Manifold Relativity: Sub-Additivity from Spectral Truncation

Preprint v15.0 — Computational Record Reconciliation

and Product-Regime Verification

Continuing the Manifold Relativity programme
(v1–v11 appeared under the historical title "Entropy Waves, Coordinate Systems,

and the Self-Referential Universe")

Developed through extended human–AI
Collaborative Augmented Consciousness (CAC)

Human Node:
Paul E. Sorvik, Alexandria, Egypt
ORCID: 0009-0008-5717-7110

AI Builder Nodes:
Claude Opus 4.6 (Anthropic) *Structural & Epistemic Logic*
Gemini 3.1 Pro (Google DeepMind) *Geometric Synthesis*

AI Referee Node:
ChatGPT (OpenAI) *CAC Referee / Analytical Reviewer*

Available: paulsorvik.wordpress.com

April 2026

Abstract

Version 12 of the Manifold Relativity programme consolidated the framework's terminology, proved three structural propositions about spectral accessibility, and stated candidate observables with falsification criteria. It identified two frontier targets: O31 (emergence of the κ-addition composition law from spectral truncation) and O1 (derivation of the off-diagonal spatial metric component $g_{IE}(I, E)$).

This edition advances O31 computationally and clarifies O1 structurally.

First, a **convention lock** fixes the spectral filter definition and its direction once, preventing sign or cutoff ambiguity in downstream calculations.

Second, we perform an explicit computation of O31 across multiple system sizes. Composite toy Dirac operators from 4×4 to 16×16 dimensions are spectrally truncated, and the resulting entropy composition is compared against κ-addition.

Result: Spectral truncation produces sub-additive entropy composition in every bipartite test case, at every system size tested (up to 16×16). The effective κ parameter increases monotonically with the fraction of spectrum retained: more accessible spectrum means weaker apparent correlation. The full-spectrum limit recovers additive composition ($\kappa \to \infty$). However, the specific κ-addition functional form (defect proportional to $S_A \cdot S_B$) is not uniquely selected by the computation: alternative sub-additive forms fit the data comparably well. The qualitative mechanism of the v7 bridge conjecture is validated; its exact quantitative form remains open.

Third, we derive an **exact analytical formula** for the mutual information created by spectral truncation of product thermal states. The mutual information is a Jensen gap measuring the non-uniformity of spectral accessibility across subsystems. It vanishes if and only if the truncation mask is a product set. This identifies the precise mechanism: the non-product geometry of the truncation mask is the sole source of apparent correlation.

Fourth, we clarify the structure of the O1 inverse problem. A critical finding is that the vacuum Einstein equations are trivially satisfied for *any* two-dimensional metric, making the naive "backward EFE recovery" constraint vacuous when applied to the (I, E) sector alone. The actual constraints on g_{IE} are identified: the three-dimensional metric construction via the replacement product (v3), the determinant condition for spatial dimensionality, the Ryu–Takayanagi geodesic structure, and the rotation map convergence.

Fifth, we prove a Space–Time Complementarity Corollary from the v12 structural propositions: spatial spectral richness and temporal processing rate are inversely related across temperature, with explicit consequences for cycle-endpoint behaviour.

Sixth, v15 reconciles the v13/v14 computational record. The 9×9 3-site-chain discrepancy reported in v13/v14 is resolved as a methodological artifact in the original verification script (Script 05): the analytical formula applied a product-basis mask, while the numerical comparison projected the thermal state onto the composite eigenbasis, which rotates arbitrarily within degenerate subspaces under floating-point diagonalisation. Forensic reconstruction (Script 07b) confirms that the two methods computed different truncated states. Under a consistent product-basis mask, Proposition 3.8 is verified to machine precision across all tested truncation levels in the product-Hamiltonian regime, including cases with composite spectral degeneracies. The v14 "domain boundary" framing is correspondingly retracted.

What changed in v15. Version 15 reconciles the v13/v14 computational record after an internal adversarial probe (Script 07b) identified the previously reported 3-site chain discrepancy as a methodological artifact rather than a failure of Proposition 3.8. The original Script 05 compared an analytical calculation (over a product-basis mask) against a numerical calculation (over the composite eigenbasis, which numpy's `eigh` rotates arbitrarily within degenerate subspaces). The two computations therefore evaluated different truncated states. Under a consistent product-basis mask, Proposition 3.8 is verified to machine precision across all tested product-Hamiltonian cases, including the 9×9 3-site-chain case. Remark 3.9 is correspondingly rewritten (Section 3.8 documents the reconciliation as a preserved evidence trail). The α/κ functional-form non-uniqueness result (O38) is unaffected and remains open.

Contents

Convention Lock

All calculations in this edition use the following fixed convention, inherited from v12 Definition 2.1.

Definition 1.1 (Spectral Filter Convention — Locked)**.** The accessible spectral sector at temperature T is:

$$U_T := \{\lambda \in \mathrm{spec}(D_{\mathcal{W}}) : |\lambda| \geq k_{\mathrm{B}}T/\hbar\}. \tag{1}$$

The spectral filter Π_T projects onto U_T. The convention is:

- Higher $T \implies$ higher threshold $\implies$ fewer eigenvalues kept $\implies$ smaller accessible sector.
- Lower $T \implies$ lower threshold $\implies$ more eigenvalues kept $\implies$ larger accessible sector.
- $T \to 0$: threshold $\to 0$, full spectrum accessible.
- $T \to T_P$: threshold $\to E_{\mathrm{P}}/\hbar$, maximal truncation.

This convention is load-bearing for the v12 monotone filtration (Proposition 3.1), the P20 measurement-floor prediction, the P23 chart-mismatch residual, and the RG-flow conjecture (O36).

Related Work and Positioning

The Manifold Relativity programme sits at the intersection of three bodies of prior work: nonextensive statistical mechanics, the quantum-information theory of thermal states, and the theory of thermodynamic relativity. This section locates the v13/v14 results within that landscape.

Nonextensive Statistical Mechanics

The κ-addition composition law tested in Section 3 belongs to the family of generalised entropy frameworks initiated by Tsallis [4], who proposed the q-entropy $S_q = (1 - \sum_i p_i^q)/(q-1)$ as a generalisation of Boltzmann–Gibbs statistics admitting non-extensive composition. The thermodynamic consistency of such frameworks — including stability, first-law structure, and uniqueness of the generalised entropy functional — was developed by Abe and collaborators [5, 6].

The present work differs in derivation posture. Nonextensive frameworks typically *postulate* a generalised entropy functional and derive composition laws from it. We instead *derive* sub-additivity as a topological consequence of spectral truncation of a candidate Dirac operator, with the specific functional form of the defect left open (O38). In this sense, the κ-addition form — if confirmed at larger operator dimensions — would appear in our framework as an emergent macroscopic law induced by coarse-graining, not as a foundational postulate.

Mutual Information in Thermal States

The analytical formula of Proposition 3.8 sits within the quantum-information literature on mutual information and correlation structure in Gibbs and thermal states. The backbone result is the strong subadditivity of von Neumann entropy [7], with related convexity and subadditivity properties of quantum entropy established in [8]. More recent work has established decay bounds on conditional mutual information in high-temperature Gibbs states [9], and structural results on matrix-product representations of thermal states [10].

The novelty of Proposition 3.8 within this context is the identification of the mutual information as a Jensen gap measuring the non-product geometry of a spectral truncation mask, with the exact closed form (14) under diagonal-product-state preconditions. The domain-of-validity boundary of that formula is discussed in Remark 3.9.

Thermodynamic Relativity and the Entropy Defect

The bridge target of the programme is the theory of thermodynamic relativity developed by Livadiotis and McComas [2], which proposes that entropy and velocity obey analogous relativistic composition laws with a shared invariant upper bound κ. The entropy-defect concept underlying that framework has precursors in Livadiotis' earlier work on κ-distribution thermodynamics and the statistical-mechanical interpretation of the κ index [11, 12].

Our sub-additivity result (Section 3) validates the qualitative mechanism of that bridge — spectral coarse-graining creates non-extensive entropy composition, and the strength of non-extensivity scales with the severity of the coarse-graining — without yet uniquely selecting the specific H-function composition law. The quantitative bridge between κ-addition and the $\mathcal{W}$-atlas remains O38.

The O31 Computation: Sub-Additivity from Spectral Truncation

Setup

Open Problem O31 (v10) asks whether the H-function composition law of thermodynamic relativity [2] emerges from spectral truncation of the candidate Dirac operator. We test this in the simplest non-trivial setting.

Subsystem operators. Following v6.1, each subsystem is a two-site graph with

Dirac operator

$$D_A = D_B = \begin{pmatrix} 0 & a \\ a & 0 \end{pmatrix}, \tag{2}$$

where $a > 0$ is the coupling strength (identified with $c_S = k_B T/\hbar$ in v6.1). The spectrum is $\{-a, +a\}$.

Composite operator. The composite system $A \otimes B$ has Hilbert space $\mathcal{H} = \mathbb{C}^2 \otimes \mathbb{C}^2 = \mathbb{C}^4$ and Dirac operator

$$D_{AB} = D_A \otimes \mathbf{1}_B + \mathbf{1}_A \otimes D_B. \tag{3}$$

The spectrum of D_{AB} is $\{\lambda_i + \mu_j\}$ where $\lambda_i \in \{-a, +a\}$ and $\mu_j \in \{-b, +b\}$.

For the symmetric case $a = b$: $\text{spec}(D_{AB}) = \{-2a, 0, 0, +2a\}$.

For the asymmetric case $a \neq b$: $\text{spec}(D_{AB}) = \{-(a+b), -(a-b), +(a-b), +(a+b)\}$.

Thermal state. The thermal density matrix at inverse temperature β is

$$\rho = \frac{e^{-\beta D_{AB}}}{\text{Tr}(e^{-\beta D_{AB}})}. \tag{4}$$

For a product Hamiltonian $D_{AB} = D_A \otimes \mathbf{1} + \mathbf{1} \otimes D_B$, the thermal state factors: $\rho = \rho_A \otimes \rho_B$. The von Neumann entropy is therefore *additive*: $S(\rho) = S(\rho_A) + S(\rho_B)$.

Spectral Truncation

Definition 3.1 (Spectral Truncation of a Thermal State)**.** Given a threshold $\theta > 0$, the spectral truncation Π_T retains eigenstates of D_{AB} with $|\lambda| \geq \theta$ and removes those with $|\lambda| < \theta$. The truncated state is the renormalised projection:

$$\rho_\theta \coloneqq \frac{\Pi_T \, \rho \, \Pi_T}{\text{Tr}(\Pi_T \, \rho \, \Pi_T)}. \tag{5}$$

For the symmetric case with $\theta > 0$ (removing the two zero-eigenvalue states), the truncated subspace is spanned by $\{|{+}{+}\rangle, |{-}{-}\rangle\}$ — a maximally correlated subspace. Even though the full state ρ is a product state with no correlations, the truncated state ρ_θ lives in a correlated subspace.

The truncation creates apparent correlation by discarding the uncorrelated sector.

The Entropy Defect

For the full (untruncated) state:

$$S(\rho) = S(\rho_A) + S(\rho_B) \qquad \text{(additive, zero defect)}. \tag{6}$$

For the truncated state, define the marginals $\rho_\theta^A := \mathrm{Tr}_B(\rho_\theta)$ and $\rho_\theta^B := \mathrm{Tr}_A(\rho_\theta)$. The entropy defect is:

$$\Delta S := S(\rho_\theta) - S(\rho_\theta^A) - S(\rho_\theta^B) < 0. \tag{7}$$

Result: Sub-Additivity Confirmed; Exact κ-Form Remains Open

Computation 3.2 (O31 Toy Result). For every parameter combination tested — symmetric and asymmetric couplings, multiple inverse temperatures β, multiple truncation thresholds θ, and system dimensions from 4 to 16 — spectral truncation of a product thermal state produces **sub-additive** entropy composition:

$$S(\rho_\theta) < S(\rho_\theta^A) + S(\rho_\theta^B) \qquad \text{(negative entropy defect).} \tag{8}$$

The sub-additivity vanishes in the full-spectrum limit ($\theta \to 0$, no truncation), recovering standard additive composition.

The defect can be expressed in κ-addition form:

$$S(\rho_\theta) = S(\rho_\theta^A) + S(\rho_\theta^B) - \frac{1}{\kappa}\, S(\rho_\theta^A)\, S(\rho_\theta^B), \tag{9}$$

where $\kappa > 0$ is defined by the defect. However, this expression is formally satisfiable for *any* bipartite entropy split with nonzero marginal entropies (one free parameter, one equation). The non-trivial content is not the formal match of the formula but the following three properties:

1. **Sub-additivity:** The defect is always negative (entropy is sub-additive) for direct spectral truncation of product states. This corresponds to $\kappa > 0$.
2. **Monotonic scaling:** The effective κ increases monotonically with the fraction of spectrum retained. More accessible spectrum means weaker apparent correlation.
3. **Additivity recovery:** In the full-spectrum limit, $\kappa \to \infty$ and additive composition is recovered exactly.

Remark 3.3 (The κ-addition form is not uniquely selected). A cross-check was performed: the alternative model $\Delta S = -\alpha(S_A + S_B)$ was fitted alongside $\Delta S = -S_A S_B/\kappa$. At fixed truncation across varying β, the coefficient of variation of α (0.36) is lower than that of κ (0.55). Neither parameter is constant across β, indicating that the true functional form of the defect is more complex than either simple model. The specific κ-addition form is **consistent** with the data but is **not uniquely determined** by the computation. The emergence of the exact Livadiotis–McComas composition law from spectral truncation remains an open question.

Table 1: O31 computation results across system sizes. dim: Hilbert space dimension. n: eigenvalues retained. κ: empirical κ-addition parameter. All listed bipartite cases exhibit negative entropy defect; empirical κ-parameterisation shown for comparison.

System	dim	$n/\dim$	S_{AB}	$S_A + S_B$	ΔS	κ
2-site ⊗ 2-site (symmetric)						
$a{=}b{=}1, \beta{=}0.5$	4	2/4	0.365	0.731	−0.365	0.37
$a{=}b{=}1, \beta{=}1.0$	4	2/4	0.090	0.180	−0.090	0.09
2-site ⊗ 2-site (asymmetric)						
$a{=}1, b{=}0.5$	4	3/4	0.714	0.758	−0.045	2.07
$a{=}1, b{=}0.3$	4	3/4	0.574	0.692	−0.118	0.90
3-site chain ⊗ 3-site chain						
	9	6/9	1.427	1.661	−0.234	2.95
	9	4/9	0.975	1.436	−0.461	1.12
	9	2/9	0.215	0.431	−0.215	0.22
4-site chain ⊗ 4-site chain						
	16	12/16	2.105	2.275	−0.171	7.58
	16	8/16	1.643	1.961	−0.317	3.03
	16	3/16	0.760	1.171	−0.411	0.83

Scaling Behaviour

The results in Table 1 reveal a robust scaling pattern:

1. **System-size robustness:** Sub-additive entropy defect appears at every dimension tested (4×4 through 16×16). It is not a small-system artifact.

2. **Monotonic κ–truncation relationship:** The empirical κ increases monotonically with the fraction of spectrum retained. For the 16-dimensional system: $n/\dim = 75\% \Rightarrow \kappa \approx 7.6$; $50\% \Rightarrow \kappa \approx 3.0$; $19\% \Rightarrow \kappa \approx 0.83$. More accessible spectrum means weaker apparent correlation.

3. **H-function not required at tested level:** Simple κ-addition ($H = \mathrm{id}$) suffices to represent the tested toy cases at the fitted-parameter level. This does not establish uniqueness of the simple κ-form.

4. **Three-subsystem tests:** For tripartite systems (8-dimensional), direct bipartite splits (e.g., A vs BC) satisfy κ-addition with positive κ. However, marginal-of-marginal splits (e.g., A vs B within the AB marginal of a truncated ABC state) can yield negative κ values, corresponding to super-additive entropy. These cases arise from inherited truncation that does not correspond to a direct spectral threshold on the subsystem Hamiltonian.

Conjecture 3.4 (Sub-Additivity from Spectral Truncation). For any product thermal state $\rho = \rho_A \otimes \rho_B$ on a composite Hilbert space $\mathcal{H}_A \otimes \mathcal{H}_B$, and any spectral truncation Π_T

defined by a threshold on the eigenvalues of the product Hamiltonian $D_A \otimes \mathbf{1} + \mathbf{1} \otimes D_B$, the truncated entropy composition is sub-additive:

$$S(\rho_\theta) \leq S(\rho_\theta^A) + S(\rho_\theta^B), \tag{10}$$

with equality if and only if no eigenvalues are truncated. Whether the specific functional form of the defect is the κ-addition product $S_A \cdot S_B/\kappa$ or a more general sub-additive function remains to be determined analytically.

Remark 3.5 (Why sub-additivity is plausible). For product Hamiltonians, the eigenstates of D_{AB} are tensor products $|i\rangle_A|j\rangle_B$. Spectral truncation selects a subset V of these product basis states. The truncated state remains diagonal in the product basis: $\rho_\theta = \sum_{(i,j)\in V}(p_i q_j/Z_\theta)\,|ij\rangle\langle ij|$. The non-product geometry of the subspace V is the sole source of the entropy defect. The subspace V is defined by the condition $|\lambda_i + \mu_j| \geq \theta$, which couples the A and B indices through the threshold, creating apparent correlation even for an initially uncorrelated state.

Remark 3.6 (Physical interpretation). The computation reveals the mechanism behind κ-addition:

1. **Full spectrum:** When all eigenvalues are accessible ($\theta = 0$), the thermal state of a product Hamiltonian is a product state. Entropy is additive. $\kappa \to \infty$.

2. **Truncated spectrum:** Spectral truncation removes eigenstates from the accessible sector. Even for an initially uncorrelated (product) state, the truncation creates apparent correlation by projecting onto a subspace that is not itself a product space.

3. **The κ parameter:** The strength of the apparent correlation is measured by κ. Stronger truncation (fewer kept eigenvalues) gives smaller κ (stronger non-additivity). This is precisely the v7/v8 interpretation: κ measures how much of the correlation structure is invisible to the observer at temperature T.

4. **The asymmetric cases are non-degenerate:** When subsystem couplings differ ($a \neq b$) and partial truncation removes some but not all mixed-sector eigenvalues, κ takes finite values that depend nontrivially on the thermal state and truncation level.

Independent κ Test (Cycle 2)

To test whether the κ-addition form has structural content beyond a one-parameter fit, we performed two additional analyses.

Cross-β stability test. At fixed truncation (keeping 4/9 eigenvalues of the 3-site chain composite), the effective κ was computed across seven values of β. An alternative model (defect $= -\alpha(S_A + S_B)$) was fitted in parallel.

Result: Neither κ nor α is constant across β. The coefficient of variation for α (0.36) is lower than for κ (0.55). This means the alternative sub-additive model fits at least as stably as κ-addition. The true functional form of the truncation-induced defect is more complex than either simple parametrisation.

Multi-system scan. We collected 28 data points across 2-site, 3-site, and 4-site chain composites (4 to 16 dimensions), varying β and truncation threshold. Key observations:

1. At the same truncation fraction (50%), κ ranges from 0.09 (2-site, $\beta = 1$) to 103.6 (4-site, $\beta = 2$). The κ parameter is **not** determined by the truncation fraction alone.
2. The direction of κ versus β **reverses** depending on truncation severity. For heavy truncation (keeping $\leq 25\%$), κ decreases with β. For light truncation (keeping $\geq 50\%$), κ increases with β. This is physically interpretable: it depends on the overlap between the truncation mask and the thermal occupation distribution.
3. No simple ratio $I(A{:}B)/f(S_A, S_B)$ is constant across β at fixed truncation. The mutual information created by truncation is a complex function of all parameters.

Remark 3.7 (What this means for the v7 bridge). The qualitative mechanism of the v7 bridge conjecture is confirmed: spectral truncation creates non-extensive composition, and the strength depends on the severity of the truncation. The specific κ-addition form (defect $\propto S_A \cdot S_B$) is consistent with the data but is not uniquely selected. Two possibilities remain:

1. The exact κ-addition law emerges in a regime not yet tested (e.g., the large-system limit, or for specific spectral geometries matching the expander graph structure).
2. The actual composition law from spectral truncation is a generalisation of κ-addition that reduces to it in specific limits.

Both possibilities are consistent with the data. Neither is confirmed. The quantitative bridge between the $\mathcal{W}$-atlas and thermodynamic relativity remains open but its qualitative mechanism is supported at the tested toy/scaling level.

Analytical Result: The Truncation Mutual Information

Proposition 3.8 (Exact Mutual Information from Spectral Truncation of Diagonal Product States). *Let $\rho = \rho_A \otimes \rho_B$ be a product thermal state with weights $p_i = e^{-\beta\lambda_i}/Z_A$ and $q_j = e^{-\beta\mu_j}/Z_B$ in the eigenbases of D_A and D_B respectively. Let $V = \{(i,j) : |\lambda_i + \mu_j| \geq \theta\}$*

be the spectral truncation mask. Define:

$$Z_V := \sum_{(i,j)\in V} p_i\, q_j \qquad \text{(retained thermal weight)}, \tag{11}$$

$$Q_i := \sum_{j\in V(i)} q_j \qquad \text{(}B\text{-weight accessible from }A\text{-state }i\text{)}, \tag{12}$$

$$P_j := \sum_{i\in V(j)} p_i \qquad \text{(}A\text{-weight accessible from }B\text{-state }j\text{)}. \tag{13}$$

Then the mutual information of the truncated state is exactly:

$$\boxed{I(A{:}B) = \ln Z_V - \frac{1}{Z_V}\sum_i p_i\, Q_i \,\ln Q_i - \frac{1}{Z_V}\sum_j q_j\, P_j \,\ln P_j.} \tag{14}$$

Proof. The truncated state $\rho_\theta = \sum_{(i,j)\in V}(p_i q_j/Z_V)\,|ij\rangle\langle ij|$ is diagonal in the product eigenbasis. Direct computation of $S(\rho_\theta)$, $S(\mathrm{Tr}_B\,\rho_\theta)$, and $S(\mathrm{Tr}_A\,\rho_\theta)$ yields (14). Verified numerically against direct product-basis computation of mutual information for all tested product-Hamiltonian cases, including 4×4, 9×9 (3-site chain), and 16×16 systems, across multiple retention levels and inverse temperatures. The verification holds to machine precision when the numerical comparison uses a consistent product-basis mask; see Remark 3.9 and the computational record note (Section 3.8). □

Remark 3.9 (Resolution of the v13/v14 3-Site Discrepancy). Versions 13 and 14 of this manuscript reported an apparent discrepancy between analytical and numerical mutual information at the 9×9 (3-site chain) truncation case ($I(A{:}B)\approx 0.607$ analytical versus ≈ 0.215 numerical, at the 4/9-retention row of Script 05 in the v13 computational addendum), and bounded Proposition 3.8 with a domain-boundary remark. A subsequent forensic probe (Script 07b; Section 3.8) has resolved this discrepancy as a methodological artifact of the original verification script. The two reported numbers are both correct — but they correspond to two different truncated states, not to two computations of the same truncated state:

- The analytical value (0.607) evaluates Proposition 3.8 on a product-basis mask $V_{2d} = \{(i,j) : |\lambda_i + \mu_j| \geq \theta\}$ defined over product-state index pairs.
- The numerical value (0.215) projects the thermal state onto a subspace of $\mathcal{H}_A \otimes \mathcal{H}_B$ selected by applying the same threshold to the flat eigenvalue spectrum of D_{AB}, using eigenvectors produced by numpy's `eigh` diagonalisation of D_{AB}.

When $D_{AB} = D_A \otimes \mathbf{1} + \mathbf{1} \otimes D_B$ has non-degenerate spectrum, the composite eigenvectors are exactly the product states $|i\rangle_A|j\rangle_B$, and the two methods coincide. When D_{AB} has degenerate eigenvalue subspaces — as the 3-site chain does, with composite eigenvalues 0 (multiplicity 3) and $\pm\sqrt{2}$ (multiplicity 2 each) — numpy's `eigh` returns an arbitrary orthonormal basis within each degenerate subspace. Stage 2 of Script 07b confirms that only

2 of the 9 composite eigenvectors are pure product states; the remaining 7 are superpositions within degenerate subspaces. Projecting the thermal state onto these non-product eigenvectors and taking partial traces produces a truncated state genuinely different from the one the analytical formula describes.

Under a consistent product-basis mask (numerical Method 2A in Script 07b, which forms the joint $p_i q_j$ distribution over (i, j) pairs in V_{2d} directly), the numerical mutual information agrees with the analytical formula to machine precision at every tested case, including the 9×9 3-site-chain case. The v14 "domain boundary" framing is therefore retracted: Proposition 3.8 is verified across all tested product-Hamiltonian truncation cases, including composites with spectral degeneracies.

Scope statement. The verification holds in the *product-Hamiltonian regime*: $D_{AB} = D_A \otimes \mathbf{1} + \mathbf{1} \otimes D_B$ with no direct A–B coupling term. Extending Proposition 3.8 to genuinely interacting composite Hamiltonians (containing an explicit H_{AB} coupling term) is a distinct open question and is not addressed by the present verification.

Corollary 3.10 (Product Truncation Criterion). *$I(A{:}B) = 0$ if and only if V is a product set ($V = V_A \times V_B$), in which case Q_i is constant across all $i \in V_A$ and P_j is constant across all $j \in V_B$.*

The mutual information measures the non-product geometry of the truncation mask: Q_i varies across i precisely when the mask couples the A and B indices through the threshold condition, creating apparent correlation in an initially uncorrelated state.

Remark 3.11 (Physical interpretation). The quantities Q_i and P_j measure **spectral accessibility asymmetry**: how much of subsystem B is "visible" to an observer in eigenstate $|i\rangle_A$, and vice versa. When the truncation threshold couples the two subsystems (as spectral truncation generically does), different eigenstates of A see different amounts of B. This asymmetry is the source of the apparent correlation, and it is quantified exactly by (14).

Remark 3.12 (Epistemic status). The computation is a **numerical result across multiple system sizes**, not a general analytical proof. It establishes that spectral truncation produces sub-additive entropy composition from 4-dimensional through 16-dimensional composite systems. Conjecture 3.4 proposes the analytical generalisation of the sub-additivity. The specific functional form of the defect (whether κ-addition or a more general sub-additive law) and the extension to non-product Hamiltonians remain open (O38).

Remark 3.13 (Connection to Livadiotis–McComas). The result confirms the v7 bridge conjecture's **qualitative mechanism**: spectral coarse-graining creates non-extensive entropy composition, and the strength of non-extensivity scales with the severity of the coarse-graining. This is consistent with the Livadiotis–McComas identification of $1/\kappa$ with the strength of inter-particle correlations.

However, the computation does not uniquely select the specific κ-addition form. The **quantitative** bridge — the emergence of the exact H-function composition law from spectral truncation — remains open (O38).

Computational Record Note: Resolution of the v13/v14 3-Site Discrepancy

This subsection documents, as part of the programme's preserved evidence trail, the forensic probe that reconciled an apparent discrepancy in the v13 computational record. We include it because the programme's discipline is to preserve failed and corrected intermediate findings alongside the verified results.

The v13 record

Script 05 of the v13 computational addendum (`05_o31_analytical_formula.py`) tested Proposition 3.8 against direct numerical computation at three representative cases: 2-site chain (4×4, verified), 3-site chain (9×9, at 4/9 retention with $\beta = 0.5$), and 4-site chain (16×16, verified). For the 3-site case Script 05 reported $I_{\text{analytical}} = 0.6069288341$ against $I_{\text{numerical}} = 0.2152715693$ and printed `Match:  False`. This was the datum that motivated Remark 3.9 in v14.

The Script 07b forensic probe

Script 07b of the v15 cycle reproduces both numbers from Script 05 and identifies the mechanism. It compares three computational methods on the same system:

1. **Method 1 (analytical, product basis):** Evaluates Proposition 3.8 on the product-basis mask $V_{2d} = \{(i, j) : |\lambda_i + \mu_j| \geq \theta\}$. Yields $I = 0.6069288341$.

2. **Method 2A (numerical, product basis):** Computes $I = S_A + S_B - S_{AB}$ directly from the joint distribution $p_i q_j$ restricted to V_{2d}, using partial sums over product-state indices. Yields $I = 0.6069288341$. This is the numerical check consistent with the analytical formula's precondition.

3. **Method 2B (numerical, composite eigenbasis; the Script 05 method):** Diagonalises D_{AB} via `numpy.linalg.eigh`, selects eigenvectors with $|\lambda| \geq \theta$, projects the thermal state onto those eigenvectors, and computes entropies by partial trace. Yields $I = 0.2152715693$.

Methods 1 and 2A agree to machine precision. Method 2B disagrees with both.

The mechanism: degenerate-subspace basis ambiguity

Stage 2 of Script 07b inspects the composite eigenvectors of the 3-site chain composite D_{AB}. The composite spectrum is $\{-2\sqrt{2}, -\sqrt{2}, -\sqrt{2}, 0, 0, 0, +\sqrt{2}, +\sqrt{2}, +2\sqrt{2}\}$ with multiplicities $1, 2, 3, 2, 1$. Only the non-degenerate eigenvectors at $\pm 2\sqrt{2}$ are pure product states $|i\rangle_A|j\rangle_B$. Within each degenerate subspace, `numpy.linalg.eigh` returns an arbitrary orthonormal basis — one that is, generically, a superposition of product states. Script 07b explicitly prints the product-state overlaps of each composite eigenvector and confirms that 7/9 are superpositions, not pure products.

When Method 2B projects the thermal state onto a subspace that partially retains a degenerate subspace, it selects a subset of these superpositions. The resulting projection is *not* the same subspace as V_{2d} under the natural product-state identification; it is a different truncated state living in a rotated basis within the degenerate sector. Partial trace of this rotated state produces entropies different from those of the product-basis truncation.

Stage 3 of Script 07b confirms that the disagreement between Method 1 and Method 2B appears across all v13 cases *exactly* where the truncation threshold partially retains a degenerate subspace, and vanishes wherever the threshold falls cleanly between distinct eigenvalue magnitudes or wherever a degenerate subspace is either fully retained or fully excluded.

Conclusion of the probe

The 3-site discrepancy reported in v13/v14 is a mask-definition inconsistency[1] within Script 05, not a failure of Proposition 3.8. Script 05's analytical and numerical calculations were computing mutual information of two different truncated states, related by an arbitrary basis rotation within degenerate subspaces. When the numerical comparison is re-run under a mask definition consistent with the analytical formula's precondition (Method 2A of Script 07b), Proposition 3.8 is verified to machine precision for every tested case, including the 9×9 3-site-chain case that v14 flagged as a domain boundary. The v14 "domain boundary" framing is retracted in Remark 3.9.

Script 07b is included in the v15 computational addendum alongside Scripts 01–05 of the v13 record, as a preserved evidence trail documenting the reconciliation. The

[1]Script 07d confirms the mechanism directly: under an asymmetric perturbation $H_A = H_{\text{chain}} + \epsilon\,\text{diag}(-1, 0, +1)$ with retention count fixed at $N = 4$, the Method 1 / Method 2B gap collapses from 3.92×10^{-1} at $\epsilon = 0$ (four-fold composite degeneracy present) to $\leq 10^{-15}$ once $\epsilon \geq 10^{-4}$ lifts the degeneracies. The discrepancy is localized to partial retention inside a degenerate subspace under a composite-eigenbasis projection. Script 07c, a fixed-threshold variant, is preserved in the computational addendum as a failed intermediate probe: its fixed-threshold design caused retention-count instability across the perturbation sweep and produced noisy gaps, and it is retained only as an evidence-trail artifact consistent with the programme's discipline of preserving corrected intermediate findings.

v13 addendum is not rewritten; the Script 05 output remains as published for historical fidelity.

Scope clarification

The reconciliation above establishes Proposition 3.8 in the product-Hamiltonian regime $D_{AB} = D_A \otimes \mathbf{1} + \mathbf{1} \otimes D_B$ across all tested retention levels and system sizes, including cases with composite spectral degeneracies. It does *not* establish Proposition 3.8 for genuinely interacting Hamiltonians (those containing an explicit H_{AB} coupling term between the subsystems). The product-Hamiltonian class is the full scope of the v13 and v15 numerical evidence; the interacting-Hamiltonian extension is a legitimate forward direction, but it is not developed in the present cycle.

The O1 Inverse Problem: Structure and a Critical Correction

The Problem

Open Problem O1 (v1) asks for the derivation of the off-diagonal metric component $g_{IE}(I, E)$ completing the spatial metric:

$$h_{ab} = \begin{pmatrix} g_{II} & g_{IE} \\ g_{IE} & g_{EE} \end{pmatrix} = \begin{pmatrix} \dfrac{\hbar^2}{4I} & g_{IE}(I, E) \\ g_{IE}(I, E) & \dfrac{4G\hbar}{k_{\mathrm{B}}^2 E^2} \end{pmatrix}. \tag{15}$$

A Critical Finding: The 2D EFE Constraint Is Vacuous

A natural attack on O1, proposed during the v13 planning cycle, was to solve "backward" from the requirement that the (I, E) Ricci tensor recovers the vacuum Einstein equations.

This approach fails for a fundamental reason.

In two dimensions, the Einstein tensor vanishes identically:

$$G_{\mu\nu} \equiv R_{\mu\nu} - \tfrac{1}{2} R\, g_{\mu\nu} \equiv 0 \qquad \text{(in 2D, for } \textit{any} \text{ metric).} \tag{16}$$

This is a standard result of differential geometry: in 2D, the Ricci tensor is fully determined by the scalar curvature, $R_{\mu\nu} = (R/2)\, g_{\mu\nu}$, making $G_{\mu\nu}$ identically zero.

Since the (I, E) sector is two-dimensional, the vacuum Einstein equations $G_{\mu\nu} = 0$ are trivially satisfied for *any* choice of $g_{IE}(I, E)$. The "backward EFE" constraint places no restriction on the unknown function.

The Actual Constraints on g_{IE}

The EFE recovery must operate on the full three-dimensional emergent spatial geometry, not the two-dimensional (I, E) base metric. The three spatial dimensions emerge through the replacement product construction of v3 (degree-3 graph from the cycle structure).

The constraints on $g_{IE}(I, E)$ are therefore:

1. **3D metric construction:** The function g_{IE} enters the three-dimensional spatial metric through the replacement product lift. The EFE recovery condition must be applied to the 3D metric, not the 2D base. This requires the explicit construction of the 3D metric from $(h_{ab}, \phi\text{-sector})$.
2. **Determinant condition:** For three spatial dimensions to emerge (v1–v2), the determinant of h_{ab} must approach zero in the classical limit $E \to E_{\max}$, $I \to I_{\max}$:
$$\det(h_{ab}) = \frac{G\hbar^3}{I\, k_{\mathrm{B}}^2\, E^2} - g_{IE}^2 \to 0 \quad \Longrightarrow \quad |g_{IE}| \to \sqrt{g_{II}\, g_{EE}}. \tag{17}$$
This constrains the asymptotic behaviour of g_{IE} but not its full functional form.
3. **Ryu–Takayanagi compatibility:** The geodesic structure on the (I, E) surface, which depends on g_{IE}, must be compatible with the RT entanglement entropy formula (v1).
4. **Rotation map convergence:** The discrete rotation map Rot_{AP_q} of v3 must converge to g_{IE} in the continuum limit (O14, v3/v5).
5. **Spectral filter compatibility:** The function g_{IE} must be compatible with the monotone filtration (Proposition 3.1 of v12): the spatial metric induced at temperature T must be consistent with the accessible spectral sector U_T.

Remark 4.1 (Status of the O1 inverse programme). The 2D vacuity finding eliminates the strongest proposed constraint from the inverse programme. The remaining constraints (items 1–5 above) are nontrivial but require the explicit 3D metric construction to be carried out. This construction depends on the replacement product geometry of v3, which is well-defined at the graph level but whose continuum limit has not been computed. O1 therefore remains open, but its constraint structure is now better understood.

Space–Time Complementarity

The Corollary

Corollary 5.1 (Space–Time Complementarity). *Let $N_{\mathrm{acc}}(T)$ denote the number of eigenvalues of $D_{\mathcal{W}}$ in the accessible sector U_T (assuming discrete spectrum). Then:*

1. *$N_{\text{acc}}(T)$ is monotonically non-increasing in T (from Proposition 3.1 of v12, Nested Accessibility).*

2. *The temporal processing rate $c_S(T) = k_{\text{B}}T/\hbar$ is monotonically increasing in T (from the v1 identification).*

Therefore, spatial spectral richness and temporal processing rate are inversely related across temperature:

$$T\uparrow \implies N_{\text{acc}}\downarrow,\ c_S\uparrow \qquad \textit{and} \qquad T\downarrow \implies N_{\text{acc}}\uparrow,\ c_S\downarrow. \tag{18}$$

Proof. Item (1) follows directly from v12 Proposition 3.1: if $T_1 < T_2$, then $U_{T_2} \subseteq U_{T_1}$, so $N_{\text{acc}}(T_2) \le N_{\text{acc}}(T_1)$. Item (2) follows from the definition $c_S(T) = k_{\text{B}}T/\hbar$, which is manifestly increasing in T. The inverse relationship is the conjunction of these two monotone behaviours. □

Consequences at Cycle Endpoints

If the framework is applied to a cosmological setting, the complementarity corollary suggests asymptotic behaviour at the endpoints:

Hypothesis 5.2 (Cycle-Endpoint Asymptotics).

1. **Hot endpoint ($T \to T_P$):** $N_{\text{acc}} \to N_{\min}$ (minimal spatial spectral richness) while c_S reaches its maximum value $k_{\text{B}}T_P/\hbar = E_{\text{P}}/\hbar$ (maximal temporal rate). The observer reconstructs minimal spatial geometry at maximum temporal resolution.

2. **Cold endpoint ($T \to 0$):** $N_{\text{acc}} \to N_{\max}$ (full spectrum accessible) while $c_S \to 0$ (temporal coordinate degenerates). The observer has access to maximal spatial richness but can no longer generate temporal flow. This is the chart failure identified in v8.

Remark 5.3 (Epistemic status). The complementarity corollary (Corollary 5.1) is **Derived** from v12 propositions. The cycle-endpoint hypothesis (Hypothesis 5.2) is **Proposed**: it requires the framework to be applied to cosmological scales, which involves additional assumptions about the global structure of $D_{\mathcal{W}}$ not yet established.

The 3+1 Emergence Hypothesis

The O31 result, the O1 constraint analysis, and the complementarity corollary together support a structured 3+1 emergence hypothesis:

Hypothesis 6.1 (3+1 Emergence from Spectral Accessibility). Standard spacetime coordinates (x, y, z, t) are emergent reconstruction variables on the observer-accessible sector of the $\mathcal{W}$-atlas:

1. **Spatial coordinates** (x, y, z) arise from the observer's reconstruction of the accessible (I, E)-sector geometry, lifted to three dimensions through the replacement product cycle structure (v3). The effective spatial scale depends on the spectral sector accessible at the observer's temperature.

2. **The temporal coordinate** t arises as a chart-dependent ordering parameter associated with the entropy processing rate $c_S(T) = k_B T/\hbar$. Its effective resolution is set by the observer's thermal state.

3. **Observer dependence:** Different observer charts (different T) reconstruct different effective 3+1 descriptions from different accessible spectral sectors. In the matched-chart limit (v12, Theorem 3.6), these descriptions agree and no chart-mismatch distortion appears.

4. **Cosmological expansion** is reinterpreted: as the universe cools, the spectral filter Π_T widens (Proposition 3.1 of v12), the accessible (I, E) sector enriches, and the reconstructed spatial geometry gains resolution and scale.

Remark 6.2 (Epistemic status). This hypothesis is **Proposed**, not Derived. Items (1) and (4) depend on O1 being resolved. Item (2) is consistent with the v1 (S, ϕ) geodesic structure but has not been given a comparably rigorous emergence derivation. Item (3) follows from v12's Matched-Chart Consistency Theorem. A full derivation of the 3+1 map from the $\mathcal{W}$-atlas remains a frontier target.

Updated Open Problems

O1. Derivation of $g_{IE}(I, E)$. The naive "backward EFE" approach is vacuous in 2D (Section 2). The constraint must operate through the 3D replacement product lift. Status: *open, constraint structure clarified.*

O31. H-function composition from spectral truncation. Status: *qualitative mechanism supported at toy/scaling level; exact composition law remains open (Section 2).*

O37. 3D metric from replacement product lift. Construct the explicit three-dimensional spatial metric from the (I, E) base metric h_{ab} and the v3 replacement product cycle structure. This is required before O1 can be constrained by EFE recovery.

O38. O31 functional form determination. Determine analytically whether the sub-additive defect from spectral truncation of product thermal states takes the specific κ-addition form ($\Delta S \propto S_A S_B$), or a different sub-additive form, or a more complex function of S_A, S_B, β, and θ. Test with non-product Hamiltonians and in the continuum limit.

O39. Temporal emergence map. Provide a rigorous construction of the temporal coordinate from the (S, ϕ) sector of the $\mathcal{W}$-atlas, comparable in precision to the spatial emergence programme for the (I, E) sector. Currently at the level of a working hypothesis (Hypothesis 6.1, item 2).

Series Arc

Table 2: The Manifold Relativity programme: series arc v1–v13.

Version	Date	Central Contribution
v1–v11	Mar–Apr 2026	(Under historical title "Entropy Waves".) $\mathcal{W}$-manifold, discrete substrate, W-atlas, candidate operator, chart matching
v12	Apr 2026	**Series renamed to Manifold Relativity.** Programme reframing; structural propositions; candidate observables; falsification criteria
v13	Apr 2026	Sub-additivity from spectral truncation (O31 qualitative validation); O1 constraint clarification (2D EFE vacuity); exact mutual information formula; Space–Time Complementarity Corollary; 3+1 emergence hypothesis
v14	Apr 2026	**Referee-hardened submission draft.** Domain-of-validity bounding for Proposition 3.8 (Remark 3.9); expanded literature positioning (Section 2); conclusion sharpened on α vs κ non-uniqueness at tested toy scale
v15	Apr 2026	**Computational record reconciliation and retraction.** v13/v14 3-site discrepancy resolved as Script 05 mask-definition artifact via forensic probe Script 07b, with causal closure confirmed by Script 07d (Section 3.8); Remark 3.9 rewritten; formal **retraction** of the v14 "domain boundary" framing; Proposition 3.8 verified to machine precision across all tested product-Hamiltonian cases, including cases with composite degeneracies

Conclusion

This edition advances the programme's frontier in two directions: one substantial positive result (with an important caveat) and one corrective finding.

The positive result is O31. Spectral truncation of product thermal states creates subadditive entropy composition across all tested system dimensions (4 through 16). The effective κ parameter increases monotonically with the fraction of spectrum retained, and the full-spectrum limit recovers standard additivity. This validates the v7 bridge conjecture's qualitative mechanism: spectral truncation creates apparent correlation by projecting product states onto non-product subspaces, and the strength of the apparent correlation scales with the severity of the truncation.

Proposition 3.8 verification status (updated in v15). The analytical mutual-information formula of Proposition 3.8 is verified to machine precision across all tested product-Hamiltonian truncation cases, including 4×4, 9×9 (3-site chain), and 16×16 composites, at multiple retention levels and inverse temperatures, when the numerical comparison uses a consistent product-basis mask. The v13/v14 "3-site discrepancy" has been resolved (Section 3.8, Remark 3.9) as a methodological inconsistency in the original Script 05 verification, not as a failure of the formula. The verification scope is the *product-Hamiltonian* regime $D_{AB} = D_A \otimes \mathbf{1} + \mathbf{1} \otimes D_B$; extension to genuinely interacting composite Hamiltonians with an explicit H_{AB} coupling term remains a legitimate forward direction not addressed by the present numerical evidence.

The caveat: The specific κ-addition functional form (defect $\propto S_A \cdot S_B$) is not uniquely selected by the computation. An alternative linear defect model (defect $\propto S_A + S_B$, parameterised by α) fits the toy-scale data more tightly: across the reported cross-β test conditions, the coefficient of variation of α is approximately 0.36, compared with approximately 0.55 for κ. Neither parameter is constant across β at fixed truncation, but at the tested toy scale the α-parameterisation exhibits lower dispersion than the κ-parameterisation. The present data therefore do not uniquely favor the κ-addition form; the quantitative bridge to thermodynamic relativity's specific composition law remains a functional conjecture (O38) awaiting larger operator dimensions and the interacting-Hamiltonian extension. The v15 reconciliation of Proposition 3.8 does not address or close O38; the functional-form non-uniqueness stands as an independent open result.

The corrective finding is on O1. The "backward EFE" approach to deriving g_{IE} is vacuous in two dimensions because the Einstein tensor vanishes identically for any 2D metric. The actual constraints on g_{IE} must operate through the three-dimensional replacement product lift (new Open Problem O37). This redirects the O1 programme from a naive inverse PDE to a more structurally informed construction problem.

The Space–Time Complementarity Corollary is a structural observation derived from v12's propositions. It is not deep mathematics, but it has physical content: it formalises the inverse relationship between spatial richness and temporal rate across temperature, with specific consequences for cycle-endpoint behaviour.

The 3+1 emergence hypothesis organises the programme's spatial and temporal claims into a unified statement. It is Proposed, not Derived. Its derivation is the long-term target of the programme.

The frontier after v13:

- O37 (3D metric construction) enables a non-vacuous O1 attack.
- O38 (O31 scaling) determines whether the κ-addition mechanism generalises.
- O39 (temporal emergence) closes the 3+1 programme on the time side.

We preserve all computational outcomes, including misleading first-pass successes, weaker alternative fits, negative uniqueness checks, and the v13/v14 3-site mask-definition artifact resolved in the v15 cycle, as supplementary addenda rather than only the final positive claims. The computational evidence trail (Scripts 01–05 from v13 plus Script 07b from v15, with raw outputs) is available as a downloadable addendum alongside this preprint.

References

[1] Sorvik, P. E. et al. (2026). Manifold Relativity v12: Observer-Dependent Spectral Accessibility. Available: https://paulsorvik.wordpress.com

[2] Livadiotis, G. & McComas, D. J. (2024). The theory of thermodynamic relativity. *Scientific Reports*, **14**, 22641.

[3] Connes, A. (1994). *Noncommutative Geometry*. Academic Press.

[4] Tsallis, C. (1988). Possible generalization of Boltzmann–Gibbs statistics. *Journal of Statistical Physics*, **52**, 479–487.

[5] Abe, S. (2004). Tsallis entropy: how unique? *Continuum Mechanics and Thermodynamics*, **16**, 237–244.

[6] Abe, S. (2001). Nonextensive thermodynamic relations. *Physics Letters A*, **281**, 126–130.

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

[8] Lieb, E. H. (1975). Some convexity and subadditivity properties of entropy. *Bulletin of the American Mathematical Society*, **81**, 1–13.

[9] Kato, K. & Brandão, F. G. S. L. (2019). Quantum approximate Markov chains are thermal. *Communications in Mathematical Physics*, **370**, 117–149.

[10] Molnár, A., Schuch, N., Verstraete, F. & Cirac, J. I. (2015). Approximating Gibbs states of local Hamiltonians efficiently with projected entangled pair states. *Physical Review B*, **91**, 045138.

[11] Livadiotis, G. (2018). Thermodynamic origin of kappa distributions. *Europhysics Letters*, **122**, 50001.

[12] Livadiotis, G. (2021). Thermodynamic definitions of temperature and kappa, and introduction of the entropy defect. *Entropy*, **23**, 1683.