The CPSV16 Unit-Weight Convention and Renormalization-Fixed-Point Scale

2026-08-15

Verdict: false-source (resolved).

1 Two scale conventions in the source

The canonical-form construction of Cirac–Pérez-García–Schuch–Verstraete first writes a tensor as

$$A^i = \bigoplus_k \mu_k A_k^i, \tag{1}$$

where every A_k is normal and its transfer map has spectral radius one [CPGSV16, lines 214–246]. Since the states are not normalized there, line 246 adopts the standard convention

$$|\mu_k| \leq 1 \qquad \text{for every } k, \qquad |\mu_{k_0}| = 1 \qquad \text{for some } k_0. \tag{2}$$

The formal BNT predicate records the second clause through `MPSTensor.IsBNTCanonicalForm.weight_unit_exists`. This is one global unit-weight witness, not one witness in every sector; the latter distinction is discussed separately in `docs/paper-gaps/cpsv16_global_vs_persector_unit_witness.tex`.

Definition 4.1 fixes the scale of the tensor in a different way. For a tensor M already in canonical form and generating MPDOs, it requires trace-preserving completely positive maps relating the one-site and two-site physical closures [CPGSV16, lines 645–660]. Trace preservation gives

$$\mathcal{T}_M^2 = \mathcal{T}_M, \tag{3}$$

where $\mathcal{T}_M = \sum_i M^{ii}$ is the physical-trace transfer; this is the calculation at lines 1333–1340. The formal statement is `MPOTensor.physTraceTransfer_sq_of_isRFPViaTS`.

These requirements concern two transfer objects with different scaling laws. For the doubled-index MPS transfer map

$$\mathcal{E}_M(X) = \sum_i M^i X (M^i)^\dagger, \tag{4}$$

one has

$$\mathcal{E}_{cM} = |c|^2 \mathcal{E}_M. \tag{5}$$

For the physical-trace transfer, linearity gives

$$\mathcal{T}_{cM} = c\mathcal{T}_M. \tag{6}$$

Equation (2) normalizes the first object by choosing normal BNT representatives and placing the overall scale in their coefficients. Equation (3) constrains the second object at the scale of the displayed tensor. The related up-to-scalar treatment of source zero correlation length is recorded in `docs/paper-gaps/cpsv16_zcl_canonical_form_normalization.tex`.

2 Fixed-representative scale pinning

By (6), replacing M by cM replaces $\mathcal{T}_M$ by $c\mathcal{T}_M$. If $\mathcal{T}_M^2 = \mathcal{T}_M \neq 0$ and the nonzero rescaling cM also satisfies the same literal idempotence equation, then

$$c^2\mathcal{T}_M = (c\mathcal{T}_M)^2 = c\mathcal{T}_M, \tag{7}$$

so $c = 1$. Thus Definition 4.1 is not invariant under an arbitrary nonzero rescaling of a fixed representative.

By contrast, the simplicity definition at lines 815–822 is phrased through the elements of a BNT: the tensor is simple when none of those elements has nilpotent ket-against-bra contraction. Multiplying one BNT element by a nonzero scalar does not change nilpotency. This observation concerns the nilpotency clause alone; it does not imply that an arbitrary complex rescaling preserves the MPDO condition or either full simplicity predicate.

Indeed, if $\rho_N(M)$ denotes the closed length-N MPO, then the exact law is

$$\rho_N(cM) = c^N \rho_N(M), \tag{8}$$

not $|c|^{2N}\rho_N(M)$. Consequently positivity is preserved in both directions by a strictly positive real rescaling rM. The formal theorems `MPOTensor.isMPDO_smul_ofReal_iff`, `MPOTensor.isSourceSimple_smul_ofReal_iff`, and `MPOTensor.isNonvanishingSourceSimple_smul_ofReal_iff` establish the corresponding invariances for the MPDO, source-simple, and nonvanishing source-simple predicates. They do not assert invariance under arbitrary nonzero complex scalars, and they do not imply that the rescaled tensor satisfies the same literal fixed-scale RFP equations of Definition 4.1.

The declaration `MPOTensor.IsSourceSimple` records the normalization-free active canonical-block reading of lines 217–246 and Definition 4.7. It asks for some positive blocking and an active BNT sector presentation whose representative physical-trace transfers are all non-nilpotent. Every representative has a positive number of copies and every copy has nonzero weight, so dormant candidates are absent. This is the local correction documented in `docs/paper-gaps/cpsv16_bnt_uniqueness_zero_coefficient.tex`, not the unrestricted literal-BNT reading. The theorem `MPOTensor.IsSourceSimple.exists_mpo_ne_zero` derives a positive length N for which

$$\rho_N(M) \neq 0, \tag{9}$$

from the active presentation; this is not an additional defining assumption. Isolated vanishing lengths remain possible. The predicate does not impose the line-246 unit-weight convention. The strengthened project predicate `MPOTensor.IsNonvanishingSourceSimple` adds the condition

$$\rho_N(M) \neq 0 \qquad \text{for every } N > 0. \tag{10}$$

By contrast, `MPOTensor.IsSimpleCanonicalForm` and `MPOTensor.IsSimple` are normalized fixed-representative predicates. The first is a predicate on one displayed canonical-form representative. The second permits positive physical blocking, but still asks the resulting blocked tensor itself to admit such a normalized witness. Neither declaration is a quotient by nonzero scalar rescaling, and no equivalence with either active simplicity predicate is asserted. The implication `MPOTensor.IsSimple.isSourceSimple` is unconditional. Its normalized sector decomposition is itself an active presentation after transport through the canonical-form gauge. Positive-length nontriviality in (9) then follows from the general source-simplicity theorem, whose proof uses eventual linear independence, a non-eventually-zero sector coefficient, and physical unblocking. The strengthened implication `MPOTensor.IsSimple.isNonvanishingSourceSimple` uses (10).

The tensor `MPOTensor.SimpleVanishingCounterexample.M`, whose sole local virtual matrix is $\mathrm{diag}(1, -1)$, satisfies

$$\rho_N(M) = \bigl(1 + (-1)^N\bigr)I. \tag{11}$$

Since $\mathrm{diag}(1, -1)^2 = I$, its two-site blocking has one sector, the scalar unit normal representative, two copies of unit weight, and total virtual dimension two. The doubled physical-trace transfer is the one-by-one matrix (1) and is non-nilpotent. Thus `MPOTensor.SimpleVanishingCounterexample.M_isSimple`, `MPOTensor.SimpleVanishingCounterexample.M_isSourceSimple`, and `MPOTensor.SimpleVanishingCounterexample.M_not_isNonvanishingSourceSimple` prove that the tensor is normalized-simple and source-simple: $\rho_2(M) = 2I \neq 0$ exhibits the positive-length nontriviality derived from its active presentation. It is not nonvanishing source-simple because $\rho_1(M) = 0$. This witness does not refute Definition 4.7, is not an RFP claim, and gives no rigidity statement for BNT structure coefficients. For a fixed tensor, a theorem may also quantify over every normalized BNT witness rather than rely on one chosen presentation; the dimer obstruction below has precisely this form.

3 The parity tensor in Example 4.12

Example 4.12 prints the two physical letters $M^{00} = I$ and $M^{11} = \sigma_z$. Its physical-trace transfer satisfies

$$\mathcal{T}_M^2 = 2\mathcal{T}_M, \tag{12}$$

so the two trace-preserving channel equations cannot hold for the printed representative. The density-normalized representative

$$\widehat{M} = \tfrac{1}{2}M \tag{13}$$

instead has idempotent physical-trace transfer. Explicit parity coarse-graining and refinement channels prove `MPOTensor.CPSVExample412NormalizedRFP.Mhat_isRFPViaTS` without any additional hypothesis; the complete calculation is recorded in `docs/paper-gaps/cpsv16_example_4_12_normalization.tex`.

The doubled-index MPS has two bond-one sectors, with physical vectors $(1, 0, 0, 1)$ and $(1, 0, 0, -1)$. Each vector has squared norm 2, so its normal representative is obtained by division by $\sqrt{2}$. Consequently the two horizontal normal-block weights of $\widehat{M}$ are both $1/\sqrt{2}$. Thus the scale selected by the channel equations does not supply the global unit-weight witness in (2). The formal theorem deliberately records the bare channel predicate and does not claim that $\widehat{M}$ satisfies every standing canonical-form convention in Definition 4.1.

4 The displayed dimer representative

The project tensor R gives a concrete instance of this tension. Its doubled-index MPS transfer map satisfies

$$\mathcal{E}_R^2 = \frac{337}{512}\mathcal{E}_R. \tag{14}$$

Consequently the displayed one-block CPSV canonical form is

$$R^{\bullet\bullet} = \mu\widehat{R}, \qquad \mu = \sqrt{\frac{337}{512}}, \tag{15}$$

where the transfer map of the retained block $\widehat{R}$ is idempotent and hence has spectral radius one. This presentation has $0 < \mu < 1$, so it does not itself supply the global unit-weight witness in (2).

On the other hand, the physical-trace transfer of the displayed tensor R is a nonzero idempotent: its $(0, 0)$ entry is $1/2$. The retained block has

$$\mathcal{T}_{\widehat{R}} = \mu^{-1}\mathcal{T}_R. \tag{16}$$

Since $\mu \neq 1$, the scale-pinning calculation above shows that this normalized representative does not satisfy the same literal physical-trace idempotence equation as R.

The obstruction is stronger than this displayed-presentation mismatch. For every positive blocking length L, the doubled-index transfer map of the blocked tensor satisfies

$$\mathcal{E}_{R^{[L]}}^2 = \left(\frac{337}{512}\right)^L \mathcal{E}_{R^{[L]}}. \tag{17}$$

Suppose that $R^{[L]}$ admitted any normalized horizontal BNT witness. The gauge to that witness transports (17). The global unit-weight condition from line 246 supplies a copy whose weight has modulus one. Restricting the transported equation to that copy removes the weight. Its normal representative is left canonical, so its transfer map is trace preserving. Taking the trace of the quasi-idempotence equation at the identity forces

$$\left(\frac{337}{512}\right)^L = 1. \tag{18}$$

This contradicts $0 < (337/512)^L < 1$. Hence no positive blocking of R admits a normalized horizontal canonical form. Since the formal simple-canonical-form predicate includes such a horizontal witness, the theorem `R_not_isSimple` in the namespace `MPOTensor.RescalingStableLengthDependentRFP` proves

$$\neg\, \mathsf{MPOTensor.IsSimple}\, R. \tag{19}$$

The argument quantifies over the arbitrary normalized BNT witness allowed by the formal predicate; it does not depend on the displayed one-block presentation.

This failure is not a nilpotency result. The retained-block theorem proves that the displayed canonical presentation has non-nilpotent physical-trace transfer. Together with positivity and the active one-site BNT refinement, it witnesses

$$\mathsf{MPOTensor.IsSourceSimple}\, R. \tag{20}$$

The explicit nonzero closed-MPO entry at every positive length further gives

$$\mathsf{MPOTensor.IsNonvanishingSourceSimple}\, R. \tag{21}$$

The source quantifier is existential and is witnessed by $L = 1$; no assertion is made about every positive blocking. The normalized obstruction instead combines quasi-idempotence of the doubled-index transfer, the line-246 unit-weight copy, and trace preservation of its normal representative. It uses none of the vertical coefficient identities discussed below.

5 Coefficient absorption with all Appendix C.2 conditions

The two-sector coefficient-absorption example realizes the normalization tension without dropping the conditions used in the simple Case-II passage [CPGSV16, lines 1628–1665 and 1740–1782]. Its nonzero physical slices are

$$K^{00} = K^{11} = \mathrm{diag}(1/2, 0), \qquad\qquad K^{22} = \mathrm{diag}(0, 1). \tag{22}$$

Thus, for $N > 0$,

$$\langle\sigma|\rho_N(K)|\tau\rangle = \delta_{\sigma,\tau}\left(2^{-N}\mathbf{1}_{\{\sigma_k\in\{0,1\}\ \forall k\}} + \mathbf{1}_{\{\sigma_k=2\ \forall k\}}\right), \quad \mathrm{tr}\,\rho_N(K) = 2 > 0. \tag{23}$$

The tensor therefore generates positive semidefinite periodic operators with nonzero normalization. With

$$\begin{aligned} P &= \mathrm{diag}(1,1,0), & Q &= I - P, \\ \omega_P &= \tfrac{1}{4}(P\otimes P), & \omega_Q &= Q\otimes Q, \end{aligned} \tag{24}$$

its physical closures are

$$M_1(X) = \tfrac{X_{00}}{2}P + X_{11}Q, \qquad\qquad M_2(X) = X_{00}\omega_P + X_{11}\omega_Q. \tag{25}$$

Partial trace and the measure-and-prepare map

$$\begin{aligned} \mathcal{S}(Y) &= \mathrm{tr}_2(Y), \\ \mathcal{R}(Y) &= \mathrm{tr}(PYP)\omega_P + \mathrm{tr}(QYQ)\omega_Q \end{aligned} \tag{26}$$

are trace-preserving completely positive maps satisfying

$$\mathcal{S}[M_2(X)] = M_1(X), \qquad\qquad \mathcal{R}[M_1(X)] = M_2(X). \tag{27}$$

These are the explicit Definition 4.1 equations at lines 638–660. They give SAL, while the slice sum gives the literal Definition 4.2 identity

$$\mathcal{T}_K = I_2, \qquad\qquad \mathcal{T}_K^2 = \mathcal{T}_K. \tag{28}$$

The displayed representatives are normal and biCF, the weights $(1/\sqrt{2}, 1)$ satisfy the global unit-weight convention, and the ambient tensor has normalized fixed-representative simple canonical form. Nevertheless,

$$\rho(\mathcal{E}_{\mu_0 A_0}) = \frac{1}{2}, \qquad\qquad \mu_0 A_0 \text{ is not normal.} \tag{29}$$

The component results and their conjunction are in `TNLean/MPS/MPDO/CaseIIAbsorptionCounterexample.lean`, culminating in `MPOTensor.CaseIIAbsorptionCounterexample.ambient_isSAL_isSimpleCanonicalForm_and_firstAbsorbed_not_isNormalTensor`. This repairs the previously missing full-condition counterexample statement and classifies the printed absorption inference as *proof-path drift*. It does not itself decide the claimed all-sector factorization. The non-Cartesian witness in `TNLean/MPS/MPDO/NonCartesianActiveSectorCounterexample.lean` refutes the implication from injectivity, SAL, literal ZCL, and scalar relation to a normal tensor. Its singleton coefficient has modulus $\sqrt{r} \approx 0.533 < 1$, so the unit clause of the line-246 convention has no witness. Rescaling to unit weight destroys literal ZCL. Thus GitHub issue #6775 remains open at the ambient source-context boundary. The alternative direct construction of the blocked channel pair, using a genuine coarsening or mixing of physical-sector labels, remains in GitHub issue #6793. The subsequent combination through the orthogonal outer physical sectors is formalized in GitHub issue #6632.

6 Unequal raw copies refute Theorem 4.9(iv)⇒(v)

The scale mismatch is decisive even when the source's global unit-weight convention is satisfied. Let $A = 1$ be the sole scalar BNT representative and take two raw copies with weights 1 and 1/2. Then the one-letter tensor has local virtual matrix

$$K^{00} = \mathrm{diag}(1, 1/2).$$

This is exactly a BNT canonical presentation: both weights have modulus at most one and the first has modulus one. Every positive-length contraction is $1 + 2^{-N} > 0$, the scalar representative is normal and nonnilpotent, and the singleton family has simultaneous one-letter span. Condition (iv) also holds: the representative generates MPDOs, the distinct-layer equation is vacuous, and the scalar physical-sector factorization has positive neighboring operator and normalized rank-one trace coefficients.

Nevertheless, the sole local matrix of $K^{[2]}$ is $(K^{00})^2$. If the blocked tensor satisfied Definition 4.1, its physical-trace transfer would be idempotent, so $(K^{00})^4 = (K^{00})^2$. Taking traces gives the false equality

$$1 + 2^{-4} = 1 + 2^{-2}.$$

The declaration in `TNLean/MPS/MPDO/Theorem49RepeatedCopyCounterexample.lean`, `MPOTensor.CaseIIAbsorptionCounterexample.printed_theorem49_iv_to_v_is_false` kernel-checks explicit component witnesses for all standing data, MPDO positivity of the blocked tensor, and this failure in one conjunction. Thus the literal implication (iv)⇒(v) is refuted and classified as an *unfaithful statement*. The printed route has *proof-path drift*: condition (iv) does not constrain repeated-copy weights, so no outer-sector assembly can repair the claim.

The source's viable route is instead (ii)⇒(v). In the earlier Case-II argument, lines 1646–1665 use literal zero correlation length to make all copy weights within each BNT sector equal and absorb their common value. The repeated-copy tensor above fails condition (ii), since $\mathrm{diag}(1, 1/2)$ is not idempotent. Consequently it does not contradict the source proposition `prop2to5`; the remaining work there is either the printed per-representative factorization or another construction of the same blocked channel pairs, followed by their projector-controlled combination after common-weight absorption.

7 The vertical one-label identity

The same scale bookkeeping is visible in the algebraic form of Theorem 4.14. Its length-one condition is

$$m_\gamma = \sum_{\alpha,\beta} c^{(1)}_{\alpha,\beta,\gamma} m_\alpha m_\beta, \tag{30}$$

where $m_\alpha = \mathrm{tr}(\mu_\alpha)$ [CPGSV16, lines 972–993]. In the displayed one-label dimer presentation,

$$c^{(1)}_{0,0,0} = \frac{32}{25}, \qquad\qquad m_0 = \frac{25}{32}, \tag{31}$$

so the nonzero solution is fixed by $(c^{(1)}_{0,0,0})^{-1}$ rather than by a unit-weight convention. This is a statement about that explicit presentation, not an invariance theorem for coefficients under changes of presentation.

8 Horizontal periodic equality does not fix the vertical normalization

The distinction between a horizontal periodic family and a chosen vertical presentation is already visible for one physical letter. Put

$$\begin{aligned} P &= \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \\ X &= \begin{pmatrix} 1 & -3/4 \\ 0 & 1 \end{pmatrix}, \\ Q &= XPX^{-1} = \begin{pmatrix} 1 & 3/4 \\ 0 & 0 \end{pmatrix}. \end{aligned} \tag{32}$$

Let M_P and M_Q be the one-letter MPO tensors with sole matrices P and Q. Since P and Q are idempotent and have trace one, for every positive length N,

$$\rho_N(M_P) = \rho_N(M_Q) = 1. \tag{33}$$

Thus both tensors generate MPDOs and have the same unnormalized periodic operator family. They are related by the nonunitary horizontal gauge X, which is part of the general MPV gauge freedom in [CPGSV16, lines 187–199].

The normalized one-label vertical representatives are different. They are

$$\begin{aligned} A_P &= P, & m_P &= 1, \\ A_Q &= \frac{4}{5}Q, & m_Q &= \frac{5}{4}. \end{aligned} \tag{34}$$

Both are normal bond-one tensors, and each gives a positive one-label algebra clause. Since

$$A_P^2 = A_P,$$
$$A_Q^2 = \frac{4}{5}A_Q,$$

their structure coefficients are

$$c_P^{(L)} = 1,$$
$$c_Q^{(L)} = \left(\frac{4}{5}\right)^L. \tag{35}$$

There is only one relabelling of a singleton label set, so the coefficients cannot agree after relabelling. Moreover, the first family is length independent and the second is not.

The declarations in `TNLean/MPS/MPDO/VerticalCoefficientPresentationCounterexample.lean` prove the gauge relation, MPDO positivity, both tensor-attached algebra clauses, the formulas in (35), and the failure of the relabelled equality. In particular, `MPOTensor.VerticalCoefficientPresentationCounterexample.sameMPVPos_and_no_relabelled_coefficient_equality` refutes the theorem proposed in GitHub issue #6395 under its stated bare positive-length MPV-equality hypothesis.

This is not a counterexample to Proposition 4.13 or Theorem 4.14. The source starts with one horizontally canonical MPDO tensor and then chooses its vertical decomposition [CPGSV16, lines 943–991]. No horizontal canonical-form assertion is made here for the sheared tensor M_Q. The source comparison in Appendix C.4 also retains the positive ratio $m_\alpha/n_{\sigma(\alpha)}$ rather than discarding the normalization scale [CPGSV16, lines 1997–2008]. Consequently, the source's fixed-presentation characterization is unchanged, as is the project's explicit length-dependent RFP example. Rigidity between two literal horizontal canonical presentations is not established by this counterexample.

The scale-covariant statement requires explicit vertical transport. Suppose there are a label equivalence e, nonzero scalars s_α, and algebra isomorphisms Φ_L such that, for every label α,

$$O'_L(\alpha) = s_\alpha^L \Phi_L\big(O_L(e\alpha)\big). \tag{36}$$

At a positive length where the target operators are linearly independent, uniqueness of the product expansion gives

$${c'}^{(L)}_{\alpha,\beta,\gamma} = \left(\frac{s_\alpha s_\beta}{s_\gamma}\right)^L c^{(L)}_{e\alpha,e\beta,e\gamma}. \tag{37}$$

Exact equality therefore requires a normalization condition such as $s_\alpha s_\beta = s_\gamma$ on every active product triple. Deriving (36) from two horizontal canonical presentations, and extending (37) to every positive length, are separate statements and are not supplied by bare horizontal periodic equality.

9 Formalization boundary

Verdict: three distinct formal predicates. The documented active canonical-block predicate `MPOTensor.IsSourceSimple` and its nonvanishing strengthening `MPOTensor.IsNonvanishingSourceSimple` are true for R, with positive blocking length $L = 1$. The normalized fixed-representative predicate `MPOTensor.IsSimple` is false for R: every positive blocking has quasi-idempotence scalar $\left(337/512\right)^L < 1$, while every normalized horizontal BNT witness would force that scalar to be one. The latter conclusion is independent of the choice of normalized BNT witness.

The normalized verdict remains representative-dependent. The line-246 field should not be deleted: it is the source's standing canonical-form convention. The obstruction does not produce a nilpotent BNT block, does not use the vertical coefficient family, and does not refute the documented active source-simplicity predicate. Conversely, the one-site source-simplicity witness does not assert that every positive blocking is source-simple, that arbitrary complex rescaling preserves the MPDO condition, that positive rescaling preserves the same fixed-scale RFP equations, or that the simplicity predicates are equivalent. The valid comparisons are the unconditional implication from normalized simplicity to source simplicity, whose normalized sector decomposition gives an active presentation after gauge transport, and the derived positive-length nontriviality theorem, proved by eventual linear independence, a non-eventually-zero sector coefficient, and unblocking, and the conditional implication to nonvanishing source simplicity using (10), together with strictly positive real rescaling invariance of the MPDO and both active simplicity predicates.[1]

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.

[1] Deletion audit: the declaration-free compatibility shell `TNLean/MPS/MPDO/RescalingStableLengthDependentRFPBNTClause.lean` was removed; the live explicit BNT declarations remain in `TNLean/MPS/MPDO/RescalingStableExplicitVerticalBNT.lean`.