Strong Subadditivity as a Standalone Axiom and Its Lieb-Concavity Elimination Route

2026-06-22

Verdict: open-gap (resolved).

1 The assertion and its current formal status

Strong subadditivity of the von Neumann entropy is the inequality

$$S(\rho_{ABC}) + S(\rho_B) \le S(\rho_{AB}) + S(\rho_{BC}) \tag{SSA}$$

for every density operator ρ_{ABC} on a tripartite finite-dimensional space $\mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_C$, where $S(\rho) = -\operatorname{tr}(\rho \log \rho)$ and the reduced operators are the corresponding partial traces [1].

The inequality (SSA) was previously asserted as a standalone axiom and has now been discharged: it is the theorem `strong_subadditivity_general` in `TNLean/Channel/Schwarz/StrongSubadditivityPosDef`, derived from Lieb concavity along the route recorded below, and the former axiom has been removed from `TNLean/Entropy/SSAEqualityCharacterization` with its single re-export `Entropy.strongSubadditivity` redirected to the derived theorem.[1] The derivation rests on the operator-concavity input that the same development already isolates for the Ando–Lieb theory, namely Lieb's joint concavity theorem.[2] The only entropy-side dependency in the closure of `strong_subadditivity_general`, formerly the axiom, is now the proved theorem `lieb_concavity_posDef`. This note records the mathematical route that connects the two, layer by layer.

[1] The entropy is the eigenvalue-based $S(\rho) = \sum_i \eta(\lambda_i)$, $\eta(x) = -x \log x$, of `TNLean/Analysis/Entropy`.

[2] Lieb concavity is the theorem `lieb_concavity_posDef` in `TNLean/Analysis/LiebConcavity` (`TNLean/Analysis/LiebConcavity.lean`), consumed by `lieb_concavity` in `TNLean/Channel/Schwarz/AndoLieb`.

2 The point at issue

The standard derivation of (SSA) factors through the quantum relative entropy and its behaviour under coarse-graining. Several of the intermediate objects are now formalized and several are still missing, so the gap is not a single missing lemma but a chain of them; the per-layer status is recorded below and summarized under *Current formal status*. In source order the chain is the following.

1. *Entropy bridge.* The identity $S(\rho) = -\operatorname{tr}(\rho \log \rho)$ relating the trace-of-logarithm form to the eigenvalue form $\sum_i \eta(\lambda_i)$ used in the formal definition. This bridge is the prerequisite that makes the trace-functional manipulations of the relative entropy meaningful for the formal entropy, and it is dispatched separately.[3]

2. *Relative entropy.* The definition

$$D(\rho \,\|\, \sigma) = \operatorname{tr}\big(\rho(\log \rho - \log \sigma)\big), \tag{1}$$

for density operators ρ, σ with $\ker \sigma \subseteq \ker \rho$.

3. *Klein's inequality.* Nonnegativity $D(\rho \,\|\, \sigma) \geq 0$, with equality exactly when $\rho = \sigma$. This is the order property that makes D a divergence. The nonnegativity half is now formalized on the full support domain: `quantumRelativeEntropy_nonneg_general` in `TNLean/Analysis/KleinInequality` (`TNLean/Analysis/KleinInequality.lean`) proves $D(\rho \,\|\, \sigma) \geq 0$ for density operators ρ, σ with $\ker \sigma \subseteq \ker \rho$, exactly the source hypothesis. The full-rank base case `quantumRelativeEntropy_nonneg` is proved first by the eigenbasis reduction to the doubly stochastic overlap matrix and the scalar Gibbs inequality $\log x \leq x - 1$; the support case follows by regularizing the reference state to $\sigma'_\varepsilon = (1 + \varepsilon N)^{-1}(\sigma + \varepsilon \mathbf{1})$, which is positive definite for $\varepsilon > 0$, and passing to the limit $\varepsilon \to 0^+$. Because σ'_ε shares σ's eigenbasis, the matrix limit reduces to scalar logarithm limits on the positive eigenvalues, while the support condition kills the diagonal weights on the zero eigenvalues (`tendsto_re_trace_mul_log_perturb`); no eigenvalue-continuity input is needed. It consumes neither Lieb concavity nor operator Jensen. The equality case is now formalized for full-rank σ: `quantumRelativeEntropy_eq_zero_iff` proves $D(\rho \,\|\, \sigma) = 0 \iff \rho = \sigma$ by sharpening the Gibbs reduction. When $D = 0$ every term of the Gibbs slack is

[3] Tracked on GitHub issue #784.

tight, so the strict tangent inequality $\log x < x-1$ for $x \neq 1$ forces the eigenvalues to match on the overlap support; the matching makes the overlap matrix intertwine the eigenvalue diagonals, and conjugating σ into ρ's eigenbasis recovers ρ, giving $\rho = \sigma$. This argument too consumes neither Lieb concavity nor operator Jensen.

4. *Joint convexity.* Convexity of $(\rho, \sigma) \mapsto D(\rho \,\|\, \sigma)$ in the pair of arguments. This is the step that consumes Lieb's joint concavity theorem: the concavity of $(A, B) \mapsto \operatorname{tr}\big(K^\dagger A^s K B^{1-s}\big)$ for $s \in [0,1]$ yields, in the limit governing the derivative at $s = 1$, the convexity of the relative entropy. The present formal Lieb theorem is stated for positive definite inputs, whereas the relative-entropy domain in layer (2) allows singular reference states provided $\ker \sigma \subseteq \ker \rho$. Therefore this layer also needs the standard extension from the full-rank case to the support-condition case: replace the pair by strictly positive regularizations, prove convexity there, and pass to the limit using continuity on the common support and lower semicontinuity at the boundary. Without this regularization step, the later support lemma would place the singular pairs in the domain of D but would not yet connect them to the positive-definite Lieb input. This layer is now formalized on the full support domain: the positive-definite case (`convexOn_quantumRelativeEntropy`) is the base case, and the singular extension under $\ker \sigma \subseteq \ker \rho$ is `convexOn_quantumRelativeEntropy_support`, obtained by regularizing both arguments through the affine trace-one perturbation $M_\varepsilon = (1+\varepsilon N)^{-1}(M + \varepsilon\,\mathbf{1})$ and passing to the limit $\varepsilon \to 0^+$.

5. *Data-processing inequality (DPI).* Monotonicity $D(\Lambda\rho \,\|\, \Lambda\sigma) \leq D(\rho \,\|\, \sigma)$ under every trace-preserving completely positive map Λ, in particular under the partial trace. Monotonicity under the partial trace is the form of DPI used below; it follows from joint convexity together with unitary invariance of D and the ancilla-additivity identity $D(\rho \otimes \tau \,\|\, \sigma \otimes \tau) = D(\rho \,\|\, \sigma)$ for a fixed ancilla state τ. The standard twirling argument writes the partial trace tr_C as an average of unitary conjugations on $\mathcal{H}_C$ applied after appending the maximally mixed ancilla $\tau = \mathbf{1}_C/d_C$, so all three ingredients — joint convexity, unitary invariance, and ancilla-additivity — are needed; joint convexity and unitary invariance alone do not suffice.

6. *Final assembly.* The inequality (SSA) read as a single instance of DPI under the partial trace over C. Take the reference state $\sigma_{ABC} = (\mathbf{1}_A/d_A) \otimes \rho_{BC}$, the tensor product of the maximally mixed state $\mathbf{1}_A/d_A$

on $\mathcal{H}_A$ ($d_A = \dim \mathcal{H}_A$) with the reduced state ρ_{BC} on $\mathcal{H}_B \otimes \mathcal{H}_C$. Being a density operator is *not* enough to place the pair $(\rho_{ABC}, \sigma_{ABC})$ in the domain fixed in layer (2): that domain is the kernel-containment condition $\ker \sigma_{ABC} \subseteq \ker \rho_{ABC}$, and as soon as ρ_{BC} fails to be full rank the reference state σ_{ABC} is singular, so the condition is a genuine hypothesis to be discharged and is also what licenses the displayed logarithm. The containment is supplied by the following *marginal support lemma*: for any density operator $\rho_{ABC} \geq 0$ with marginal $\rho_{BC} = \mathrm{tr}_A\, \rho_{ABC}$,

$$\mathcal{H}_A \otimes \ker \rho_{BC} \subseteq \ker \rho_{ABC}, \qquad \text{equivalently} \qquad \operatorname{supp} \rho_{ABC} \subseteq \mathcal{H}_A \otimes \operatorname{supp} \rho_{BC}. \tag{2}$$

To see (2), let P_{BC} be the projection onto $\operatorname{supp} \rho_{BC}$ and $Q = \mathbf{1}_A \otimes (\mathbf{1}_{BC} - P_{BC})$ the projection onto $\mathcal{H}_A \otimes \ker \rho_{BC}$. Using $\mathrm{tr}_A\, \rho_{ABC} = \rho_{BC}$ and $(\mathbf{1}_{BC} - P_{BC})\, \rho_{BC} = 0$,

$$\mathrm{tr}\big(Q\, \rho_{ABC}\big) = \mathrm{tr}\big((\mathbf{1}_{BC} - P_{BC})\ \mathrm{tr}_A\, \rho_{ABC}\big) = \mathrm{tr}\big((\mathbf{1}_{BC} - P_{BC})\, \rho_{BC}\big) = 0. \tag{3}$$

Since $\rho_{ABC} \geq 0$ and Q is a projection, $\mathrm{tr}(Q\, \rho_{ABC}) = \|Q\, \rho_{ABC}^{1/2}\|_2^2 = 0$ forces $Q\, \rho_{ABC}^{1/2} = 0$, hence $Q\, \rho_{ABC} = 0$, so the range of Q — namely $\mathcal{H}_A \otimes \ker \rho_{BC}$ — lies in $\ker \rho_{ABC}$. Because $\mathbf{1}_A/d_A$ is full rank we have $\ker \sigma_{ABC} = \mathcal{H}_A \otimes \ker \rho_{BC}$, so (2) gives exactly $\ker \sigma_{ABC} \subseteq \ker \rho_{ABC}$ and places the pair in the domain of layer (2). On $\operatorname{supp} \sigma_{ABC}$, where the relative entropy is evaluated, the logarithm splits as $\log\big((\mathbf{1}_A/d_A) \otimes \rho_{BC}\big) = -(\log d_A)\, \mathbf{1}_{ABC} + \mathbf{1}_A \otimes \log \rho_{BC}$ (with $\log \rho_{BC}$ taken on $\operatorname{supp} \rho_{BC}$), and the support condition guarantees ρ_{ABC} carries no weight off this support, so the relative entropy evaluates to

$$D\big(\rho_{ABC} \,\|\, (\mathbf{1}_A/d_A) \otimes \rho_{BC}\big) = \log d_A + S(\rho_{BC}) - S(\rho_{ABC}). \tag{4}$$

Applying the partial trace $\Lambda = \mathrm{tr}_C$ sends $\rho_{ABC} \mapsto \rho_{AB}$ and $(\mathbf{1}_A/d_A) \otimes \rho_{BC} \mapsto (\mathbf{1}_A/d_A) \otimes \rho_B$. The image pair inherits its own support condition from (2) applied to ρ_{AB} with marginal $\rho_B = \mathrm{tr}_A\, \rho_{AB}$, so it again lies in the domain of layer (2) and the image relative entropy is

$$D\big(\rho_{AB} \,\|\, (\mathbf{1}_A/d_A) \otimes \rho_B\big) = \log d_A + S(\rho_B) - S(\rho_{AB}). \tag{5}$$

Monotonicity under Λ, namely $D(\Lambda\rho \,\|\, \Lambda\sigma) \leq D(\rho \,\|\, \sigma)$, then reads $\log d_A + S(\rho_B) - S(\rho_{AB}) \leq \log d_A + S(\rho_{BC}) - S(\rho_{ABC})$; the common

$\log d_A$ cancels, leaving $S(\rho_B) - S(\rho_{AB}) \leq S(\rho_{BC}) - S(\rho_{ABC})$, which on rearrangement is exactly (SSA); equivalently it is the nonnegativity of the conditional mutual information $I(A : C \mid B) \geq 0$. The discarded subsystem is C and the reference state carries the maximally mixed state on the *retained* factor A; discarding A instead against $\rho_A \otimes \rho_{BC}$ would only yield subadditivity $I(A : BC) \geq 0$, not strong subadditivity.

The mathematical question is therefore not whether (SSA) is true — it is a theorem of Lieb and Ruskai — but which of the layers (2)–(6) must be built so that the standalone axiom can be deleted and replaced by a proof whose only entropy-side input is the Lieb-concavity theorem already present.

3 The elimination route

The route deletes the axiom `strong_subadditivity` once the chain above is in place. Concretely:

- Layer (1) is the entropy bridge and is independent of Lieb concavity; it is the lemma that lets the eigenvalue-based formal entropy be treated as a trace of a matrix logarithm.

- Layers (2) and (3) define the divergence and establish its nonnegativity; they need only the matrix logarithm and its operator monotonicity, not Lieb concavity. Layer (2) is in place (`quantumRelativeEntropy`) and layer (3), Klein's inequality, is now formalized on the full support domain $\ker \sigma \subseteq \ker \rho$ (`quantumRelativeEntropy_nonneg_general`), via the full-rank eigenbasis reduction and scalar Gibbs inequality together with the support-respecting regularization, rather than operator monotonicity of the logarithm.

- Layer (4) is the single consumer of `lieb_concavity_posDef`: joint convexity of D is the entropy-theoretic content of Lieb's theorem. This layer is now formalized on the positive-definite domain: `convexOn_quantumRelativeEntropy` in `TNLean/Channel/Schwarz/RelativeEntropyConvexity` (`TNLean/Channel/Schwarz/RelativeEntropyConvexity.lean`) proves that $(\rho, \sigma) \mapsto D(\rho \,\|\, \sigma)$ is jointly convex on pairs of positive definite matrices, via the Lindblad/Uhlmann approximant $g_s(\rho, \sigma) = (\operatorname{tr} \rho - \operatorname{Re} \operatorname{tr}(\rho^s \sigma^{1-s}))/(1 - s)$: each g_s is jointly convex for $s \in [0, 1)$ (convexOn_relativeEntropyApprox, the consumer of lieb_concavity_posDef) and converges pointwise to D as $s \to 1^-$ (tendsto_relativeEntropyApprox), and the pointwise limit of convex functions is convex. The two-calculus trace double-sum identity re_trace_cfc_mul_cfc_eq_double_sum underlying the limit consumes no Lieb input. The extension from the positive-definite case to singular reference states under $\ker \sigma \subseteq \ker \rho$ is now formalized as convexOn_quantumRelativeEntropy_support in the same file: both arguments are regularized through the affine trace-one perturbation, the positive-definite convexity applies to the four regularized endpoints and the regularized mixture, and the inequality passes to the limit using the both-arguments limit lemma tendsto_relativeEntropyPerturb. The domain is convex because the kernel of a positively weighted sum of positive semidefinite matrices is the intersection of the kernels (mulVec_eq_zero_of_smul_add). This extension consumes no further Lieb input beyond the positive-definite base case.

- Layer (5) combines joint convexity, unitary invariance, and ancilla-additivity $D(\rho \otimes \tau \,\|\, \sigma \otimes \tau) = D(\rho \,\|\, \sigma)$ into the partial-trace monotonicity statement via the twirling representation of tr_C. The unitary-invariance ingredient is now formalized for arbitrary Hermitian arguments: quantumRelativeEntropy_conj_unitary in TNLean/Channel/Schwarz/RelativeEntropyUnitaryInvariance (TNLean/Channel/Schwarz/RelativeEntropyUnitaryInvariance.lean) proves $D(U\rho U^\dagger \,\|\, U\sigma U^\dagger) = D(\rho \,\|\, \sigma)$ for every unitary U, by the covariance of the matrix logarithm under unitary conjugation (Matrix.log_conj_unitary, itself the $f = \log$ case of the functional-calculus covariance Matrix.cfc_conj_unitary) followed by trace cyclicity. This ingredient needs no Lieb input. The ancilla-additivity ingredient is formalized on the positive-definite domain: quantumRelativeEntropy_kronecker in TNLean/Channel/Schwarz/RelativeEntropyAncillaAdditivity (TNLean/Channel/Schwarz/RelativeEntropyAncillaAdditivity.lean) proves $D(\rho\otimes\tau \,\|\, \sigma\otimes\tau) = D(\rho \,\|\, \sigma)$ for positive definite ρ, σ and a positive definite τ of unit trace. It reduces to the logarithm of a tensor product $\log(\rho\otimes\tau) = \log\rho\otimes\mathbf{1}+\mathbf{1}\otimes\log\tau$ (Matrix.log_kronecker): the one-sided splits $\log(A\otimes\mathbf{1}) = \log A\otimes\mathbf{1}$ (Matrix.cfc_kronecker_one) and $\log(\mathbf{1}\otimes B) = \mathbf{1}\otimes\log B$ (Matrix.cfc_one_kronecker) follow from the functional calculus through the unital tensor embeddings, and recombining them uses the logarithm of a commuting product $\log(XY) = \log X + \log Y$ for commuting positive definite X, Y

(Matrix.PosDef.cfc_log_mul in TNLean/Analysis/CfcLogAdditive, TNLean/Analysis/CfcLogAdditive.lean). The difference of tensor logarithms then cancels the $\mathbf{1} \otimes \log \tau$ terms, the trace of the Kronecker product factors as a product of traces, and the unit trace of τ leaves $D(\rho \,\|\, \sigma)$. The corresponding singular-support identity is quantumRelativeEntropy_kronecker_support in TNLean/Channel/Schwarz/PetzEqualityPrerequisites (TNLean/Channel/Schwarz/PetzEqualityPrerequisites.lean): it assumes positive semidefiniteness, $\ker \sigma \subseteq \ker \rho$, and unit trace of τ, but no normalization of ρ or σ. Its proof uses the tensor-logarithm formula on support projections and support absorption. These results depend only on the standard Lean axioms and need no Lieb input. The third ingredient of layer (5) is the twirling identity, which writes $\operatorname{tr}_C \rho \otimes (\mathbf{1}_C/d_C)$ as a uniform average of conjugations by a unitary 1-design on C. Because the current dependency carries no Haar measure on the matrix unitary group, the average is taken over the finite Heisenberg–Weyl 1-design. Its design property is now formalized: Matrix.sum_weyl_conj in TNLean/Channel/Schwarz/WeylTwirl (TNLean/Channel/Schwarz/WeylTwirl.lean) proves that for a primitive d-th root of unity the uniform average of the conjugations by the d^2 Weyl operators $W(a,b) = X^a Z^b$ equals the completely depolarizing channel $M \mapsto (\operatorname{tr} M/d)\,\mathbf{1}$, via the root-of-unity character-sum orthogonality (Matrix.sum_rootOfUnity_pow_eq). This result depends only on the standard Lean axioms and needs no Lieb input. Lifting that design property to the ancilla factor presents tr_C through the average: Matrix.sum_kronecker_one_weyl_conj in TNLean/Channel/Schwarz/RelativeEntropyDataProcessing (TNLean/Channel/Schwarz/RelativeEntropyDataProcessing.lean) proves that the uniform average of the conjugations by the d_C^2 unitaries $\mathbf{1}_S \otimes W(a,b)$ equals $(\operatorname{tr}_C M) \otimes (\mathbf{1}_C/d_C)$. With all three ingredients in place, layer (5) is now formalized on the positive-definite domain: quantumRelativeEntropy_partialTraceRight_le in the same file proves $D(\operatorname{tr}_C \rho \,\|\, \operatorname{tr}_C \sigma) \le D(\rho \,\|\, \sigma)$ for positive definite ρ, σ, via the ancilla-additivity expansion of the reduced-state relative entropy, the twirling identity, joint convexity over the d_C^2 equally weighted conjugated pairs, and unitary invariance of each term. The transport of joint convexity to the product index uses the reindexing invariance quantumRelativeEntropy_submatrix_equiv. The extension to the singular support domain $\ker \sigma \subseteq \ker \rho$ is now formalized as quantumRelativeEntropy_partialTraceRight_le_support in the same file: both arguments are regularized through the affine trace-shrinking perturbation $M_\varepsilon = (1+N_{SC}\varepsilon)^{-1}(M+\varepsilon\,\mathbf{1})$, which is positive definite for $\varepsilon > 0$, so the positive-definite case applies at each ε. The right-hand side converges to $D(\rho\,\|\,\sigma)$ by the both-arguments limit. The left-hand side is the relative entropy of the marginals, which are themselves affine regularizations of $\operatorname{tr}_C \rho, \operatorname{tr}_C \sigma$ with the same scaling rate N_{SC} but shift rate d_C (`partialTraceRight_regPerturbAffine`); the support condition transfers to the marginals (`Matrix.partialTraceRight_support`), so the same limit applies once it is generalized to an arbitrary positive-affine regularization $M \mapsto (1+a\varepsilon)^{-1}(M+b\varepsilon\,\mathbf{1})$ (`TNLean.RelativeEntropyConvexity.tendsto_relativeEntropyPerturbAffine`). Passing the inequality through both limits gives the support-domain bound. This extension consumes no further Lieb input beyond the positive-definite base case.

- Layer (6) instantiates layer (5) at the tripartite splitting to recover (SSA); it additionally needs the marginal support lemma (2), which puts the singular reference state $(\mathbf{1}_A/d_A)\otimes\rho_{BC}$ into the domain of layer (2) and so is a prerequisite of the instantiation rather than a consequence of normalization. The marginal support lemma is now formalized (`Matrix.marginal_support` in `TNLean/Analysis/MarginalSupport`, `TNLean/Analysis/MarginalSupport.lean`): for a positive semidefinite ρ_{ABC} a Hermitian idempotent on $B\otimes C$ annihilating the marginal $\operatorname{tr}_A \rho_{ABC}$ has its identity-on-A lift annihilate ρ_{ABC}. Its operator core — a projection of zero expectation against a positive semidefinite state annihilates it (`Matrix.PosSemidef.proj_mul_eq_zero_of_trace_eq_zero`) — and the partial-trace adjoint identity it uses (`Matrix.traceA_ABC_lift_trace`) are independent of Lieb concavity. These three results depend only on the standard Lean axioms. On the positive-definite domain the instantiation is now formalized: `strong_subadditivity_posDef` in `TNLean/Channel/Schwarz/StrongSubadditivityPosDef` (`TNLean/Channel/Schwarz/StrongSubadditivityPosDef.lean`) proves the inequality (SSA) for a positive definite ρ_{ABC}. There the reference state $(\mathbf{1}_A/d_A)\otimes\rho_{BC}$ is itself positive definite, so the marginal support lemma is not invoked and the layer-(2) domain condition holds outright; the relative entropy of each pair evaluates to the displayed entropy difference through the tensor logarithm split and the partial-trace adjoint, and layer (5) data processing between the two pairs yields (SSA). The singular-reference extension to the general density-matrix domain is now formalized as well: `strong_subadditivity_general` in the same file proves the inequality (SSA) for every positive semidefinite unit-trace ρ_{ABC}. The marginal support lemma places the singular reference state $(\mathbf{1}_A/d_A) \otimes \rho_{BC}$ in the domain of layer (2) through the kernel projection of ρ_{BC}, after which the singular-domain data processing of layer (5) applies. The two relative entropies evaluate to the displayed entropy differences by the singular-reference evaluation `rel_entropy_eval_support`, which regularizes the reduced state through the affine perturbation $(1 + d_R\varepsilon)^{-1}(\rho_R + \varepsilon\mathbf{1})$ and passes to the limit, the tensor logarithm split being unavailable for a singular reference. Its only entropy-side dependency, formerly the axiom, is now the proved theorem `lieb_concavity_posDef`, so this theorem discharges the content of the standalone strong subadditivity axiom on the full density-matrix domain.

Current formal status. Layers (1)–(6) are formalized, with layers (3), (4), (5), and the layer-(6) instantiation on the full density-matrix domain. Layer (4) joint convexity is formalized on the singular support domain $\ker\sigma \subseteq \ker\rho$ (`convexOn_quantumRelativeEntropy_support`), extending the positive-definite base case by regularizing both arguments and passing to the limit. Layer (5), data-processing (monotonicity of D under the partial trace), combines the layer-(4) joint convexity with unitary invariance and ancilla-additivity $D(\rho\otimes\tau \,\|\, \sigma\otimes\tau) = D(\rho \,\|\, \sigma)$ via the twirling representation of tr_C. Joint convexity (layer 4, on the singular support domain), unitary invariance (`quantumRelativeEntropy_conj_unitary`, for arbitrary Hermitian arguments), ancilla-additivity (`quantumRelativeEntropy_kronecker`, on the positive-definite domain), the finite Heisenberg–Weyl 1-design twirl (`Matrix.sum_weyl_conj`), and its ancilla-factor lift presenting tr_C through the twirl (`Matrix.sum_kronecker_one_weyl_conj`) combine into the partial-trace monotonicity `quantumRelativeEntropy_partialTraceRight_le` on the positive-definite domain, which is then extended to the singular support domain $\ker\sigma \subseteq \ker\rho$ as `quantumRelativeEntropy_partialTraceRight_le_support` by regularizing both arguments, transferring the support condition to the marginals (`Matrix.partialTraceRight_support`), and passing to the limit through the arbitrary-rate affine regularization (`TNLean.RelativeEntropyConvexity.tendsto_relativeEntropyPerturbAffine`). The layer-(6) instantiation is formalized on the full density-matrix domain. The positive-definite case (`strong_subadditivity_posDef`) is the base case, and the singular-reference extension to every positive semidefinite unit-trace ρ_{ABC} is `strong_subadditivity_general`, whose only entropy-side dependency, formerly the axiom, is now the proved theorem `lieb_concavity_posDef`. The marginal support lemma places the singular reference state $(\mathbf{1}_A/d_A)\otimes\rho_{BC}$ in the layer-(2) domain through the kernel projection of ρ_{BC}; the two relative entropies evaluate to the displayed entropy differences by the singular-reference evaluation `rel_entropy_eval_support` (which regularizes the reduced state and passes to the limit, the tensor logarithm split being unavailable for a singular reference); and the singular-domain data processing of layer (5) between the two pairs yields the inequality. Because `strong_subadditivity_general` carries exactly the hypotheses of the standalone axiom `strong_subadditivity`, it discharges the content of that axiom on the full density-matrix domain.

The inequality is now a theorem on the full density-matrix domain whose only entropy-side dependency on this route, formerly the axiom, is now the proved theorem `lieb_concavity_posDef`; the standalone axiom can be removed and its consumers redirected to the derived theorem. The downstream consumers in the stable entropy interface [4] require no change because the statement they import is unchanged; only its justification changes from axiom to theorem.

4 Verdict

The strong-subadditivity inequality is now a theorem connected to the Lieb-concavity theorem (formerly asserted as an axiom, now proved as `lieb_concavity_posDef`) available in the same repository; the former standalone axiom has been removed. The elimination route is the classical relative-entropy chain: entropy bridge, definition of relative entropy, Klein's inequality, joint convexity (the step consuming Lieb concavity), data-processing under the partial trace, and final assembly at the tripartite splitting. Layers (2) and (3) are formalized (the latter on the full support domain $\ker\sigma \subseteq \ker\rho$, `quantumRelativeEntropy_nonneg_general`), and layer (4) is formalized on the same full support domain: the positive-definite base case (`convexOn_quantumRelativeEntropy`) is extended to $\ker\sigma \subseteq \ker\rho$ by `convexOn_quantumRelativeEntropy_support` through the both-arguments regularization, so the Lieb-concavity theorem is now connected to the joint convexity of the relative entropy on the full support domain.

[4] The stable interface re-exports the inequality from `TNLean/Entropy/StrongSubadditivity`; consumers import it from there and not from the axiom module.

Layer (5), data-processing, is formalized on the full support domain: the positive-definite case (`quantumRelativeEntropy_partialTraceRight_le`) combines unitary invariance (`quantumRelativeEntropy_conj_unitary`), ancilla additivity (`quantumRelativeEntropy_kronecker`), the finite Heisenberg–Weyl 1-design twirl (`Matrix.sum_weyl_conj`), and the ancilla-factor lift presenting tr_C through the twirl (`Matrix.sum_kronecker_one_weyl_conj`) through joint convexity into the partial-trace monotonicity $D(\mathrm{tr}_C\,\rho \,\|\, \mathrm{tr}_C\,\sigma) \le D(\rho \,\|\, \sigma)$, and the singular extension under $\ker\sigma \subseteq \ker\rho$ (`quantumRelativeEntropy_partialTraceRight_le_support`) follows by regularizing both arguments, transferring the support condition to the marginals (`Matrix.partialTraceRight_support`), and passing to the limit through the arbitrary-rate affine regularization. The layer-(6) instantiation is formalized on the full density-matrix domain. The positive-definite case (`strong_subadditivity_posDef`), where the positive-definite reference state $(\mathbf{1}_A/d_A) \otimes \rho_{BC}$ sidesteps the marginal support lemma, is the base case; the singular-reference extension to every positive semidefinite unit-trace ρ_{ABC} is `strong_subadditivity_general`, which places the singular reference state in the layer-(2) domain through the marginal support lemma and evaluates the two relative entropies by the singular-reference evaluation `rel_entropy_eval_support`. The inequality is therefore a theorem on the full density-matrix domain whose only entropy-side dependency, formerly the axiom, is now the proved theorem `lieb_concavity_posDef`, and it carries exactly the hypotheses of the former standalone axiom, so that axiom has been removed and its single re-export `Entropy.strongSubadditivity` redirected to the derived theorem.

References

[1] E. H. Lieb and M. B. Ruskai, *Proof of the strong subadditivity of quantum-mechanical entropy*, Journal of Mathematical Physics **14** (1973), 1938–1941.

[2] E. H. Lieb, *Convex trace functions and the Wigner–Yanase–Dyson conjecture*, Advances in Mathematics **11** (1973), 267–288.