

Causal Topology Vacuum Model: A Statistics-Resolved, Matroidal, and Horizon-Regulated Framework Integrated v15.5 Draft

Russell S. Clark II

Brian L. Swift, Ph.D.

May 2026

Abstract

This integrated draft is a research-program white paper that harmonizes three strands of the Causal Topology Vacuum Model (CTVM): the broad v7 dark-sector program, the shorter v8 corrective scaffold, and the Volovik/statistical-sector revision. The central claim is no longer that CTVM has been derived as a theorem of algebraic quantum field theory (AQFT). The corrected claim is more precise: CTVM is a layered, testable framework in which regulated vacuum correlations are separated into operational response sectors, encoded by response-weighted representable matroids, and conditionally related to gravitational sourcing through a separate semiclassical or holographic dictionary. Version 15.5 retains the referee-facing viability improvements and critical-path improvements: a tighter dependency schema, a modular-compatibility status note, a compact SK-versus-CPTP worked example, an explicit matroid falsifier, clearer Volovik analogy boundaries, and more visible falsifier/non-claim summaries. First, the transport–teleportation diagnostic is rewritten as operational prose rather than as a numbered pseudo-equation, while preserving the distinction between path-dependent geometric response and resource-mediated algebraic relabeling. Second, the near-degenerate horizon-edge assumption is strengthened into a spectral-shell hypothesis with an explicit expansion parameter for the approximation $\gamma_+ \simeq \gamma_-$. These are deliberately clarifying revisions; they sharpen presentation, operational falsifiability, and horizon-edge kinetics without adding new load-bearing claims. Version 15.5 further tightens the architectural-tradeoff statement and quarantines the conditional residual-sector construction: a thresholded binary reduction of the residual response block, a matroid-to-CSS code ansatz, a quantified addressability/leakage trade-off, a small three-detector CSS illustration, and an analytic longitudinal-mode asymptotic in Appendix H.12. These additions are framed as construction targets and calculable hypotheses, not as completed proofs of residual gravitational inertness. Version 15.5 also retains a modular-equivalence consistency check motivated by the companion Entropic Inertia analysis: if inertia is a Rindler modular work response, then any sector excluded from gravitational sourcing must also be absent from the inertial modular-work functional, or else the Einstein equivalence principle would be violated. This is stated as a conditional sector-matching principle, not as a proof of the equivalence principle.

The key repair is a separation of maps that earlier drafts conflated. The Schwinger-Keldysh retarded projection is a linear response-functional filter on correlation cumulants, not automatically a completely-positive trace-preserving state channel. The split-property construction supplies type-I tensor bookkeeping only after a split inclusion and a compatible vacuum-preserving conditional expectation are chosen. The response matroid records support, rank, and dependence structure of regulated detector data; it does not by itself prove gravitational coupling. Its role is instead diagnostic: exact sector separation requires a direct-sum decomposition with no cross-circuits across the proposed split, so the matroid supplies a discrete consistency test and possible falsifier for detector/split compatibility. The Volovik analogy is incorporated as a controlled universality-class analogy: ${}^4\text{He}$ -like bosonic collective media motivate Bose-enhanced

pairwise response channels and emergent metric kinematics, while ${}^3\text{He-A}/{}^3\text{He-B}$ -like fermionic topological media motivate Pauli-blocked residual structure, Weyl-point topology, and protected boundary modes.

The dark-matter benchmark is now stated conservatively as a response target range, $f_{\text{eq}} \in [0.20, 0.30]$, with the historical point $27/106.75 \simeq 0.253$ retained only as a reference value. It becomes a physical claim only if the species- and polarization-resolved response integral $f_{\text{vec}}^{\text{eff}}(T_{\text{EW}}; \mathcal{P})$ is evaluated across the electroweak transition and flows into that range in a regulator-stable, gauge-consistent protocol window. The dark-energy number $\Omega_\Lambda \simeq 2/3$ is retained only as a conditional normalization target: it requires a chosen null-boundary convention, an edge-action/vacuum-energy dictionary, a cosmological-horizon holography extension, and a late-time blocking history with $f_{\text{edge}} \rightarrow 1$. CTVM in this form does not solve the magnitude/cancellation face of the cosmological-constant problem; it gives a conditional horizon-scale and coincidence framework whose bulk-vacuum cancellation or residual-decoupling mechanism remains an open target. The central deliverable for any future cancellation or sequestration claim is the coarse-grained edge theory of Gate E1, which must produce the edge-action/vacuum-energy dictionary rather than assume it. The tracking behavior $\rho_\Lambda \propto f_{\text{edge}} H^2$ is elevated to a leading viability gate because constant f_{edge} would behave as early tracking dark energy rather than a cosmological constant.

The result is a longer and more coherent manuscript that preserves the ambition of v7, the organizational clarity of v8, and the discipline of the Volovik/statistical revision. It states exactly what is proved, what is conditional, what remains conjectural, and what would falsify the framework. Near-term falsifiers include failure of detector/split block compatibility, absence of reproducible non-uniform matroid strata in physically realizable protocols, failure of a vector-channel response calculation to support the $[0.20, 0.30]$ benchmark range, strong null-boundary normalization sensitivity of κ_{edge} , or absence of the predicted Bose/Pauli bias in analogue-vacuum harvesting experiments.

Reader’s Guide: the framework in one page

This white paper is a dependency map, not a completed derivation of dark matter or dark energy. The strong CTVM gravity interpretation requires several logically independent gates to close. The durable core that remains if those gates fail is the detector-response/SK-filter/matroid toolkit.

Conditional CTVM chain. If a modular-compatible split or type-II substitute is constructed; if detector response data block-diagonalize relative to the proposed split/residual decomposition; if the response matroid has no stable circuits crossing the operational split; if the split-sector data define a regulator-stable pairwise network satisfying the required entropy-cone inequalities; and if a holographic or semiclassical sourcing dictionary plus residual decoupling are supplied, then the split-sector response may be assigned the gravitational source functional of the model. If the edge-action/vacuum-energy dictionary and late-time blocking history are also supplied, the dark-energy number becomes a conditional normalization target. If the vector response integral is stable across the electroweak transition, the dark-matter benchmark range becomes a dynamical target.

If any of these gates fail, the strong CTVM gravity interpretation fails at that link. What remains is more modest but still well-defined: regulated detector response, Schwinger–Keldysh response filtering, and response-matroid diagnostics for testing whether vacuum correlation data admit operational sector decompositions.

Typed dependency diagram. The arrows below change mathematical type; this is a dependency diagram, not a commutative diagram in a single category.

$$\begin{array}{ccccccc}
\mathcal{A}_{\text{reg}} & \xrightarrow{E_{\mathcal{N}}} & \mathcal{N}_{\text{reg}} & \xrightarrow{\mathcal{P}} & X & \xrightarrow{[\mathfrak{R}|\mathfrak{S}]} & \hat{X} & \xrightarrow{M[-]} & M[\hat{X}] \\
& & & & & & \downarrow \text{operational partition} & & \\
& & & & & & (E_{\text{split}}, E_{\text{res}}) & & \\
& & & & & & \downarrow \text{Hyp. 9.1} & & \\
& & & & & & G_{\text{split}} & \xrightarrow{\text{cuts}} & S_{\text{net}} \\
& & & & & & & & \downarrow \text{Hyp. 9.2/9.3} \\
& & & & & & & & S_{\text{hol}} \xrightarrow{\text{Hyp. 9.5}} \langle T_{\mu\nu} \rangle_{\text{source}}
\end{array}$$

Fastest falsifiers. The cheapest decisive tests are entropy-cone admissibility for a small network, the $f_{\text{edge}}(z)$ tracking fit, the vector-channel response integral, and the certified matroid split test. The type-II horizon algebra, de Sitter holography, and the coarse-grained edge theory are deeper architectural gates and should not be treated as solved.

Contents

Reader’s Guide: the framework in one page	2
1 Integration map and governing discipline	7
1.1 Status labels	8
1.2 Master dependency schema	8
1.3 Key notation and dictionary objects	10
1.4 Regulator stability convention	11
1.5 Architectural trade-off: from immediate cancellation to falsifiable sequestration . . .	12
2 Logical strata of the revised model	13
3 Causal topology in AQFT	14
3.1 Haag-Kastler net, vacuum, and microcausality	14
3.2 Geometric modular flow	14
3.3 Modular curvature and retarded response	14
4 Two distinct filters: split expectation and Schwinger-Keldysh response	15
4.1 The split-relative conditional expectation	15
4.2 Modular-compatibility status	16
4.3 Construction target for cosmological-horizon split expectations	16
4.4 Candidate strategies for constructing $\mathcal{N}_{\epsilon,\Lambda}$	17
4.5 The Schwinger-Keldysh response filter	20
4.6 Model-level causal topology map	22
4.7 Modular Equivalence Principle and inertial sequestration	22
5 Operational EPR sector	24
6 Statistics-resolved detector response and Volovik analogy	25
6.1 Bosonic and fermionic response kernels	25

6.2	Volovik dictionary	26
7	Response-weighted matroids	27
7.1	Realified response representation	27
7.2	Loops, coloops, and what they mean	29
7.3	Regulator dependence and continuum-limit bookkeeping	30
7.4	Matroids as validators and falsifiers of the split hypothesis	31
7.4.1	Continuous response data and discrete rank strata	31
7.4.2	Operational column partition	31
7.4.3	Sector separation as a direct-sum condition	32
7.4.4	Algorithmic split test	33
7.4.5	What the three-detector example does and does not validate	33
7.4.6	Matroid falsifiers	34
7.4.7	Non-EPR is not the same claim as non-gravitating	35
8	Corrected set-theoretic vacuum structure	35
9	Conditional Gravity Dictionary: From Response Data to Holographic Entropy and Sourcing	36
9.1	Network assignment	36
9.2	Entropy dictionary	37
9.3	Sourcing dictionary	38
9.4	Residual-sector decoupling	38
9.5	Dictionary ladder	40
9.6	Residual matroid-to-CSS construction target and addressability trade-off	40
10	Boundary modes, nuclearity, and area scaling	43
11	Dark matter from vector-mediated causal harvesting	44
11.1	Channel-counting benchmark	44
11.2	Response-integral target replacing the bare count	45
11.3	Boltzmann–CTVM kinetic spine	46
11.4	Role of longitudinal modes	46
12	Horizon-regulated dark energy	46
12.1	Horizon scale and edge density	47
12.2	Conditional null-boundary normalization	47
12.3	KMS detailed balance and fermionic stabilization	48
13	Cosmological dynamics and equation of state	50
14	QCD, topological sectors, and self-sustained vacuum analogies	51
15	Operational analogies and analogue experiments	51
15.1	Transport versus teleportation-like relabeling	52
15.2	BTZ and image-sum harvesting	52
15.3	Analogue-vacuum experimental program	52

16 Predictions, benchmarks, and falsifiers	53
16.1 Benchmarks retained as open targets	53
16.2 Near-term falsifiers	54
17 Open derivational targets and architectural gates	55
17.1 Critical path: cheapest decisive tests first	55
17.2 Architectural gates	56
17.3 Dark-matter phenomenology gates	58
17.4 Dark-energy phenomenology gates	59
17.5 Technical and computational targets	59
18 Revised axiom ledger	60
19 What CTVM does not claim	61
20 Conclusion	62
Glossary of abbreviations	62
A Schwinger-Keldysh identities	63
B Detector density matrix bookkeeping	63
C Minimal worked detector model: massive scalar split test	63
D Constant-parameter edge kinetic solution	64
E Why the dark-matter number is a target, not evidence	64
F Three-detector matroid example: physical interpretation and scope	65
F.1 Codimension correction	65
F.2 Physical interpretation	66
F.3 What the example shows	66
F.4 What the example does not show	66
F.5 Falsifier sharpened by the example	66
F.6 Verification by exhaustive pair enumeration	67
F.7 Three-detector CSS miniature	68
G Null-boundary normalization sensitivity	69
H Vector-channel response calculation: from Proca toy model to $f_{\text{vec}}^{\text{eff}}$	69
H.1 Protocol gate: detector-current model	70
H.2 Massive-vector spectral kernel	71
H.3 Detector response matrix	71
H.4 Species-resolved response fraction	72
H.5 Regulator and protocol scan	72
H.6 Longitudinal and gauge-consistency audit	72
H.7 Kramers-Kronig and causality diagnostics	73
H.8 Numerical performance hint	73

H.9	Matroid split diagnostic	73
H.10	Species-resolved denominator plan	74
H.11	Reporting rule	74
H.12	Longitudinal asymptotic near the electroweak transition	75
I	Theorem and proposition audit table	75
J	Executable critical-path protocols	78
J.1	Certified matroid split test	79
J.2	Entropy-cone admissibility test	79
J.3	$f_{\text{edge}(z)}$ tracking pre-filter	80
J.4	Vector-channel response protocol	80

1 Integration map and governing discipline

The preceding drafts contained several good ideas entangled with several over-strong derivational moves. Version 7 supplied the broad physical program: a causal-topology filter, dark-matter channel counting, horizon-edge dark energy, Schwinger-Keldysh response language, a response matroid, KMS detailed balance, and operational signatures. Version 8 compressed this program into a shorter roadmap and attempted to repair the presentation, while some claims still needed to be demoted to explicit hypotheses. The Volovik revision imposed the right discipline: use superfluid vacuum analogies as mechanism-building guidance, not as proof.

This integrated draft keeps the following items from v7:

1. the causal-topology diagnostic based on geometric modular flow, modular curvature, and retarded response;
2. the two-component effective vacuum fluid with an exchange term;
3. the Schwinger-Keldysh distinction between spectral/retarded response and symmetric Keldysh noise;
4. the detector-harvesting response matrix and response-matroid bookkeeping;
5. the vector-channel benchmark range $[0.20, 0.30]$, with $27/106.75$ retained as a historical reference point;
6. the horizon-edge kinetic spine and KMS detailed-balance argument;
7. the null-boundary normalization argument for κ_{edge} , but only conditionally;
8. the operational signatures involving harvesting, teleportation, and curved-spacetime image sums.

It keeps the following from v8:

1. the cleaner roadmap and statement of falsifiability;
2. the insistence on self-contained definitions in each technical section;
3. the emphasis on Pauli blocking and emergent pairing as statistics-driven mechanisms;
4. the desire to restore full document length and technical depth rather than merely patching a nine-page summary.

It keeps the following from the Volovik/statistical revision:

1. the layered logical structure;
2. the single representable linear matroid, augmented by a response-weight layer;
3. the corrected statement that Bose enhancement and Pauli blocking normally change weights, not matroid rank;
4. the explicit warning that Volovik's ${}^3\text{He}$ and ${}^4\text{He}$ systems are analogue universality classes, not literal components of the physical vacuum;

- the insistence that a holographic or semiclassical dictionary is an additional hypothesis, not an AQFT theorem.

The governing rule is simple: no theorem language is used unless the hypotheses actually imply the conclusion. Where the argument requires an extra compatibility condition, that condition is stated as a hypothesis. Where a calculation gives the correct numerical coefficient only after a convention or normalization is chosen, the convention is stated explicitly. Where a map acts on correlation functions rather than density operators, it is not called a CPTP state channel.

A second rule is equally important. The integrated document is a program map, not a completed solution of the cosmological constant problem. It conditionally addresses how a horizon-scale contribution of order $M_{\text{Pl}}^2 H^2$ might arise and why the present value could be an $O(1)$ fraction of critical density. It does not yet derive a cancellation or sequestration of the microscopic bulk vacuum energy. That magnitude/cancellation problem is isolated as an open target rather than treated as solved.

The matroid construction in earlier drafts of this program used a Gram-factorization representation; the present draft replaces it with the realified column-doubling representation. The null-GHY edge contribution, previously presented as a derivation of $\kappa_{\text{edge}} = 1/(4\pi)$, is now stated as a labeled normalization hypothesis with an explicit convention catalogue in Appendix G.

1.1 Status labels

We use five status labels throughout.

Label	Meaning
Definition	A choice of terminology or construction. It is neither true nor false until applied.
Theorem/proposition	A mathematical statement following from stated hypotheses.
Conditional mechanism	A physically motivated implication that holds only if a listed dictionary or compatibility assumption is imposed.
Benchmark	A numerical target or coincidence that must be derived by a deeper calculation before becoming a prediction.
Conjecture/open target	A claim not established here, but sharpened into a testable or calculable form.

This status discipline is the main difference between the present draft and the stronger derivational language of earlier versions.

1.2 Master dependency schema

The whole paper can be read as one conditional implication, with several independent gates. Making the gates explicit prevents later sections from silently promoting benchmarks into derivations.

Dependency Schema 1.1 (Conditional CTVM dependency schema). *Assume the following data and hypotheses.*

- A regulated AQFT detector setup with a chosen modular-compatible split inclusion: after the regulator has been imposed, the chosen type-I split factor is invariant under the reference-state modular flow and therefore admits the required state-preserving conditional expectation.

2. A detector protocol whose realified response matrix is compatible with the split-relative image/kernel decomposition.
3. A response-weighted matroid decomposition with no circuits crossing the proposed split/residual sector boundary.
4. A regulator-stable pairwise network assignment from the split-sector response block.
5. A holographic entropy dictionary in a regime where RT/HRT or bit-thread methods are justified, including the required entropy-cone admissibility conditions.
6. An additional cosmological-horizon extension of that dictionary when the application is de Sitter-like rather than AdS-like.
7. A semiclassical sourcing dictionary converting the corresponding entropy or relative-entropy variation into a conserved gravitational source.
8. Residual-sector decoupling from the source functional.
9. An edge-action/vacuum-energy dictionary dictionary if the dark-energy benchmark is invoked.

Then, within the selected dictionary, the split-sector pairwise response block supplies the gravitational source functional assigned by the model, while the residual block is inert with respect to that functional. If the edge normalization and late-time blocking hypotheses also hold, the dark-energy benchmark becomes a convention-dependent target. If the vector-channel response calculation reproduces the effective vector fraction, the dark-matter benchmark becomes a dynamical target rather than a bare count.

Justification. This is a dependency schema, not a new physical derivation. Items 1–3 define the regulated split and matroidal sector bookkeeping. Item 4 constructs a capacity-bearing network from the split block. Items 5–6 permit that network cut functional to be interpreted as holographic entropy in the relevant setting. Item 7 turns entropy variation into a gravitational source. Item 8 removes residual-coordinate dependence from that source functional. Item 9 is needed only for the horizon-edge numerical normalization. The conclusion is therefore exactly the conjunction of the stated dictionaries and decoupling assumptions.

Remark 1.1 (No optional links). *No item in Dependency Schema 1.1 is optional for the strong CTVM gravity claim. Failure of any item breaks the corresponding strong conclusion, although the detector-response construction, the SK response filter, and the matroidal bookkeeping may remain meaningful as independent tools.*

Remark 1.2 (What survives if the gravity gates fail). *The strong CTVM gravity claim requires the full dictionary stack. If Gates 5–9 fail, the manuscript still leaves a smaller but well-defined toolkit: the split/SK type separation, the regulated detector-response construction, the realified response-matroid falsifier, and the analogue-vacuum statistics tests. Those tools can be studied without claiming gravitational sourcing or residual-sector inertness. Thus the dependency schema should be read as a map of what must be added to the standalone response apparatus, not as evidence that the gravity dictionary has already been established.*

Gate	Enables	Downstream conclusion blocked if absent
1	Regulated split expectation or type-II replacement	Algebraic split/residual bookkeeping.
2	Detector/split compatibility	Response matrix can represent the proposed split.
3	Matroid direct-sum compatibility	Exact sector separation is not falsified by cross-circuits.
4	Pairwise response network	Cut/capacity language can be assigned to split-sector data.
5	Controlled holographic entropy dictionary	Network cuts may be interpreted as RT/HRT or bit-thread entropy.
6	Cosmological-horizon extension	Dark-energy application can use the entropy dictionary outside AdS-like settings.
7	Semiclassical sourcing dictionary	Entropy or response data can be read as gravitational source data.
8	Residual decoupling	Residual sector is inert with respect to the selected source functional.
9	Edge-action/vacuum-energy dictionary	Horizon-edge number $\Omega_\Lambda \simeq 2/3$ becomes a conditional target.

Independence clarification. The gates are logically independent unless an explicit theorem in the text connects them. They are also empirically independent in the conservative sense: success of one gate, such as a non-uniform response matroid, does not supply evidence for another gate, such as a cosmological-horizon holographic dictionary.

1.3 Key notation and dictionary objects

The model intentionally uses several mathematical layers. The following table is included to keep the types of objects separate.

Symbol	Type of object	Role in the framework
ρ_Ω	QFT state	Vacuum or reference state before detector restriction.
\mathcal{A}_{reg}	regulated local algebra	Algebraic input after cutoff, smearing, or split regularisation.
Π_{split}	conditional-expectation map, when it exists	Split-relative algebraic projection onto an intermediate type-I factor; not a gravity map.
P_{SK}	linear filter on cumulants	Keeps retarded/spectral Schwinger-Keldysh cumulants; not a CPTP state channel.
X	detector-response matrix	Protocol-dependent matrix built from smeared Wightman or response kernels.
\widehat{X}	realified response representation	Column-doubled real matrix $[\text{Re } X \mid \text{Im } X]$.
$M[\widehat{X}]$	representable matroid	Records support, rank, circuits, direct sums, and possible falsifiers of sector compatibility.
w	response-weight layer	Records amplitudes, Bose enhancement, Pauli blocking, and EPR-yield proxies.
G_{split}	weighted pairwise network	Additional construction from split-sector data; carries capacities, not matroid rank alone.
$S_{\text{net}}, S_{\text{RT}}$	entropy functionals	Network cut entropy and its conditional holographic/RT interpretation.
$\langle T_{\mu\nu} \rangle_{\text{grav}}$	gravitational source assignment	Requires a separate semiclassical sourcing dictionary.
κ_{edge}	dimensionless edge coefficient	Conditional normalization in $\rho_\Lambda = \kappa_{\text{edge}} f_{\text{edge}} m_{\text{pl}}^2 H^2$.
$f_{\text{sec}}(a)$	time-dependent fraction	Active-sector fraction $\rho_1/(\rho_1 + \rho_2)$ governing the dark-matter locking dynamics of §11.3; formerly written f .
f_{eq}	static benchmark range	Conservative response target range $[0.20, 0.30]$ to which f_{sec} is supposed to lock at the electroweak epoch if the response integral supports it; $27/106.75 \simeq 0.253$ is retained as the historical reference point.
$f_{\text{vec}}^{\text{eff}}(T; \mathcal{P})$	static response fraction	Effective vector-channel response fraction defined in §11.2 and computed in Appendix H; the dynamical target that must flow to f_{eq} for the benchmark to be supported.
$f_{\text{edge}}(N)$	time-dependent fraction	Active horizon-edge fraction governing the dark-energy kinetic spine of §12.3; appears as $\Omega_\Lambda = (8\pi/3)\kappa_{\text{edge}} f_{\text{edge}}$; formerly written f .

1.4 Regulator stability convention

The phrase “regulator-stable” is used throughout the white paper, so we fix a numerical convention for the execution program.

Definition 1.1 (Regulator stability). Let $Q(\epsilon, \Lambda; \mathcal{P})$ be a scalar, vector norm, matrix norm, response fraction, or entropy functional computed from a regulated detector protocol \mathcal{P} . Fix a regulator window

$$\epsilon \in [\epsilon_{\min}, \epsilon_{\max}], \quad \Lambda \in [\Lambda_{\min}, \Lambda_{\max}],$$

and hold all non-regulator protocol parameters fixed. We call Q regulator-stable at tolerance δ_{reg} if

$$\frac{\max_{\text{window}} Q - \min_{\text{window}} Q}{\text{median}_{\text{window}} |Q| + \eta} < \delta_{\text{reg}},$$

where η is a small numerical floor used only to avoid division by zero. Unless otherwise stated, the default white-paper tolerance is

$$\delta_{\text{reg}} = 0.1.$$

For dimensionful matrix data the same definition is applied to dimensionless ratios such as $\|\widehat{X}_{\text{off}}\|/\|\widehat{X}\|$, entropy-cone margins normalized by total capacity, or the response fraction $f_{\text{vec}}^{\text{eff}}$.

A result may be called *robust* when the same test passes with $\delta_{\text{reg}} < 0.05$, *qualified* when it passes with $0.05 \leq \delta_{\text{reg}} < 0.20$, and *unsupported* when the spread exceeds 0.20 or when a required gauge, Goldstone, Ward-identity, or entropy-cone consistency condition fails.

1.5 Architectural trade-off: from immediate cancellation to falsifiable sequestration

The present white paper chooses an explicit architectural trade-off. Earlier versions gestured toward a direct microscopic cancellation of vacuum energy. Version 15.5 does not claim such a derivation. Instead, it treats the cosmological-constant problem as a staged program whose immediate mathematical target is *structural sequestration*: identifying response sectors that remain accessible to retarded, source-carrying probes and residual sectors that are protected, weakly addressable, or absent from the selected source dictionary.

This trade-off protects the durable part of the framework. The split/SK distinction, detector-response construction, response-matroid falsifier, and analogue-vacuum statistics tests can survive even if the later gravity-dictionary gates fail. The strong cosmological claim does not survive those failures. In particular, the numbers 27/106.75 and 2/3 are retained as open targets whose evidential status depends on future response-integral, edge-theory, and expansion-history calculations.

In this framing the EPR-accessible split sector is operationally defined by detector protocols and computable bipartite witnesses. It is the sector that can enter pairwise response networks and, under the additional gravity dictionary of Section 9, a candidate source functional. The residual sector is not merely a bin for discarded correlations. It is the candidate home of multipartite, code-like, or topologically protected response structure. The construction in Section 9.6 is deliberately weaker than a proof: it gives a thresholded binary reduction and a possible route toward CSS-code protection only after an explicit valid CSS datum is supplied.

This architecture buys rigor and falsifiability at the cost of immediate numerical closure. The framework no longer says that the value of Λ has been computed. It says that a future cancellation, sequestration, or edge-conversion claim must close definite gates: residual decoupling, a coarse-grained edge theory, perturbation-order stability, and a viable $f_{\text{edge}}(z)$ history. The payoff is that failure is now well-defined. Cross-sector circuits, nonconvergent detector blocks, entropy-cone violations, gauge-unstable vector response, or an excluded tracking history each kills a specified arm of the program.

The Entropic Inertia companion sharpens this trade-off. If inertial response is the modular work cost of displacing energy relative to a Rindler screen, then residual gravitational inertness is not enough: a residual sector excluded from the source dictionary must also be excluded from the inertial modular-work functional. Otherwise CTVM would predict an equivalence-principle violation. Section 4.7 therefore adds a modular-equivalence consistency condition tying gravitational sourcing and inertial response to the same source-carrying sector.

The architectural thesis is therefore modest but operational: CTVM is a conditional framework for converting the old slogan “only the causal sector gravitates” into a sequence of calculable tests. The residual-code material below should be read as one proposed construction target inside that sequence, not as an established topological realization of the residual sector.

2 Logical strata of the revised model

The revised CTVM is not one derivation. It is a stack of constructions.

- (I) **Regulated AQFT structure.** A Haag-Kastler net with a cyclic and separating vacuum, microcausality, and suitable split inclusions supplies local algebras and type-I split factors. This yields tensor-product-like regulated bookkeeping. It does not by itself select a gravitational sector.
- (II) **Response-functional structure.** The Schwinger-Keldysh formalism separates retarded/spectral response from symmetric Keldysh fluctuations. Linear response of $\langle T_{\mu\nu} \rangle$ to metric perturbations is governed by retarded kernels. This supports a response filter on cumulants. It is not automatically a positive map on states.
- (III) **Operational detector structure.** A finite family of smeared detectors converts vacuum correlators into a regulated detector response matrix. An EPR-extractable sector is defined operationally by a post-harvesting LOCC distillation functional. This prevents the phrase “EPR sector” from becoming an unsupported ontology.
- (IV) **Statistics-resolved Volovik layer.** Bosonic collective channels carry Bose factors $1 + n_B$ and can amplify pairwise harvestability. Fermionic channels carry Pauli factors $1 - n_F$ and can suppress final-state access while preserving boundary or defect modes. Volovik’s ${}^4\text{He}$ and ${}^3\text{He}$ systems motivate this split as a universality-class analogy.
- (V) **Matroidal encoding.** The detector response matrix defines a representable linear matroid after realification. The matroid records rank, circuits, loops, and possible direct sums. A separate weight or valuation layer records response amplitudes. Magnitude effects do not generally change the matroid.
- (VI) **Conditional gravity dictionary.** A pairwise weighted network extracted from the split-factor sector may be assigned a bit-thread/RT interpretation only if a holographic or semiclassical dictionary is imposed. This step is outside pure Minkowski AQFT.
- (VII) **Cosmological phenomenology.** The dark-matter and dark-energy numbers are benchmark mechanisms built on the previous layers. They are retained because they are suggestive and falsifiable; they are not presented as completed first-principles derivations.

3 Causal topology in AQFT

3.1 Haag-Kastler net, vacuum, and microcausality

Let $O \mapsto \mathcal{A}(O)$ be a Haag-Kastler net of von Neumann algebras on a globally hyperbolic spacetime, represented on a Hilbert space \mathcal{H} with vacuum vector Ω . For causally disjoint regions $O_1 \perp O_2$, microcausality states

$$[A_1, A_2] = 0, \quad A_i \in \mathcal{A}(O_i). \quad (1)$$

For a scalar field this implies that the Pauli-Jordan commutator function vanishes at spacelike separation. With the convention

$$[\phi(x), \phi(y)] = i\Delta_{\text{PJ}}(x, y), \quad (2)$$

the Wightman function decomposes as

$$D^+(x, y) = \langle \Omega | \phi(x) \phi(y) | \Omega \rangle = \frac{1}{2} G^{(1)}(x, y) + \frac{i}{2} \Delta_{\text{PJ}}(x, y), \quad (3)$$

where $G^{(1)}(x, y) = \langle \{\phi(x), \phi(y)\} \rangle$ is the Hadamard function. This convention avoids the double- i error that appeared in earlier drafts.

Microcausality constrains directed response. It does not eliminate spacelike vacuum correlations. Reeh-Schlieder-type behavior and the nonzero Wightman function mean that spacelike-separated detectors can harvest correlations even when the commutator vanishes. This distinction between symmetric correlation and retarded influence is the starting point of the causal-topology diagnostic.

3.2 Geometric modular flow

For a region R such that Ω is cyclic and separating for $\mathcal{A}(R)$, Tomita-Takesaki theory defines the modular automorphism group

$$\sigma_t^R(A) = \Delta_R^{it} A \Delta_R^{-it}, \quad A \in \mathcal{A}(R). \quad (4)$$

Definition 3.1 (Geometric modular flow). *The pair $(\mathcal{A}(R), \Omega)$ exhibits geometric modular flow if there exists a one-parameter family of spacetime automorphisms φ_t^R preserving R such that σ_t^R implements φ_t^R on local operators.*

The Bisognano-Wichmann theorem gives this property for wedge regions in relativistic QFT, and conformal symmetry extends related local modular Hamiltonians to balls in CFT vacua. In these cases the modular Hamiltonian has a local stress-tensor representation. In generic regions modular flow is nonlocal; then it cannot directly define a local geometric transport structure.

3.3 Modular curvature and retarded response

Let λ^i parametrize a smooth family of regions $R(\lambda)$, with reduced state ρ_λ and modular Hamiltonian $K_\lambda = -\log \rho_\lambda$. Let P_0 denote the projection onto the centralizer or modular zero-mode sector. The modular connection and curvature may be written schematically as

$$A_i(\lambda) = P_0[\partial_i K_\lambda], \quad F_{ij} = \partial_i A_j - \partial_j A_i + i[A_i, A_j]. \quad (5)$$

In geometric modular-flow cases, shape variations of K_λ can be represented by stress-tensor integrals. The commutator term in F_{ij} is then tied to Kubo response. Thus nonzero modular curvature is a diagnostic of nontrivial transport/response structure under deformations.

Definition 3.2 (Causal-topology diagnostic). *A state-region pair $(\mathcal{A}(R), \Omega)$ has causal topology, relative to a specified deformation family, if it satisfies:*

1. *geometric modular flow or a controlled horizon-regulated substitute;*
2. *nonzero modular curvature or holonomy for the deformation family;*
3. *nonzero retarded/spectral support in the channel relevant for the proposed stress response.*

The diagnostic deliberately refers to response, not merely to entanglement. A spacelike Bell pair is highly entangled but cannot be used for directed local influence. Conversely, a retarded response channel may be weakly entangled but causally active. CTVM is therefore not simply “bipartite entanglement gravitates.” Its intended criterion is response-compatible causal topology.

Remark 3.1 (Terminology). *The phrase “causal topology” is a diagnostic name in this manuscript, not yet a claim that a genuine topological invariant has been constructed for every response sector. Actual topological content enters only when a protected invariant—for example a Chern number, winding number, \mathbb{Z}_2 index, or K -theory class—is supplied by a boundary or residual-sector model. Until then, the diagnostic should be read as a response-geometric criterion rather than as a completed topological classification.*

4 Two distinct filters: split expectation and Schwinger-Keldysh response

Earlier drafts blurred two maps. They must be separated.

4.1 The split-relative conditional expectation

Let $O_1 \Subset O_2$ be a split inclusion. The split property gives an intermediate type-I factor

$$\mathcal{A}(O_1) \subset \mathcal{N} \subset \mathcal{A}(O_2). \quad (6)$$

The type-I factor supplies a tensor-product-like description for regulated local degrees of freedom. It does not imply a unique canonical factor, and it does not imply that all entanglement represented there is bipartite. It also does not automatically supply a vacuum-preserving conditional expectation onto every chosen subalgebra.

Hypothesis 4.1 (Modular-compatible split expectation). *For the chosen split inclusion and reference state ω , the regulated type-I factor is assumed to satisfy*

$$\sigma_t^\omega(\mathcal{N}_{\text{reg}}) = \mathcal{N}_{\text{reg}}, \quad \forall t \in \mathbb{R}, \quad (7)$$

where σ_t^ω is the modular automorphism group of $(\mathcal{A}_{\text{reg}}, \omega)$. Under this modular-compatibility condition we assume the existence of a normal ω -preserving conditional expectation

$$E_{\mathcal{N}} : \mathcal{A}_{\text{reg}} \rightarrow \mathcal{N}_{\text{reg}}, \quad \omega \circ E_{\mathcal{N}} = \omega. \quad (8)$$

Remark 4.1 (Scope of the split expectation). *The Doplicher-Longo split property gives intermediate type-I factors. It does not, by itself, guarantee that an arbitrary intermediate factor is invariant under the modular flow of the ambient regulated algebra. By Takesaki's criterion, this modular invariance is the relevant condition for a state-preserving normal conditional expectation. The expectation $E_{\mathcal{N}}$ is therefore a regulated modular-compatibility assumption, not a generic consequence of the type-III split property. Clean geometric control is available in wedge/Bisognano-Wichmann settings and in conformal-ball settings in CFT; causal diamonds, light-cone tubes, and cosmological horizons require an explicitly constructed modular-compatible regulator or an independent proof of modular invariance.*

4.2 Modular-compatibility status

For referee clarity we separate what is known, what is conjectural, and what would have to be constructed.

- **Wedges and conformal balls.** In Rindler wedges and CFT balls the modular flow is geometrically controlled. These cases are the cleanest laboratories for testing the split-expectation logic.
- **Cosmological horizons.** For de Sitter or cosmological-horizon applications, a KMS state and a static-patch Killing flow do not by themselves give a state-preserving conditional expectation onto a chosen split factor. The horizon split remains a hypothesis or construction target.
- **Constructive path to proof.** A proof requires either a regulated modular-compatible type-I split factor as in Construction 4.1, or an observer-inclusive type-II crossed-product route as in Construction 4.2, together with regulator-stable detector and stress-response data. The massive-scalar detector model in Appendix C gives only a controlled approximate detector/split test; Appendix G separately records the normalization sensitivity of the horizon-edge application.

4.3 Construction target for cosmological-horizon split expectations

Hypothesis 4.1 is strongest in wedge and conformal-ball settings, where the modular flow is known geometrically. The cosmological-horizon application is more delicate. The Bunch-Davies or Gibbons-Hawking state restricted to a de Sitter static patch is KMS with respect to the static-patch Killing flow, but this does not by itself construct a state-preserving conditional expectation onto a chosen intermediate type-I split factor.

The required construction is the following.

Construction 4.1 (Regulated modular-compatible horizon split). *Fix a static patch \mathcal{D}_H with Killing flow α_s generated by the static-patch horizon time. Let ω_H be the reference KMS state. Introduce a finite-resolution horizon regulator (ϵ, Λ) , for example by replacing the null horizon by a stretched horizon at proper distance ϵ and restricting to modes below energy cutoff Λ . Let*

$$\mathcal{A}_{\epsilon, \Lambda}(\mathcal{D}_H) \tag{9}$$

be the resulting regulated algebra.

Choose nested regulated regions or collars

$$\mathcal{O}_{\text{in}}^{\epsilon} \Subset \mathcal{O}_{\text{out}}^{\epsilon} \subset \mathcal{D}_H \tag{10}$$

and an intermediate type-I factor

$$\mathcal{A}_{\epsilon,\Lambda}(\mathcal{O}_{\text{in}}) \subset \mathcal{N}_{\epsilon,\Lambda} \subset \mathcal{A}_{\epsilon,\Lambda}(\mathcal{O}_{\text{out}}). \quad (11)$$

The construction is admissible for CTVM only if the following conditions hold:

1. **Regulated type-I property.** $\mathcal{N}_{\epsilon,\Lambda}$ is type-I or finite-dimensional at fixed regulator.

2. **Modular covariance.** The regulated modular flow satisfies

$$\sigma_s^{\omega_H, \epsilon, \Lambda}(\mathcal{N}_{\epsilon, \Lambda}) = \mathcal{N}_{\epsilon, \Lambda} \quad \forall s \quad (12)$$

at fixed regulator.

3. **State-preserving expectation.** There exists a normal conditional expectation

$$E_{\epsilon, \Lambda} : \mathcal{A}_{\epsilon, \Lambda}(\mathcal{D}_H) \longrightarrow \mathcal{N}_{\epsilon, \Lambda} \quad (13)$$

such that

$$\omega_H \circ E_{\epsilon, \Lambda} = \omega_H. \quad (14)$$

4. **Detector compatibility.** For the detector protocols used to build \widehat{X} , the correlation functions computed after applying $E_{\epsilon, \Lambda}$ converge, or at least stabilize, in the regulated detector window.

5. **Stress-tensor compatibility.** The renormalised stress tensor obtained after the split expectation preserves the Ward identity up to controlled regulator errors:

$$\nabla^\mu \langle T_{\mu\nu} \rangle_{E_{\epsilon, \Lambda}} = O(\epsilon, \Lambda^{-1}). \quad (15)$$

6. **Horizon limit.** The limit

$$E_H = \lim_{\epsilon \rightarrow 0, \Lambda \rightarrow \infty} E_{\epsilon, \Lambda} \quad (16)$$

need not exist as a normal expectation on the type-III continuum algebra, but the detector response data and source functionals used in the CTVM dictionary must have a regulator-stable limit.

The key point is that the conditional expectation is a regulated object. The continuum type-III algebra need not admit the same finite-resolution expectation as a normal map. CTVM only requires the regulated expectation to be stable on the detector-response and stress-response observables that enter the dictionary.

Remark 4.2 (Why the cosmological-horizon case is open). *The static-patch Killing flow provides a natural candidate for the geometric part of the modular flow, and the Gibbons-Hawking state is thermal with respect to that flow. However, the existence of a Killing flow and a KMS state does not imply that an arbitrary intermediate split factor is invariant under the reference modular automorphism group. Establishing*

$$\sigma_s^{\omega_H, \epsilon, \Lambda}(\mathcal{N}_{\epsilon, \Lambda}) = \mathcal{N}_{\epsilon, \Lambda} \quad (17)$$

is the missing step. This is why the dark-energy application remains conditional on Construction 4.1.

4.4 Candidate strategies for constructing $\mathcal{N}_{\epsilon, \Lambda}$

There are four plausible routes.

Strategy I: static-patch mode algebra. Diagonalise the regulated static-patch modular Hamiltonian $K_H^{\epsilon,\Lambda}$, and build $\mathcal{N}_{\epsilon,\Lambda}$ from spectral bands that are invariant under

$$\sigma_s^{\omega_H,\epsilon,\Lambda}(A) = e^{isK_H^{\epsilon,\Lambda}} A e^{-isK_H^{\epsilon,\Lambda}}. \quad (18)$$

In a finite-dimensional cutoff model, a subalgebra generated by blocks of the modular-energy decomposition is automatically invariant under this flow. The price is that locality is no longer automatic: a modular-energy block algebra may fail to approximate a local horizon collar or detector region. A successful Strategy I construction must therefore prove both modular covariance and approximate locality in the detector-response window.

Strategy II: stretched-horizon lattice algebra. Replace the horizon collar by a finite lattice of stretched-horizon cells. Choose the reference state

$$\rho_{\epsilon,\Lambda} \propto e^{-\beta_H K_{\epsilon,\Lambda}} \quad (19)$$

and choose $\mathcal{N}_{\epsilon,\Lambda}$ to be a tensor factor or block algebra commuting with the modular block decomposition of $\rho_{\epsilon,\Lambda}$. Then a finite-dimensional $\rho_{\epsilon,\Lambda}$ -preserving conditional expectation can be constructed by the corresponding block conditional expectation. The hard part is proving regulator stability and approximate covariance under static-patch time evolution.

Strategy III: modularly averaged split factor. Start with any split type-I factor \mathcal{N}_0 , then define its modular orbit

$$\mathcal{N}_s = \sigma_s^\omega(\mathcal{N}_0). \quad (20)$$

One can attempt to construct a modularly invariant effective factor by coarse-graining over the orbit,

$$\mathcal{N}_{\text{av}} \sim \text{vN} \left(\bigcup_{|s| \leq s_*} \mathcal{N}_s \right) \quad (21)$$

followed by a finite-resolution projection back to a type-I algebra. This route is conceptually natural but technically fragile. The modular average may become too large, lose approximate locality, or cease to be type-I unless the regulator and projection are part of the definition. A successful Strategy III construction must show that the averaged-and-projected algebra still gives stable detector-response data and does not wash out the split/residual distinction.

Strategy IV: observer-inclusive crossed-product/type-II construction. A fourth route is to avoid demanding a type-I split factor directly inside the continuum static-patch type-III algebra. Instead, one first passes to an observer-inclusive algebra in which the modular flow has been incorporated by a crossed-product construction or by gravitational dressing to an observer worldline. Motivated by the type-II₁ static-patch algebra of observables for de Sitter space, define a crossed-product horizon algebra schematically by

$$\mathfrak{A}_H = \mathcal{A}(\mathcal{D}_H) \rtimes_{\sigma^{\omega_H}} \mathbb{R}, \quad (22)$$

or, in the regulated setting,

$$\mathfrak{A}_H^{\epsilon,\Lambda} = p_{\epsilon,\Lambda} \left(\mathcal{A}_{\epsilon,\Lambda}(\mathcal{D}_H) \rtimes_{\sigma^{\omega_H,\epsilon,\Lambda}} \mathbb{R} \right) p_{\epsilon,\Lambda}, \quad (23)$$

where $p_{\epsilon,\Lambda}$ is a finite-trace or finite-resolution projection selecting the observer energy window and stretched-horizon cutoff.

Umegaki expectation in the finite algebra. The main advantage of the observer-inclusive type-II route is that the conditional expectation problem becomes mathematically cleaner. If the regulated static-patch algebra is a finite von Neumann algebra \mathfrak{A}_H equipped with a faithful normal trace τ_H , then for any unital von Neumann subalgebra $\mathfrak{B} \subset \mathfrak{A}_H$ there is a canonical τ_H -preserving conditional expectation

$$\mathbb{E}_{\mathfrak{B}} : \mathfrak{A}_H \rightarrow \mathfrak{B}, \quad \tau_H \circ \mathbb{E}_{\mathfrak{B}} = \tau_H. \quad (24)$$

Concretely, $\mathbb{E}_{\mathfrak{B}}$ is the $L^2(\mathfrak{A}_H, \tau_H)$ -orthogonal projection onto $L^2(\mathfrak{B}, \tau_H)$, restricted back to \mathfrak{A}_H . This is the finite-algebra/Umegaki form of conditional expectation [49]. Thus, if the observer-inclusive crossed-product construction is adopted, the existence of the trace-preserving expectation is no longer the hard part. The remaining hard problems are physical and regulator-theoretic: identifying the correct subalgebra \mathfrak{B} corresponding to the split-factor response sector and showing that the resulting detector-response and stress-response data are stable under the horizon regulator.

The Umegaki expectation is canonically τ_H -preserving. If the CTVM source dictionary requires preservation of a nonracial physical state ω_H , then an additional compatibility condition between \mathfrak{B} and the modular flow of ω_H is required. In the type-II route, the trace-preserving expectation is therefore algebraically guaranteed once the finite algebra and unital subalgebra are fixed, while physical-state preservation remains part of the subalgebra-selection and regulator-stability problem.

Construction 4.2 (Observer-inclusive crossed-product horizon algebra). *A CTVM type-II horizon construction consists of the following data:*

1. a static-patch algebra $\mathcal{A}(\mathcal{D}_H)$ and reference KMS state ω_H ;
2. the modular automorphism group $\sigma_s^{\omega_H}$ or its regulated version;
3. a crossed-product or observer-dressed algebra \mathfrak{A}_H that is type II, or becomes type II after a controlled finite-trace projection;
4. a finite-resolution projection $p_{\epsilon,\Lambda}$ and a unital von Neumann subalgebra $\mathfrak{N}_{\epsilon,\Lambda} \subset p_{\epsilon,\Lambda} \mathfrak{A}_H p_{\epsilon,\Lambda}$; the Umegaki expectation below preserves the finite trace, while preservation of any nonracial detector reference state is an additional compatibility condition;
5. the Umegaki trace-preserving conditional expectation

$$\mathbb{E}_{\epsilon,\Lambda} : p_{\epsilon,\Lambda} \mathfrak{A}_H p_{\epsilon,\Lambda} \longrightarrow \mathfrak{N}_{\epsilon,\Lambda}, \quad \tau_H^{\epsilon,\Lambda} \circ \mathbb{E}_{\epsilon,\Lambda} = \tau_H^{\epsilon,\Lambda}; \quad (25)$$

6. a comparison map from $\mathfrak{N}_{\epsilon,\Lambda}$ back to the finite detector-response algebra used to build \widehat{X} .

The construction is successful only if the induced detector response data, response matroid, and source functional are stable under the allowed horizon regulators.

Remark 4.3 (Why the type-II route is a preferred algebraic target). *The type-I split route must separately prove modular compatibility of the chosen intermediate factor. The observer-inclusive type-II route supplies a finite trace, and hence a trace-preserving conditional expectation onto any chosen unital von Neumann subalgebra by the finite-algebra/Umegaki construction. This is a powerful algebraic simplification, but it does not solve the physical-state problem: the Umegaki expectation*

preserves the trace τ_H , not automatically the nontracial static-patch or KMS state ω_H that may enter the source dictionary. If ω_H -preservation is required, the subalgebra must still be compatible with the ω_H modular flow. Thus the type-II route relocates the modular-compatibility obstruction from existence of an expectation to the selection of the physically correct subalgebra and regulator-stable state.

Remark 4.4 (Caveats for the type-II route). *The price is interpretive and technical. The type-II construction is observer-inclusive, not a state-independent local algebra of the original AQFT net. It may not satisfy the same locality, covariance, or split-inclusion properties as the original net. It also shifts the CTVM split from a type-I split factor in the field algebra to a trace-preserving expectation inside an observer-dressed horizon algebra. Readers who do not accept the observer-inclusive interpretation may prefer Strategies I–III, but then the modular-compatible type-I split factor must be constructed directly.*

Minimal success criterion. For CTVM, the minimal success criterion is not a continuum theorem on the full type-III algebra. It is the weaker, operational condition that for the chosen detector family \mathcal{P} ,

$$\lim_{(\epsilon, \Lambda), (\epsilon', \Lambda')} \left| X_{ij}^{E_{\epsilon, \Lambda}} - X_{ij}^{E_{\epsilon', \Lambda'}} \right| = 0 \quad (26)$$

inside a stable regulator window, and that the resulting \hat{X} , $M[\hat{X}]$, and source functional are stable under allowed regulator changes.

Current priority. Construct either a regulated modular-compatible type-I horizon split factor $\mathcal{N}_{\epsilon, \Lambda}$ or, preferably, an observer-inclusive type-II crossed-product horizon algebra \mathfrak{A}_H . In the type-II route, the trace-preserving expectation exists algebraically by the finite-algebra/Umegaki construction once the finite algebra and unital subalgebra are fixed; the remaining tasks are to identify the physically correct split-sector subalgebra, address physical-state preservation if a nontracial state is used, and prove regulator stability of detector-response and stress-response consequences.

Given a detector restriction/coarse-graining map \mathcal{R}_{det} , define

$$\Pi_{\text{split}} = E_{\mathcal{N}} \circ \mathcal{R}_{\text{det}}. \quad (27)$$

This is the split-relative algebraic projection used for finite-regulated tensor bookkeeping. Its image and kernel define

$$\mathcal{S}_{\text{reg}} = \text{im } \Pi_{\text{split}} \oplus \ker \Pi_{\text{split}} \quad (28)$$

only on the chosen regulated vector space of states or cumulants. The CTVM interpretation of $\text{im } \Pi_{\text{split}}$ as the EPR-accessible sector and $\ker \Pi_{\text{split}}$ as residual/topological structure is an additional modeling identification.

4.5 The Schwinger-Keldysh response filter

The closed-time-path generating functional is

$$Z[J_+, J_-] = \int \mathcal{D}\phi_+ \mathcal{D}\phi_- \exp \{ iS[\phi_+] - iS[\phi_-] + iJ_+ \phi_+ - iJ_- \phi_- \}. \quad (29)$$

After the Keldysh rotation

$$\phi_c = \frac{1}{\sqrt{2}}(\phi_+ + \phi_-), \quad \phi_q = \frac{1}{\sqrt{2}}(\phi_+ - \phi_-), \quad (30)$$

the free two-point structure may be organized as

$$G = \begin{pmatrix} G^K & G^{\text{ret}} \\ G^{\text{adv}} & 0 \end{pmatrix}, \quad (31)$$

where G^K is the symmetric/statistical correlator and G^{ret} is the retarded propagator.

A metric perturbation $h_{\alpha\beta}$ changes the stress-tensor one-point function by

$$\delta\langle T_{\mu\nu}(x) \rangle = \frac{1}{2} \int d^d y \chi_{\mu\nu,\alpha\beta}^{\text{ret}}(x, y) h^{\alpha\beta}(y). \quad (32)$$

The retarded kernel, not the symmetric noise kernel, controls linear response.

Definition 4.1 (Schwinger-Keldysh response filter). *On the two-point cumulant block define*

$$P_{\text{SK}}^{(2)} : \begin{pmatrix} G^K & G^{\text{ret}} \\ G^{\text{adv}} & 0 \end{pmatrix} \mapsto \begin{pmatrix} 0 & G^{\text{ret}} \\ G^{\text{adv}} & 0 \end{pmatrix}. \quad (33)$$

At higher order, $P_{\text{SK}}^{(n)}$ retains only those cumulants containing response legs sufficient to generate causal variation of the relevant one-point function. The direct sum over cumulant order is denoted P_{SK} .

Proposition 4.1 (Conditional Proposition 4.1). *P_{SK} is a linear idempotent map on the chosen correlation-functional space. If the full renormalized retarded stress response satisfies the Ward identity, then the retained retarded part satisfies the corresponding linearized conservation law.*

Proof. Linearity and idempotence follow from the block projection. The Ward statement follows because the divergence of the retained retarded response kernel is unchanged by the projection, once the renormalized Ward identity is assumed as input from the locally covariant stress-tensor construction. \square

Remark 4.5 (Ward identity input). *The Ward identity is an **input** from the renormalized locally covariant stress-tensor construction, such as the Hollands–Wald framework; it is not proved in this subsection. The proposition states only that the SK response filter does not disturb the retained retarded Ward identity once that identity is available. The trace anomaly and contact-term bookkeeping belong to the renormalized stress-tensor construction, not to the cumulant projection itself.*

Remark 4.6 (Not a CPTP theorem). *P_{SK} is not automatically a completely-positive trace-preserving map on density operators. Retarded Green functions are response kernels, not density matrices. A CPTP map may exist in a separate open-system construction, but it is not obtained merely by setting G^K to zero. This is one of the central repairs relative to v7.*

Example 4.1 (A cumulant filter that cannot be a state channel). *Consider a single canonically normalised smeared mode with quadratures q, p and $[q, p] = i$. A physical state must satisfy the Robertson bound*

$$\langle q^2 \rangle \langle p^2 \rangle - \frac{1}{4} \langle qp + pq \rangle^2 \geq \frac{1}{4}. \quad (34)$$

The two-point SK filter may consistently retain the commutator/retarded kernel while deleting the symmetric fluctuation block, which in this one-mode caricature corresponds to setting the symmetric covariances $\langle q^2 \rangle$, $\langle p^2 \rangle$, and $\langle qp + pq \rangle$ to zero while keeping $[q, p] = i$. The resulting cumulant data violate the uncertainty inequality and therefore cannot be the covariance matrix of any density operator. The retained retarded kernel is still the correct causal response function, but the filter is a cumulant-level linear projection, not a state channel.

4.6 Model-level causal topology map

The model-level causal-topology map should not be written as an ordinary composition of maps of the same type. The split-relative conditional expectation acts on the regulated algebraic state/correlation space; the Schwinger-Keldysh filter acts on response cumulants; and the sourcing dictionary acts on stress-tensor or entropy variations. The logical order is therefore this: first a split-relative algebraic projection regulates the local algebra and supplies tensor-product bookkeeping; then a response-functional filter selects the retarded cumulant content; then an additional semiclassical or holographic sourcing dictionary maps the retained response data to gravitational stress-energy. Only the first piece is a conditional expectation under the split/Takesaki hypotheses. Only the second piece is a linear filter on cumulants. Only the third piece speaks about gravity.

The type discipline is more transparent in table form:

Input object	Operation	Output object
state on \mathcal{A}_{reg}	split-relative conditional expectation $E_{\mathcal{N}}$	state on \mathcal{N}_{reg}
state on \mathcal{N}_{reg}	detector restriction/coarse-graining	detector state $\rho_D^{(\mathcal{P})}$
detector state or generating functional	connected-cumulant expansion	CTP cumulants $C_{s_1 \dots s_n}^{(n)}$
CTP cumulants	linear SK response filter P_{SK}	retarded/spectral cumulants
retarded/spectral cumulants	semiclassical or holographic sourcing dictionary	candidate $\langle T_{\mu\nu} \rangle_{\text{grav}}$

The rows are logical processing steps. They are not a single composition of same-type maps.

4.7 Modular Equivalence Principle and inertial sequestration

The Entropic Inertia companion argues that, in a Rindler wedge, inertia can be read as a modular thermodynamic work response rather than as a primitive force law. In the Bisognano–Wichmann setting the wedge modular Hamiltonian is local, the restricted vacuum is KMS at the Unruh temperature, and the entanglement first law gives the chain

$$\delta\langle K_W \rangle \longrightarrow \delta S_W \longrightarrow \delta W = T_U \delta S_W = \frac{E}{c^2} a \delta x. \quad (35)$$

For rest energy $E = mc^2$ this gives $\delta W = ma \delta x$, hence $F = ma$ in the quasi-static Rindler-screen limit [41]. The input is not Newton’s second law but the local modular/KMS structure of the accelerated observer’s accessible algebra.

This observation changes how CTVM should formulate the equivalence principle. The equality of inertial and gravitational mass is not taken here as a separate dark-sector postulate. It is treated as a sector-matching consistency condition: the same causally accessible, retarded-response sector must control both the modular work cost of acceleration and the gravitational source assignment. Schematically,

$$(\mathcal{A}_{\text{inert}}, \omega_{\text{inert}}, K_{\text{inert}}, \chi_{\text{inert}}^{\text{ret}}) \simeq (\mathcal{A}_{\text{grav}}, \omega_{\text{grav}}, K_{\text{grav}}, \chi_{\text{grav}}^{\text{ret}}) \quad (36)$$

inside the source-carrying split sector. In this form the Einstein equivalence principle becomes a *Modular Equivalence Principle*: inertial response and gravitational response agree when they are computed from the same accessible algebra, restricted state, modular flow, and retarded stress-response kernel.

The residual sector therefore faces a stronger consistency requirement than gravitational decoupling alone. If a residual sector contributed to the inertial modular work functional but not to the gravitational source functional, then CTVM would predict an equivalence-principle violation. Thus a residual sector that is sequestered from gravity must also be sequestered from ordinary inertial response, at least to current experimental bounds.

A useful algebraic way to state the condition is in terms of the spatial momentum generator normal to the local Rindler screen. Let P^1 generate physical translations in the x^1 direction and let \mathcal{A}_{res} be the residual regulated algebra. We do *not* require the microscopic commutator $[P^1, A]$ to vanish for every residual operator A ; that would be too strong and generally false. What CTVM requires is that the commutator vanish after projection to the inertial/source-response quotient. One possible defect functional is

$$D_{P^1}^{\text{res}} = \sup_{\substack{A \in \mathcal{A}_{\text{res}} \\ \|A\| \leq 1}} \left\| \Pi_{\text{src}}^{\text{ret}}([P^1, A]) \right\|, \quad (37)$$

where $\Pi_{\text{src}}^{\text{ret}}$ denotes the combined operation of retaining the retarded response component and restricting to the source-carrying sector. Exact inertial sequestration is the condition

$$D_{P^1}^{\text{res}} = 0, \quad (38)$$

or, in a finite-resolution model, $D_{P^1}^{\text{res}} = O(F_\nu^2)$ with the same residual-suppression factor used for leakage.

Equivalently, for the residual modular Hamiltonian K_{res} , spatial displacement should not change the residual contribution to the modular work functional:

$$\left. \frac{d}{ds} \omega \left(e^{isP^1} K_{\text{res}} e^{-isP^1} \right) \right|_{s=0} = i \omega([P^1, K_{\text{res}}]) = 0 \quad (39)$$

inside the same source-response quotient, again up to controlled $O(F_\nu^2)$ corrections if leakage is allowed. The corresponding retarded statement is that residual operators have no source-relevant response to the translation generator,

$$\chi_{P^1, \mathcal{A}_{\text{res}}}^{\text{ret}}(t) = i\theta(t) \omega([P^1(t), A_{\text{res}}(0)]) \mapsto 0 \quad (40)$$

after the split/SK/source projection. This retarded kernel is the response-function version of the commutator defect $D_{P^1}^{\text{res}}$ above: $D_{P^1}^{\text{res}}$ controls the projected operator response, while $\chi_{P^1, \mathcal{A}_{\text{res}}}^{\text{ret}}$ controls the corresponding linear-response signal.

Equivalence-principle bound. If a fraction ϵ_{res} of a bound system’s inertial energy is stored in residual-sector degrees of freedom that do not enter the gravitational source functional, then the Eötvös parameter between two test bodies A and B receives a contribution of order

$$\eta_{AB} \sim \left| \epsilon_{\text{res}}^{(A)} - \epsilon_{\text{res}}^{(B)} \right|. \quad (41)$$

Equivalence-principle experiments therefore constrain the same residual participation parameter that appears in inertial sequestration. In the topological-uniformity picture one expects

$$\epsilon_{\text{res}} = O(F_\nu^2), \quad (42)$$

so Eötvös bounds become direct constraints on any sector that is gravitationally sequestered but might still contribute to ordinary inertial response. The present manuscript does not compute a numerical bound; it records the required consistency condition.

This is the mathematical answer to how \mathcal{A}_{res} should decouple from P^1 in the next commutator-algebra setup. The residual algebra may carry nonlocal or code-like internal modular flow, but its response to physical translations must be invisible to the same local modular work and retarded stress-response functionals that define inertia and gravitational sourcing. Establishing this is an open derivational target, not a theorem of the present manuscript.

5 Operational EPR sector

The phrase “EPR sector” must be made operational. Otherwise it functions as decorative language.

Let \mathcal{P} denote the full harvesting protocol: detector trajectories, spacetime smearing, switching functions, energy gaps, couplings, regulator, and any allowed post-selection convention. Let $\mathcal{D} = \{D_i\}_{i=1}^N$ be the associated finite detector family, initially in a product state $\rho_{D,0}$. The post-interaction detector state is

$$\rho_D^{(\mathcal{P})} = \text{Tr}_\Phi \left[U_{\mathcal{P}}(\rho_{D,0} \otimes |\Omega\rangle\langle\Omega|) U_{\mathcal{P}}^\dagger \right]. \quad (43)$$

The unitary $U_{\mathcal{P}}$ is local in spacetime, but the field is a shared quantum resource. Therefore this harvesting map is not LOCC on the detector state alone. Entanglement can appear between initially separable detectors because the vacuum itself carries correlations.

Definition 5.1 (Protocol-dependent computable EPR witness). *For each detector pair (i, j) , let*

$$\rho_{ij}^{(\mathcal{P})} = \text{Tr}_{\overline{ij}} \rho_D^{(\mathcal{P})}$$

be the two-detector reduced state produced by the harvesting protocol \mathcal{P} . The operational witness of record in this paper is the logarithmic negativity

$$E_{\mathcal{N}}^{(\mathcal{P})}(i, j) = \log_2 \left\| \left(\rho_{ij}^{(\mathcal{P})} \right)^{T_i} \right\|_1.$$

Given a threshold $\epsilon_E \geq 0$, define the thresholded pair support

$$\mathcal{E}_{\text{pair}}^{\mathcal{P}, \epsilon} = \left\{ (i, j) : E_{\mathcal{N}}^{(\mathcal{P})}(i, j) > \epsilon_E \right\}.$$

A detector site is EPR-active if it participates in at least one pair in $\mathcal{E}_{\text{pair}}^{\mathcal{P}, \epsilon}$.

The CTVM EPR-accessible sector is therefore the part of the regulated detector response detected by the chosen bipartite witness. It is a property of the vacuum together with a specified harvesting protocol, threshold, and witness, not a detector-independent property of the vacuum alone. The residual sector is correlation structure not captured by this bipartite witness. This definition does not imply that the residual sector is gravitationally inert; it only defines the operational split that later hypotheses may act on.

For network and matroid tests we use the same thresholded pair support and define, when needed, the capacity proxy

$$c_{ij}^{\mathcal{P}} = \alpha_{\mathcal{P}} E_{\mathcal{N}}^{(\mathcal{P})}(i, j) \quad \text{or} \quad c_{ij}^{\mathcal{P}} = \alpha_{\mathcal{P}} \mathcal{C} \left(\frac{|X_{ij}^{(\mathcal{P})}|}{\sqrt{P_i^{(\mathcal{P})} P_j^{(\mathcal{P})} + \epsilon^2}} \right). \quad (44)$$

Thus ‘‘EPR-accessible sector’’ may denote a scalar witness value, a thresholded pair support, or a capacity layer, depending on context. It never denotes an intrinsic detector-independent subspace of the vacuum unless an additional protocol-independence theorem is supplied.

Remark 5.1 (Relation to optimal LOCC distillation). *The optimal LOCC-distillable yield remains the conceptual target, but it is not the operational primitive used in the present falsifier tests. All matroid-sector partitions and network-capacity diagnostics in this paper use the computable logarithmic-negativity witness unless explicitly stated otherwise. For the first executable detector tests, where each detector is a qubit and each pair is a 2×2 system, PPT is equivalent to separability; hence nonzero logarithmic negativity is a faithful entanglement witness for the pair. For higher-dimensional detectors or grouped multiparty bipartitions, logarithmic negativity is not complete: it is an upper bound on distillable entanglement and can be blind to PPT-bound entanglement. In those settings the witness can underpopulate the EPR-active support. This limitation is part of the protocol definition and must be tracked when interpreting the split.*

Remark 5.2 (Thesis-ladenness of the witness). *The witness is deliberately bipartite. It therefore encodes the CTVM choice to use bipartite EPR accessibility as the operational proxy for the split-sector response. A passing matroid split test is consequently weaker evidence than a theorem of vacuum-sector selection; it says that the data are compatible with the chosen bipartite proxy. A failing test, by contrast, is stronger: a stable cross-sector circuit is an ordinary linear-algebraic obstruction to the proposed exact sector assignment.*

Remark 5.3 (Filtering versus transformation). *After the detectors have harvested a state ρ_D , subsequent EPR distillation is LOCC and cannot increase the chosen entanglement monotone on average. The field-coupling stage is different: it consumes the vacuum as an entangled resource. Hence the correct statement is not ‘‘harvesting is LOCC.’’ The correct statement is: post-harvesting distillation filters pairwise extractable correlations from a detector state generated by local couplings to a pre-correlated vacuum.*

6 Statistics-resolved detector response and Volovik analogy

6.1 Bosonic and fermionic response kernels

For scalar-like bosonic channels, a leading-order detector response has the schematic form

$$X_{ij}^{(B)} = \lambda_i \lambda_j \int d\tau d\tau' \chi_i(\tau) \chi_j(\tau') e^{-i(\Omega_i \tau - \Omega_j \tau')} W^{(B)}(x_i(\tau), x_j(\tau')). \quad (45)$$

For fermionic harvesting one must specify the detector coupling. A detector may couple to a fermion bilinear, to a non-Hermitian field with an internal detector degree of freedom, or to an effective composite response. We write the corresponding contribution abstractly as

$$X_{ij}^{(F)} = \lambda_i \lambda_j \int d\tau d\tau' \chi_i(\tau) \chi_j(\tau') e^{-i(\Omega_i \tau - \Omega_j \tau')} W_{\text{eff}}^{(F)}(x_i(\tau), x_j(\tau')). \quad (46)$$

The statistics-resolved matrix is

$$X_{\text{stat}} = \begin{pmatrix} X^{(B)} & X^{(BF)} \\ X^{(FB)} & X^{(F)} \end{pmatrix}. \quad (47)$$

The mixed blocks are present only when the chosen detector protocol couples the bosonic and fermionic response channels.

In a spectral representation the kernels may contain occupation factors

$$\mathcal{F}_B(\omega) = 1 + n_B(\omega), \quad \mathcal{F}_F(\omega) = 1 - n_F(\omega), \quad (48)$$

representing Bose stimulation and Pauli blocking. These factors affect response amplitudes and distillation yield. They do not by themselves prove gravitational coupling, superselection, or dark-sector abundances.

6.2 Volovik dictionary

Volovik's superfluid-vacuum program is used here as a disciplined analogue. The physical vacuum is not claimed to be literal helium. The point is that helium systems and related quantum liquids exhibit universality classes in which low-energy relativistic quasiparticles, effective metrics, topology, gauge-like fields, and protected boundary modes arise from an underlying many-body medium.

Analogue system	Statistical/topological mechanism	CTVM role
^4He -like bosonic superfluid or BEC	Bose coherence, hydrodynamic collective modes, phonon/acoustic metric behavior	Motivates a bosonic collective channel in which pairwise detector response may be amplified and metric-like kinematics may emerge.
^3He -A-like fermionic topological phase	Weyl points, chiral quasiparticles, emergent tetrads/gauge fields, Pauli exclusion	Motivates a fermionic topological response channel with suppressed bulk final-state access and protected spectral structure.
^3He -B-like gapped topological phase	Gapped bulk with protected boundary Majorana modes	Motivates a boundary/defect residual sector and topological protection, but does not derive CTVM edge charge.
Self-sustained quantum vacuum	Equilibrium thermodynamic identity for vacuum pressure	Supplies the structural lesson that microscopic vacuum energy need not equal gravitating response; CTVM does not inherit the equilibrium zero-pressure conclusion.
Analogy used for		Not claimed
Universality-class motivation for Bose enhancement and Pauli blocking in detector response.		The physical vacuum is not claimed to be literal $^3\text{He}/^4\text{He}$ or a two-component helium mixture.
Protected boundary/defect modes in topological media.		Volovik invariants are not imported as CTVM invariants without an independent edge or response theory.
Self-sustained-vacuum thermodynamic lessons.		No Volovik equilibrium cancellation theorem is assumed to solve the cosmological constant problem.
Analogue-vacuum experimental design.		Analogue success would not by itself prove gravitational inertness or dark-sector abundances.

The Volovik analogy is used only for mechanism building and analogue tests, not as a derivation of CTVM.

Hypothesis 6.1 (Statistics-resolved sector bias). *For detector protocols that separately probe bosonic collective and fermionic topological channels, the bosonic block generically has enhanced weighted pairwise response in occupied collective regimes, while the fermionic block generically has suppressed bulk pairwise final-state access and enhanced sensitivity to boundary or defect-supported modes.*

This is a hypothesis to be tested in analogue systems. It is not a theorem of AQFT.

7 Response-weighted matroids

7.1 Realified response representation

Given a complex detector response matrix $X \in M_N(\mathbb{C})$, we fix the column-doubling realification convention

$$\hat{X} = [\text{Re } X \mid \text{Im } X] \in M_{N \times 2N}(\mathbb{R}). \quad (49)$$

Thus the ground set of $M[\widehat{X}]$ consists of the $2N$ real response-feature columns obtained from the real and imaginary parts of the detector response. We do not use the row-stacking convention $\begin{pmatrix} \text{Re } X \\ \text{Im } X \end{pmatrix}$ in this manuscript; it generally defines a different matroid with a different ground set. The earlier conflation of a pseudograph ground set with a matrix-column ground set is not allowed.

Remark 7.1 (Operational justification of column-doubling). *The choice of*

$$\widehat{X} = [\text{Re } X \mid \text{Im } X]$$

rather than the row-stacking convention

$$[\text{Re } X; \text{Im } X]$$

is dictated by the operational role of the real and imaginary quadratures of the detector response.

In the canonical real-smearing convention, with no additional detector phase convention obscuring the decomposition, the Wightman function splits as

$$D^+(x, y) = \frac{1}{2}G^{(1)}(x, y) + \frac{i}{2}\Delta_{\text{PJ}}(x, y),$$

where $G^{(1)}$ is the symmetric Hadamard kernel and Δ_{PJ} is the Pauli–Jordan commutator function. In that convention, the real part of the smeared response inherits the symmetric/statistical kernel, while the imaginary part inherits the antisymmetric/retarded-response content.

For general detector protocols, switching functions, detector gaps, and phase conventions can mix these kernels at the level of individual matrix entries. Nevertheless, the two real quadratures $\text{Re } X$ and $\text{Im } X$ remain operationally distinct features of the complex response. The purpose of column-doubling is to make those two quadratures separate ground-set elements of the matroid. A linear dependence among real-response columns has a different diagnostic meaning from a linear dependence among imaginary-response columns, and cross-dependencies between the two are precisely the structural relations the causal-topology diagnostic is meant to expose.

The row-stacking convention treats real and imaginary parts as two coordinates of a single feature: its ground set has size N , not $2N$, and it cannot assign independent matroid elements to the symmetric-response and antisymmetric-response quadratures. Column doubling is therefore not merely cosmetic. It is the convention compatible with using the matroid layer as a causal-topology diagnostic.

Definition 7.1 (Response matroid). *The response matroid is the representable linear matroid*

$$M[X] = M[\widehat{X}] \tag{50}$$

whose ground set is the chosen column set of \widehat{X} and whose independent sets are linearly independent column subsets.

Definition 7.2 (Response-weighted matroid). *A response-weighted matroid is a pair*

$$\mathfrak{M}_X = (M[X], w), \tag{51}$$

where $M[X]$ records support, rank, circuits, and direct sums, while w records response magnitude.

Examples include edge-like weights

$$w_{ij} = \log \left(1 + \frac{|X_{ij}|}{\sqrt{P_i P_j + \epsilon^2}} \right) \quad (52)$$

or basis weights

$$w(B) = \log \left(\epsilon + \sqrt{\det(\widehat{X}_B^T \widehat{X}_B)} \right), \quad (53)$$

for a basis-sized column submatrix \widehat{X}_B . The object is called a valuated matroid only if the valuated-matroid axioms, including the tropical Plücker relations, are separately verified.

Remark 7.2 (Weighted matroid terminology). *The pair $(M[\widehat{X}], w)$ is called a response-weighted matroid in this paper. It should not be called a valuated matroid unless the basis weights satisfy the valuated-matroid exchange axiom, equivalently the relevant tropical Pluecker relations. This distinction matters because arbitrary response weights, including Bose-enhanced and Pauli-suppressed amplitudes, need not define a valuated matroid.*

Proposition 7.1 (Statistics usually changes weights, not matroid rank). *Suppose Bose or Pauli factors rescale nonzero columns of \widehat{X} by nonzero scalars. Then $M[\widehat{X}]$ is unchanged. The matroid changes only when the statistical dressing produces exact zero columns, exact block decomposition, or exact rank-changing degeneracies.*

Proof. Representable matroids are invariant under nonzero rescaling of columns. The dependence relations change only when a column is sent to zero, when a block decomposition removes all cross-block support, or when scalings enforce an exact linear relation that was absent before. \square

7.2 Loops, coloops, and what they mean

In a linear matroid an element is a loop if and only if its representing column is the zero vector. Therefore a detector site or feature is a loop only if the chosen representation gives a zero column. It is not a loop merely because it is “diagonal noise,” and it is not non-loop merely because it labels an off-diagonal pair.

A coloop is an element contained in every basis. Nuclearity does not guarantee coloops. A coloop statement requires a structural theorem showing that a boundary feature is indispensable in every maximal independent set.

Proposition 7.2 (Exact direct sum). *If the realified response representation has block form*

$$\widehat{X} \sim \begin{pmatrix} \widehat{X}_{\text{split}} & 0 \\ 0 & \widehat{X}_{\text{res}} \end{pmatrix}, \quad (54)$$

then

$$M[\widehat{X}] \cong M[\widehat{X}_{\text{split}}] \oplus M[\widehat{X}_{\text{res}}]. \quad (55)$$

No circuit crosses the two blocks.

Proof. For a block-diagonal matrix, a column set is independent exactly when its projections onto each block are independent. Minimal dependencies therefore live entirely in one block or the other, which is the standard direct-sum criterion. \square

Hypothesis 7.1 (Detector/split compatibility). *The detector feature map used to build \widehat{X} is compatible with the split-relative decomposition $\text{im } \Pi_{\text{split}} \oplus \ker \Pi_{\text{split}}$ in the sense that a regulator-dependent change of basis produces the block form above.*

Microcausality constrains commutators. It does not itself imply detector/split block diagonalization. Pauli blocking can suppress mixing weights. It does not prove exact matroid superselection unless it enforces exact zeros or an exact selection rule.

Example 7.1 (Massive scalar support for approximate compatibility). *A simple case in which the compatibility hypothesis holds approximately is a free massive scalar in $1+1$ dimensions with two detector clusters separated by a spacelike distance d much larger than the Compton scale. At equal times the massive Wightman/Hadamard kernel has the large-distance form*

$$G^+(d) \propto K_0(md) \sim \sqrt{\frac{\pi}{2md}} e^{-md}, \quad md \gg 1. \quad (56)$$

For smeared detectors of support radius $R \ll d$, cross-cluster matrix entries are suppressed as $O(e^{-m(d-2R)})$ up to polynomial factors. The representing matrix is therefore block diagonal up to exponentially small off-block terms. This does not prove exact superselection, but it gives a concrete model in which detector/split compatibility is realized to controlled exponential accuracy.

A second toy calculation in Appendix F shows that the matroid layer can carry information beyond rank: for a three-detector Hermitian response matrix, a non-uniform realified matroid with a parallel pair appears on a codimension-one protocol stratum. The example is deliberately modest. It demonstrates non-vacuity of the matroid bookkeeping, not a vacuum-sector theorem.

7.3 Regulator dependence and continuum-limit bookkeeping

Every object in this section is protocol dependent. A detector protocol \mathcal{P} includes smearing radius σ , switching time T , detector gaps Ω_i , ultraviolet cutoff or Hadamard subtraction prescription Λ , and a choice of feature map used to build \widehat{X} . The exact matroid $M[\widehat{X}]$ can change when a minor crosses zero; the weight layer w changes continuously under generic regulator changes.

For a capacity assignment used later in the gravity dictionary, a minimal stability requirement is

$$c_{ij}^{(\mathcal{P})} = \alpha_{\mathcal{P}} \mathcal{C} \left(\frac{|X_{ij}^{(\mathcal{P})}|}{\sqrt{P_i^{(\mathcal{P})} P_j^{(\mathcal{P})} + \epsilon^2}} \right), \quad \lim_{\sigma \rightarrow 0, T \rightarrow \infty, \Lambda \rightarrow \infty}^{\text{ren}} c_{ij}^{(\mathcal{P})} = c_{ij}^{\text{ren}}, \quad (57)$$

where the limit is meant only after the chosen renormalisation/subtraction scheme is fixed. Without such regulator stability, the network layer of Section 9 is not physically meaningful.

The massive-scalar example above gives one template: the off-block norm scales as

$$\|X_{\text{off}}\| = O(e^{-m(d-2R)}) \quad (58)$$

for fixed smearing support radius R and separation $d \gg R, m^{-1}$. This gives controlled approximate decoupling, not exact superselection. A submission-ready completion of the matroid program should repeat this kind of estimate for the vector, fermionic, and analogue-vacuum protocols used in the physical sections.

7.4 Matroids as validators and falsifiers of the split hypothesis

The matroid layer has an epistemic role: it gives a discrete test of whether a chosen detector response is compatible with the conjectured split between a split-factor, EPR-accessible sector and a residual sector. This role must be stated carefully. The matroid is not a protocol-independent invariant of the vacuum. It is a discrete invariant associated with a specified regulated detector protocol. For that protocol, however, it is stable under continuous changes of detector positions, switching functions, and couplings until the response data cross a rank-degeneracy wall where a minor vanishes or ceases to vanish.

7.4.1 Continuous response data and discrete rank strata

The response matrix X is built from smeared Wightman or response kernels and therefore varies continuously with the detector protocol \mathcal{P} . The realified matroid $M[\widehat{X}]$, whose ground set is the column set of the column-doubled matrix, changes only when exact linear dependencies appear or disappear. Equivalently, the protocol space is stratified by the vanishing pattern of minors of \widehat{X} .

For a matrix of rank r on $m = 2N$ realified columns, the generic stratum is often the uniform matroid $U_{r,m}$, meaning that every subset of at most r columns is independent. A non-uniform matroid, such as one containing a parallel-pair circuit, records an exact algebraic relation among response features. Such a relation is not automatically a vacuum-sector invariant, but it is precisely the sort of discrete structure that a sector-selection hypothesis must be able to explain or rule out.

7.4.2 Operational column partition

The split falsifier in §7.4.3 is conditional on an assignment of columns of \widehat{X} to

$$E_{\text{split}} \sqcup E_{\text{res}}.$$

We now make one such assignment explicit by linking it to the operational EPR sector of §5. This assignment is external to the matroid rank computation: it is obtained from the harvested detector state and an entanglement monotone, not from the circuit structure itself.

Definition 7.3 (Operational column partition). *Fix a detector protocol \mathcal{P} , a bipartite entanglement monotone E , and a threshold $\epsilon_E \geq 0$. Define the thresholded pair support*

$$\mathcal{E}_{\text{pair}}^{\mathcal{P},\epsilon} = \left\{ (i, j) : E\left(\rho_{ij}^{(\mathcal{P})}\right) > \epsilon_E \right\},$$

where

$$\rho_{ij}^{(\mathcal{P})} = \text{Tr}_{\bar{i}\bar{j}} \rho_D^{(\mathcal{P})}$$

is the two-detector reduced state on detectors i and j .

A detector site $j \in \{1, \dots, N\}$ is called EPR-active at tolerance ϵ_E if there exists $i \neq j$ such that

$$(i, j) \in \mathcal{E}_{\text{pair}}^{\mathcal{P},\epsilon}.$$

Let

$$S_{\text{act}} \subseteq \{1, \dots, N\}$$

be the set of EPR-active detector sites. Under the column-doubling convention

$$\widehat{X} = [\text{Re } X \mid \text{Im } X],$$

the operational column partition of the ground set

$$E = \{1, \dots, 2N\}$$

is

$$E_{\text{split}}^{(\mathcal{P}, \epsilon)} = S_{\text{act}} \cup \{j + N : j \in S_{\text{act}}\},$$

$$E_{\text{res}}^{(\mathcal{P}, \epsilon)} = E \setminus E_{\text{split}}^{(\mathcal{P}, \epsilon)}.$$

Hypothesis 7.2 (Operational/matroid compatibility). *For the detector protocol \mathcal{P} , the operational column partition of Definition 7.3 is matroid-compatible: no circuit of $M[\widehat{X}]$ has elements in both*

$$E_{\text{split}}^{(\mathcal{P}, \epsilon)} \quad \text{and} \quad E_{\text{res}}^{(\mathcal{P}, \epsilon)}.$$

The falsifier of §7.4.3 now has a definite target. If the operational column partition is nontrivial, meaning both $E_{\text{split}}^{(\mathcal{P}, \epsilon)}$ and $E_{\text{res}}^{(\mathcal{P}, \epsilon)}$ are nonempty, then the proposed exact sector assignment fails for that protocol if a stable circuit of $M[\widehat{X}]$ crosses the two sectors.

If no choice of ϵ_E in a physically reasonable threshold window produces a nontrivial partition, then the protocol is non-diagnostic for this split test. If this happens for an entire physically relevant protocol class, then the matroid layer carries no sector-separation information beyond rank for that class.

Remark 7.3 (Coarseness of the operational bridge). *Definition 7.3 is deliberately coarse. The operational EPR yield of §5 is a scalar or thresholded support condition built from the full two-detector state $\rho_{ij}^{(\mathcal{P})}$. It does not, by itself, distinguish the symmetric-correlation part of the harvested response from the retarded-response part.*

The bridge therefore assigns the real and imaginary columns associated with a given detector site to the same sector, treating the EPR activity of detector j as a single operational bit. A finer partition that independently assigns

$$\text{Re } X_{.j} \quad \text{and} \quad \text{Im } X_{.j}$$

would require an operational distinguishability test sensitive to the symmetric-correlation versus retarded-response content of the harvested state. The present bipartite distillation functional does not supply that refinement.

Constructing such a refined test would let the matroid layer probe causal-topology structure at the level of the Hadamard/Pauli–Jordan or response/noise decomposition, rather than only at the level of detector sites. This is added as an open derivational target in §17.

7.4.3 Sector separation as a direct-sum condition

The detector/split compatibility hypothesis says that, after a regulator-dependent change of basis, the realified response representation can be brought to block form,

$$\widehat{X} \sim \begin{pmatrix} \widehat{X}_{\text{split}} & 0 \\ 0 & \widehat{X}_{\text{res}} \end{pmatrix}, \quad (59)$$

where $\widehat{X}_{\text{split}}$ is supported on the split-factor image and \widehat{X}_{res} is supported on the residual kernel. In matroid language this implies

$$M[\widehat{X}] \cong M[\widehat{X}_{\text{split}}] \oplus M[\widehat{X}_{\text{res}}], \quad (60)$$

and therefore no circuit crosses the proposed sector boundary.

This gives a necessary condition for exact sector separation. If an independent split calculation assigns columns to $E_{\text{split}} \sqcup E_{\text{res}}$, and the measured or computed response matroid contains a circuit with elements in both parts, then that exact sector assignment is falsified for the protocol under study. Conversely, if the matroid decomposes as a direct sum with circuits confined to one block, then the response data are compatible with the split. Compatibility is not yet identification: one still needs an independent argument that the blocks correspond to $\text{im } \Pi_{\text{split}}$ and $\text{ker } \Pi_{\text{split}}$.

7.4.4 Algorithmic split test

Given a computed realified response matrix \widehat{X} and a proposed sector partition $E = E_{\text{split}} \sqcup E_{\text{res}}$, the exact split can be tested directly.

1. Compute the represented matroid $M[\widehat{X}]$ by rank tests on column subsets, or equivalently by the vanishing pattern of minors up to the relevant rank.
2. Enumerate circuits up to the maximum size accessible in the finite detector model. For small examples this can be done exhaustively; for larger examples one may use incremental rank-oracle algorithms.
3. Build the circuit graph: its vertices are the columns of \widehat{X} , and two vertices are connected if they occur in a common circuit.
4. The connected components of this circuit graph are the connected components of the represented matroid.
5. Exact detector/split compatibility requires E_{split} and E_{res} to be unions of these components. Equivalently, no circuit may intersect both proposed sectors.

Thus a single cross-sector circuit falsifies the exact block-diagonal sector assignment for that protocol. If all off-block minors are merely small rather than zero, the appropriate conclusion is approximate decoupling, not exact superselection.

7.4.5 What the three-detector example does and does not validate

The three-detector example in Appendix F should not be oversold. The codimension-one condition $\text{Im } z_{23} = 0$ produces a parallel pair $\{5, 6\}$ in the column-doubled realification. This shows that the matroid layer is non-vacuous: non-uniform matroid strata occur in the space of Hermitian response matrices and survive nonzero rescalings of the parallel columns.

It does not, by itself, prove a direct-sum sector split. A parallel pair is not automatically a separate matroid component. In general a parallel extension of a rank-three uniform matroid can still have circuits involving the parallel pair and other columns. Thus the example is evidence that matroidal structure can carry information beyond rank; it is not evidence that the split-factor image and residual kernel have been identified. To turn such an example into sector evidence, one would have to show three further facts: first, that a physically realizable detector protocol lies on the relevant stratum; second, that a split-relative conditional expectation assigns the parallel-pair

columns to the residual block or to the split block; and third, that no circuit crosses the independently assigned sector boundary.

7.4.6 Matroid falsifiers

The matroid layer gives sharp falsifiers of the detector/split compatibility hypothesis.

1. If a proposed split assignment produces cross-sector circuits, then the exact direct-sum version of detector/split compatibility fails for that protocol.
2. If every physically realizable protocol in a proposed experimental class produces only generic uniform matroids, then the matroid layer carries no sector information beyond rank for that class.
3. If non-uniform matroids occur but their circuits do not correlate with independently identified split or residual structures, then the matroid remains useful bookkeeping but does not validate CTVM sector selection.
4. If non-uniform matroids occur reproducibly in topological regimes, for example near Weyl points or boundary-mode regimes in a Volovik-style analogue system, and their circuits are confined to independently identified residual blocks, then the matroid supplies operational evidence for sector separation.

Minimal crossing-circuit falsifier. Consider the realified response submatrix

$$\widehat{X}_{\text{toy}} = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 0 \end{pmatrix}. \quad (61)$$

The column matroid has rank 2 and the unique minimal circuit $\{1, 2, 3\}$, since $c_3 = c_1 + c_2$. If an independent operational split assigns $E_{\text{split}} = \{1, 2\}$ and $E_{\text{res}} = \{3\}$, then the circuit crosses the proposed split and the exact direct-sum assignment is falsified. The operational test is concrete: compute the three response columns for the chosen detector protocol, verify $\text{rk}\{1, 2, 3\} = 2$ while every two-column subset is independent, and check whether the EPR-active column partition places column 3 in the opposite sector from columns 1, 2.

Step	Experimental or numerical split test
1	Build X from the specified detector protocol and realify to $\widehat{X} = [\text{Re } X \mid \text{Im } X]$.
2	Use the operational EPR partition of Definition 7.3 to propose $E_{\text{split}} \sqcup E_{\text{res}}$.
3	Compute ranks of candidate subsets using a stable tolerance window or exact arithmetic when available.
4	Enumerate circuits and mark any circuit intersecting both proposed sectors.
5	Report: pass, fail by cross-circuit, or non-diagnostic if no nontrivial operational partition exists.

7.4.7 Non-EPR is not the same claim as non-gravitating

The matroid can test whether response features separate into independent linear-dependence components. It cannot prove that one component is gravitationally inert. The assignment of a component to “non-EPR” requires the operational distillation and split-compatibility analysis of Sections 5 and 7. The assignment of that same component to “non-gravitating” requires the gravity dictionary of Section 9. Matroid decomposability is therefore a necessary consistency condition for the exact split, not a proof of gravitational decoupling.

Remark 7.4 (Epistemic role of the matroid). *The response matroid validates the CTVM split only in the limited sense of testing its linear-algebraic compatibility with detector data. Exact sector separation requires a direct-sum decomposition after an independently specified sector assignment. Indecomposable matroids, or matroids with cross-sector circuits, falsify that exact assignment. Decomposable matroids with circuits confined to one block support the assignment only when an independent split calculation identifies that block with $\text{im } \Pi_{\text{split}}$ or $\text{ker } \Pi_{\text{split}}$. The matroid supplies the algebraic scaffold on which the split hypothesis can be tested; it does not decide the gravitational interpretation by itself.*

8 Corrected set-theoretic vacuum structure

The v7 statement $C_N \subset E_N \subset F_N^{\text{pairwise}}$ was too strong as written. Mutual causal contact does not automatically imply harvestable entanglement above threshold, and if the regions are assumed mutually spacelike then the original definition of C_N may be empty. The corrected structure is threshold-dependent.

Fix detector regions O_i and a detector protocol \mathcal{P} .

Definition 8.1 (Pairwise Wightman support). $F_N^{\text{pair}}(\mathcal{P}, \epsilon_W)$ is the set of N -tuples for which the smeared pairwise Wightman kernels relevant to \mathcal{P} exceed the chosen support tolerance ϵ_W for the required pairs. The limiting case $\epsilon_W = 0$ is the exact-support idealization.

Definition 8.2 (Retarded causal support). $C_N(\mathcal{P}, \epsilon_R)$ is the set of N -tuples for which the retarded/commutator graph of the protocol has response connectivity above tolerance ϵ_R in the channel required by the causal-topology diagnostic.

Definition 8.3 (Harvestable entanglement support). $E_N(\mathcal{P}, \epsilon_E)$ is the set of N -tuples for which the post-interaction detector state has specified entanglement yield above threshold ϵ_E , either bipartite yield $\mathcal{E}_2^{\mathcal{P}} > \epsilon_E$ or a specified multipartite monotone above threshold.

Proposition 8.1 (Corrected inclusions). *For leading-order pairwise harvesting protocols, one has the thresholded implication*

$$E_N(\mathcal{P}, \epsilon_E) \subseteq F_N^{\text{pair}}(\mathcal{P}, \epsilon_W) \tag{62}$$

only when the chosen entanglement witness requires pairwise kernels above the corresponding tolerance. In general neither $C_N(\mathcal{P}, \epsilon_R) \subseteq E_N(\mathcal{P}, \epsilon_E)$ nor $E_N(\mathcal{P}, \epsilon_E) \subseteq C_N(\mathcal{P}, \epsilon_R)$ follows without additional threshold and geometry assumptions.

Proof. At leading order the detector density matrix entries used by the witness are functions of smeared two-point kernels. If those kernels vanish in all required channels, the witness cannot be positive. Conversely, nonzero kernels may be too weak relative to local noise to produce entanglement, so membership in F_N^{pair} does not imply membership in E_N . Retarded causal support

and entanglement support are distinct diagnostics, so neither contains the other without extra hypotheses. \square

Corollary 8.1 (Noncausal harvestable sector). *The operationally interesting residual region is*

$$E_N(\mathcal{P}, \epsilon_E) \setminus C_N(\mathcal{P}, \epsilon_R), \quad (63)$$

namely detector-harvestable entanglement without retarded causal support in the chosen protocol. This is the clean version of the CTVM “non-EPR” or topological-residual sector.

For a free Gaussian scalar field, connected N -point functions vanish for $N > 2$ in the usual cumulant sense, while nonzero detector entanglement can still arise from pairwise Wightman kernels. In interacting fields connected higher-point functions can contribute, and the hierarchy must be recomputed.

9 Conditional Gravity Dictionary: From Response Data to Holographic Entropy and Sourcing

The response matroid does not prove gravity. The representable linear matroid $M[\widehat{X}]$ records rank, circuits, and direct-sum structure of the regulated detector-response representation. It contains neither edge capacities nor a bulk geometry nor an extremal-surface prescription nor a stress-tensor sourcing rule. Therefore the chain

$$\lambda_M = \text{min-cut} = \text{bit-thread flux} = S_{\text{RT}} \quad (64)$$

is not valid for $M[\widehat{X}]$ by itself. The gravitational step requires additional structure.

9.1 Network assignment

Let $\widehat{X}_{\text{split}}$ denote the detector-response block associated with the split-factor sector, and let w_{split} denote the corresponding response-weight layer. We introduce a protocol-dependent network assignment

$$\mathcal{N}_{\mathcal{P}} : (\widehat{X}_{\text{split}}, w_{\text{split}}) \mapsto G_{\text{split}} = (V, E, c), \quad (65)$$

where \mathcal{P} encodes the regulator, detector smearing, switching functions, and post-processing convention.

Hypothesis 9.1 (Pairwise response-network dictionary). *There exists a network assignment $\mathcal{N}_{\mathcal{P}}$ such that:*

1. *vertices V represent the regulated detector cells or boundary subregions in the split-factor sector;*
2. *an edge $e_{ij} \in E$ is present only when the split-sector two-detector response passes the chosen operational pairwise-response criterion;*
3. *each edge carries a nonnegative capacity*

$$c_{ij} = \alpha_{\mathcal{P}} \mathcal{C} \left(\frac{|\widehat{X}_{ij}^{\text{split}}|}{\sqrt{P_i P_j}} \right), \quad (66)$$

or, more operationally,

$$c_{ij} = \alpha_{\mathcal{P}} E_{\mathcal{N}}(\rho_{ij}^{\text{split}}), \quad (67)$$

where $E_{\mathcal{N}}$ is the chosen bipartite entanglement monotone;

4. the capacity assignment is regulator-stable in the continuum or large- N limit;
5. residual-sector coordinates \widehat{X}_{res} are not used in defining G_{split} .

The network G_{split} is not the matroid $M[\widehat{X}]$. It is a weighted pairwise object constructed from the same regulated response data. The matroid records dependence structure; the network records capacities.

9.2 Entropy dictionary

Given G_{split} , define the network cut functional

$$S_{\text{net}}(A) = \min_{\gamma_A} \sum_{e \in \gamma_A} c(e), \quad (68)$$

where γ_A ranges over cuts separating A from its complement. This is a network entropy candidate, not yet gravitational entropy.

Hypothesis 9.2 (Controlled AdS-like holographic entropy dictionary). *The pair $(G_{\text{split}}, S_{\text{net}})$ lies in a regime admitting a controlled semiclassical holographic interpretation, such as an AdS/RT/HRT-like setting with a well-defined boundary Hilbert-space interpretation. In this regime there exists a continuum or large- N limit in which*

$$S_{\text{net}}(A) = \max_v \int_A v \cdot n \, dA = \frac{\text{Area}(\gamma_A)}{4G_N}, \quad (69)$$

with v satisfying the Freedman-Headrick bit-thread constraints and γ_A the corresponding Ryu-Takayanagi or HRT surface. This hypothesis also requires that S_{net} is finite, regulator-stable, submodular, and compatible with the relevant holographic entropy inequalities in the approximation being used.

Hypothesis 9.3 (Cosmological-horizon holography extension). *For dark-energy applications, the controlled entropy dictionary must be extended to a de Sitter or cosmological-horizon setting. This is a separate conjecture. A viable completion must specify whether the relevant dictionary is static-patch, stretched-horizon, bilayer-screen, DSSYK/JT-de Sitter, crossed-product/type-II, or some other de Sitter-inspired construction, and must state what replaces the usual AdS boundary Hilbert-space interpretation. Relevant but still conjectural guideposts include dS/CFT, static-patch/stretched-horizon proposals, DSSYK/JT-de Sitter toy models, crossed-product/type-II observer algebras, and bilayer static-patch entanglement prescriptions [34, 35, 36, 37, 38, 44, 45, 46, 47].*

Under Hypothesis 9.2 we may rename S_{net} as S_{hol} . For cosmological applications Hypothesis 9.3 is additionally required. Without these hypotheses, S_{net} remains only a weighted network cut functional.

Hypothesis 9.4 (Entropy-cone admissibility). *The network entropy vector*

$$A \longmapsto S_{\text{net}}(A) \quad (70)$$

satisfies the entropy inequalities required of the intended holographic regime. At minimum this includes positivity and strong subadditivity. In static classical AdS/RT-like regimes it also includes monogamy of mutual information (MMI), meaning the holographic inequality $I_3(A : B : C) \leq 0$, and the appropriate holographic entropy-cone inequalities. If S_{net} violates these constraints, the pairwise response network cannot be interpreted as an RT/HRT entropy in that regime.

9.3 Sourcing dictionary

Even a holographic entropy is not yet a gravitational field equation. To obtain a source term, one must assume a semiclassical response principle relating variations of the entropy functional to variations of modular energy or stress-energy.

Hypothesis 9.5 (Semiclassical sourcing dictionary). *For the class of states and regions under consideration, variations of the split-sector holographic entropy satisfy an entanglement-first-law, relative-entropy, or semiclassical response identity of the schematic form*

$$\delta S_{\text{hol}}(A) = \delta \langle K_A^{\text{split}} \rangle \implies \delta G_{\mu\nu} = 8\pi G \delta \langle T_{\mu\nu} \rangle_{\text{split}}, \quad (71)$$

or an equivalent local gravitational response relation, with the stress-tensor response conserved:

$$\nabla^\mu \delta \langle T_{\mu\nu} \rangle_{\text{split}} = 0. \quad (72)$$

This is the step at which the CTVM entropy dictionary becomes a gravity dictionary. The step is not supplied by matroid theory and is not a theorem of Minkowski AQFT alone.

9.4 Residual-sector decoupling

The residual sector is absent from the ordinary pairwise network G_{split} by construction. This is a representation statement. Gravitational inertness requires an additional decoupling condition.

Hypothesis 9.6 (Residual decoupling within the chosen dictionary). *The gravitational response functional depends only on the split-sector capacities:*

$$S_{\text{source}} = S_{\text{source}}(c_{ij}^{\text{split}}). \quad (73)$$

Equivalently, for any smooth parametrisation of the regulated response data,

$$D_{\widehat{X}_{\text{res}}} S_{\text{source}} = 0, \quad D_{\widehat{X}_{\text{res}}} \langle T_{\mu\nu} \rangle_{\text{source}} = 0. \quad (74)$$

Lemma 9.1 (Dependency Lemma 9.1 – bookkeeping implication inside the selected gravity dictionary). *Assume the detector/split compatibility hypothesis, the pairwise response-network dictionary, the controlled holographic entropy dictionary including entropy-cone admissibility, the semiclassical sourcing dictionary, and residual decoupling (Hypothesis 9.6). For cosmological-horizon applications also assume Hypothesis 9.3.*

Then, within the chosen dictionary, the split-factor pairwise response sector supplies the gravitational source functional assigned by the model, while the residual response sector is inert with respect to that functional:

$$D_{\widehat{X}_{\text{res}}} S_{\text{source}} = 0, \quad D_{\widehat{X}_{\text{res}}} \langle T_{\mu\nu} \rangle_{\text{source}} = 0. \quad (75)$$

Proof. By the pairwise response-network dictionary, the network G_{split} and its capacities are functions only of split-sector data. By the holographic entropy dictionary together with entropy-cone admissibility, the cut functional is identified with a bit-thread/RT entropy in the stated regime. By the semiclassical sourcing dictionary, variations of that entropy define the gravitational response functional. By residual decoupling, this functional has no dependence on residual-sector coordinates. Hence its directional derivative with respect to residual response data vanishes within the chosen dictionary. \square

Remark 9.1 (No independent evidential force). *This lemma records the logical consequence of imposing residual decoupling together with the network, entropy, and sourcing dictionaries. It has **no independent evidential force** for gravitational inertness of the residual sector. If any of those dictionary gates fail, what remains is not the CTVM gravity claim but the standalone detector-response, SK-filter, and matroid-diagnostic apparatus.*

Remark 9.2 (Failure modes). *The conditional implication fails if the network assignment is not regulator-stable, if the cut function violates the required entropy inequalities, if the weighted network has no semiclassical holographic limit, if the residual sector contributes to the same stress-response functional after coarse-graining, or if the residual sector admits an independent hypergraph, tensor-network, or quantum-error-correcting representation with its own geometric dual. In those cases the CTVM sector split may remain useful as detector-response bookkeeping, but the strong gravitational-inertness interpretation would fail.*

Open-system route to source conservation. If the split sector exchanges response with a residual topological sector, the split-sector stress tensor need not be conserved by itself. The conservation law should instead hold for the combined regulated system. One possible completion is to formulate the split-sector dynamics as an open quantum system, with a leakage dissipator $\mathcal{D}_{\text{leak}}^{(\nu_{\text{res}})}$ generated by boundary operators coupling the split factor to residual topological modes. In that formulation the exchange current Q_ν appearing in the two-fluid phenomenology would be the stress-energy bookkeeping of the same leakage channel. This construction is left as an open target and is not assumed in the present dependency schema.

Equivalence-principle consistency of residual decoupling. The residual-decoupling hypothesis must be compatible with the entropic-inertia reading of local inertial response. If the residual block is absent from the gravitational source functional but contributes to the Rindler modular work functional, then the inertial and gravitational response coefficients would not agree. A strengthened consistency condition is therefore

$$D_{\widehat{X}_{\text{res}}} \delta W_{\text{inert}} = 0, \quad D_{\widehat{X}_{\text{res}}} \delta \langle T_{\mu\nu} \rangle_{\text{source}} = 0, \quad (76)$$

with both derivatives understood inside the same regulated split/SK/source dictionary. In the commutator language of Section 4.7, this requires the residual algebra to have vanishing source-relevant response to the local translation generator P^1 , up to controlled leakage corrections. This is not an additional proof of residual inertness; it is an equivalence-principle consistency condition that any future residual-sequestration mechanism must satisfy.

9.5 Dictionary ladder

The conditional chain can be summarized as Figure 1. The figure is a type-labelled dependency diagram, not a commutative diagram of same-category maps. Algebraic restriction, detector response extraction, realification, matroid encoding, network assignment, entropy interpretation, and semiclassical sourcing are different operations with different domains and codomains.

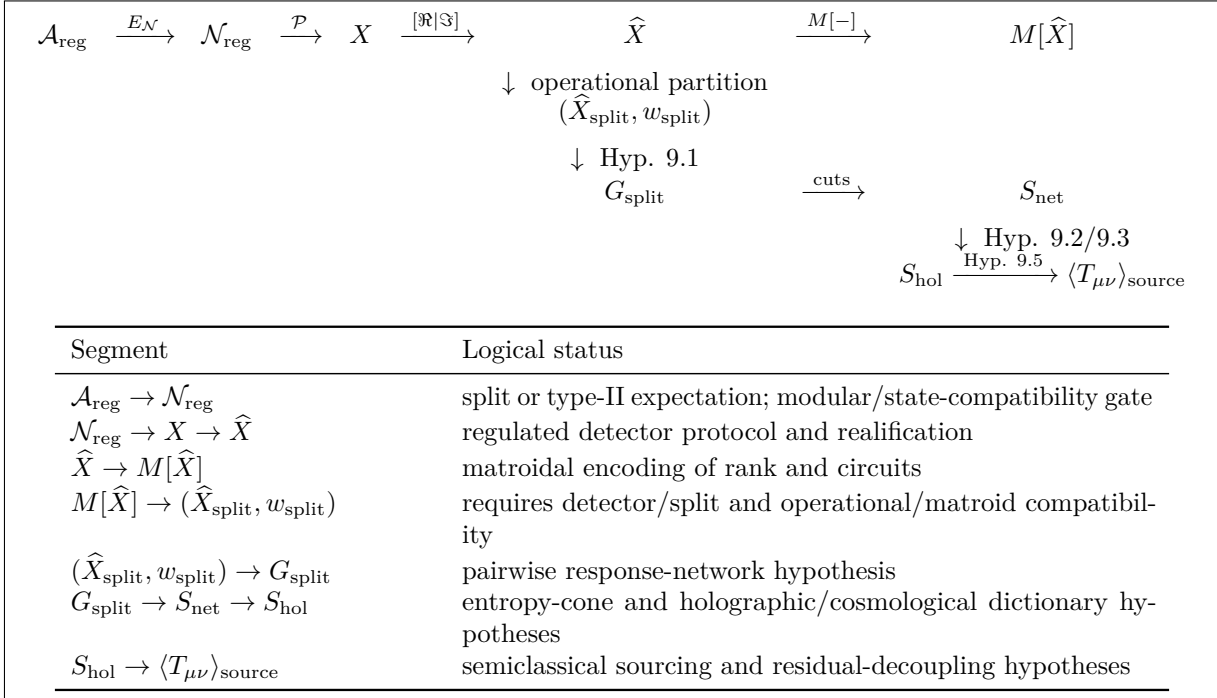


Figure 1: The CTVM typed dependency ladder. The arrows are not maps in a single category; they move between algebras, detector data, real matrices, matroids, weighted networks, entropy candidates, and source assignments. Only the detector/AQFT/matroid arrows are internal constructions. Network, holographic, cosmological-horizon, and sourcing arrows are additional hypotheses.

9.6 Residual matroid-to-CSS construction target and addressability trade-off

The preceding subsections define residual decoupling as a condition on the source functional. This subsection outlines a possible route to model the residual sector as a protected, weakly addressable code-like structure. **The construction is deliberately conditional and is presented as a target, not as a completed construction.** A realized residual response matrix does not automatically define a quantum error-correcting code, and a binary matroid alone does not automatically supply valid CSS stabilizers.

Remark 9.3 (CSS status warning). *A binary matroid does not by itself define a CSS code. A valid CSS stabilizer code requires two binary parity-check matrices H_X and H_Z satisfying the orthogonality condition*

$$H_X H_Z^T = 0. \quad (77)$$

Circuits of the binary residual matroid may suggest candidate stabilizer checks, but they are not stabilizers until the orthogonality condition has been verified and a consistent logical-operator structure

has been constructed. The material in this subsection is therefore a **construction target**, not a completed code.

Let X_{res} be the complex residual response block selected by the detector/split decomposition, and form the column-doubled realification

$$\widehat{X}_{\text{res}} = [\text{Re } X_{\text{res}} \mid \text{Im } X_{\text{res}}]. \quad (78)$$

Given a regulator-dependent threshold $\epsilon(\sigma, \Lambda)$, define the binary support matrix

$$(B_{\text{res}})_{ab} = \begin{cases} 1, & |(\widehat{X}_{\text{res}})_{ab}| > \epsilon(\sigma, \Lambda), \\ 0, & |(\widehat{X}_{\text{res}})_{ab}| \leq \epsilon(\sigma, \Lambda). \end{cases} \quad (79)$$

The associated binary matroid is the representable matroid

$$M_{\text{bin}} = M_{\mathbb{F}_2}[B_{\text{res}}], \quad (80)$$

whose ground set is the column set of B_{res} . This binary reduction is a lossy diagnostic: it discards amplitude information and preserves only thresholded support and parity-dependence structure. It is useful only inside a stable regulator window.

Definition 9.1 (CSS-compatible residual code data). *A residual CSS code datum consists of two binary parity-check matrices H_X and H_Z , extracted from the real and imaginary residual blocks or from selected circuit-incidence data of M_{bin} , such that the CSS orthogonality condition holds:*

$$H_X H_Z^T = 0 \quad \text{over } \mathbb{F}_2. \quad (81)$$

The corresponding stabilizer code is generated by the X-type checks from rows of H_X and the Z-type checks from rows of H_Z , as in the standard CSS construction [51, 52].

Remark 9.4 (Circuit data versus stabilizer data). *In a parity-check representation, circuits of the column matroid are minimal dependent column sets; equivalently, they are supports of low-weight codewords or undetected-error candidates. Stabilizer generators are rows of H_X and H_Z . Circuit-incidence vectors become check generators only after a choice of check set has been made and the CSS orthogonality condition and independence of the chosen checks have been verified. This distinction prevents the matroid-to-CSS target from silently becoming another decorative theorem.*

For a binary matrix H , let $d_{\text{circ}}(H)$ denote the girth of the represented column matroid, i.e. the minimum size of a dependent column set. If a valid CSS datum (H_X, H_Z) can be constructed from the circuit and cocircuit data of the residual matroid, then the minimum circuit weight supplies a candidate lower bound on the corresponding code distance:

$$d_{\text{CSS}} \gtrsim \min\{d_{\text{circ}}(H_X), d_{\text{circ}}(H_Z)\}, \quad (82)$$

up to the usual exclusion of stabilizer-supported codewords in the exact CSS distance. Without an explicit valid CSS datum, this remains a heuristic guide rather than a theorem. Suppressing low-weight circuits is therefore evidence for improved code-like protection only after the CSS datum has been constructed.

Hypothesis 9.7 (Residual circuit suppression). *For the residual response representation and the chosen binary reduction, the density of low-weight residual circuits decreases with the residual topological index ν_{res} . Equivalently, for some regulator-stable window and some constants $C, a > 0$,*

$$\Pr(\text{residual circuit of weight} \leq w) \leq Ce^{-a\nu_{\text{res}}} \quad (83)$$

for fixed low w , or the residual circuit girth grows with system size. This is a coding-theoretic strengthening of residual decoupling, not a consequence of matroid theory alone.

Under Hypothesis 9.7, the residual block may support a code-like sector whose logical operators require extended support. The intended interpretation is that such logical operators would be natural candidates for data in a DHR-like superselection structure, because they are not generated by local EPR-accessible operations in the split block. This identification is not established here. It would require an additional AQFT construction of the relevant representation category [54]; the CSS model supplies only a regulated finite-resolution proxy.

The same topological index that suppresses residual leakage in the open-system route may be used to parametrize code addressability. Let $\rho_{\text{cross}}(w_{\text{max}})$ denote the density of cross-sector circuits of weight at most w_{max} , and let r_{cross} be the total rank carried by those cross-sector circuit supports. A minimal addressability/error ansatz is

$$p_L \lesssim \alpha r_{\text{cross}} \rho_{\text{cross}}(w_{\text{max}}) \left(\frac{w_{\text{max}}}{N} \right)^\beta, \quad \beta \geq 1, \quad (84)$$

where α is set by the controlled interface coupling and N is the effective number of residual features. The formula expresses the operational trade-off: a residual code must have enough sparse cross-sector support to be addressable, but too much low-weight cross-sector support increases logical leakage.

The topological suppression factor used in the open-system route may be written schematically as

$$F_\nu^2 \sim (1 + \chi\nu_{\text{res}})^{-1}. \quad (85)$$

When the same invariant controls the leakage dissipator, the off-block response norm, and the cross-sector circuit density, one obtains the topological-uniformity picture:

$$\|\widehat{X}_{\text{off}}\|, \quad \rho_{\text{cross}}(w_{\text{max}}), \quad \|\mathcal{D}_{\text{leak}}^{(\nu_{\text{res}})}\| \quad \text{are suppressed by the same residual index.} \quad (86)$$

This is a falsifiable correlation, not a theorem. If analogue experiments or numerical response matrices show that these quantities scale differently with ν_{res} , then the topological-uniformity version of residual decoupling fails even if the weaker detector/matroid bookkeeping remains useful.

Construction 9.1 (Minimal leakage-model target: SSH boundary chain). *A minimal microscopic test of the residual-leakage hypothesis is a one-dimensional Su–Schrieffer–Heeger (SSH) boundary chain or an equivalent topological interface model. Let ν denote the number or index of protected boundary modes and let b_ν denote the boundary mode coupled to a split-sector interface operator Φ_{split} . A toy interaction is*

$$H_{\text{int}} = g_\nu \Phi_{\text{split}} \otimes (b_\nu + b_\nu^\dagger), \quad g_\nu = g_0 F_\nu, \quad (87)$$

with either algebraic or tunneling suppression,

$$F_\nu^2 = (1 + \chi\nu)^{-1} \quad \text{or} \quad F_\nu^2 = e^{-c\nu}. \quad (88)$$

The falsifiable test is to compute, as functions of ν ,

$$\|\widehat{X}_{\text{off}}\|, \quad \rho_{\text{cross}}(w_{\text{max}}), \quad p_L, \quad (89)$$

and check whether they share the same suppression law. Failure of this common scaling falsifies the topological-uniformity version of residual decoupling. This construction target is not used as an input elsewhere in the manuscript.

Remark 9.5 (Monogamy and residual coding). *Monogamy explains why maximal EPR pairs are concentrated and fragile, making them natural probes of pairwise geometry but poor carriers of robust multipartite memory. Residual multipartite entanglement can instead support high-distance code structure. Monogamy alone does not decide what gravitates; gravitational status is still assigned only by the response filter and the sourcing dictionary.*

10 Boundary modes, nuclearity, and area scaling

Nuclearity controls the phase-space size of localized states and is closely related to the split property. A typical nuclearity statement bounds a map

$$\Theta_{\beta,O}(A) = e^{-\beta H} A \Omega \quad (90)$$

by a nuclear norm whose logarithm grows no faster than a regulated boundary/volume expression. In favorable settings this supports an upper bound of the form

$$\text{rk } M_{\partial H} \leq C(\Lambda, \Delta, \mathcal{F}) \frac{A_H}{\ell^{d-2}}, \quad (91)$$

where ℓ is the cellulation scale, Λ represents energy cutoff data, Δ is a bounded-degree/cell-shape control, and \mathcal{F} is the detector feature map.

The reverse inequality is not supplied by nuclearity. A lower bound requires a positive density of independent boundary generators. Protected edge modes require more: a theorem that certain elements are present in every basis, or a topological invariant that prevents their removal.

Hypothesis 10.1 (Topological boundary invariant). *There exists a regulated boundary response operator or effective edge Hamiltonian K_{edge} whose low-energy sector carries a topological invariant*

$$\nu_{\partial H} \in \mathbb{Z}, \quad \mathbb{Z}_2, \quad \text{or an appropriate } K\text{-theory class}, \quad (92)$$

such that nonzero $\nu_{\partial H}$ protects a family of boundary response modes against all local perturbations preserving the relevant symmetry and gap or near-gap condition. A positive density of protected matroid generators follows only after this invariant has been identified and shown to act faithfully on the detector-response representation.

Proposition 10.1 (Conditional area-rank scaling). *Assume the nuclearity upper bound and the topological boundary-sector hypothesis. Then*

$$\text{rk } M_{\partial H} = \Theta \left(\frac{A_H}{\ell^{d-2}} \right). \quad (93)$$

If, in addition, the physical edge charge is proportional to boundary rank, $Q_{\text{edge}} = \alpha \text{rk } M_{\partial H}$, then $Q_{\text{edge}} \propto A_H$ at fixed cutoff.

Volovik’s ${}^3\text{He-B}$ analogue supports this hypothesis as a model-building guide: a gapped topological bulk can support protected boundary modes. It does not prove the CTVM edge theorem. The proof must be supplied inside the CTVM response representation.

11 Dark matter from vector-mediated causal harvesting

11.1 Channel-counting benchmark

The dark-matter mechanism posits that Standard Model gauge bosons are the connection-mediating channels that convert delocalized vacuum correlations into causally topologized response. Because the channel count has unresolved epoch and Goldstone bookkeeping issues, the conservative white-paper benchmark is a range, $f_{\text{eq}} \in [0.20, 0.30]$. The historical reference fraction is defined by

$$f_{\text{hist}} = \frac{g_{\text{conn}}}{g_*}. \quad (94)$$

At temperatures near electroweak symmetry breaking the Standard Model relativistic count is $g_* = 106.75$. The proposed connection-channel count after electroweak symmetry breaking is

$$g_{\text{conn}} = 8 \times 2 + 1 \times 2 + 3 \times 3 = 16 + 2 + 9 = 27, \quad (95)$$

where the terms correspond respectively to gluons, the photon, and the three massive weak vector bosons with longitudinal polarizations. This gives the historical reference point

$$f_{\text{hist}} = \frac{27}{106.75} \simeq 0.253. \quad (96)$$

Counting scheme	Vector numerator	Comment
Above EWSB, unbroken gauge basis	$12 \times 2 = 24$	Consistent with the usual high-temperature $g_* = 106.75$.
Below EWSB, broken vector polarizations	$8 \times 2 + 1 \times 2 + 3 \times 3 = 27$	Includes W^\pm and Z longitudinal modes.
Mixed benchmark	$27/106.75 \simeq 0.253$	Combines a below-EWSB numerator with an above-EWSB denominator; kept as a response-integral target, not a derivation.

The specific value $27/106.75$ is therefore best read as a historical reference point inside the broader response target range

$$f_{\text{eq}} \in [0.20, 0.30]. \quad (97)$$

The quantity that can make this range physical is not the integer ratio by itself but the response integral $f_{\text{vec}}^{\text{eff}}(T_{\text{EW}}; \mathcal{P})$ defined below.

Remark 11.1 (Status of the number). *The agreement with the observed dark-matter fraction is suggestive but the count is not a first-principles degeneracy derivation. It is a model-dependent identification of connection-mediating vector channels, subject to two distinct caveats. First, the numerator uses a below-electroweak-symmetry-breaking polarization count ($8 \times 2 + 1 \times 2 + 3 \times 3 = 27$, including the longitudinal modes of W^\pm and Z), while the denominator $g_* = 106.75$ is the usual*

above-electroweak thermodynamic count. Resolving this epoch mismatch is part of the field-theoretic response calculation set out in Appendix H. A further bookkeeping issue is that the broken-phase longitudinal W/Z modes are the eaten Goldstone modes of the Higgs sector, while the high-temperature denominator already counts the unbroken electroweak scalar content; the response calculation must decide how this transition is counted without double-counting. Until that calculation flows into the target range in a regulator-stable, gauge-consistent protocol window, the ratio $27/106.75$ stands only as a historical reference point, not a prediction. Second, numerical proximity alone carries limited evidential weight. As Appendix E records, a conservative look-elsewhere estimate over integer numerators of the form n/g_* already admits eleven integers in the broad window $[0.20, 0.30]$ at $g_* = 106.75$, and allowing weighted ratios, fermionic $7/8$ factors, or alternative denominators at different epochs broadens the comparison set further. The model should therefore stand or fall on the full package – vector-channel response dynamics, freeze-out behavior, growth history, and falsifiable cosmological deviations – not on the numerical coincidence alone.

For orientation, in Planck 2018 base- Λ CDM one has $\Omega_c h^2 \simeq 0.120$ and $h \simeq 0.674$, giving $\Omega_c \simeq 0.264$ [19]. The benchmark $27/106.75 \simeq 0.253$ is therefore roughly 4% low relative to that reference value. This comparison is dataset- and model-dependent and is included only to keep the numerical target honest; it does not constitute evidence for the model.

Dark-matter benchmark logic.

- (1) The bare channel count gives only the historical reference point $f_{\text{hist}} = 27/106.75$ inside the target range $f_{\text{eq}} \in [0.20, 0.30]$.
- (2) The response integral $f_{\text{vec}}^{\text{eff}}(T; \mathcal{P})$ is the dynamical quantity that must enter that target range in a regulator-stable and gauge-consistent calculation.
- (3) The Boltzmann–CTVM locking equation is a conditional mechanism that would make the target an attractor if the response calculation and locking rate support it.

Appendix H lists the numerical checks that would falsify this benchmark: regulator instability, gauge or longitudinal-mode contamination, or failure of $f_{\text{vec}}^{\text{eff}}(T_{\text{EW}}; \mathcal{P})$ to enter the target range.

11.2 Response-integral target replacing the bare count

The calculation that would upgrade the benchmark into a derivation is a polarization- and species-resolved response integral. For a field species a , let $\mathcal{K}_a^{\text{ret}}(\omega, \mathbf{k}; T)$ denote the retarded harvesting or stress-response kernel in the chosen detector protocol, and let $\mathcal{W}_a^{\mathcal{P}}(\omega, \mathbf{k}; T)$ denote the protocol weight, including smearing, switching, occupation factors, and polarization projectors. Define

$$f_{\text{vec}}^{\text{eff}}(T; \mathcal{P}) = \frac{\sum_{a \in \text{vector}} \int d\Pi \mathcal{K}_a^{\text{ret}}(\omega, \mathbf{k}; T) \mathcal{W}_a^{\mathcal{P}}(\omega, \mathbf{k}; T)}{\sum_{b \in \text{all}} \int d\Pi \mathcal{K}_b^{\text{ret}}(\omega, \mathbf{k}; T) \mathcal{W}_b^{\mathcal{P}}(\omega, \mathbf{k}; T)}. \quad (98)$$

The actual dynamical target is therefore

$$f_{\text{vec}}^{\text{eff}}(T_{\text{EW}}; \mathcal{P}) \xrightarrow{?} [0.20, 0.30] \quad (99)$$

in a regulator-stable and epoch-consistent calculation, with $27/106.75$ retained only as the historical reference point. Until this response integral is evaluated, the integer ratio remains a mnemonic

target, not a prediction.

11.3 Boltzmann–CTVM kinetic spine

Let ρ_1 be the causally active sector and ρ_2 the residual sector. The effective two-fluid equations are

$$\dot{\rho}_1 + 3H(1 + w_1)\rho_1 = +Q, \quad (100)$$

$$\dot{\rho}_2 + 3H(1 + w_2)\rho_2 = -Q. \quad (101)$$

For the active-sector fraction

$$f_{\text{sec}} = \frac{\rho_1}{\rho_1 + \rho_2},$$

a minimal locking equation is

$$\frac{df_{\text{sec}}}{d \ln a} = \frac{\Gamma_{\text{lock}}}{H} (f_{\text{eq}} - f_{\text{sec}}) - 3f_{\text{sec}}(1 - f_{\text{sec}})(w_1 - w_2),$$

with $f_{\text{eq}} \in [0.20, 0.30]$ as the conservative channel-counting benchmark range (§11.1), conditional on the response-integral calculation of Appendix H flowing into that range. During a narrow locking epoch in which the first term dominates,

$$f_{\text{sec}}(a) = f_{\text{eq}} + (f_{\text{sec},i} - f_{\text{eq}}) \exp\left[-\int_{a_i}^a \frac{\Gamma_{\text{lock}}}{H} d \ln a'\right],$$

where $f_{\text{sec},i}$ is the initial active-sector fraction at a_i . After freeze-out the active component is assumed to redshift as pressureless matter. That assumption is a phenomenological requirement; it must be tested against structure formation.

11.4 Role of longitudinal modes

The longitudinal modes of the massive W^\pm and Z are taken to enable stable topology locking. At high energies the Equivalence Theorem relates longitudinal vector amplitudes to Goldstone amplitudes schematically as

$$\mathcal{M}(V_{L1} \cdots V_{Ln} + X) = i^n \mathcal{M}(\phi_1 \cdots \phi_n + X) + O(M_V/E). \quad (102)$$

In the present framework this is not a proof of locking. It is the motivation for a calculation of polarization-dependent harvesting kernels. The open target is

$$K_V^{\text{long}}(T) \quad \text{versus} \quad K_V^{\text{trans}}(T), \quad (103)$$

and whether the longitudinal kernels really dominate the causal-topology transition.

12 Horizon-regulated dark energy

The dark-energy application uses the gravity dictionary in a cosmological-horizon setting. This is an additional step beyond the AdS/CFT regime in which RT/HRT and bit-thread methods are best controlled. The de Sitter or cosmological-horizon version of the dictionary is therefore not part of the AQFT/matroid construction; it is the extra cosmological-holography conjecture listed separately in Dependency Schema 1.1.

12.1 Horizon scale and edge density

At a cosmological horizon the macroscopic scale is H^{-1} . With the unreduced Planck convention $m_{\text{Pl}}^2 = 1/G_N$, write the horizon-regulated energy density as

$$\rho_\Lambda(t) = \kappa_{\text{edge}} f_{\text{edge}}(t) m_{\text{Pl}}^2 H^2(t). \quad (104)$$

The critical density is

$$\rho_{\text{crit}} = \frac{3H^2}{8\pi G_N} = \frac{3}{8\pi} m_{\text{Pl}}^2 H^2, \quad (105)$$

so

$$\Omega_\Lambda(t) = \frac{\rho_\Lambda}{\rho_{\text{crit}}} = \frac{8\pi}{3} \kappa_{\text{edge}} f_{\text{edge}}(t). \quad (106)$$

If $\kappa_{\text{edge}} = 1/(4\pi)$ and $f_{\text{edge}} \rightarrow 1$, then $\Omega_\Lambda \rightarrow 2/3$.

For orientation, Planck 2018 base- Λ CDM gives $\Omega_m \simeq 0.315$, hence $\Omega_\Lambda \simeq 0.685$ in a flat model [19]. The conditional CTVM target $2/3 \simeq 0.667$ is roughly 3% low relative to that reference value. The comparison does not supply evidence for the model; it only calibrates the target.

12.2 Conditional null-boundary normalization

For a null boundary \mathcal{N} with generator k^a and non-affinity κ_N defined by $k^b \nabla_b k^a = \kappa_N k^a$, the null boundary term can be written

$$I_N = \frac{1}{8\pi G_N} \int_{\mathcal{N}} d\lambda d^{d-2} y \sqrt{\gamma} \kappa_N, \quad (107)$$

up to the usual normalization and joint-term subtleties of null boundaries. The key caveat is that null-boundary terms depend on the normalization of the null generators; if generators are affinely parametrized then $\kappa_N = 0$ on that segment, and joint terms carry part of the information.

Hypothesis 12.1 (Causal-diamond edge normalization). *For a de Sitter-like causal diamond, choose the null normalization associated with the horizon Killing flow, so that each of the two null sheets has effective non-affinity H over a parameter range of order H^{-1} .*

Under this hypothesis the two null sheets give an action per horizon area

$$\frac{I_{N^+ \cup N^-}}{A_H} \sim \frac{2H}{8\pi G_N} H^{-1} = \frac{1}{4\pi G_N}. \quad (108)$$

This quantity has the dimensions of $1/G_N$. It is not yet the dimensionless coefficient κ_{edge} in $\rho_\Lambda = \kappa_{\text{edge}} f_{\text{edge}} m_{\text{Pl}}^2 H^2$.

Hypothesis 12.2 (Edge-action/vacuum-energy dictionary). *There exists a coarse-grained edge theory, derived from the CTVM split-factor and residual sectors, whose effective boundary action density I_{edge}/A_H is converted into a horizon-regulated bulk vacuum density by the edge dictionary*

$$\rho_\Lambda = f_{\text{edge}} H^2 \frac{I_{\text{edge}}}{A_H}. \quad (109)$$

Equivalently, in the parametrization $\rho_\Lambda = \kappa_{\text{edge}} f_{\text{edge}} m_{\text{Pl}}^2 H^2$, the dimensionless coefficient is

$$\kappa_{\text{edge}} = \frac{1}{m_{\text{Pl}}^2} \frac{I_{\text{edge}}}{A_H}. \quad (110)$$

Thus $\kappa_{\text{edge}} = 1/(4\pi)$ follows only after the causal-diamond normalization of the null boundary action and this edge-action/vacuum-energy dictionary are both imposed. The existence of such a coarse-grained edge theory is not proved here; it is a load-bearing assumption of the CTVM dark-energy mechanism.

With this additional normalization convention,

$$\kappa_{\text{edge}} = \frac{1}{4\pi}. \quad (111)$$

Proposition 12.1 (Conditional κ_{edge}). *Given the causal-diamond edge normalization and the edge-action/vacuum-energy dictionary hypotheses, $\kappa_{\text{edge}} = 1/(4\pi)$ within the adopted CTVM conventions.*

It is useful to display the convention dependence explicitly. Let

$$\kappa_{\text{edge}} = \kappa_{\text{edge}}[\nu, \lambda, j, c_{\text{ct}}], \quad (112)$$

where ν denotes the null-generator normalization, λ the affine/Killing parameter convention, j the joint/corner prescription, and c_{ct} the null-boundary counterterm convention. The CTVM causal-diamond choice is the particular point

$$\kappa_{\text{edge}}^{\text{CTVM}} = \kappa_{\text{edge}}[\nu_{\text{CD}}, \lambda_{\text{CD}}, j_{\text{CD}}, c_{\text{ct}}^{\text{CD}}] = \frac{1}{4\pi}. \quad (113)$$

Alternative conventions define alternative model variants unless a microscopic edge theory shows that they flow to the same infrared coefficient.

Remark 12.1. *This is a normalization argument, not an unconditional derivation from GR. The remaining work is to show that the CTVM edge sector really couples to the null-boundary action with this normalization, that the parametrization/joint-term ambiguities are fixed in the proposed way, and that the dimensional bridge from edge action density to effective vacuum energy is produced by an actual coarse-grained edge theory. Appendix G shows explicitly that κ_{edge} is normalization-sensitive; $1/(4\pi)$ is a convention-dependent CTVM choice, not a universal constant derived from the null action alone.*

The spectral-shell hypothesis below is also part of this conditional package. Changing the null normalization changes the horizon temperature and rate units used in the kinetic spine, so stability of the edge mechanism requires both a stable normalization convention and a stable low-energy transition shell.

12.3 KMS detailed balance and fermionic stabilization

Let $\mathcal{A}_{\text{edge}}$ be the horizon edge algebra in a Gibbons-Hawking/KMS state with temperature

$$T_H = \frac{H}{2\pi}. \quad (114)$$

Let V be an operator mediating transitions between inactive and active edge states. The KMS condition gives detailed balance:

$$\frac{|A_+(\alpha \rightarrow \beta)|^2}{|A_-(\beta \rightarrow \alpha)|^2} = e^{-(E_\beta - E_\alpha)/T_H}. \quad (115)$$

Hypothesis 12.3 (Near-degenerate edge shell). *Let $W_{\text{edge}}(\omega)$ denote the effective transition weight for the edge activation/deactivation operator. This weight includes the edge density of states, matrix-element factors, and the spectral weight of the coupling operator. Assume:*

1. **Low-energy concentration.** *The transition weight $W_{\text{edge}}(\omega)$ is concentrated on a shell of transitions with energy transfer*

$$0 \leq \omega = E_\beta - E_\alpha \ll T_H, \quad (116)$$

with contributions from $\omega \gtrsim T_H$ subleading.

2. **Smoothness on the shell.** *The transition weight varies slowly across the shell:*

$$W_{\text{edge}}(\omega) = W_{\text{edge}}(0) [1 + O(\omega/T_H)]. \quad (117)$$

Under these assumptions, the intrinsic transition rates per available state pair, defined by integrating the KMS-related weight over the near-degenerate shell, satisfy

$$\frac{\gamma_+}{\gamma_-} = \frac{\int_{\text{shell}} d\omega W_{\text{edge}}(\omega) e^{-\omega/T_H}}{\int_{\text{shell}} d\omega W_{\text{edge}}(\omega)} = 1 - \frac{\langle \omega \rangle_W}{T_H} + O\left(\frac{\langle \omega^2 \rangle_W}{T_H^2}\right), \quad (118)$$

where $\langle \cdot \rangle_W$ denotes the shell average weighted by W_{edge} . Hence $\gamma_+ \simeq \gamma_-$ to leading order in the shell thickness $\langle \omega \rangle_W/T_H$.

Remark 12.2 (Meaning of the rate equality). *The equality $\gamma_+ \simeq \gamma_-$ is not a statement that the macroscopic activation and deactivation fluxes are equal. It concerns the intrinsic transition rates per available state pair. Occupation, blocking, and final-state availability are accounted for separately by b_{edge} . This separation is what allows the kinetic fixed point*

$$f_{\text{edge},*} = \frac{\gamma_+}{\gamma_+ + \gamma_- b_{\text{edge}}} \quad (119)$$

to move from 1/2 in the unblocked case to 1 as $b_{\text{edge}} \rightarrow 0$.

Remark 12.3 (Why near-degeneracy is natural for horizon edge modes). *Two structural features make low-energy concentration plausible without proving it. First, gapless or nearly gapless edge excitations—such as soft horizon modes or stretched-horizon Goldstone-like modes—can put substantial transition weight at small ω . Second, the Gibbons-Hawking temperature $T_H = H/(2\pi)$ is the intrinsic thermal scale, so transitions with $\omega \gtrsim T_H$ are thermally suppressed in the KMS weight. This is an assumption about the edge spectrum, not a theorem; a microphysical edge theory is required to establish it.*

The edge fraction evolves according to

$$\frac{df_{\text{edge}}}{dN} = \gamma_+ (1 - f_{\text{edge}}) - \gamma_- b_{\text{edge}}(N) f_{\text{edge}},$$

where $N = \ln a$ and b_{edge} is the availability of deactivation final states. The instantaneous fixed point is

$$f_{\text{edge},*} = \frac{\gamma_+}{\gamma_+ + \gamma_- b_{\text{edge}}}.$$

If $\gamma_+ \simeq \gamma_-$, then unblocked deactivation $b_{\text{edge}} = 1$ gives $f_{\text{edge},*} = 1/2$, while full Pauli blocking $b_{\text{edge}} \rightarrow 0$ gives $f_{\text{edge},*} \rightarrow 1$. This yields a factor-of-two increase in the active edge fraction.

A minimal edge model making the blocking claim calculable is

$$\mathcal{H}_{\text{edge}} = \bigotimes_{\alpha=1}^{N_{\text{edge}}} \mathcal{H}_{\alpha}, \quad K_{\text{edge}} = \sum_{\alpha} \epsilon_{\alpha} n_{\alpha} + \sum_{\alpha\beta} J_{\alpha\beta} n_{\alpha} n_{\beta} + \dots$$

For deactivation channels with rate density $\Gamma_-(E)$, define

$$b_{\text{edge}}(T, \mu) = \frac{\int dE \rho_{\text{edge}}(E) [1 - n_F(E; \mu, T)] \Gamma_-(E)}{\int dE \rho_{\text{edge}}(E) \Gamma_-(E)}.$$

Then $\mu - E_{\text{edge}} \gg T$ gives $b_{\text{edge}} \rightarrow 0$ in the ideal blocked limit. This formula turns the factor-of-two enhancement into a microphysical target rather than a verbal Pauli-blocking slogan.

A microscopic blocking mechanism requires a conserved edge charge, a nonzero chemical potential, or topological saturation. For a Fermi–Dirac occupation

$$n_F(E; T, \mu) = \frac{1}{e^{(E-\mu)/T} + 1},$$

blocking becomes strong when $n_F \rightarrow 1$ for the deactivation final states. That is plausible in protected edge sectors, but it is not yet derived.

13 Cosmological dynamics and equation of state

Because $\rho_{\Lambda} \propto f_{\text{edge}} H^2$, a constant f_{edge} would track the critical density rather than behave as a strict cosmological constant. The model therefore needs nontrivial late-time dynamics of f_{edge} .

This is the leading viability issue for the dark-energy arm. If an $O(1)$ value of f_{edge} were present during radiation or matter domination, the horizon component would behave as early tracking dark energy and would alter the expansion rate during BBN, CMB, and structure-formation epochs. Thus CTVM needs $f_{\text{edge}} \ll 1$ at early times and a late transition or freeze-out to $f_{\text{edge}} = O(1)$ near the onset of cosmic acceleration. The tanh profile below is only a phenomenological template for that transition, not a microphysical history.

The Friedmann equation may be written

$$H^2(z) = \frac{8\pi G_N}{3} \frac{\rho_m(z) + \rho_r(z)}{1 - \Omega_{\Lambda}(z)}, \quad \Omega_{\Lambda}(z) = \frac{8\pi}{3} \kappa_{\text{edge}} f_{\text{edge}}(z).$$

A useful phenomenological template is

$$b_{\text{edge}}(N) = \frac{1 - \tanh[(N - N_t)/\Delta]}{2},$$

which moves from unblocked to blocked deactivation as the universe evolves.

From $\rho_{\Lambda} = \kappa_{\text{edge}} f_{\text{edge}} m_{\text{Pl}}^2 H^2$ and the conservation identity

$$\frac{d \ln \rho_{\Lambda}}{d \ln(1+z)} = 3(1 + w_{\text{eff}}),$$

one obtains

$$w_{\text{eff}}(z) = -1 + \frac{1}{3} \frac{d \ln f_{\text{edge}}}{d \ln(1+z)} + \frac{2}{3} \frac{d \ln H}{d \ln(1+z)}.$$

This is a key diagnostic: the theory is not identical to Λ CDM unless $f_{\text{edge}}(z)$ evolves in just the right way.

For the dark-matter-like active sector the linear perturbation equation is assumed to reduce to

$$\ddot{\delta}_1 + 2H \dot{\delta}_1 - 4\pi G_N (\rho_b + \rho_1) \delta_1 = 0,$$

with negligible pressure and anisotropic stress. The density ρ_1 is the background active-sector density, related to the sector fraction by $\rho_1 = f_{\text{sec}}(\rho_1 + \rho_2)$. This equation is a phenomenological constraint, not a consequence of the sector split alone: it requires the residual-decoupling hypothesis 9.6 to hold not just at background level but at linearized perturbation order, with the active sector gravitating exactly like CDM at that order. Failure of this stronger form of residual decoupling would generically produce scale-dependent growth deviations from Λ CDM.

14 QCD, topological sectors, and self-sustained vacuum analogies

QCD topological sectors and contact terms have been proposed in other approaches to vacuum energy. They are relevant to CTVM for a limited reason: they may supply microscopic examples of conserved or slowly relaxing topological edge data. A dimensional relation such as

$$\rho_\Lambda \sim H \Lambda_{\text{QCD}}^3 \tag{120}$$

is suggestive but not a derivation. A CTVM completion would require an effective action matching between QCD topological susceptibility, horizon edge modes, and the CTVM edge fraction f_{edge} . The QCD and Volovik analogies in this section are mechanism-building analogies only; they do not import topological invariants or vacuum-cancellation theorems into CTVM without an explicit matching calculation.

Volovik’s self-sustained vacuum analogy is also relevant but limited. In controlled equilibrium quantum liquids, large microscopic contributions to vacuum energy can cancel in the macroscopic pressure through a thermodynamic identity relating microscopic energy density, chemical potential, and pressure. This is stronger than the generic effective-field-theory slogan that microscopic vacuum energy need not equal gravitating response: it shows, in a real many-body system, how an equilibrium response functional can differ from the naive microscopic energy sum. CTVM borrows the structure of that lesson—the possibility that a sector-balance or edge-equilibrium law, rather than the microscopic zero-point sum, controls gravitational response—but it does not inherit Volovik’s equilibrium cancellation theorem, nor does it derive the observed value of Ω_Λ .

15 Operational analogies and analogue experiments

Among the operational proposals, the analogue-vacuum statistics tests constitute the most concrete near-term experimental face of the program. The transport/teleportation contrast is retained as a conceptual diagnostic whose experimental implementation requires a specified interferometric protocol and careful carrier-phase calibration.

15.1 Transport versus teleportation-like relabeling

Transport of a quantum system along a continuous worldline is a continuous dynamical process and can accumulate geometric or Berry phase. Teleportation using a harvested resource and classical communication is not transport of the quantum state along a spacetime path: it consumes entanglement and classical information to reprepare the state at a distant location.

A continuous worldline $\gamma : [0, 1] \rightarrow M$ carries topological data—winding number, holonomy, and geometric phase accumulated around a closed loop,

$$\phi = \oint A_\mu dx^\mu.$$

A teleported state is not continuously carried along such a worldline. It is reconstructed using a shared entanglement resource and classical communication. The laboratory devices and classical signals used to implement the protocol still propagate through spacetime and may acquire their own phases, but those carrier phases must be calibrated separately. The CTVM-relevant distinction is that the teleported quantum state itself has no continuous trajectory γ to which a spacetime holonomy can be assigned.

Earlier drafts described transport and teleportation as continuous modular flow, $\sigma_t = \Delta^{it}$, versus discrete modular conjugation, J . That description should be read as a structural analogy, not as a proven modular-theoretic identity. Establishing it rigorously would require showing that the harvested teleportation channel implements the modular conjugation of the relevant local algebra exactly, which is not done here. What survives without that identification is the operational contrast:

*Transport probes path-dependent geometric response;
teleportation probes resource-mediated algebraic relabeling.*

This contrast is operationally diagnostic. In a controlled interferometric comparison, a quantum system transported continuously around a closed loop can acquire the geometric phase $\phi = \oint A_\mu dx^\mu$. A teleported state is reconstructed rather than transported around the loop; after carrier-phase calibration, it should not acquire the same path-dependent holonomy. A measurable phase difference between continuous transport and teleportation-assisted round trips is therefore a candidate operational signature of the transport/teleportation distinction, independent of any claim about modular conjugation.

15.2 BTZ and image-sum harvesting

In quotient spacetimes such as BTZ, the Wightman function can be written as an image sum

$$D_{\text{BTZ}}^+(x, x') = \sum_n D_{\text{AdS}}^+(x, \gamma^n x'). \quad (121)$$

The image terms can enhance detector correlations near horizon scales. This is a legitimate curved-spacetime harvesting signature. It is not unique to CTVM, but CTVM predicts that such enhancements should appear specifically in the pairwise response sector selected by the model's network dictionary.

15.3 Analogue-vacuum experimental program

The Volovik/statistical layer suggests two experimental arms.

Bosonic arm. Use detector-like impurities or localized modes coupled to a BEC or ^4He -like collective medium. Measure a harvestability proxy such as negativity or mutual information after subtracting direct signalling. The diagnostic ratio is

$$\eta_B(\omega) = \frac{\mathcal{N}_{\text{harv}}^{(B)}(n_B)}{\mathcal{N}_{\text{harv}}^{(B)}(0)}. \quad (122)$$

A CTVM-compatible outcome is enhanced pairwise harvestability in regimes where bosonic occupation enhances final-state access.

Fermionic arm. Use degenerate Fermi gases, synthetic Weyl systems, or ^3He -like topological media. Measure

$$\eta_F(\omega) = \frac{\mathcal{N}_{\text{harv}}^{(F)}(n_F)}{\mathcal{N}_{\text{harv}}^{(F)}(0)}. \quad (123)$$

A CTVM-compatible outcome is suppressed bulk pairwise harvesting together with persistent boundary or defect correlations.

Positive analogue results would not prove dark matter or dark energy. They would show that quantum statistics can operationally bias harvested correlations into bosonic pairwise and fermionic topological patterns, making CTVM a measurable statistical-filter program rather than a purely formal projection postulate.

16 Predictions, benchmarks, and falsifiers

16.1 Benchmarks retained as open targets

The dark-matter benchmark is included because it is falsifiable, not because its numerical proximity currently carries substantial evidential weight. The dark-energy benchmark is included for the same reason: it is a compact target that only becomes physically meaningful if the edge-normalization, blocking-history, and cosmological-horizon dictionary gates are closed.

1. **Dark matter.** $f_{\text{eq}} \in [0.20, 0.30]$ as a conservative response target range, with $27/106.75 \simeq 0.253$ retained as the historical reference point; contingent on the vector-channel response integral and the freeze-out mapping from causal fraction to matter abundance.
2. **Dark energy.** $\Omega_\Lambda = (8\pi/3)\kappa_{\text{edge}}f_{\text{edge}} \rightarrow 2/3$ if $\kappa_{\text{edge}} = 1/(4\pi)$ and $f_{\text{edge}} \rightarrow 1$.
3. **No standard particle signal.** If the dark components are vacuum-sector effects, direct-detection experiments looking for a new stable particle may remain null.
4. **Finite-redshift equation-of-state transition.** The blocking transition generically gives a nontrivial $w_{\text{eff}}(z)$ rather than an exact constant -1 .
5. **Response-matroid structure.** Specified detector protocols should show the block, rank, or weight patterns required by the model. Failure of these patterns falsifies the corresponding sector-selection claim.
6. **Analogue statistics.** Bosonic media should enhance pairwise harvestability; fermionic topological media should suppress bulk final-state access while preserving boundary/defect correlations.

16.2 Near-term falsifiers

1. A rigorous Standard Model response calculation yields an effective vector fraction outside the target range $[0.20, 0.30]$.
2. Longitudinal vector modes fail to enhance or lock the proposed causal topology across the electroweak transition.
3. The response representation cannot be block-diagonalized, even approximately, in any physically natural detector/split basis.
4. The holographic pairwise network dictionary cannot be constructed for the split-factor sector.
5. Null-boundary and modular-edge calculations fix κ_{edge} to a value incompatible with $1/(4\pi)$ under the CTVM conventions.
6. No plausible horizon edge microphysics drives $b_{\text{edge}} \rightarrow 0$ at late times.
7. The induced $w_{\text{eff}}(z)$ or growth history conflicts with combined distance, CMB, BAO, weak-lensing, and structure data.
8. Analogue experiments show no statistics-resolved bias of the type required by the Volovik-CTVM mechanism.

Layer	Falsifier	Effort	If it fails
Split algebra	No modular-compatible split or type-II substitute with stable detector data.	High	The strong gravity dictionary has no regulated algebraic input for the relevant horizon.
Matroid diagnostic	Stable circuits cross the proposed operational split.	Low	The exact detector/split assignment is falsified for that protocol.
Vector response	$f_{\text{vec}}^{\text{eff}}$ does not enter $[0.20, 0.30]$ in a stable protocol window.	Low–medium	The dark-matter benchmark remains numerology rather than a response target.
Entropy cone	A small CTVM network violates required entropy inequalities.	Low	The proposed network cannot be given the intended RT/HRT interpretation.
Edge normalization/history	κ_{edge} or $f_{\text{edge}}(z)$ is convention- or data-incompatible.	Medium	The dark-energy normalization and expansion-history arm fails.
Statistics layer	No Bose/Pauli response bias appears in analogue harvesting.	Medium	The Volovik/statistical mechanism loses empirical support.
Residual decoupling	Residual response alters background or perturbation-order source terms.	High	The residual-inertness and CDM-like growth assumptions fail.

17 Open derivational targets and architectural gates

The preceding sections deliberately separate definitions, conditional mechanisms, benchmarks, and conjectures. This section collects the remaining derivational debt. The items are ordered by how much downstream phenomenology they block, not by where they appear in the paper.

The gate numbering follows Dependency Schema 1.1: that schema is the logical dependency map, while the present section is the execution plan for closing, testing, or falsifying those gates. We use the status tags DERIVED, CONDITIONAL, and OPEN as textual markers: DERIVED means the item follows from stated assumptions; CONDITIONAL means it holds only if an explicit hypothesis or dictionary is imposed; OPEN means no derivation is attempted here and the item is a target for future calculation.

17.1 Critical path: cheapest decisive tests first

The list below is ordered by decisiveness per unit effort rather than by conceptual depth. Rank 1, the entropy-cone check, is the cheapest general kill test for the holographic/network arm. Rank 2 attacks the leading dark-energy viability risk, the $f_{\text{edge}}(z)$ tracking history. Rank 3, the vector-response integral, is medium effort but attacks the dark-matter arm directly. These tests should be attempted before the frontier gates such as the full de Sitter holography dictionary or the complete type-II horizon-algebra construction.

Rank	Task	Effort	Decisive failure mode
1	Run an entropy-cone linear-programming check on a four- or five-cell candidate response network.	Low	Violation of strong subadditivity, MMI, or the relevant entropy-cone inequalities kills the holographic interpretation for that network.
2	Fit the $f_{\text{edge}}(z)$ template against early-dark-energy and late-time distance/growth constraints.	Medium	If no viable transition is allowed, the dark-energy arm fails independently of the edge-normalization arithmetic.
3	Execute the Appendix H response integral in a Proca/vector toy plus representative scalar/fermion denominator.	Low–medium	If $f_{\text{vec}}^{\text{eff}}$ is not in the target range in any stable, gauge-consistent window, the dark-matter benchmark loses support.
4	Compute δ_{block} and run the certified matroid split test on the massive-scalar example and at least one vector protocol, using the logarithmic-negativity witness of Definition 5.1.	Low	If all physical protocols produce uniform matroids, unstable operational partitions, or robust cross-circuits, the matroid layer carries no sector information for those protocols.
5	Quantify witness dependence by comparing logarithmic negativity with any available entanglement witnesses in the same detector protocol.	Low	If the operational column partition is highly witness- or threshold-sensitive, the split test is non-diagnostic.
6	Catalogue the convention-spread interval for κ_{edge} under the null-boundary choices of Appendix G.	Low–medium	If the spread is $O(1)$ and no microscopic edge theory selects a convention, the Ω_{Λ} normalization has no predictive content.
7	Only after the above, attack the frontier gates: type-II subalgebra selection, cosmological-horizon holography, and the coarse-grained edge theory.	High	These are architecture-building tasks, not cheap falsifiers.

17.2 Architectural gates

A1. Modular-compatible split inclusion or observer-inclusive type-II horizon algebra.

Hypothesis 4.1 is the deepest architectural gate in the type-I split formulation. The split property supplies intermediate type-I factors, but it does not by itself imply that the chosen regulated type-I factor is invariant under the reference-state modular flow:

$$\sigma_t^\omega(\mathcal{N}_{\text{reg}}) = \mathcal{N}_{\text{reg}}.$$

By Takesaki's criterion, such modular invariance is the condition that allows a normal ω -preserving conditional expectation onto \mathcal{N}_{reg} . Wedge and conformal-ball geometries are the cleanest cases because the modular flow is geometrically controlled. For causal diamonds, light-cone tubes, stretched horizons, and cosmological horizons, the required modular-compatible split factor must be explicitly constructed or the dark-energy application remains conditional.

The observer-inclusive crossed-product/type-II route offers a cleaner alternative. If the static-patch horizon algebra is replaced by a finite type-II algebra \mathfrak{A}_H with faithful normal trace τ_H , then a τ_H -preserving Umegaki conditional expectation exists onto any chosen unital von Neumann subalgebra $\mathfrak{B} \subset \mathfrak{A}_H$. This trace-preserving expectation should not be conflated with preservation of an arbitrary nontracial physical state ω_H ; if the source dictionary requires ω_H -preservation, additional modular compatibility of \mathfrak{B} with ω_H is still required. In this route the hard problem is therefore not the existence of the tracial expectation; it is identifying the physically correct split-sector subalgebra, addressing physical-state preservation when needed, and proving regulator stability of the detector-response and stress-response data.

A2. Detector/split compatibility for physical protocols. Hypothesis 7.1 requires the detector feature map to be compatible with the split-relative image/kernel decomposition. The massive scalar example gives an exponentially accurate approximate block structure in one controlled case, but not an exact direct sum. Vector, fermionic, and analogue-vacuum protocols must still be tested.

A3. Operational/matroid compatibility. Hypothesis 7.2 links the operational EPR-active column partition to the matroid split test. It has not yet been verified for any physical detector protocol. The task is to compute $\rho_{ij}^{(P)}$, determine $E_{\text{split}}^{(P,\epsilon)} \sqcup E_{\text{res}}^{(P,\epsilon)}$, and check whether stable circuits of $M[\widehat{X}]$ cross the proposed partition.

A4. Pairwise response-network dictionary. Hypothesis 9.1 postulates a regulator-stable assignment

$$(\widehat{X}_{\text{split}}, w_{\text{split}}) \mapsto G_{\text{split}} = (V, E, c).$$

The capacities c_{ij} must be derived from a specified detector protocol, entanglement monotone, subtraction scheme, and regulator window. Regulator stability of G_{split} is required before the network can be interpreted physically.

A5. Holographic entropy dictionary in controlled AdS-like settings. Hypothesis 9.2 is now restricted to the controlled AdS/RT/HRT-like part. The task is to show that the response network obtained from the split block has a cut function that can be embedded in a regime where RT/HRT or bit-thread methods are actually justified.

A6. Entropy-cone admissibility for CTVM networks. Hypothesis 9.4 requires the entropy vector generated by the CTVM network to satisfy the inequalities appropriate to the intended

holographic regime, including strong subadditivity and, when applicable, monogamy of mutual information and other holographic entropy inequalities. No concrete CTVM network has yet passed this test.

A7. Cosmological-horizon holography dictionary. Hypothesis 9.3 is an additional conjectural step beyond the better-controlled AdS/RT/HRT setting. A viable completion must specify whether the relevant dictionary is a static-patch, stretched-horizon, bilayer-screen, DSSYK-like, crossed-product/type-II, or other de Sitter-inspired construction, and must state what replaces the usual AdS boundary Hilbert-space interpretation.

A8. Semiclassical sourcing dictionary. Hypothesis 9.5 is the step from entropy or relative-entropy data to a local source term,

$$\delta S_{\text{hol}} \longrightarrow \delta \langle T_{\mu\nu} \rangle_{\text{source}} \longrightarrow \delta G_{\mu\nu}.$$

The present text states this schematically. A completion must specify the state class, the modular Hamiltonian, the renormalised stress tensor, contact terms, anomalies, and conservation scheme.

A leading candidate for the nonperturbative completion is a regulated local open-system formulation. In such a formulation the split-sector state evolves under a GKSL/Lindblad generator,

$$\frac{d\rho_{\text{split}}}{ds} = -i[K_{\text{split}} + H_{\text{LS}}, \rho_{\text{split}}] + \mathcal{D}_{\text{leak}}^{(\nu_{\text{res}})}(\rho_{\text{split}}), \quad (124)$$

where the dissipator is derived from boundary/interface operators coupling the split sector to residual topological modes [55, 56]. Conservation of the total stress tensor would then be imposed on the combined split-plus-residual system, while the split-sector source alone may exchange energy-momentum through the leakage current. This is a construction target, not yet part of the main derivation.

A9. Residual decoupling at background and perturbation order. Hypothesis 9.6 is needed at background level. Section 13 additionally requires a stronger linearised version,

$$D_{\widehat{X}_{\text{res}}} \delta \langle T_{\mu\nu} \rangle_{\text{source}} = 0,$$

to obtain CDM-like growth without residual-sector scale-dependent corrections. This perturbation-order residual decoupling is not derived.

In the open-system route, residual decoupling becomes a quantitative statement about the leakage dissipator. One seeks a regulated interface model in which

$$\mathcal{D}_{\text{leak}}^{(\nu_{\text{res}})}$$

is suppressed by a residual topological index, for example through a factor

$$F_{\nu}^2 \sim (1 + \chi\nu_{\text{res}})^{-1},$$

so that residual-to-split leakage, off-block response norms, and tolerance-certified cross-sector circuits are suppressed together. Failure of this common scaling would falsify the topological-uniformity version of residual decoupling, even if the weaker detector/matroid bookkeeping remains valid.

A10. Modular equivalence and inertial sequestration. The Entropic Inertia interpretation makes residual decoupling sharper. A sector removed from gravitational sourcing must also be removed from the modular work functional that defines ordinary inertial response, or CTVM would generically violate the Einstein equivalence principle. A completion must show that the residual algebra has vanishing source-relevant commutator with the physical translation generator,

$$D_{P^1}^{\text{res}} = \sup_{A \in \mathcal{A}_{\text{res}}, \|A\| \leq 1} \left\| \Pi_{\text{src}}^{\text{ret}}([P^1, A]) \right\| = 0,$$

or is suppressed by the same leakage factor that controls residual-to-split communication. Failure of this condition would mean that the residual sector contributes to inertial mass without contributing to gravitational mass.

A11. Topological boundary invariant and faithful action on response data. Hypothesis 10.1 requires an edge or boundary response operator carrying a genuine topological invariant, such as a Chern number, winding number, \mathbb{Z}_2 index, or K -theory class. The invariant must act faithfully on the detector-response representation. This is stronger than simply assuming protected generators exist.

17.3 Dark-matter phenomenology gates

D1. Vector-channel response integral. Appendix H defines

$$f_{\text{vec}}^{\text{eff}}(T; \mathcal{P}) = \frac{\sum_{a \in \text{vectors}} R_a^{\mathcal{P}}(T)}{\sum_{b \in \text{all channels}} R_b^{\mathcal{P}}(T)}.$$

This calculation must be carried out before the benchmark range can be promoted from target to prediction; 27/106.75 remains the historical reference point. It must resolve the below-EWSB numerator versus above-EWSB denominator mismatch, regulator dependence, longitudinal modes, gauge consistency, and species weighting.

D2. Two-fluid exchange term Q and locking rate Γ_{lock} . The two-fluid equations use an exchange term

$$\dot{\rho}_1 + 3H(1 + w_1)\rho_1 = +Q, \quad \dot{\rho}_2 + 3H(1 + w_2)\rho_2 = -Q.$$

The locking equation introduces Γ_{lock}/H as a phenomenological rate. A microphysical derivation must connect Q and Γ_{lock} to the response kernels of Appendix H.

D3. Longitudinal-versus-transverse vector kernels. The claim that longitudinal vector modes drive causal-topology locking requires polarization-resolved kernels

$$K_V^{\text{long}}(T), \quad K_V^{\text{trans}}(T).$$

The Equivalence Theorem motivates the longitudinal channel but does not compute its detector-response weight. The required calculation must include Goldstone contributions and gauge-parameter cancellation in an electroweak-like setting.

D4. Statistics-resolved sector bias. Hypothesis 6.1 states that bosonic channels enhance pairwise harvestability while fermionic/topological channels suppress or redirect bulk response toward residual or boundary modes. This is motivated by the Volovik/statistical layer and

analogue experiments, but no AQFT-level calculation has yet established it for the CTVM detector protocols.

17.4 Dark-energy phenomenology gates

E1. Coarse-grained edge theory for Hypothesis 12.2. This is the deepest unfilled dark-energy gap. Hypothesis 12.2 requires a coarse-grained edge theory, derived from the split-factor and residual sectors, whose effective boundary action density I_{edge}/A_H is converted into the horizon-regulated bulk vacuum density

$$\rho_\Lambda = \kappa_{\text{edge}} f_{\text{edge}} M_{\text{Pl}}^2 H^2.$$

Without such an edge theory, $\kappa_{\text{edge}} = 1/(4\pi)$ remains a conditional normalization benchmark.

E2. Causal-diamond null normalization robustness. Hypothesis 12.1 fixes a causal-diamond null normalization. Appendix G catalogs how κ_{edge} varies under null-generator normalization, affine/Killing parametrization, joint terms, and counterterm choices. A microscopic edge theory is needed to decide which convention is physically selected.

E3. Fermionic blocking microphysics. The late-time limit $b_{\text{edge}} \rightarrow 0$ requires a conserved edge charge, nonzero chemical potential, or topological saturation mechanism. The minimal formula for $b_{\text{edge}}(T, \mu)$ gives a target but not a derivation.

E4. Late-time history $b_{\text{edge}}(N)$. The tanh profile used in Section 13 is a phenomenological template. A microphysical model must determine the redshift, width, and monotonicity of the transition into the blocked regime.

E5. Boundary-rank/edge-charge proportionality. The proportionality

$$Q_{\text{edge}} = \alpha \text{rk } M_{\partial H}$$

is an additional physical input. It must be derived from a boundary response theory or replaced by a different edge-charge functional.

E6. QCD topological-susceptibility to horizon-edge matching. Section 14 notes a possible connection between QCD topological sectors and the CTVM edge fraction f_{edge} . A completion would require an effective-action matching between QCD topological susceptibility, horizon edge modes, and the CTVM edge-sector variables. Dimensional numerology is not enough.

17.5 Technical and computational targets

T1. Ward identity, anomalies, and contact terms. Proposition 4.1 should be expanded or supported by an explicit curved-spacetime renormalisation reference. The proof must address trace anomalies, contact terms, and renormalisation-scheme dependence of the stress tensor.

T2. Geometric modular flow beyond wedges and conformal balls. Bisognano-Wichmann controls wedges, and conformal symmetry controls balls in CFT vacua. The controlled horizon-regulated substitute used for cosmological applications is not yet constructed.

T3. Fermionic detector kernels. The effective fermionic kernels $W_{\text{eff}}^{(F)}$ must be tabulated for Dirac, Majorana, and fermion-bilinear couplings. The chiral-Majorana stress-tensor-like operator with $J_{\nu}^F(\omega) \sim \omega^3$ should appear as one explicit entry.

T4. Continuum limit of capacity assignment. The renormalised capacities

$$c_{ij}^{\text{ren}} = \lim_{\epsilon \rightarrow 0} c_{ij}^{(\mathcal{P}, \epsilon)}$$

exist only after a subtraction or renormalisation scheme is fixed. That scheme must be specified and tested for stability.

T5. Operational EPR-witness robustness. Definition 5.1 fixes logarithmic negativity as the computable witness of record for near-term tests. The remaining task is to quantify how the operational column partition changes under threshold variation, detector dimension, and alternative witnesses. For two-qubit detector pairs the witness is faithful for entanglement detection, but for higher-dimensional bipartitions it can miss PPT-bound entanglement.

T6. Refined operational column partition. Construct an operational distinguishability test, sensitive to symmetric-correlation versus retarded-response content of the harvested detector state, that can independently assign $\text{Re } X_{.j}$ and $\text{Im } X_{.j}$ to sectors and thereby refine Definition 7.3.

T7. Nonlinear structure formation. The linear perturbation equation is constrained to reproduce CDM-like growth. Nonlinear scales, cluster environments, voids, and possible screening or environmental effects remain uncomputed.

18 Revised axiom ledger

Table 1: Status of the original CTVM axioms in the integrated draft.

Axiom	Earlier claim	Integrated status
S1	A canonical causal-topology projection selects the gravitating sector.	Replaced by two distinct maps: a split-relative conditional expectation, when it exists, and an SK response-functional filter. The model-level Π_{CT} requires an additional sourcing dictionary.
S2	Diagonal terms are classical loops; off-diagonal terms are quantum non-loops.	Replaced by a response matroid built from a fixed realified matrix representation. Loops are zero columns. Physical labels require protocol-specific interpretation.
S3	Only the projected pairwise sector gravitates.	Not derived from AQFT or matroid theory. Conditional on detector/split compatibility, a pairwise response-network assignment, a holographic entropy dictionary, entropy-cone admissibility, a semiclassical sourcing dictionary, and residual decoupling of the source functional.

Axiom	Earlier claim	Integrated status
T1	The non-EPR sector is super-selected.	Conditional on exact block decomposition or an exact selection rule. Microcausality and Pauli blocking alone are insufficient.
T2	Horizon boundaries carry protected edge modes.	Motivated by topological media and ${}^3\text{He-B}$ analogues. Requires an actual topological boundary theorem in the CTVM representation.
T3	Edge charge scales with area.	Nuclearity supplies an upper bound; lower area scaling requires a positive-density protected boundary sector and an edge-charge/rank proportionality.
κ_{edge}	$1/(4\pi)$ derived absolutely.	Convention-dependent conditional normalization: null-generator normalization, affine/Killing parameter choice, joint terms, counterterms, and the edge-action/vacuum-energy dictionary must be fixed.
Ω_{DM}	Parameter-free prediction.	Model benchmark from a vector-channel count. Needs response-weight derivation.

19 What CTVM does not claim

For clarity, the integrated framework does not claim the following:

1. It does not derive gravity from AQFT alone. The sourcing step requires an additional semiclassical or holographic dictionary.
2. It does not identify entanglement with gravity. Retarded response, operational harvestability, and gravitational sourcing are distinct layers.
3. It does not claim the physical vacuum is literally a Volovik helium-like superfluid. The Volovik layer is a universality-class analogy for mechanism building.
4. It does not provide a unique or canonical split factor for generic regions. Modular-compatible split factors or type-II substitutes remain construction targets.
5. It does not complete a microscopic derivation of the dark-matter or dark-energy abundances. The numerical values remain conditional benchmark targets until the response and edge calculations are performed.
6. It does not solve the magnitude/cancellation face of the cosmological-constant problem. The present framework conditionally addresses the horizon-scale/coincidence face; a mechanism that cancels or sequesters bulk vacuum energy remains an open target.
7. It does not yet supply a topological invariant behind the phrase “causal topology” in every sector. The phrase names a response-geometric diagnostic until a protected invariant is constructed.
8. It does not derive the Einstein equivalence principle from CTVM. The Modular Equivalence Principle introduced here is a consistency condition: the same sector must control inertial modular work and gravitational sourcing, or the residual sector must be sequestered from both.

20 Conclusion

The integrated CTVM is more coherent once three distinctions are enforced. First, retarded response is not the same as entanglement. Second, a linear matroid is not the same as a weighted network or a holographic entropy functional. Third, Volovik’s superfluid vacuum analogies are not proofs of cosmology; they are controlled examples of how bosonic collective response and fermionic topology can coexist in quantum media.

With these distinctions in place, the model becomes sharper rather than weaker. The AQFT layer supplies split inclusions, modular flow, and response diagnostics. The detector layer defines what it means to extract EPR-type correlations operationally. The Volovik layer supplies a plausible statistical mechanism: Bose enhancement biases collective pairwise channels, while Pauli blocking and topology stabilize residual boundary/defect structure. The matroid layer records rank and support, while the weight layer records response strength. The holographic layer, when imposed, turns a pairwise weighted network into an entropy functional. The cosmological layer then proposes two conditional targets: a vector-channel response fraction near 0.253 and a horizon-edge normalization near $2/3$. These targets become predictive only after the response integral, edge dictionary, and expansion-history gates are discharged.

The result is not a completed derivation of dark matter or dark energy. It is a conditional framework with exposed load-bearing assumptions. That is the right form for the next stage: the theory can now be attacked by concrete calculations, analogue experiments, and cosmological fits.

Glossary of abbreviations

AQFT	Algebraic quantum field theory.
BEC	Bose–Einstein condensate.
BW	Bisognano–Wichmann theorem.
CSS	Calderbank–Shor–Steane quantum error-correcting code.
CTP	Closed-time path.
CTVM	Causal Topology Vacuum Model.
DHR	Doplicher–Haag–Roberts superselection theory.
DSSYK	Double-scaled Sachdev–Ye–Kitaev model.
EEP	Einstein equivalence principle.
EPR	Einstein–Podolsky–Rosen pairwise entanglement.
GHY	Gibbons–Hawking–York boundary term.
GKSL	Gorini–Kossakowski–Sudarshan–Lindblad master equation.
HRT	Hubeny–Rangamani–Takayanagi formula.
KMS	Kubo–Martin–Schwinger thermal condition.
LOCC	Local operations and classical communication.

MEP	Modular Equivalence Principle.
RT	Ryu–Takayanagi formula.
SK	Schwinger–Keldysh formalism.
UDW	Unruh–DeWitt detector.

A Schwinger-Keldysh identities

With $\phi_c = (\phi_+ + \phi_-)/\sqrt{2}$ and $\phi_q = (\phi_+ - \phi_-)/\sqrt{2}$, the two-point functions can be organized as

$$G^{cc}(x, y) = \frac{i}{2} \langle \{\phi(x), \phi(y)\} \rangle, \quad (125)$$

$$G^{cq}(x, y) = G^{\text{ret}}(x, y), \quad (126)$$

$$G^{qc}(x, y) = G^{\text{adv}}(x, y), \quad (127)$$

$$G^{qq}(x, y) = 0. \quad (128)$$

The response filter sets G^{cc} to zero and retains G^{cq} and G^{qc} . This is safe as a statement about a chosen cumulant representation. It is not, without additional construction, a physical operation on a density matrix.

B Detector density matrix bookkeeping

For two identical UDW detectors in the usual weak-coupling expansion, the detector density matrix contains local noise terms P_i , exchange/correlation terms, and nonlocal coherence terms. A schematic two-detector block has entries of the form

$$P_i \sim \lambda_i^2 \int d\tau d\tau' \chi_i(\tau) \chi_i(\tau') e^{-i\Omega_i(\tau-\tau')} D^+(x_i(\tau), x_i(\tau')), \quad (129)$$

$$X_{ij} \sim \lambda_i \lambda_j \int d\tau d\tau' \chi_i(\tau) \chi_j(\tau') e^{-i(\Omega_i\tau + \Omega_j\tau')} D^+(x_i(\tau), x_j(\tau')), \quad (130)$$

with precise signs and phases depending on detector conventions. The important repair is the double integral with switching, smearing, and gap phases. A single integral of $\langle \phi^2 \rangle$ is not a correct general detector response.

C Minimal worked detector model: massive scalar split test

This appendix gives the template calculation behind the approximate detector/split compatibility example in Section 7. Consider a free massive scalar field in 1 + 1 Minkowski spacetime. Let two detector clusters A and B have compact spatial smearings supported in intervals of radius R , with separation $d > 2R$, and smooth Gaussian switching

$$\chi(t) = e^{-t^2/(2T^2)}. \quad (131)$$

For a detector with gap Ω , the leading response matrix element has the form

$$X_{ij} = \lambda_i \lambda_j \int dt dt' dx dx' \chi(t) \chi(t') s_i(x) s_j(x') e^{-i\Omega(t-t')} D^+(t, x; t', x'), \quad (132)$$

with s_i the spatial smearing. At equal times and spacelike separation r , the massive two-point kernel obeys

$$D^+(r) \propto K_0(mr) \sim \sqrt{\frac{\pi}{2mr}} e^{-mr}, \quad mr \gg 1. \quad (133)$$

Hence every cross-cluster entry satisfies a bound of the schematic form

$$|X_{ij}^{AB}| \leq C(\lambda, T, \Omega, s_i, s_j) e^{-m(d-2R)} \quad (134)$$

up to polynomial factors in $m(d-2R)$. In block form,

$$\hat{X} = \begin{pmatrix} \hat{X}_A & O(e^{-m(d-2R)}) \\ O(e^{-m(d-2R)}) & \hat{X}_B \end{pmatrix}. \quad (135)$$

If the smallest nonzero singular values of the diagonal blocks are bounded below by γ , and $e^{-m(d-2R)} \ll \gamma$, then the split is stable as an approximate decoupling. It is not an exact matroid direct sum unless the off-block entries vanish exactly. This example shows what a controlled regulator-dependent split estimate should look like; analogous estimates are still required for vector, fermionic, and analogue-vacuum protocols.

D Constant-parameter edge kinetic solution

If b_{edge} , γ_+ , and γ_- are constant, then

$$f_{\text{edge}}(N) = f_{\text{edge},*} + (f_{\text{edge}}(N_i) - f_{\text{edge},*}) e^{-(\gamma_+ + \gamma_- b_{\text{edge}})(N - N_i)},$$

where

$$f_{\text{edge},*} = \frac{\gamma_+}{\gamma_+ + \gamma_- b_{\text{edge}}}.$$

This is the minimal solution used in the horizon-edge phenomenology.

E Why the dark-matter number is a target, not evidence

This appendix is included to prevent numerical proximity from being mistaken for evidence. The value 27/106.75 is retained because it defines a falsifiable response-integral target, not because nearby integer ratios are rare.

The numerical proximity of 27/106.75 to the observed dark-matter fraction is interesting, but numerical proximity alone is weak evidence. The correct evidential weight depends on how many ratios were considered, how independently the channel-counting rule was motivated, and whether the same rule predicts other observables. As a crude operational check, allow ratios of the form n/g_* with an integer channel count $1 \leq n \leq 40$ and ask how many lie in the broad interval $[0.20, 0.30]$ when $g_* = 106.75$. This interval corresponds to $21.35 \leq n \leq 32.025$, so $n = 22, \dots, 32$ gives eleven integer numerators. This is a conservative look-elsewhere estimate: allowing weighted ratios, fermionic 7/8 factors, alternative denominators at different epochs, or different natural partitions

of Standard Model degrees of freedom broadens the comparison set further. The model should therefore stand or fall on the full package: vector-channel dynamics, freeze-out behavior, growth history, and falsifiable deviations, not on the coincidence alone.

F Three-detector matroid example: physical interpretation and scope

This appendix records a small but useful three-detector example. Its purpose is not to prove the CTVM sector split. Its purpose is to show that the column-doubled realified matroid can be non-uniform, can stratify protocol parameter space, and can distinguish structural linear dependence from accidental magnitude choices. The boxed example in §7.4.6 gives the complementary minimal falsifier: a circuit that explicitly crosses a proposed operational split.

F.1 Codimension correction

For a 3×3 Hermitian response matrix with off-diagonal entries $(z_{12}, z_{13}, z_{23}) \in \mathbb{C}^3$, the off-diagonal parameter space is real six-dimensional. Under the column-doubling convention

$$\widehat{X} = [\operatorname{Re} X \mid \operatorname{Im} X], \quad (136)$$

the imaginary-part columns relevant to the example may be written

$$v_5 = (\operatorname{Im} z_{12}, 0, -\operatorname{Im} z_{23})^T, \quad v_6 = (\operatorname{Im} z_{13}, \operatorname{Im} z_{23}, 0)^T. \quad (137)$$

The previously suggested codimension-two description of the parallel-pair stratum was too strong. The analytic check is immediate from the 2×2 minors of the two-column matrix $(v_5 \ v_6)$. Writing

$$a = \operatorname{Im} z_{12}, \quad b = \operatorname{Im} z_{13}, \quad c = \operatorname{Im} z_{23}, \quad (138)$$

we have

$$v_5 = (a, 0, -c)^T, \quad v_6 = (b, c, 0)^T. \quad (139)$$

The three 2×2 minors are

$$ac, \quad bc, \quad c^2 \quad (140)$$

up to signs. Under the open conditions $a \neq 0$ and $b \neq 0$, these minors vanish simultaneously if and only if $c = 0$.

The $\{5, 6\}$ parallel-pair stratum is therefore defined by the single equation

$$\operatorname{Im} z_{23} = 0, \quad (141)$$

together with the open conditions $\operatorname{Im} z_{12} \neq 0$ and $\operatorname{Im} z_{13} \neq 0$. When $\operatorname{Im} z_{23} = 0$, both v_5 and v_6 collapse onto the line $\operatorname{span}\{e_1\} \subset \mathbb{R}^3$ and are automatically proportional whenever both are nonzero. The ratio $\operatorname{Im} z_{13}/\operatorname{Im} z_{12}$ selects a representative inside the stratum; it is not part of the defining equation of the stratum. Thus the codimension of this parallel-pair wall is one, not two.

F.2 Physical interpretation

The parallel pair $\{5, 6\}$ records an exact algebraic relation among detector response coefficients, not a magnitude relation. A representative example may fix $\text{Im } z_{13}/\text{Im } z_{12} = 2$, but any nonzero real value of this ratio gives the same parallel pair as long as $\text{Im } z_{23} = 0$. Similarly, multiplication of v_5 or v_6 by a nonzero Bose stimulation factor $1 + n_B$ or a nonzero Pauli factor $1 - n_F$ preserves the represented matroid. Only exact zeros, exact block decompositions, or exact new linear relations modify the matroid.

In this sense the example illustrates the distinction between the matroid layer and the weight layer. The matroid detects exact support and dependence structure. The response weights detect amplitudes, enhancement, and suppression.

F.3 What the example shows

- (i) **The matroid layer is non-vacuous.** Without an example, one might worry that a full-rank 3×3 Hermitian response matrix always yields the uniform matroid $U_{3,6}$ after column-doubling. The example gives a counterexample: the realified matroid can be a parallel extension of $U_{3,5}$, distinguished from $U_{3,6}$ by the parallel pair $\{5, 6\}$.
- (ii) **The matroid stratifies protocol space.** The $\{5, 6\}$ parallel pair persists under all perturbations preserving $\text{Im } z_{23} = 0$ and the nonzero open conditions. It defines a codimension-one wall in the six-dimensional space of off-diagonal complex entries. Extra relations, such as an accidental relation of the form $v_1 + v_4 = v_5$ for a special parameter choice, define smaller strata. The matroid distinguishes the generic wall from such non-generic intersections.

F.4 What the example does not show

- (i) **It does not show that physically realizable protocols produce vacuum-sector circuits.** The conditions above are protocol-level conditions on detector phases, switching profiles, and field-coupling smearings. They are not automatically properties of the vacuum. The example shows what the matroid would record if such protocol conditions held; it does not show that a physically realizable protocol satisfies them, nor that satisfying them reflects a vacuum-level topological invariant.
- (ii) **It does not assign the circuit to $\text{im } \Pi_{\text{split}}$ or $\ker \Pi_{\text{split}}$.** The example operates at the response-matrix level. It does not construct a split inclusion, a conditional expectation, or a sector assignment. Establishing such an assignment would require constructing the compatible split expectation for a concrete detector protocol.

F.5 Falsifier sharpened by the example

The matroid layer of CTVM carries information beyond rank only if physically realizable detector protocols can produce non-uniform matroids. The example shows that non-uniform matroids exist as mathematical objects in the space of possible response matrices. It therefore sharpens the following falsifier:

If every physically realizable detector protocol, including configurations accessible to the bosonic and fermionic arms of the analogue-vacuum experimental program, produces only uniform matroids of type $U_{3,N}$ or their obvious higher-rank analogues, then the

matroidal sector-bookkeeping layer carries no information beyond rank and is physically empty.

Conversely, observation of stable non-uniform matroid structures—parallel pairs, larger circuits, or sector-correlated circuit strata—in topologically nontrivial regimes such as near Weyl points in ${}^3\text{He-A}$ or in boundary harvesting from a ${}^3\text{He-B}$ -like gapped phase would be evidence that vacuum causal topology can imprint itself on the matroid layer. The example establishes neither outcome. It establishes only that the question is substantive.

F.6 Verification by exhaustive pair enumeration

A simple numerical check confirms the codimension-one statement. Keeping $\text{Im } z_{23} = 0$ while varying the nonzero ratio $\text{Im } z_{13} / \text{Im } z_{12}$ leaves the parallel pair $\{5, 6\}$ intact. Moving $\text{Im } z_{23}$ away from zero breaks the parallel pair. Varying the real parts while holding $\text{Im } z_{23} = 0$ also leaves the pair intact. Figure 2 gives a finite-tolerance view of this codimension-one wall.



Figure 2: Numerical visualization of the $\{5, 6\}$ parallel-pair stratum in the three-detector example. Yellow points mark the finite-tolerance neighbourhood of the $\{5, 6\}$ parallel-pair stratum; dark points mark parameter values where the realified response matroid remains the generic uniform matroid $U_{3,6}$. The exact stratum is the codimension-one wall $\text{Im } z_{23} = 0$ with $\text{Im } z_{12}$ and $\text{Im } z_{13}$ nonzero. The finite width is a plotting tolerance, not an additional physical parameter.

The following script implements the check.

```

def matroid_parallel_pairs(z12, z13, z23):
    a = b = c = 1.0
    X = np.array([[a, z12, z13],
                  [np.conj(z12), b, z23],
                  [np.conj(z13), np.conj(z23), c]], dtype=complex)
    Xhat = np.hstack([X.real, X.imag])
    pairs = []
    for i in range(6):
        for j in range(i+1, 6):
            vi, vj = Xhat[:, i], Xhat[:, j]
            if (np.linalg.norm(vi) > 1e-10
                and np.linalg.norm(vj) > 1e-10
                and np.linalg.matrix_rank(np.column_stack([vi, vj]),
                                          tol=1e-10) < 2):
                pairs.append((i+1, j+1))
    return pairs

```

The verification is modest but useful: it corrects the codimension count and gives a concrete check that the matroidal type stratifies response-parameter space.

F.7 Three-detector CSS miniature

The following binary matrices are illustrative placeholders. A physical derivation would threshold a computed response matrix from, for example, three Unruh–DeWitt detectors coupled to a scalar field with a specified smearing, switching, gap, and regulator protocol. That calculation is left as a future executable target, together with the refined operational column-partition problem of Open Target T6.

The three-detector example can also be read as a minimal warning about residual coding. In the realified response matrix, the parallel pair $\{5, 6\}$ is a weight-two circuit. If this circuit is assigned to the residual block and then thresholded into a binary parity-check representation, it produces the lowest possible nontrivial dependency. In a CSS interpretation this is dangerous: a weight-two circuit can support a low-weight undetected pattern or a low-distance check, depending on how the parity-check matrices are chosen.

A toy binary reduction of the two parallel columns has the schematic form

$$B_{\parallel} = \begin{pmatrix} 1 & 1 \\ 0 & 0 \\ 0 & 0 \end{pmatrix}, \quad (142)$$

whose column matroid has the circuit $\{1, 2\}$. Without additional residual bias, the corresponding girth is two. If a controlled residual-sector bias removes one of the two columns from the residual binary support or lifts the exact parallelism by moving the protocol off the wall $\text{Im } z_{23} = 0$, the weight-two circuit disappears. In the smallest possible three-column toy replacement,

$$B_3 = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{pmatrix}, \quad (143)$$

all two-column subsets are independent and the minimal circuit has weight three. Thus the girth

increases from two to three.

This example should not be read as a physical derivation of a CSS code. It shows only the local mechanism by which residual circuit suppression can raise the minimum circuit weight. To become a genuine CSS code, the binary residual data must be embedded into matrices H_X, H_Z satisfying $H_X H_Z^T = 0$. The example illustrates the escape from the ‘‘GHZ trap’’ in the following limited sense: a fragile low-weight dependency can be replaced by a higher-weight dependency only when the residual protocol supplies an exact selection rule or a stable bias in the binary support.

G Null-boundary normalization sensitivity

The coefficient $\kappa_{\text{edge}} = 1/(4\pi)$ is not a standalone theorem of the null-boundary action. It is the value obtained after choosing the causal-diamond normalization of Hypothesis 12.1 and the edge-action/vacuum-energy dictionary of Hypothesis 12.2. The following catalogue makes the sensitivity explicit.

Choice	Effect on boundary term	Consequence for CTVM
Killing/horizon normalization	$\kappa_N \sim H$ over parameter range $\Delta\lambda \sim H^{-1}$ on two null sheets	gives $I/A_H \sim 1/(4\pi G_N)$ under the adopted causal-diamond convention
Affine normalization on a null segment	$\kappa_N = 0$ on that segment	shifts information into joint/corner terms; cannot by itself fix κ_{edge}
Rescaling $k^a \mapsto e^\beta k^a$	changes κ_N by a derivative of β along generators	changes boundary and joint bookkeeping unless the counterterm prescription is fixed
Different diamond cutoff or tip regularisation	changes the effective integration range and joint contribution	can shift the numerical coefficient unless controlled by a physical edge theory

Thus the falsifiable object is not merely the algebraic factor $1/(8\pi G_N)$ in the gravitational action. It is the complete edge dictionary: null normalization, joint-term convention, coarse-grained edge action, and conversion to the bulk density parameterization $\rho_\Lambda = \kappa_{\text{edge}} f_{\text{edge}} m_{\text{Pl}}^2 H^2$. A robust microscopic completion should show that reasonable regulator choices flow to the same dimensionless coefficient, or else quantify the allowed spread of κ_{edge} .

H Vector-channel response calculation: from Proca toy model to

$f_{\text{vec}}^{\text{eff}}$

The benchmark

$$f_{\text{eq}} \simeq \frac{27}{106.75} \tag{144}$$

is not derived by polarization counting alone in the present framework. The required calculation is a response-weighted, species-resolved field-theory calculation across the electroweak transition. This appendix states the calculation in a form suitable for analytic and numerical implementation.

The central object is an effective vector response fraction

$$f_{\text{vec}}^{\text{eff}}(T; \mathcal{P}) = \frac{\sum_{a \in \text{vectors}} R_a^{\mathcal{P}}(T)}{\sum_{b \in \text{all channels}} R_b^{\mathcal{P}}(T)}, \quad (145)$$

where \mathcal{P} denotes the detector protocol: current coupling, smearing, switching, detector gaps, regulator choices, and post-processing convention. The CTVM benchmark becomes the test

$$f_{\text{vec}}^{\text{eff}}(T_{\text{EW}}; \mathcal{P}) \xrightarrow{?} \frac{27}{106.75}. \quad (146)$$

A Proca calculation supplies only a controlled vector-channel toy model. It does not by itself compute the full denominator R_{tot} , and therefore cannot by itself establish the Standard Model benchmark.

H.1 Protocol gate: detector-current model

The vector propagator is not itself the detector response. A response kernel is obtained only after contracting the vector spectral density with a detector current or form factor. Before any numerical calculation, one must choose the detector-current model.

Three useful choices are:

1. Conserved-current coupling

$$H_{\text{int}} = g \int d\tau \mu(\tau) J_\mu(\tau) A^\mu(x(\tau)), \quad p_\mu J^\mu(p) = 0. \quad (147)$$

This is the cleanest option for avoiding unphysical longitudinal enhancement.

2. Field-strength or dipole coupling

$$H_{\text{int}} = g \int d\tau \mu(\tau) d^{\mu\nu}(\tau) F_{\mu\nu}(x(\tau)). \quad (148)$$

This coupling is gauge-invariant and is the preferred route for electroweak-like extensions.

3. Proca polarization probe

$$H_{\text{int}} = g \int d\tau \mu(\tau) e_\mu A^\mu(x(\tau)). \quad (149)$$

This is acceptable for a pure massive Proca toy model. The longitudinal mode is then physical. The $m \rightarrow 0$ limit is singular unless the current is conserved, so this option should not be interpreted as Maxwell theory.

All numerical values below are protocol-dependent until this choice is fixed. For the first electroweak-relevant executable calculation, the default choice should be the conserved-current or field-strength coupling, not the unconstrained Proca polarization probe. The Proca polarization probe remains useful as a toy model for physical longitudinal massive-vector modes, but it cannot define the gauge-invariant Standard Model response fraction by itself.

H.2 Massive-vector spectral kernel

As a first controlled toy model, take a massive Proca field in flat spacetime with

$$\Omega_k = \sqrt{k^2 + m^2}, \quad P_T^{ij} = \delta^{ij} - \hat{k}^i \hat{k}^j, \quad P_L^{ij} = \hat{k}^i \hat{k}^j. \quad (150)$$

With the metric and sign conventions fixed in the main text, the spatial retarded kernel may be represented schematically as

$$G_R^{ij}(\omega, \mathbf{k}) = \frac{P_T^{ij} + (\Omega_k^2/m^2)P_L^{ij}}{(\omega + i0^+)^2 - \Omega_k^2}. \quad (151)$$

The associated spectral density is

$$\rho^{ij}(\omega, \mathbf{k}) = -2 \operatorname{Im} G_R^{ij}(\omega, \mathbf{k}). \quad (152)$$

Equivalently, using a pole prescription,

$$\rho^{ij}(\omega, \mathbf{k}) = \frac{\pi}{\Omega_k} \left[P_T^{ij} + \frac{\Omega_k^2}{m^2} P_L^{ij} \right] [\delta(\omega - \Omega_k) - \delta(\omega + \Omega_k)], \quad (153)$$

up to the global sign convention used for G_R . Numerical work should replace the delta function by a Lorentzian regulator

$$\delta_\eta(x) = \frac{1}{\pi} \frac{\eta}{x^2 + \eta^2} \quad (154)$$

and verify stability as $\eta \rightarrow 0$.

H.3 Detector response matrix

The vector contribution to the detector response matrix is

$$X_{ij}^{(V)} = g_i g_j \int \frac{d\omega}{2\pi} \frac{d^3k}{(2\pi)^3} \mathcal{J}_i^\mu(\omega, \mathbf{k}) \rho_{\mu\nu}^{(V)}(\omega, \mathbf{k}) \mathcal{J}_j^\nu(\omega, \mathbf{k})^* \mathcal{F}_B(\omega; T), \quad (155)$$

where

$$\mathcal{F}_B(\omega; T) = 1 + n_B(\omega; T) \quad (156)$$

for a bosonic channel. The detector-current form factor contains the current orientation, switching, smearing, detector gap, and detector position. For example,

$$\mathcal{J}_i^\mu(\omega, \mathbf{k}) = e_i^\mu \tilde{\chi}_i(\omega - \Delta_i) \tilde{F}_i(\mathbf{k}) e^{i\mathbf{k}\cdot\mathbf{x}_i}. \quad (157)$$

For Gaussian switching and smearing,

$$\tilde{\chi}_i(\omega - \Delta_i) = T_i \sqrt{\pi} \exp \left[-\frac{T_i^2}{4} (\omega - \Delta_i)^2 \right], \quad (158)$$

$$\tilde{F}_i(\mathbf{k}) = \exp \left[-\frac{r_{s,i}^2 k^2}{2} \right]. \quad (159)$$

The response matrix $X^{(V)}$ is then realified using the fixed column-doubling convention

$$\hat{X}^{(V)} = [\text{Re } X^{(V)} \mid \text{Im } X^{(V)}]. \quad (160)$$

H.4 Species-resolved response fraction

For a channel a , define the protocol-weighted response

$$R_a^{\mathcal{P}}(T) = \sum_{\lambda \in a} g_{a,\lambda} \int \frac{d^3 k}{(2\pi)^3} \mathcal{W}_{a,\lambda}^{\mathcal{P}}(\Omega_{a,k}, k) \rho_{a,\lambda}(\Omega_{a,k}, k; T) \mathcal{F}_a(\Omega_{a,k}; T). \quad (161)$$

Here a labels field species or effective channels, λ labels polarization or spin states, $g_{a,\lambda}$ is the internal degeneracy, $\rho_{a,\lambda}$ is the relevant spectral density, and $\mathcal{W}_{a,\lambda}^{\mathcal{P}}$ is the detector/protocol window function. The statistics factor is $\mathcal{F}_B = 1 + n_B$ for bosonic stimulation and, for fermionic final-state access, $\mathcal{F}_F = 1 - n_F$.

The effective vector fraction is

$$f_{\text{vec}}^{\text{eff}}(T; \mathcal{P}) = \frac{\sum_{a \in \text{vectors}} R_a^{\mathcal{P}}(T)}{\sum_{b \in \text{all channels}} R_b^{\mathcal{P}}(T)}. \quad (162)$$

The comparison with the CTVM benchmark is therefore a dynamical test:

$$f_{\text{vec}}^{\text{eff}}(T_{\text{EW}}; \mathcal{P}) \stackrel{?}{\simeq} \frac{27}{106.75}. \quad (163)$$

Bare polarization counting gives a target. The response integral decides whether that target is dynamically supported.

H.5 Regulator and protocol scan

The smearing radius r_s and switching time T_{sw} should be treated as physical protocol scales, not merely disposable regulators. The ultraviolet cutoff Λ may be removed only at fixed r_s, T_{sw} if the response is stable.

A minimal scan is

$$\Lambda/\Delta \in \{10, 20, 50, 100\}, \quad r_s \Delta \in \{0.1, 0.3, 1.0\}, \quad T_{\text{sw}} \Delta \in \{0.5, 1.0, 2.0\}. \quad (164)$$

Report

$$f_{\text{vec}}^{\text{eff}} = \bar{f} \pm \sigma_{\Lambda} \pm \sigma_{r_s} \pm \sigma_{T_{\text{sw}}} \pm \sigma_{\text{gauge}}. \quad (165)$$

The benchmark is classified as robust if total systematic variation is below 5%, qualified if the variation is between 5% and 20%, and not supported if variation exceeds 20% or if gauge or longitudinal sensitivity dominates.

H.6 Longitudinal and gauge-consistency audit

For a pure Proca toy model, the longitudinal mode is physical. Compute

$$R_T, \quad R_L, \quad \frac{R_L}{R_T + R_L}. \quad (166)$$

If R_L diverges or dominates as $m/\Delta \rightarrow 0$, check whether the chosen detector current violates current conservation.

For an electroweak-like extension, one should repeat the calculation in a covariant R_ξ gauge and include the Goldstone sector. The physical response must be independent of the gauge parameter:

$$\left| \frac{\partial f_{\text{vec}}^{\text{eff}}}{\partial \xi} \right| \ll 1 \quad (167)$$

within numerical tolerance after all physical contributions are included. If this fails, the calculation has not isolated a physical vector-channel response.

H.7 Kramers-Kronig and causality diagnostics

Finite grids and UV cutoffs require subtracted Kramers-Kronig tests. Use

$$\text{Re } G_R(\omega) - \text{Re } G_R(0) = \frac{2\omega^2}{\pi} \text{P} \int_0^\infty d\omega' \frac{\text{Im } G_R(\omega')}{\omega'(\omega'^2 - \omega^2)}. \quad (168)$$

Define the residual

$$\epsilon_{\text{KK}} = \frac{\|G_R^{\text{direct}} - G_R^{\text{KK}}\|}{\|G_R^{\text{direct}}\|}. \quad (169)$$

The target is $\epsilon_{\text{KK}} < 10^{-3}$ away from broadened pole neighborhoods. In the time domain, the retarded transform should vanish outside the retarded domain within numerical tolerance.

H.8 Numerical performance hint

Nested adaptive quadrature in ω and k against a Lorentzian-regulated near-delta spectral density is often slow and unreliable. Two safer strategies are recommended. First, near each pole use the change of variables

$$\omega = \Omega_k + \eta \tan \theta \quad (170)$$

and apply fixed Gauss-Legendre quadrature in θ over a finite window. Second, when Gaussian switching is used, perform the ω integral analytically: the Lorentzian regulator convolved with the Gaussian switching profile gives a Voigt-type profile, leaving only the k integral for numerical evaluation. The second route is preferred for production scans.

H.9 Matroid split diagnostic

For each computed response matrix, form

$$\widehat{X} = [\text{Re } X \mid \text{Im } X]. \quad (171)$$

Numerical rank tests define an ϵ -diagnostic, not an exact matroid unless symbolic or certified arithmetic is available. For a candidate column subset S , let

$$\sigma_1(S) \geq \dots \geq \sigma_r(S) \quad (172)$$

be its singular values and define

$$\text{rk}_\epsilon(S) = \#\{a : \sigma_a(S) > \epsilon \sigma_1(\widehat{X})\}. \quad (173)$$

A subset C is a circuit candidate if

$$\text{rk}_\epsilon(C) < |C| \quad (174)$$

while

$$\text{rk}_\epsilon(C \setminus \{e\}) = |C| - 1 \quad \forall e \in C. \quad (175)$$

Use two thresholds,

$$\epsilon_{\text{low}} = 10^{-8}, \quad \epsilon_{\text{high}} = 10^{-5}, \quad (176)$$

and classify any candidate with singular values in the gray zone as unresolved. Exact detector/split compatibility passes only if no stable circuit candidate intersects both the proposed split block and the proposed residual block.

If all physically realizable protocols produce only uniform matroids, the matroid layer carries little information beyond rank. Conversely, stable non-uniform strata, such as parallel pairs or larger circuit structures, provide operational evidence that detector protocols can imprint nontrivial vacuum-response structure on the matroid layer.

H.10 Species-resolved denominator plan

The Proca toy gives a controlled vector numerator. The denominator requires a staged program. At the toy level one compares a representative scalar channel, fermionic channel, and vector channel:

$$f_{\text{vec, toy}}^{\text{eff}} = \frac{R_V^{\mathcal{P}}}{R_S^{\mathcal{P}} + R_F^{\mathcal{P}} + R_V^{\mathcal{P}}}. \quad (177)$$

Below electroweak symmetry breaking the vector numerator is schematically

$$R_{\text{vec}}^{\text{broken}} = 16R_g + 2R_\gamma + \sum_{V=W^+, W^-, Z} (2R_{V,T} + R_{V,L}). \quad (178)$$

Across the electroweak transition the calculation must use a temperature-dependent interpolation of spectral densities, masses, occupation factors, and detector/protocol weights:

$$R_a^{\mathcal{P}}(T) = \sum_\lambda g_{a,\lambda}(T) \int \frac{d^3k}{(2\pi)^3} \mathcal{W}_{a,\lambda}^{\mathcal{P}} \rho_{a,\lambda}(\Omega_{a,k}, k; T) \mathcal{F}_a(\Omega_{a,k}; T). \quad (179)$$

The test is whether $f_{\text{vec}}^{\text{eff}}(T; \mathcal{P})$ develops a stable plateau or freeze-out value in the target range $[0.20, 0.30]$, with 27/106.75 retained only as the historical reference point. A true electroweak calculation must include Goldstone and Higgs-sector contributions consistently and verify gauge-parameter independence.

H.11 Reporting rule

The Proca calculation tests whether massive vector channels are dynamically favored by the chosen detector protocol. It does not derive the Standard Model benchmark by itself. The benchmark range $[0.20, 0.30]$ is supported only if the full species-resolved response integral flows into that range across a regulator-stable, gauge-consistent protocol window; 27/106.75 remains the historical reference point.

H.12 Longitudinal asymptotic near the electroweak transition

This subsection records a toy asymptotic used to diagnose longitudinal-mode contamination in the vector-channel response calculation. It is not a Standard Model electroweak computation. It is a Proca-like saddle estimate showing why gauge-consistent current or field-strength couplings are required before any vector-channel benchmark can be trusted.

Assume a Gaussian switching window of width σ , detector gap E , and a temperature-dependent vector mass with mean-field scaling

$$m(T) \simeq m_0 \sqrt{1 - (T/T_{\text{EW}})^2}, \quad T \rightarrow T_{\text{EW}}^- \quad (180)$$

For a Proca polarization probe that couples directly to the longitudinal polarization, the spectral kernel carries a factor of order $k^2/m(T)^2$ or $\omega_k^2/m(T)^2$. The Gaussian window localizes the integral near $\omega_k \simeq E$. A saddle evaluation then gives the leading behavior

$$\mathcal{F}_L(T \rightarrow T_{\text{EW}}^-) \simeq \frac{\sqrt{\pi}|c|^2}{3} \frac{E^3 \sigma}{m(T)^2} \frac{1}{1 - e^{-E/T}}, \quad (181)$$

up to protocol-dependent normalization and subleading terms in the switching width.

The $1/m(T)^2$ factor is the finite-time footprint of the Goldstone-equivalence behavior in a nonconserved Proca probe. In a gauge-consistent electroweak calculation, longitudinal vector and Goldstone contributions must combine so that gauge-parameter dependence cancels. Therefore a large value of \mathcal{F}_L in the Proca toy is not evidence for the dark-matter benchmark. It is a warning that the detector-current model must be fixed before interpreting $f_{\text{vec}}^{\text{eff}}$.

If the longitudinal/Goldstone leakage is assigned to a residual topological channel with suppression factor

$$F_\nu^2 \sim (1 + \chi \nu_{\text{res}})^{-1}, \quad (182)$$

then the effective contribution is schematically $F_\nu^2 \mathcal{F}_L$. This is a calculable hypothesis linking Appendix H to the residual-code ansatz of Section 9.6. It must be checked by the same gauge-consistency audit that tests the vector-channel response integral.

I Theorem and proposition audit table

Appendix I is an audit table, not a second theorem list. Its purpose is to identify which claims are definitions, which are conditional implications, which are benchmarks, and which remain open targets. The table also separates several items that are easy to conflate: AdS-like holographic entropy versus cosmological-horizon holography, background residual decoupling versus perturbation-order decoupling, and the null-boundary normalization convention versus the edge-action/vacuum-energy dictionary.

Claim	Depends on	Status	Failure mode
SK response idempotence	Chosen CTP cumulant representation	Proven as linear algebra on cumulants	Fails only if promoted to a state map.

Claim	Depends on	Status	Failure mode
Hypothesis 4.1: modular-compatible split expectation	Split inclusion plus $\sigma_t^\omega(\mathcal{N}_{\text{reg}}) = \mathcal{N}_{\text{reg}}$ after regulator	Hypothesis; deepest architectural gate	No ω -preserving conditional expectation exists for the chosen regulated split factor.
Construction 4.1: regulated cosmological-horizon split expectation	Static-patch KMS state, stretched-horizon regulator, type-I split factor, modular covariance	Open construction; gates dark-energy application	Candidate split factor is not invariant under regulated modular flow, or detector/source data fail regulator-stability.
Construction 4.2: observer-inclusive type-II crossed-product algebra	Static-patch KMS state, crossed product or observer dressing, finite type-II trace, Umegaki conditional expectation onto a unital subalgebra	Preferred construction target; the trace-preserving expectation exists algebraically once the finite trace algebra and unital subalgebra are fixed	No physically correct split-sector subalgebra is identified, a required nonracial physical state is not preserved, or detector/stress-response data fail regulator stability.
Hypothesis 7.1: detector/split compatibility	Regulated detector feature map, split image/kernel decomposition	Hypothesis; partially modeled only in free massive scalar example	Off-block response remains large or cross-block circuits persist for physical protocols.
Hypothesis 7.2: operational/matroid compatibility	Operational EPR-active column partition and realified response matroid	Hypothesis; diagnostic	A nontrivial operational partition has stable cross-sector circuits.
Operational EPR witness	Detector protocol, logarithmic negativity, threshold	Definition; computable witness of record	Different thresholds or higher-dimensional detectors can change the operational partition; PPT-bound entanglement may be missed.
Residual binary reduction and CSS code datum	Thresholded residual response block, binary support matrix, CSS orthogonality $H_X H_Z^T = 0$	Conditional construction target	Binary support is regulator-unstable or no CSS-compatible pair of check matrices exists.
Hypothesis 9.7: residual circuit suppression	Binary residual matroid, residual topological index, circuit-density scaling	Hypothesis; coding-theoretic strengthening of residual decoupling	Low-weight residual or cross-sector circuits persist at finite density, preventing code-distance growth.

Claim	Depends on	Status	Failure mode
Addressability trade-off	Cross-sector circuit density, cross-circuit rank, interface coupling	Conditional error-rate ansatz	Logical error rate cannot be made small while retaining controlled addressability.
Matroid direct sum	Exact block-diagonal realified representation	Conditional proposition	Any cross-sector circuit falsifies exact split assignment.
Statistics changes weights, not rank	Nonzero column rescalings	Proven for representable matroids	Exact zeros, exact degeneracies, or exact block decoupling can change the matroid.
Hypothesis 9.1: pairwise response-network dictionary	Split response block, weight layer, capacity rule, regulator window	Hypothesis	No regulator-stable capacity assignment exists.
Hypothesis 9.2: AdS-like holographic entropy dictionary	Response network and controlled RT/HRT or bit-thread regime	Conditional hypothesis	Network cut function cannot be embedded in a controlled holographic entropy setting.
Hypothesis 9.3: Static-patch, stretched-cosmological-horizon holography extension	Static-patch, stretched-horizon, bilayer, DSSYK-like, type-II/crossed-product, or other de Sitter-inspired dictionary	Additional conjecture	No well-defined cosmological-horizon entropy dictionary supports the CTVM edge source.
Hypothesis 9.4: entropy-cone admissibility	Network entropy vector $A \mapsto S_{\text{net}}(A)$	Testable consistency condition	Strong subadditivity, MMI, or other required entropy inequalities fail.
Hypothesis 9.5: semiclassical sourcing dictionary	Entropy/relative-entropy variation, modular Hamiltonian, renormalised stress tensor	Conditional hypothesis	No conserved local stress response follows from the entropy functional.
Hypothesis 9.6: residual decoupling at background order	Source functional and residual response coordinates	Conditional hypothesis	Residual response changes the background source functional.
Hypothesis 9.6: residual decoupling at perturbation order	Linearised source response and residual perturbations	Additional condition needed for CDM-like growth	Residual sector produces scale-dependent or anisotropic growth deviations.
Hypothesis 10.1: topological boundary invariant	Boundary response operator or edge Hamiltonian carrying \mathbb{Z} , \mathbb{Z}_2 , winding, Chern, or K -theory class	Hypothesis	No protected invariant acts faithfully on the detector-response representation.

Claim	Depends on	Status	Failure mode
Hypothesis 12.2: coarse-grained edge-action to vacuum-energy dictionary	Boundary action density I_{edge}/A_H , horizon-regulated bulk density	Hypothesis; deepest dark-energy normalization gate	No edge theory produces the dimensional bridge to $\rho_\Lambda = \kappa_{\text{edge}} f_{\text{edge}} M_{\text{Pl}}^2 H^2$.
Hypothesis 12.1: causal-diamond null normalization	Null generator normalization, affine/Killing parameter, joint terms, counterterms	Convention-dependent hypothesis	Reasonable conventions shift κ_{edge} enough to spoil the benchmark.
Hypothesis 6.1: statistics-resolved sector bias	Bosonic and fermionic response kernels, Bose/Pauli factors, analogue-vacuum protocols	Hypothesis; experimentally testable	Bosonic/fermionic channels do not show the predicted response bias.
Appendix H: vector-channel response integral	Species-resolved retarded kernels and detector form factors	Open calculation	$f_{\text{vec}}^{\text{eff}}(T_{\text{EW}}; \mathcal{P})$ does not enter $[0.20, 0.30]$ in a stable protocol window.
Section 11.3: two-fluid exchange term and locking rate	Q, Γ_{lock} , vector-response kernels	Phenomenological input	No microscopic response calculation yields the assumed locking dynamics.
Section 12.3: fermionic edge blocking	$b_{\text{edge}}(T, \mu)$, edge charge, chemical potential, topological saturation	Conditional mechanism	No edge microphysics drives $b_{\text{edge}} \rightarrow 0$.
Section 13: late-time $b_{\text{edge}}(N)$ history	Transition template, edge kinetics, background expansion	Phenomenological template	No viable $w_{\text{eff}}(z)$ history compatible with data.
Section 14: QCD-to-horizon-edge matching	Topological susceptibility, horizon edge modes, effective action	Speculative correspondence	No effective-action matching exists beyond dimensional numerology.
Appendix E: look-elsewhere caution for 27/106.75	Integer numerators, alternative denominators, weighted ratios	Caution, not derivation	Numerical proximity loses evidential value absent response dynamics.

J Executable critical-path protocols

This appendix converts the cheapest decisive tests of Section 17.1 into compact protocols. They are meant to be run before investing effort in the frontier gates.

J.1 Certified matroid split test

Input: a harvested detector state $\rho_D^{(\mathcal{P})}$, a response matrix $X^{(\mathcal{P})}$, an EPR-witness threshold ϵ_E , and rank tolerances

$$\epsilon_{\text{low}} = 10^{-8}, \quad \epsilon_{\text{high}} = 10^{-5}.$$

Algorithm.

1. Compute the pairwise logarithmic negativities

$$E_{\mathcal{N}}^{(\mathcal{P})}(i, j) = \log_2 \left\| \left(\rho_{ij}^{(\mathcal{P})} \right)^{T_i} \right\|_1.$$

2. Construct the pair support $\mathcal{E}_{\text{pair}}^{\mathcal{P}, \epsilon}$ and the operational column partition $E_{\text{split}}^{(\mathcal{P}, \epsilon)} \sqcup E_{\text{res}}^{(\mathcal{P}, \epsilon)}$.
3. Realify X by column-doubling, $\hat{X} = [\Re X \mid \Im X]$.
4. Enumerate candidate circuits C up to size $r + 1$, where $r = \text{rank}(\hat{X})$.
5. Certify C as a stable circuit only if

$$\sigma_{\min}(\hat{X}_C) < \epsilon_{\text{low}} \sigma_1(\hat{X})$$

and every proper deletion satisfies

$$\sigma_{\min}(\hat{X}_{C \setminus e}) > \epsilon_{\text{high}} \sigma_1(\hat{X}), \quad e \in C.$$

6. Output PASS if no certified circuit crosses the partition, FAIL if a certified crossing circuit exists, and NON-DIAGNOSTIC if no nontrivial operational partition exists.

A FAIL result falsifies the exact sector assignment for that protocol. A PASS result does not prove gravitational inertness; it only says the data do not contradict the proposed split.

J.2 Entropy-cone admissibility test

Input: a weighted graph $G_{\text{split}} = (V, E, c)$ built from capacities

$$c_{ij} = \alpha_{\mathcal{P}} E_{\mathcal{N}}^{(\mathcal{P})}(i, j), \quad \alpha_{\mathcal{P}} > 0.$$

Algorithm.

1. For every boundary subset $A \subseteq \{1, \dots, n\}$ compute the minimum cut

$$S_{\text{net}}(A) = \min_{\gamma_A} \sum_{e \in \gamma_A} c(e).$$

2. Check subadditivity and strong subadditivity.
3. In an AdS-like static holographic interpretation, also check monogamy of mutual information (MMI, $I_3 \leq 0$),

$$S(AB) + S(AC) + S(BC) \geq S(A) + S(B) + S(C) + S(ABC).$$

4. Record every inequality margin, normalized by total network capacity.

Failure of the entropy-cone test kills the RT/HRT interpretation of that network. Passing is necessary but not sufficient for a holographic dual.

Toy numerical demonstration. As a smoke test, take a four-boundary complete graph with capacities

$$c_{12} = c_{34} = 1, \quad c_{13} = c_{24} = 0.5, \quad c_{14} = c_{23} = 0.2.$$

Using the ordinary graph-cut entropy $S(A) = \sum_{e \in \partial A} c(e)$, the singleton entropies are all 1.7, the pair entropies are

$$S(12) = S(34) = 1.4, \quad S(13) = S(24) = 2.4, \quad S(14) = S(23) = 3.0,$$

and the triplet entropies again equal 1.7. The subadditivity, strong-subadditivity, and MMI margins are nonnegative, with the smallest margins zero up to numerical roundoff, as expected for a positive undirected cut function.

To confirm that the checker is non-vacuous, feed it instead the artificial entropy vector

$$S(A) = S(B) = S(C) = 1, \quad S(AB) = S(AC) = S(BC) = 1.2, \quad S(ABC) = 1.2.$$

The strong-subadditivity margins are positive (minimum 0.2), but the MMI margin is

$$S(AB) + S(AC) + S(BC) - S(A) - S(B) - S(C) - S(ABC) = -0.6,$$

so this vector fails the static holographic entropy-cone test. The example shows that the test is a real falsifier: a proposed network entropy can pass ordinary submodularity while still failing the additional holographic inequality.

J.3 $f_{\text{edge}}(z)$ tracking pre-filter

Use the template

$$f_{\text{edge}}(N) = \frac{f_{\infty}}{2} \left[1 + \tanh \left(\frac{N - N_t}{\Delta} \right) \right], \quad \Omega_{\Lambda}(z) = \frac{8\pi}{3} \kappa_{\text{edge}} f_{\text{edge}}(z).$$

Discard any parameter set for which

$$\Omega_{\Lambda}(z \gtrsim 10^3) > \Omega_{\text{EDE}}^{\text{max}}, \quad \Omega_{\text{EDE}}^{\text{max}} = 0.02$$

in the default pre-filter. Surviving parameter sets must then be tested against CMB distance priors or a full likelihood, BAO distances, supernova distance moduli, and growth observables.

J.4 Vector-channel response protocol

The vector-response calculation is pre-registered as follows.

1. The physically meaningful vector calculation uses conserved-current or field-strength coupling. A Proca polarization probe is retained only as a toy diagnostic.

2. Compute species- and polarization-resolved responses $R_a^{\mathcal{P}}(T)$ and form

$$f_{\text{vec}}^{\text{eff}}(T; \mathcal{P}) = \frac{\sum_{a \in \text{vectors}} R_a^{\mathcal{P}}(T)}{\sum_{b \in \text{all channels}} R_b^{\mathcal{P}}(T)}.$$

3. Scan Λ/Δ , $r_s\Delta$, and $T_{\text{sw}}\Delta$.
4. Apply Definition 1.1.
5. Check Goldstone consistency: the longitudinal W/Z response must not double-count the Higgs-sector Goldstones already present in the unbroken-phase denominator.

If $f_{\text{vec}}^{\text{eff}}$ does not enter the target range $[0.20, 0.30]$ in a gauge-consistent, regulator-stable window, the dark-matter benchmark is unsupported.

References

- [1] R. Haag, *Local Quantum Physics: Fields, Particles, Algebras*, 2nd ed. Springer, 1996.
- [2] H. Araki, *Mathematical Theory of Quantum Fields*. Oxford University Press, 1999.
- [3] J. J. Bisognano and E. H. Wichmann, “On the duality condition for a Hermitian scalar field,” *J. Math. Phys.* 16, 985 (1975).
- [4] J. J. Bisognano and E. H. Wichmann, “On the duality condition for quantum fields,” *J. Math. Phys.* 17, 303 (1976).
- [5] S. Doplicher and R. Longo, “Standard and split inclusions of von Neumann algebras,” *Invent. Math.* 75, 493–536 (1984).
- [6] D. Buchholz and E. H. Wichmann, “Causal independence and the energy-level density of states in local quantum field theory,” *Commun. Math. Phys.* 106, 321–344 (1986).
- [7] M. Takesaki, “Conditional expectations in von Neumann algebras,” *J. Funct. Anal.* 9, 306–321 (1972).
- [8] E. Witten, “APS Medal for Exceptional Achievement in Research: Invited article on entanglement properties of quantum field theory,” *Rev. Mod. Phys.* 90, 045003 (2018).
- [9] L. V. Keldysh, “Diagram technique for nonequilibrium processes,” *Sov. Phys. JETP* 20, 1018 (1965).
- [10] B. L. Hu and E. Verdaguer, “Stochastic gravity: Theory and applications,” *Living Rev. Relativ.* 11, 3 (2008).
- [11] B. Reznik, “Entanglement from the vacuum,” *Found. Phys.* 33, 167 (2003).
- [12] B. Reznik, A. Retzker, and J. Silman, “Violating Bell’s inequalities in vacuum,” *Phys. Rev. A* 71, 042104 (2005).
- [13] A. Pozas-Kerstjens and E. Martin-Martinez, “Harvesting correlations from the quantum vacuum,” *Phys. Rev. D* 92, 064042 (2015).

- [14] T. R. Perche, C. Lima, and E. Martin-Martinez, “Harvesting entanglement from complex scalar and fermionic fields with linearly coupled particle detectors,” *Phys. Rev. D* 105, 065016 (2022).
- [15] S. Ryu and T. Takayanagi, “Holographic derivation of entanglement entropy from AdS/CFT,” *Phys. Rev. Lett.* 96, 181602 (2006).
- [16] M. Freedman and M. Headrick, “Bit threads and holographic entanglement,” *Commun. Math. Phys.* 352, 407–438 (2017).
- [17] L. Lehner, R. C. Myers, E. Poisson, and R. D. Sorkin, “Gravitational action with null boundaries,” *Phys. Rev. D* 94, 084046 (2016).
- [18] T. Faulkner, R. G. Leigh, O. Parrikar, and H. Wang, “Modular Hamiltonians for deformed half-spaces and the averaged null energy condition,” *JHEP* 09, 038 (2016).
- [19] Planck Collaboration, N. Aghanim et al., “Planck 2018 results. VI. Cosmological parameters,” *Astron. Astrophys.* 641, A6 (2020).
- [20] E. W. Kolb and M. S. Turner, *The Early Universe*. Addison-Wesley, 1990.
- [21] J. M. Cornwall, D. N. Levin, and G. Tiktopoulos, “Derivation of gauge invariance from high-energy unitarity bounds on the S matrix,” *Phys. Rev. D* 10, 1145 (1974).
- [22] B. W. Lee, C. Quigg, and H. B. Thacker, “Weak interactions at very high energies: The role of the Higgs-boson mass,” *Phys. Rev. D* 16, 1519 (1977).
- [23] M. S. Chanowitz and M. K. Gaillard, “The TeV physics of strongly interacting W’s and Z’s,” *Nucl. Phys. B* 261, 379 (1985).
- [24] G. E. Volovik, *The Universe in a Helium Droplet*. Oxford University Press, 2003.
- [25] G. E. Volovik, “The Superfluid Universe,” arXiv:1004.0597 (2010; revised 2012).
- [26] G. E. Volovik, “³He Universe 2020,” *J. Low Temp. Phys.* 202, 11–33 (2021).
- [27] G. E. Volovik and M. A. Zubkov, “Standard model as the topological material,” *New J. Phys.* 19, 015009 (2017).
- [28] Y.-K. Lu, Y. Margalit, and W. Ketterle, “Observation of bosonic stimulation in light scattering,” arXiv:2204.06639 (2022).
- [29] Y. Margalit, Y.-K. Lu, F. C. Top, and W. Ketterle, “Pauli blocking of light scattering in degenerate fermions,” *Science* 374, 976–979 (2021).
- [30] R. Jannin, Y. van der Werf, K. Steinebach, H. L. Bethlem, and K. S. E. Eikema, “Pauli blocking of stimulated emission in a degenerate Fermi gas,” *Nat. Commun.* 13, 6479 (2022).
- [31] J. G. Oxley, *Matroid Theory*, 2nd ed. Oxford University Press, 2011.
- [32] A. R. Zhitnitsky, “The vacuum energy density in QCD and the cosmological constant problem,” arXiv:1905.08180.
- [33] J. M. Maldacena, “The large N limit of superconformal field theories and supergravity,” *Adv. Theor. Math. Phys.* 2, 231–252 (1998).

- [34] A. Strominger, “The dS/CFT correspondence,” *JHEP* 10, 034 (2001).
- [35] H. Lin and L. Susskind, “Infinite Temperature’s Not So Hot,” arXiv:2206.01083 (2022).
- [36] L. Susskind, “De Sitter Holography: Fluctuations, Anomalous Symmetry, and Wormholes,” arXiv:2106.03964 (2021).
- [37] L. Susskind, “Entanglement and Chaos in De Sitter Space Holography,” arXiv:2201.03603 (2022).
- [38] V. Franken, H. Partouche, F. Rondeau, and N. Toumbas, “Bridging the static patches: de Sitter holography and entanglement,” *JHEP* 08, 074 (2023).
- [39] T. Padmanabhan, “Gravity and the thermodynamics of horizons,” *Phys. Rept.* 406, 49–125 (2005).
- [40] E. Verlinde, “On the origin of gravity and the laws of Newton,” *JHEP* 04, 029 (2011).
- [41] R. S. Clark II, “Entropic Inertia from Modular Flow: $F = ma$ as a First-Law Identity,” companion white paper, in preparation (2026).
- [42] H. Casini, M. Huerta, and R. C. Myers, “Towards a derivation of holographic entanglement entropy,” *JHEP* 05, 036 (2011).
- [43] S. Hollands and R. M. Wald, “Conservation of the stress tensor in perturbative interacting quantum field theory in curved spacetimes,” *Rev. Math. Phys.* 17, 227–312 (2005).
- [44] V. Chandrasekaran, R. Longo, G. Penington, and E. Witten, “An algebra of observables for de Sitter space,” *JHEP* 02, 082 (2023).
- [45] E. Witten, “Gravity and the crossed product,” *JHEP* 10, 008 (2022).
- [46] V. Narovlansky and H. Verlinde, “Double-scaled SYK and de Sitter holography,” *JHEP* 05, 032 (2025).
- [47] A. Irakleous, F. Rondeau, and N. Toumbas, “Holography for de Sitter bubble geometries,” *JHEP* 03, 240 (2026).
- [48] M. Ali, “Local generalized second law in crossed product constructions,” *Phys. Rev. D* 111, 024015 (2025).
- [49] H. Umegaki, “Conditional expectation in an operator algebra. IV. Entropy and information,” *Kodai Math. Sem. Rep.* 14, 59–85 (1962).
- [50] W. P. Su, J. R. Schrieffer, and A. J. Heeger, “Solitons in polyacetylene,” *Phys. Rev. Lett.* 42, 1698–1701 (1979).
- [51] A. R. Calderbank and P. W. Shor, “Good quantum error-correcting codes exist,” *Phys. Rev. A* 54, 1098–1105 (1996).
- [52] A. M. Steane, “Error correcting codes in quantum theory,” *Phys. Rev. Lett.* 77, 793–797 (1996).
- [53] A. Y. Kitaev, “Fault-tolerant quantum computation by anyons,” *Annals Phys.* 303, 2–30 (2003).

- [54] S. Doplicher, R. Haag, and J. E. Roberts, “Fields, observables and gauge transformations. I,” *Commun. Math. Phys.* 13, 1–23 (1969).
- [55] V. Gorini, A. Kossakowski, and E. C. G. Sudarshan, “Completely positive dynamical semigroups of N-level systems,” *J. Math. Phys.* 17, 821 (1976).
- [56] G. Lindblad, “On the generators of quantum dynamical semigroups,” *Commun. Math. Phys.* 48, 119–130 (1976).