isolated here, once the BNT linear-independence input supplies a witness length, the comparison can be made at a single chosen length $N = N_0$ up to the expected phase/root-of-unity ambiguity. The general multi-multiplicity case remains the fixed finite family of values $N \leq \max(r_{a,j}, r_{b,j})$ used by the source separation lemma and by the 2026-02-22 memo.

6 Cross-check with the blueprint

The blueprint entries that formerly paralleled this convergence-and-nonzero-limit surface have been demoted or removed from the source-labelled theorem route. They must not be marked as formalizations of the source proportional Fundamental Theorem or its equal-MPV corollary. The active blueprint now treats the source comparison as an open block-matching step whose proof must come from the BNT power-sum argument, not from external coefficient limits.

The blueprint BNT definitions are symmetric with the formal definition: the strict-modulus ordering is recorded; the normalization $|\mu_1| = 1$ is not.

7 What a corrected formalization would say

Two changes are needed to bring the formal development into agreement with the source.

The first change is to record the source convention $|\mu_1| = 1$. Once this normalization is present, the implicit bounds $|\mu_k| \leq 1$ on the remaining blocks follow from the strict modulus ordering.[6] This should be a field of `IsNormalCanonicalFormBNT` and of the corresponding blueprint definition. It is a transcription of the source definition, not an additional assumption.

The second change is to reformulate the proportional comparison around finite power-sum data rather than limits. In the general coefficient-multiplicity setting of the source proof, the blockwise input should be the finite family

$$\sum_{q=1}^{r_{a,j}} \mu_{j,q}^N = \sum_{q=1}^{r_{b,j}} (\nu_{j,q} e^{i\phi_j})^N, \qquad N = 1, \ldots, \max(r_{a,j}, r_{b,j}), \tag{8}$$

or the corresponding finite collection of witness lengths produced after the BNT span and linear-independence input has been made available. The source separation lemma then gives $r_{a,j} = r_{b,j}$ and the multiset equality of the weights.

[6] The formal field carrying this ordering is `mu_strict_anti` in `IsNormalCanonicalFormBNT`; see `TNLean/MPS/BNT/Construction.lean`, lines 184–185.

This finite-family route is what the 2026-02-22 proof-plan memo planned for the general multi-multiplicity case; it is not a single-witness argument in that generality.

In the strict-weight, one-representative specialization discussed above, the same finite-data recipe collapses to comparing one chosen witness length N_0. The mathematical input is the BNT linear independence at that length, supplied by the direct-sum span theorem on a basis of normal tensors, and the blockwise output is

$$(\mu_j^A)^{N_0} = c^{N_0}\,(\mu_j^B)^{N_0}, \qquad j = 1, \dots, g, \tag{9}$$

which fixes μ_j^A in terms of μ_j^B and c up to an N_0th root of unity. The phase ambiguity is then absorbed into the per-block phase ϕ_j already present in the source statement $B_j = e^{i\phi_j} Y_j A_j Y_j^{-1}$. No convergence statement is needed; in particular no nonzero limit hypothesis on $(\mu_j^A)^N$ is required.

Thus the corrected route is source-faithful at finite length: the multi-multiplicity version follows the finite power-sum identities from the proof of the equal-MPV corollary [CPGSV16], while the single-witness discussion is only the strict-weight specialization of that route. The circularity worry that the memo flagged in connection with the span-equality bridge does not arise for this finite-length route once the BNT linear-independence input is supplied from the direct-sum span theorem rather than extracted from the proportional comparison itself.

8 Conclusion

The verdict is that the canonical-form definition and the proportional MPV comparison originally differed from [CPGSV16] in two coupled ways. The canonical-form definition omitted the source convention $|\mu_1| = 1$. The proportional comparison replaced the source's fixed-length power-sum identity by a limit hypothesis on $(\mu_j^A)^N$ that fails on the sub-dominant blocks under the very normalization the source imposes. The proof-plan memo of 2026-02-22 already proposed the source-faithful route; the implementation diverged from both the memo and the source in a single direction, by introducing a convergence-to-nonzero-limits hypothesis that has no counterpart in either record.

Status: supplier-side normalization. The arbitrary-input supplier no longer leaves the line-246 weight choice to the caller. `MPSTensor.exists_isBNTCanonicalForm_afterBlocking_pos_normalized` rescales the prepared weights by their maximal modulus $m = \max_k |\mu_k|$ (`MPSTensor.exists_weight_normalization`), which satisfies both conditional premises of the earlier supplier, and tracks the cost as the per-site scalar m^N in the positive-length matrix-product identity. The positive-length statement is the canonical source formulation; no length-zero coefficient is part of this statement. The normalized supplier assumes exactly that the positive-length matrix-product family is nonzero: there are $N \geq 1$ and a word w of length N for which $\operatorname{tr}(A^w) \neq 0$. This condition persists under every blocking. For each $p \geq 1$, some $q \geq 1$ satisfies $\operatorname{tr}((A^w)^{pq}) \neq 0$; otherwise the Newton–Girard trace identities would make $(A^w)^p$, and hence A^w, nilpotent, contrary to its nonzero trace. The repeated word of length pqN therefore gives a nonzero coefficient after blocking by p. Thus this hypothesis excludes precisely the zero positive-length family, for which the unit-modulus weight required at line 246 is impossible. The subsequent proposition nevertheless says "for any tensor." Taken literally, its zero case is incompatible with the preceding global unit-modulus convention and with the formal canonical-form predicate. The present theorem therefore formalizes the nonzero-family case. Eliminating the restriction requires a separate zero-family clause or a canonical-form predicate whose normalization condition is conditional on the presence of a sector.

Status: structural realignment. The structural realignment has been done. The canonical-form definition now carries the source convention $|\mu_1| = 1$ as a field of `IsNormalCanonicalFormBNT`, propagated through every constructor. The convergence-to-nonzero-limit fields and the later per-N coefficient package have been removed from the formal development. The surviving boundary keeps the block-permutation and gauge-phase comparison as an explicit input until the source-faithful residual comparison theorem is available.

The earlier one-shot nonzero-overlap theorems and the matching wrappers built on them have now been deleted. Their statements asserted a full "for every block" nondecay conclusion from arbitrary bounded coefficient arrays, but that conclusion is stronger than the exposed hypotheses: a non-dominant coefficient can vanish identically unless the source's residual peeling argument or the power-sum sector structure has already been encoded. The remaining Lean interfaces keep the proportional matching conclusion as an explicit input until the source-faithful residual comparison theorem is available.

What remains is the proof itself. The paper-faithful proof follows the block-selection step of the proportional Fundamental Theorem [CPGSV16, line 1182], followed in the equal-MPV corollary by the finite power-sum comparison and global gauge assembly [CPGSV16, lines 1184–1192]. It will consume the per-N hypothesis set together with the missing residual comparison data: the lower-bound argument uses $\|b_{N,1}\| = 1$ together with $\|b_{N,k}\| \leq 1$ and $c_N \neq 0$ for the currently dominant surviving sector, and the finite power-sum identification of weights either at one witness length (the strict-weight specialization) or at finitely many witness lengths (the multi-multiplicity case). That work is tracked as Stage C of issue #1559.

The corresponding tracked items are issue #1554 for the canonical-form normalization (structurally addressed by the paper-realignment PR; closes when the proofs land) and issue #1555 for the fixed-length reformulation of the proportional comparison (structurally addressed; closes with Stage C).

References

[CPGSV16] J. Ignacio Cirac, David Pérez-García, Norbert Schuch, and Frank Verstraete. Matrix product density operators: Renormalization fixed points and boundary theories. Preprint; published version: Annals of Physics 378 (2017), 100–149, 2016.