

Timeless Partition-Strain on a Structured Algebraic Universe: A Unisound Substrate–Observer Reconstruction

Daniel de Souza Casali

May 2026

Abstract

This paper develops a timeless partition-strain framework motivated by a basic tension in present foundations. General relativity and quantum mechanics both work with extraordinary precision, yet they organize reality through incompatible treatments of time, locality, and information. In quantum mechanics, time is normally treated as an external background parameter against which states evolve. In general relativity, time is part of the dynamical geometry itself, altered by gravitational structure. Quantum theory permits nonseparable correlations across distant systems, while relativistic gravity protects local causal structure through finite light-cone propagation. Black-hole evaporation sharpens the conflict: general relativity appears to allow information loss behind horizons, while quantum mechanics requires unitary recoverability.

The framework proposed here changes the starting point. Instead of treating spacetime, time evolution, gravity, quantum state update, and thermodynamic irreversibility as separate primitives, it begins from a tenseless algebraic substrate equipped with regional structure, conditional expectations, connection data, and a paired regional disturbance bookkeeping ($K_{\text{rec}}, M_{\text{rec}}$). The first member is the universal record stiffness: the substrate’s resistance to record-supporting deformations. The second member is the universal record susceptibility: the substrate’s loading or distinguishability response under the same deformations. Both are positive operators, neither is defined from the other, and the observable record cone is controlled not by either operator separately but by the generalized spectrum

$$G_{\text{acc}} = M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2},$$

which is the geometry the framework requires to converge, under refinement, to the substrate’s intrinsic access spectrum. Physical phenomena are then read through a second layer: a memory-bearing observer whose experience is determined by record-ordered access to the substrate. In this view, time is not a primitive coordinate of the universe. Time appears at the observer junction, where a nonracial observer representation cuts a tracial, timeless substrate and generates modular record flow.

The central object is a regional partition-strain functional measuring what an observer-access algebra fails to recover from the full substrate. Entropy is the scalar unrecoverability of this disturbance. Inertia is resistance to algebraic support deformation, read through the stiffness operator K_{rec} . Gravity-like response is variation of the same disturbance with respect to accessible connection or metric-like data. Record susceptibility, read through M_{rec} , sets the entropy/distinguishability cost of those deformations. Gauge and field equations are stationarity conditions of the same regional disturbance under admissible connection variations. The framework therefore seeks a simpler organization: not many disconnected physical mechanisms, but many observer-effective readouts of one underlying paired access-defect geometry.

Thermodynamics becomes the guiding analogy. The arrow of time is not imposed as a fundamental direction in the substrate. It is the experienced direction of increasing record restriction for a Type III observer reading a Type II-like reality. The past feels fixed because its records have already been compressed into stable but incomplete access-cuts; the

underlying substrate information is no longer fully recoverable from the observer’s present algebra. The future feels open because those cuts have not yet been made. The Second Law is interpreted not as a primitive law of universal machinery, but as the geometric limit of localized ignorance: the monotone growth of unrecoverability along record-ordered access.

On this reading, the apparent conflicts between general relativity, quantum mechanics, thermodynamics, and information recovery are relocated. They are not treated as contradictions inside one spacetime ontology, but as different limits of substrate information seen through observer access. Local causal cones arise from finite record recoverability of the paired operator $(K_{\text{rec}}, M_{\text{rec}})$, not from a primitive substrate speed limit. Quantum nonseparation reflects correlations in the substrate that need not be reducible to observer-local records. Horizon information loss becomes loss relative to an exterior access algebra, not destruction in the full substrate. Gravity becomes the macroscopic connection response of regional disturbance, rather than a separate force added to quantum theory.

The aim is to formulate a mathematically disciplined reconstruction program: a timeless algebraic substrate with a paired stiffness–susceptibility bookkeeping, a record-bearing observer layer, and a junction where time, thermodynamics, locality, inertia, gravity-like response, and information loss appear as coupled consequences of partition strain and finite recoverability through the generalized record geometry G_{acc} .

Copyright © 2026 Daniel de Souza Casali. All rights reserved.

Contents

1	Introduction	7
2	Unified access-defect spine	9
2.1	Metric-compatible paired access normal form	9
2.2	Absorbing access-cut boundary principle	10
2.3	Minimal observer-access dimension: three projection directions plus one record direction	10
2.4	Sector readout dictionary	11
3	The substrate layer	11
3.1	Paired regional disturbance bookkeeping	11
3.2	Structural decomposition of the paired operators	12
3.3	Refinement compatibility of paired bookkeeping	13
3.4	Pre-oriented inclusion-boundary symmetry	16
3.5	Regional defect and master disturbance	16
3.6	Substrate-level identifications	17
3.7	Non-Abelian selection and the Wilson-defect criterion	17
4	Type II₁ trace-quadratic normal form and cone recovery	18
4.1	Conditional-expectation defects and the Type II ₁ isotropic quadratic disturbance theorem	18
4.2	Cone-recovery trace lemma	21
5	Projective-access multiplicity rules and the mass-shape sector	24
5.1	Visible–hidden Schur architecture	24
5.2	Projective-access ladder lemma (provisional form)	25
5.3	Second-order projective closure	25
5.4	Sector multiplicity and sign rules	26
5.5	Core lemma: tau and baryonic mean	27
5.6	Numerical diagnostics	27

5.7	Provisional extension: muon access ratio	27
5.8	Falsifier and predictive scope	28
5.9	Open problems and named proof obligations	28
5.10	Section status	29
6	The observer layer	30
6.1	Two primitive operations and one stability criterion	30
6.2	Local entropy-maximizing observer state	30
6.3	Inclusion index	30
6.4	Stability criterion	31
6.5	Minimal stable observer-access dimension	31
6.6	One-sided record-boundary trace	35
6.7	Access resolution and the observer-layer conjugate structure	35
6.8	Schur-complement reduction of observer-junction quadratic forms	36
6.9	Access–recovery composition for observer-effective couplings	38
6.10	Bare trace from observer-side duality: $T_{\text{bare}} = 21/2$ derived from $d = 4$ and $b = 3$	39
6.11	Trace weights at the minimal stable observer-access cut	40
6.12	Wavefunction and collapse from accessible-algebra restriction	44
6.13	Observer covariance through invariant record relations	44
6.14	Projective observer readout	45
6.15	Observer-level identifications	45
6.16	Order and arrow	45
7	Observer metrological completion	46
7.1	Level hierarchy and the dimensionless baseline	46
7.2	The seven metrological scales as a coordinate basis	47
7.3	Canonical metrology-block normalization	48
7.4	Coordinate transformation of $(K_{\text{rec}}, M_{\text{rec}})$ and invariance of G_{acc}	48
7.5	Metrology–physics direct-sum decomposition	49
7.6	Gauge conditions from the modern SI	50
7.7	Derived dimensional scales	52
7.8	Gravitational coupling: G as metrological, α_G as physical	52
7.9	Universal gas constant as a derived bridge	53
7.10	Vacuum electromagnetic constants as α -dependent reconstructions	53
7.11	Dimensional-weight predictive criterion	53
7.12	Full Level-2 metrological map	54
7.13	Observer metrological completion theorem	54
8	Recovered thermodynamics from substrate–observer coupling	55
8.1	Cuts, entropy, and disturbance flux	56
8.2	Recovered Clausius relation	56
8.3	Relation to Jacobson	56
8.4	Why the placement matters	57
9	Shared-Connection principle from one paired regional disturbance and Type II₁ trace uniqueness	57
9.1	Alternative motivation: support re-embedding as paired disturbance variation	58
9.2	Composition independence: the binding-energy question	59

10 Scope, modularity, and claim status	60
10.1 Consistency claims	60
10.2 Reconstruction claims	60
10.3 Locational claims	60
10.4 Scope of the claims	60
10.5 Modularity	61
11 Junction recovery principle for effective laws	61
11.1 Recovered Newtonian mechanics as least-disturbance record continuation	61
11.2 Finite record recoverability and the observer-junction light limit	62
11.2.1 Paired record bandwidth from $(K_{\text{rec}}, M_{\text{rec}})$	63
11.2.2 Compatibility without metric coincidence	64
11.2.3 Extreme regimes	65
11.2.4 Finite generalized-eigenvalue benchmark	65
11.2.5 Stability envelope and null record modes	66
11.2.6 Light as the saturating record sector	66
11.3 Recovered Lorentz kinematics from invariant modular-access cones	67
11.4 Wavefunction and collapse: observer-side reading	68
11.5 Recovered Maxwell-type equations from a paired holonomy readout	68
11.6 Recovered weak-field gravity from stationary accessible connection response	69
12 Algebraic necessity program: paired disturbance, access, shell growth, non-Abelian selection, and observer cuts	70
12.1 Spectral covariance of defect modes from observer coarse-graining	71
12.2 Continuum Einstein-translation of the junction variation	74
12.3 Proof obligations for an airtight Einstein reduction	76
12.4 Regional disturbance stationarity as the source of field actions	78
12.5 Algebraic shell growth instead of coordinate distance	79
12.6 Observer access from modular spectral data	79
12.7 Modularly stable observer cuts	79
12.8 Hierarchy of necessity statements	80
13 Type II₁ to Type III₁ observer junction	80
13.1 Negative constraint: a normal cut of Type II ₁ is not enough	80
13.2 Positive bridge: non-tracial observer states and scaling limits	80
13.3 Access-cut as conditional expectation and modular embedding	81
13.4 Crossed-product observer core	81
13.5 Support stiffness as a noncommutative derivation	81
13.6 Junction recovery loss and record susceptibility	81
13.7 Temperature as a scaling-limit observer parameter	82
13.8 Powers-type product states as the model example	82
13.9 Junction principle	82
13.10 Scaling-limit observer-recoverability speed: proposal status	82
13.11 Block-cut diagnostic on the Powers family	82
13.12 Correlated observer-state record-bandwidth diagnostic	83
13.13 Finite diagnostic checklist	83
14 Horizon-cut thermodynamics and black-hole radiation	83
14.1 High-disturbance recoverability traps	84
14.2 Horizons as recoverability boundaries	84
14.3 Radiation as exterior disturbance flux	84
14.4 Cut temperature	84

14.5	Thermality and modular appearance	85
14.6	Information loss and Page-like recovery	85
14.7	Summary of the framework-native reading	85
15	Internal charged sector: identification, minimality, and observer-junction diagnostics	86
15.1	Charge as minimal observer-stable internal holonomy	86
15.2	From discrete branches to continuous mixing	87
15.3	Minimal charged-sector selection principle	88
15.4	Minimal chiral content and hypercharge fixing	89
15.5	Hypercharge normalization and conditional weak angle	89
15.6	Conditional identification with the Maxwell readout and the lepton ladder	90
15.7	Schur-complement and access-recovery application to the charged sector	90
15.8	Projective access-cover operator and polarization-order expansion	91
15.9	Fine-structure diagnostic as a single Schur-trace identity	93
15.10	Neutral electroweak visibility and the SLD/LEP- <i>b</i> channel split	99
15.11	Strong-sector adjoint holonomy and electroweak-resolution overlap	114
15.12	Observer-access dimension ablation and substrate-observer duality	117
15.13	Observer-side bare-trace duality applied to the charged sector	118
15.14	Bare-trace content realization: minimal chiral content saturates the duality value	119
15.15	Independent consequences of the projective-readout axiom	120
15.16	Charged-lepton ladder	121
15.17	Charge structure consequences	121
15.18	What is not recovered	121
15.19	Status and remaining problems	122
16	Layer allocation and status	123
17	Worked finite trace diagnostics: Petz recovery as paired-regional-disturbance diagnostics	124
17.1	Exact threshold as a sanity check	125
17.2	Perturbative Petz-paired-disturbance normal form	125
17.3	Sector splits: useful diagnostic, limited analogy	125
17.4	What the example establishes	126
17.5	Approximate recovery away from the exact-erasure surface	126
18	Paired record-pair audit: universal stiffness-susceptibility consistency across sectors	127
18.1	Meaning of the susceptibility operator	127
18.2	The identity-susceptibility toy and its limitation	128
18.3	First nontrivial universal susceptibility attempt	128
18.4	Paired universal stiffness	129
18.5	Fine-structure and lepton-ladder audit after pairing	130
18.6	Audit table	130
18.7	Status of the paired audit	131
19	Relation to neighboring frameworks	131
20	Modular flow, sector visibility, and the gravitational coupling	133
20.1	Modular access-flow principle	133
20.2	Modular sector-visibility functions	135
20.3	Gravitational access and the unified modular-flow architecture	138

20.4	Boundary-stiffness Planck coefficient theorem	140
20.5	Boundary-deformation reduction and the Einstein–Hilbert coefficient	145
20.6	Substructure of the Einstein–Hilbert reduction hypotheses	149
21	Open questions	153
21.1	Type II ₁ to Type III ₁ observer junction	153
21.2	Einstein-reduction proof obligations	154
21.3	Horizon-cut radiation	154
21.4	Cosmological-constant dictionary entry	154
21.5	Multi-cut readings at distinct access depths	154
21.6	Continuation work flagged by the relocated structural results	155
21.7	Projective-access proof obligations	155
21.8	Relation to perturbative electroweak corrections	156
21.9	Falsification criteria	156
22	Conclusion	158
A	Validation material	159
	Acknowledgements	160
	Copyright	160
	References	160

1 Introduction

Standard physics writes its laws on a manifold. A spacetime is assumed, fields are placed on it, equations evolve those fields in a time coordinate, and observers measure what evolves. The tension this produces with background-independent quantum gravity is the long-discussed problem of time [1, 2, 3]. This paper takes a different organizational stance. It treats the totality of physical structure as a tenseless algebraic object and treats observers — and therefore time — as a separate layer that cuts that object and accumulates records.

The contribution is a unified reconstruction architecture. It treats established physical structures as observer-effective readouts of one substrate–junction geometry, while separating what is derived inside the framework, what is conditional on an explicit principle, and what remains open. The organizing principle is compressed bookkeeping. Familiar structures — locality, mass, inertia, gravitational response, gauge phase, time, motion, collapse, field dynamics, and thermal-ity — are not introduced as independent modules. They are organized as readouts of a paired regional disturbance functional relative to observer-access algebras.

The framework rests on two ideas. The first is that physical structure can be represented as static compatibility data of an algebraic totality without reference to a clock. The second is that physical experience — what is felt as temporal order, the arrow of time, motion, and measurement — can be organized as features of a memory-bearing access structure that cuts the totality and accumulates records. The two ideas are modular. They become a unified reconstruction framework when applied together.

The compression principle. The primitive object is not entropy, inertial mass, gravitational charge, or gauge energy separately. The primitive object is the paired regional disturbance bookkeeping $(K_{\text{rec}}, M_{\text{rec}})$. The first member, K_{rec} , is the universal record stiffness; the second member, M_{rec} , is the universal record susceptibility. The intrinsic regional disturbance is

$$\mathcal{K}_R(\rho, \nabla, g) = \tau[\mathfrak{D}_R(\rho)^* K_{\text{rec}} \mathfrak{D}_R(\rho)], \quad \mathfrak{D}_R(\rho) = \rho_R - E_R(\rho),$$

weighted by the substrate stiffness, while the record-distinguishability geometry of the same defect is carried by M_{rec} . Entropy is the scalar unrecoverability of this disturbance. Inertia is the Hessian of this disturbance along algebraic support deformations, weighted by K_{rec} . Gravity-like response is the variation of this disturbance with respect to the accessible connection or metric-like embedding. Gauge and field equations are stationarity conditions of the same disturbance under internal connection variations. Crucially, the observable record cone is not controlled by K_{rec} or M_{rec} separately; it is controlled by the generalized record operator

$$G_{\text{acc}} = M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2},$$

whose spectrum the framework requires to converge, under refinement, to the substrate’s intrinsic access spectrum. Thus the paper does not attempt to assign each phenomenon to a separate phenomenological term. It asks how many effective laws can be read from one paired regional defect geometry.

Not a quantization, not a geometrization. This is not a proposal to quantize gravity as a separate sector, nor to geometrize quantum mechanics by imposing a spacetime interpretation on operators. The statistical or quantum readout is the defect covariance, recovery loss, and modular state induced by a cut; the gravity-like readout is the connection or metric variation of the same disturbance; the gauge readout is its internal holonomy stationarity. Existing continuum theories remain reduction targets, but the gravity target is made explicit: the question is whether the spectral continuum limit of the generalized record operator G_{acc} yields an observer-effective action whose metric variation gives an Einstein tensor and whose defect variation gives

a stress-energy tensor. In this sense quantum/statistical behavior and gravity-like behavior are not two primitive languages to be merged after the fact; they are two observer readouts of one paired trace-quadratic disturbance, one through defect covariance and state restriction, the other through connection and metric-like variation.

Three classes of claim. This paper makes three classes of claim. The first is a *substrate claim*: locality, non-separation, persistence, support-stiffness, gravity-like response, gauge phase, and record-stable cuts can be represented as relations among algebraic regions, conditional expectations, and the paired master disturbance functional. The second is a *reconstruction claim*: given a memory-bearing observer layer, familiar kinematic, quantum, gauge, and weak-gravity descriptions can be read as effective descriptions of recorded experience. The third is a *normal-form claim*: in a Type II₁ factor [4, 5, 7], the local disturbance associated with conditional-expectation defects has a forced quadratic leading term under explicit covariance and isotropy hypotheses, and the paired $(K_{\text{rec}}, M_{\text{rec}})$ bookkeeping inherits the canonical trace through the generalized record operator.

Equivalence principle, recast. The equivalence-principle content is therefore not a literal identity between two separately postulated forces. It is the claim that inertial resistance and gravity-like sourcing are two variational readouts of the same paired regional disturbance functional, evaluated along different observer directions. In weak, isotropic, low-disturbance limits these readouts reduce to the same trace-normalized quadratic coefficient, giving the observer-effective cancellation associated with composition-independent free fall. The same compression also fixes the interpretation of the observer speed bound: c_* is not an independently chosen number but a record-bandwidth threshold obtained from the lowest generalized eigenvalue of $G_{\text{acc}} = M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2}$.

Plan of the paper. The text proceeds as follows. Section 2 fixes the unified access-defect spine that runs through every later sector. Sections 3–4 define the substrate layer and the Type II₁ trace-quadratic normal form, including the refinement-compatibility theorem (Section 3.3) that governs how the paired bookkeeping behaves under substrate refinement by branch-neutral labels and supplies the architectural foundation for resolved retrieval algebras used throughout the paper. Section 5 applies the substrate normal form for the first time as a predictive tool, identifying the mass-shape sector and stating the two-mode projective-access core lemma that fixes the tau and baryonic mean access ratios under named selection rules. Section 6 develops the observer layer. Section 7 introduces the observer metrological completion, the gauge-fixing layer that maps the dimensionless substrate–observer record geometry into laboratory SI form and isolates the framework’s strict predictive content as dimensionless invariants. Section 8 recovers thermodynamics at the substrate–observer junction. Section 9 states the Shared-Connection principle as a shared-generator result. Section 10 fixes scope. Section 11 states the junction recovery principle for effective laws and gives the observer-effective reconstructions: Newtonian mechanics, the finite record recoverability principle and observer-junction light limit, Lorentz kinematics, wavefunction and collapse, Maxwell-type equations, and weak-field gravity. Section 12 develops the algebraic necessity program, including the Type II₁ quadratic disturbance theorem, spectral covariance of defect modes from observer coarse-graining, the continuum Einstein-translation schema, disturbance-stationary actions, algebraic shell growth, modular observer access, modularly stable observer cuts, and non-Abelian selection. Section 13 formulates the Type II₁ → Type III₁ observer junction as a modular embedding, including the conditional-expectation access cut, derivational support stiffness, recovery-loss susceptibility, and crossed-product observer core. Section 14 gives the framework-native reading of black-hole radiation. Section 15 develops the internal charged-sector identification and its observer-junction diagnostics built from the paired $(K_{\text{rec}}, M_{\text{rec}})$ operator, including the single Schur-trace identity for the fine-structure

access-polarization diagnostic and the resolved Drell–Yan retrieval algebra that specializes the refinement-compatibility theorem of Section 3.3 to the precision-electroweak channel. Section 16 tabulates the layer allocation and status of each claim. Section 17 presents the worked finite trace diagnostics. Section 18 runs a paired-pair audit on the universal $(K_{\text{rec}}, M_{\text{rec}})$ pair across record-cone, Maxwell-type stationarity, fine-structure, and lepton-ladder sectors at finite N . Section 19 compares the framework to neighboring foundation programs and Section 21 lists the remaining mathematical tasks.

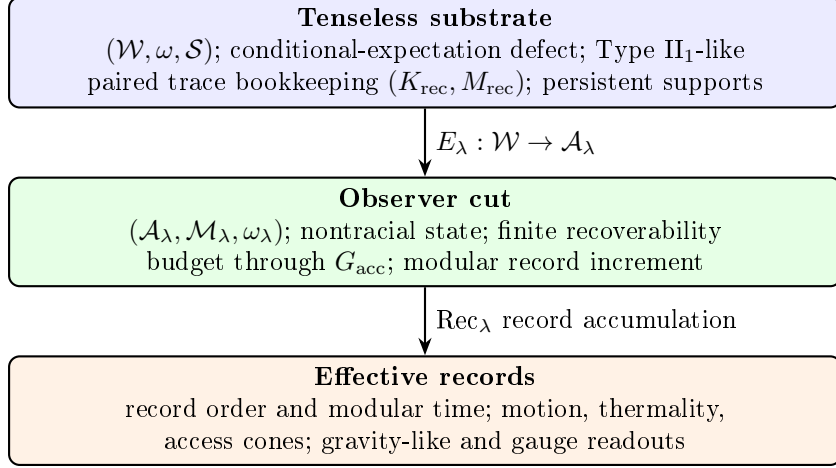


Figure 1: Layer map for the partition-strain framework. Static substrate data — including the paired stiffness–susceptibility bookkeeping $(K_{\text{rec}}, M_{\text{rec}})$ — define regional disturbance and persistent supports; a memory-bearing observer cut restricts that data; effective physics is reconstructed from stable, monotonically recoverable records.

2 Unified access-defect spine

The document is organized around one substrate–observer spine. The substrate carries a tracial regional algebraic totality and a paired record geometry $(K_{\text{rec}}, M_{\text{rec}})$. The observer does not add a second physics; it supplies a memory-bearing access cut through which the paired geometry is read. Every effective sector is therefore written as a readout of

$$(K_{\text{rec}}, M_{\text{rec}}), \quad G_{\text{acc}} = M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2},$$

or of the corresponding access operator obtained at the observer junction. This section fixes the notation used throughout the manuscript before the individual physical readouts are discussed.

2.1 Metric-compatible paired access normal form

The paired record geometry is represented by a positive stiffness operator K_{rec} and a positive susceptibility operator M_{rec} . The universal record cone is not a property of either operator alone. It is the generalized spectrum of

$$G_{\text{acc}} = M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2}.$$

The metric-compatible access normal form used throughout the reconstruction is

$$K_{\text{rec}} = M_{\text{rec}}^{1/2} L_{\text{acc}} M_{\text{rec}}^{1/2}, \quad G_{\text{acc}} = L_{\text{acc}}.$$

This form does not erase the distinction between stiffness and susceptibility. It states that, after the susceptibility loading has been accounted for, the invariant observer-access spectrum is

carried by a single intrinsic access operator L_{acc} . This is the operator whose spectrum controls record-cone recovery, Maxwell-type propagation, access dressing in the charged sector, and the finite diagnostic audits.

2.2 Absorbing access-cut boundary principle

The primitive one-dimensional access interval is the normalized record-access coordinate $u \in [0, 1]$. The endpoints are not ordinary spatial boundaries. They represent the inaccessible complement at one end and the fully absorbed record cut at the other. The minimal self-adjoint access operator compatible with this absorbing cut is

$$L_{\text{acc}} = -\frac{d^2}{du^2}, \quad f(0) = f(1) = 0.$$

Thus

$$\lambda_n = n^2 \pi^2, \quad f_n(u) = \sqrt{2} \sin(n\pi u).$$

The appearance of π in the later fine-structure diagnostic is therefore not introduced as a numerical ornament. It is the spectral constant of the normalized absorbing access interval. The charged-sector calculation uses the same L_{acc} that appears in the record-cone and access-dressing statements, so the fine-structure diagnostic is tied to the same access spectrum rather than to a separate numerical ansatz.

2.3 Minimal observer-access dimension: three projection directions plus one record direction

The observer layer is fixed as a $3 + 1$ access structure, but not as primitive substrate spacetime. The three spatial directions belong to a projective localization layer carrying a non-removable cyclic obstruction on its branch cover; the fourth direction is the record-order coordinate required for a memory-bearing cut. The full algebraic statement is developed as Theorem 6.3 and its Corollary 6.9 in Section 6.5, where the minimality of the spatial branch count $b = 3$ is proved as a Čech-cohomological obstruction theorem on a finite projective cover. The present subsection records the resulting access structure

$$O_{\text{acc}} = \Sigma_{\text{obs}}^3 \oplus \tau_{\text{obs}}, \quad b := \dim \Sigma_{\text{obs}}^3 = 3, \quad d := \dim O_{\text{acc}} = 4,$$

which is used consistently throughout the rest of the paper: in the Lorentz-recovery target (Section 11.3), in the observer-access dimension ablation and substrate–observer duality (Section 15.12), in the Clifford/access combinatorics, and in the charged-sector trace duality. The substrate itself is not thereby made into a $3 + 1$ spacetime; only the observer-access interface has this local organization.

The asymmetry between the three spatial access directions and the single record-ordering direction is structural. Projective closure lives on Σ_{obs}^3 , because the branch cover of the spatial-access layer is the support of the non-removable cyclic obstruction Ω_{ijk} of Definition 6.1. The record-ordering direction τ_{obs} supplies sequencing along the inclusion-index \preceq and contributes a one-sided boundary trace through the one-sided record-boundary normalization of Section 6.6, but it does not generate independent spatial projective sectors. The fine-structure calculation therefore depends on the pair $(b, d) = (3, 4)$ for a definite reason: $b = 3$ is the minimal non-removable projective closure of Theorem 6.3, while $d = 4$ is that cover plus the one record-ordering coordinate required by memory, with an independent algebraic closure check supplied by the duality conditions of Section 15.12.

2.4 Sector readout dictionary

With the access spine fixed, the physical sectors use the following dictionary:

Effective phenomenon	Substrate–observer readout
record order	monotone access to stable record supports
entropy	unrecoverability of conditional-expectation defect
inertia	support-deformation Hessian through K_{rec}
gravity-like response	connection/metric variation of paired disturbance
Maxwell-type field	stationary internal holonomy readout
light cone	$c_*^2 = \lambda_{\min}(G_{\text{acc}})$
collapse-like update	restriction to accessible record algebra
horizon information loss	loss relative to an exterior access algebra
fine structure	Schur trace of the charged substrate–cut operator
lepton ladder	charged support-deformation diagnostic

The purpose of the following sections is to keep these readouts in one language: paired disturbance, access spectrum, observer cut, and Schur reduction. Standard references on operator-algebraic modular structure, conditional expectation, quantum recovery, gauge unification, anomaly cancellation, and black-hole thermodynamics are cited in the sections where each technical interface is used [4, 5, 7, 12, 49, 50, 51, 54, 55, 33, 34].

3 The substrate layer

The substrate layer represents the structural side of physics. It proposes that what observers call locality, persistence, stiffness, curvature response, and gauge phase are features of a tenseless algebraic totality

$$\mathcal{S} = (\mathcal{W}, \omega, S),$$

where \mathcal{W} is a von Neumann algebra modeling all observables of the totality, ω is a faithful normal state, and S is declared additional structure — a regional decomposition, an admissible class of conditional expectations, an internal automorphism sector, and connection or metric-like data. Nothing in S refers to a clock or a direction of time.

3.1 Paired regional disturbance bookkeeping

The compressed substrate object is not a single stiffness or a single covariance. It is a paired bookkeeping

$$(K_{\text{rec}}, M_{\text{rec}}) = (\text{record stiffness, record susceptibility})$$

both positive operators on the relevant defect Hilbert space, both trace-intrinsic, and neither defined from the other. The first member, K_{rec} , encodes the substrate’s resistance to record-supporting deformations. The second member, M_{rec} , encodes the substrate’s loading or entropy-distinguishability response to those same deformations. The observable geometry of the substrate is then carried by the generalized record operator

$$G_{\text{acc}} := M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2},$$

which the framework requires to converge, under refinement of the substrate, to the intrinsic substrate access spectrum. This is a weaker compatibility condition than identifying K_{rec} with a function of M_{rec} or vice versa; it only requires

$$\text{spec}\left(M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2}\right) \rightarrow \text{spec}(G_{\text{acc}}),$$

without imposing a literal metric-compatible realization.

Invertibility of M_{rec} on the relevant support. The expression $M_{\text{rec}}^{-1/2}$ requires M_{rec} to be invertible on the support carrying the diagnostic. The framework’s paired bookkeeping construction supplies M_{rec} as a positive operator on the defect Hilbert space with faithful substrate trace, and invertibility on the trace-positive support is guaranteed by the trace’s faithfulness: any kernel of M_{rec} would correspond to a trace-zero projection and is excluded by the framework’s normal-form requirement that M_{rec} enter as a record susceptibility (a positive form rather than a degenerate one). The square root $M_{\text{rec}}^{1/2}$ and its inverse $M_{\text{rec}}^{-1/2}$ are then defined through measurable functional calculus on the positive spectrum, with the inversion operating only on the trace-positive support. Every diagnostic in this paper projects onto the support of K_{rec} , which lies within the trace-positive support of M_{rec} by the pairing condition, so the off-spectrum behavior of M_{rec} is irrelevant to the readouts the framework computes.

3.2 Structural decomposition of the paired operators

Both K_{rec} and M_{rec} admit a structural decomposition into substrate contributions. The susceptibility operator has the form

$$M_{\text{rec}} = M_0 + \eta_{\text{acc}} M_{\text{acc}} + \eta_{\text{str}} M_{\text{str}} + \eta_{\text{ns}} M_{\text{ns}} + \eta_{\text{sec}} M_{\text{sec}}$$

and the stiffness operator has the matching form

$$K_{\text{rec}} = K_0 + \kappa_{\text{acc}} K_{\text{acc}} + \kappa_{\text{str}} K_{\text{str}} + \kappa_{\text{ns}} K_{\text{ns}} + \kappa_{\text{sec}} K_{\text{sec}}$$

with the positivity requirements

$$M_{\text{rec}} > 0, \quad K_{\text{rec}} > 0.$$

The ingredients have a uniform physical interpretation. On the susceptibility side, M_0 is a baseline record-loading metric (the irreducible record-distinguishability scale); it is taken to be sectorwise isotropic, that is

$$M_0 = \bigoplus_s m_s I_s = m_{\text{loc}} I_{\text{loc}} \oplus m_{\text{int}} I_{\text{int}} \oplus m_{\text{ns}} I_{\text{ns}},$$

with a single positive susceptibility scale m_s per substrate sector (local, internal, nonseparable). The sectorwise-isotropic form expresses the physical statement that, in the absence of access compliance, strain, nonseparable, or sector-organization corrections, no internal direction of a given sector is privileged: each sector contributes one record-loading scale and not a directionally biased loading. Anisotropy in M_{rec} is then carried entirely by the access, strain, nonseparable, and sector contributions through their respective η_{\bullet} weights. The same convention is used throughout the charged-sector diagnostics below, so that nontrivial α -like structure arises from internal projection on the charged access channel rather than from anisotropy in the baseline susceptibility itself.

M_{acc} is the record-loading contribution from access compliance, capturing how recoverable record information changes under the most basic access deformations; M_{str} is the susceptibility contribution from local substrate strain; M_{ns} is the susceptibility contribution from nonseparable record structure (record-loading carried by collective non-factorizing sectors); and M_{sec} is the susceptibility contribution from sector organization (charged, inertial, field-like, or entropy-carrying channels).

On the stiffness side, K_0 is the irreducible record-stiffness baseline; K_{acc} is the access stiffness that resists changes in the accessible record configuration (in standard substrate models it specializes to the access Laplacian L_{acc}); K_{str} is the stiffness contribution from local substrate strain; K_{ns} is the stiffness contribution from nonseparable substrate structure; and K_{sec} is the stiffness contribution from sector organization.

The coefficients $\eta_{\bullet}, \kappa_{\bullet} \geq 0$ are admissible weights; they determine the relative substrate-internal balance among the contributions and are not chosen from experiment.

The point of the decomposition is conceptual. It expresses that record stiffness and record susceptibility are not generic positive matrices to be chosen; they are built up from the same admissible substrate contributions (baseline, access, strain, nonseparable, sector), and each sector contributes independently to both members of the pair. In particular, $M_{\text{rec}} \neq I$ in general, and $K_{\text{rec}} \neq L_{\text{acc}}$ in general: both operators are nontrivial, and the trivial normalization $(K_{\text{rec}}, M_{\text{rec}}) = (L_{\text{acc}}, I)$ is a degenerate special case obtained by turning off every $\eta_{\bullet}, \kappa_{\bullet}$ except the access contributions and replacing the baseline susceptibility by the identity.

3.3 Refinement compatibility of paired bookkeeping

The framework requires the access spectrum of $G_{\text{acc}} = M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2}$ to converge, under refinement of the substrate, to the intrinsic substrate access spectrum. The refinement compatibility condition was stated informally in Section 3.1; this subsection makes it precise. The result is a structural lemma about how the paired bookkeeping interacts with admissible refinements of the substrate algebra, and is used throughout the paper whenever a substrate-level algebra is refined by attaching observer-resolvable labels — flavor sectors, kinematic supports, internal-sector resolutions, access-depth refinements, and the resolved retrieval algebra of Section 15.10.

Admissible refinements

A substrate retrieval algebra \mathcal{R} with paired bookkeeping $(K_{\text{rec}}^{\mathcal{R}}, M_{\text{rec}}^{\mathcal{R}})$ admits a *refinement* when a richer algebra $\tilde{\mathcal{R}}$ is constructed by attaching branch-neutral labels to the original. The labels are carried by a W^* -algebra \mathcal{F} with a normalized faithful trace $\tau_{\mathcal{F}}$, and the refinement is intended to capture observer-resolvable structure that does not modify the substrate-level retrieval primitives.

Definition 3.1 (Admissible substrate refinement). An *admissible refinement* of a substrate retrieval algebra \mathcal{R} by a labeling algebra \mathcal{F} is a W^* -algebra $\tilde{\mathcal{R}}$ together with inclusions

$$\mathcal{R} \hookrightarrow \tilde{\mathcal{R}}, \quad \mathcal{F} \hookrightarrow \tilde{\mathcal{R}},$$

under which the following four structural conditions hold.

(A1) Untwisted extension.

The primitive support projections $e \in \mathcal{R}$ lift trivially: $e \mapsto e \otimes \mathbf{1}_{\mathcal{F}}$. No twist, crossed product, or branch-dependent action on \mathcal{F} is permitted; the refinement attaches labels to the support index, but does not alter or rotate it.

(A2) Trace-product compatibility.

The faithful normal trace on $\tilde{\mathcal{R}}$ factorizes: $\tau_{\tilde{\mathcal{R}}} = \tau_{\mathcal{R}} \otimes \tau_{\mathcal{F}}$, with $\tau_{\tilde{\mathcal{R}}}(\mathbf{1}) = 1$.

(A3) Equivariance.

Any group action $G \curvearrowright \mathcal{R}$ of substrate symmetries extending to $\tilde{\mathcal{R}}$ acts on the first tensor factor only: $g \cdot (r \otimes f) = (g \cdot r) \otimes f$ for $g \in G$, $r \in \mathcal{R}$, $f \in \mathcal{F}$.

(A4) Access-geometry preservation.

The substrate-level paired bookkeeping $(K_{\text{rec}}^{\mathcal{R}}, M_{\text{rec}}^{\mathcal{R}})$ lifts so that the generalized record operator factorizes trivially:

$$\tilde{G}_{\text{acc}} := \tilde{M}_{\text{rec}}^{-1/2} \tilde{K}_{\text{rec}} \tilde{M}_{\text{rec}}^{-1/2} = G_{\text{acc}}^{\mathcal{R}} \otimes \mathbf{1}_{\mathcal{F}}.$$

The substantive content of (A1) is the untwisted form of the extension; (A2) is a trace-normalization condition on \mathcal{F} ; (A3) is the equivariance statement that observer-resolvable labels in \mathcal{F} are inert under substrate symmetries; (A4) is the canonical preservation condition that the spectrum of G_{acc} is left invariant by the refinement. (A3) is technically implied by (A1) when \mathcal{F} carries no inherent action of G , but it is convenient to state separately to make the equivariance explicit. (A4) is the operative content and constrains the form of the lifted bookkeeping.

Canonical lift forced by access-geometry preservation

Proposition 3.2 (Common-multiplicative lift forced by (A4)). *Let $(\tilde{K}_{\text{rec}}, \tilde{M}_{\text{rec}})$ be positive operators on $\tilde{\mathcal{R}}$ of the elementary tensor form*

$$\tilde{K}_{\text{rec}} = K_{\text{rec}}^{\mathcal{R}} \otimes H_K, \quad \tilde{M}_{\text{rec}} = M_{\text{rec}}^{\mathcal{R}} \otimes H_M,$$

with H_K, H_M positive in \mathcal{F} . Then assumption (A4) holds if and only if $H_K = H_M$ up to scalar normalization.

Proof. By functional calculus on the second tensor factor,

$$\tilde{M}_{\text{rec}}^{-1/2} = (M_{\text{rec}}^{\mathcal{R}})^{-1/2} \otimes H_M^{-1/2},$$

so

$$\tilde{G}_{\text{acc}} = G_{\text{acc}}^{\mathcal{R}} \otimes \left[H_M^{-1/2} H_K H_M^{-1/2} \right].$$

This equals $G_{\text{acc}}^{\mathcal{R}} \otimes \mathbf{1}_{\mathcal{F}}$ if and only if $H_M^{-1/2} H_K H_M^{-1/2} = \mathbf{1}_{\mathcal{F}}$, that is $H_K = H_M$. Normalization is fixed by $\tau_{\mathcal{F}}(H_K) = \tau_{\mathcal{F}}(H_M) = 1$. \square

Proposition 3.2 forces a specific structural form: independent additive lifts $\tilde{K}_{\text{rec}} = K^{\mathcal{R}} \otimes \mathbf{1} + \mathbf{1} \otimes K^{\mathcal{F}}$ are ruled out, mixed lifts with $H_K \neq H_M$ are ruled out, and only common-multiplicative lifts $\tilde{K}_{\text{rec}} = K^{\mathcal{R}} \otimes H$, $\tilde{M}_{\text{rec}} = M^{\mathcal{R}} \otimes H$ for a single positive $H \in \mathcal{F}$ are admissible. The simplest representative within this family is the strict form $H = \mathbf{1}_{\mathcal{F}}$.

Definition 3.3 (Canonical lift). The *canonical lift* of $(K_{\text{rec}}, M_{\text{rec}})$ to an admissible refinement $\tilde{\mathcal{R}}$ is the strict form

$$\tilde{K}_{\text{rec}} = K_{\text{rec}}^{\mathcal{R}} \otimes \mathbf{1}_{\mathcal{F}}, \quad \tilde{M}_{\text{rec}} = M_{\text{rec}}^{\mathcal{R}} \otimes \mathbf{1}_{\mathcal{F}}.$$

Definition 3.3 commits to \mathcal{F} being a pure labeling layer, carrying neither independent record stiffness nor independent record susceptibility. Observer-resolvable labels exist algebraically as projections in $\tilde{\mathcal{R}}$, but the substrate-level paired bookkeeping is flat across them. This is the architectural division of labor between substrate primitives and observer-resolvable refinement: the substrate carries the paired bookkeeping; the observer-resolvable layer carries the labels.

Refinement factorization theorem

Theorem 3.4 (Refinement factorization). *Let $\tilde{\mathcal{R}}$ be an admissible refinement of \mathcal{R} by \mathcal{F} in the sense of Definition 3.1. Then*

$$\tilde{\mathcal{R}} \cong \mathcal{R} \otimes \mathcal{F}$$

as W^* -algebras, with:

(i) *Trace factorization:* $\tau_{\tilde{\mathcal{R}}} = \tau_{\mathcal{R}} \otimes \tau_{\mathcal{F}}$.

(ii) *Equivariant structure:* for any substrate symmetry $G \curvearrowright \mathcal{R}$ extending to $\tilde{\mathcal{R}}$, the symmetric expectation E_G acts only on the first tensor factor.

(iii) *Access-geometry factorization* (Proposition 3.2, Definition 3.3): $\tilde{G}_{\text{acc}} = G_{\text{acc}}^{\mathcal{R}} \otimes \mathbf{1}_{\mathcal{F}}$, and the metric-compatible normal form $K_{\text{rec}} = M_{\text{rec}}^{1/2} L_{\text{acc}} M_{\text{rec}}^{1/2}$ lifts as $\tilde{K}_{\text{rec}} = \tilde{M}_{\text{rec}}^{1/2} \tilde{L}_{\text{acc}} \tilde{M}_{\text{rec}}^{1/2}$ with $\tilde{L}_{\text{acc}} = L_{\text{acc}}^{\mathcal{R}} \otimes \mathbf{1}_{\mathcal{F}}$.

Proof. By (A1), the lifted support projections $\hat{e} := e \otimes \mathbf{1}_{\mathcal{F}}$ commute with every element of $\mathbf{1}_{\mathcal{R}} \otimes \mathcal{F}$ inside $\tilde{\mathcal{R}}$: for any $f \in \mathcal{F}$, $(e \otimes \mathbf{1}_{\mathcal{F}})(\mathbf{1}_{\mathcal{R}} \otimes f) = e \otimes f = (\mathbf{1}_{\mathcal{R}} \otimes f)(e \otimes \mathbf{1}_{\mathcal{F}})$. The subalgebras $\mathcal{R} \otimes \mathbf{1}_{\mathcal{F}}$ and $\mathbf{1}_{\mathcal{R}} \otimes \mathcal{F}$ therefore mutually commute, and the algebra they generate inside $\tilde{\mathcal{R}}$ is isomorphic to the algebraic tensor product $\mathcal{R} \otimes \mathcal{F}$. Completion in the natural W^* -topology gives the spatial tensor product. (i) is (A2). (ii) follows from (A3): trivial action on \mathcal{F} gives $E_G(r \otimes f) = E_{G, \mathcal{R}}(r) \otimes f$. (iii) is Proposition 3.2 together with the canonical lift of Definition 3.3, plus functional calculus on the metric-compatible normal form. \square

Two distinct objects on a refinement

Theorem 3.4 produces two structurally different operators on $\tilde{\mathcal{R}}$ that play different roles whenever a refinement is used in this paper. The distinction is essential: it is the architectural reason that measurement readouts can carry refined structure that the substrate itself does not.

The first object is the *substrate-level paired bookkeeping* $(\tilde{K}_{\text{rec}}, \tilde{M}_{\text{rec}})$, governed by (A4) and realized in the canonical lift. This pair is flat on the \mathcal{F} factor by construction: the substrate carries no independent record stiffness or susceptibility along the refinement labels.

The second object is any *measurement kernel* \tilde{K}^{meas} that the framework uses to express a sector-specific readout in terms of trace-quadratic structure on $\tilde{\mathcal{R}}$. Such kernels may carry non-trivial structure on the \mathcal{F} factor through projections that resolve the label content. They are derived measurement objects, not substrate-level paired-bookkeeping objects: their non-flatness reflects how measurement supports project against the algebraic labels, not a substrate-level loading of those labels.

Both object types live on the same algebra $\tilde{\mathcal{R}}$ but play different structural roles. The substrate stays branch-and-label neutral; measurement kernels may acquire label-resolved structure.

Recovery of the unrefined readout

Corollary 3.5 (Unrefined-limit recovery). *Let $\mathcal{E}_{\mathcal{F}} := \text{id}_{\mathcal{R}} \otimes \tau_{\mathcal{F}} : \tilde{\mathcal{R}} \rightarrow \mathcal{R}$ be the conditional expectation onto the unrefined retrieval algebra. Then for any element of the form $\tilde{K} = K_{\text{iso}} + K_{\mathcal{F}}$ with K_{iso} a scalar multiple of $\mathbf{1}_{\tilde{\mathcal{R}}}$ and $\tau_{\mathcal{F}}(K_{\mathcal{F}}) = 0$,*

$$\mathcal{E}_{\mathcal{F}}(\tilde{K}) = K_{\text{iso}}.$$

In particular, the canonical-lift access geometry coarse-grains exactly:

$$\mathcal{E}_{\mathcal{F}}(\tilde{G}_{\text{acc}}) = G_{\text{acc}}^{\mathcal{R}},$$

with no asymptotic error term.

Proof. Trace-product compatibility (Theorem 3.4(i)) gives $\mathcal{E}_{\mathcal{F}}(K_{\text{iso}}) = K_{\text{iso}}$ since K_{iso} is proportional to the identity on $\tilde{\mathcal{R}}$ with $\tau_{\mathcal{F}}(\mathbf{1}_{\mathcal{F}}) = 1$. The trace-centered condition $\tau_{\mathcal{F}}(K_{\mathcal{F}}) = 0$ gives $\mathcal{E}_{\mathcal{F}}(K_{\mathcal{F}}) = 0$. The access-geometry statement is immediate from Theorem 3.4(iii). \square

The corollary upgrades the informal convergence condition of Section 3.1 to an exact identity for canonical-lift refinements: the substrate-level access spectrum is recovered without error under coarse-graining over the refinement labels. Every spectral identity built on $G_{\text{acc}}^{\mathcal{R}}$ transfers to \tilde{G}_{acc} unchanged.

Layer-protection content and falsifier

The factorization theorem is the algebraic enforcement of the substrate–observer layer separation that runs through the entire paper. The substrate carries the paired bookkeeping; observer-resolvable refinements carry labels; the two layers do not mix at the level of substrate primitives.

Remark 3.6 (Layer-protection falsifier). A measurement or analysis requiring direct coupling between an observer-resolvable refinement label $f \in \mathcal{F}$ and the substrate-level support index in \mathcal{R} — i.e. a readout whose support projection cannot be factored as $P_{\mathcal{R}} \otimes P_{\mathcal{F}}$ within experimental or analytic uncertainty — would falsify the untwisted extension (A1) of Definition 3.1. The framework permits all observer-resolvable structure to enter through \mathcal{F} -side support weights; it does not permit such structure to modify substrate-level support primitives. Concrete instances of this falsifier appear in the Drell–Yan extension of the charged-sector readout (Section 15.10), in the access-depth refinement of the Schur expansion (Section 6.7), in the Type II₁ to Type III₁ modular embedding (Section 13), and in the algebraic shell-growth structure of Section 12.5.

3.4 Pre-oriented inclusion-boundary symmetry

The sectorwise-isotropic convention also fixes how an inclusion boundary is read before an observer has selected a record orientation. The substrate itself is tenseless: it contains no primitive clock, no built-in arrow, and no preferred side of a prospective access boundary. Thus an elementary inclusion boundary has two formal sides only after it is viewed as an interface between a smaller accessible algebra and a larger accessible algebra; before a memory-bearing observer chooses an orientation, neither side is distinguished by the substrate.

Let an unoriented inclusion interface be represented abstractly by the two formal sides of a normalized elementary boundary. The baseline isotropy and covariance hypotheses imply that these two sides are equivalent as substrate boundary data: any assignment that weights one side differently from the other would introduce a preferred record direction before the observer cut. In the Type II₁ setting the normalized trace is unique, so trace-equivalent boundary sides must receive the same scalar trace assignment. This substrate statement fixes equality of the two sides; the oriented numerical weight is produced only after the observer layer selects one side as the recorded side of the inclusion.

The full elementary boundary is therefore treated as the scalar unit of the inclusion interface, while the substrate-level symmetry says only that its two pre-oriented sides are trace-equivalent. The half-boundary value used below is not an additional substrate postulate; it is the observer-side consequence of combining this equality with the normalized unit boundary and then selecting one oriented side as record order.

3.5 Regional defect and master disturbance

For an admissible observer-access region R , let

$$E_R : \mathcal{W} \rightarrow \mathcal{A}_R$$

be the conditional expectation [12, 13] onto the accessible regional algebra, and let ρ_R denote the regional representative of the state in a finite approximant or the corresponding state restriction in the algebraic setting. The regional defect is

$$\mathfrak{D}_R(\rho) = \rho_R - E_R(\rho),$$

understood abstractly as the part of the substrate state or support data that is not losslessly represented by the chosen regional algebra. The master disturbance functional is

$$\mathcal{K}_R(\rho, \nabla, g) = \tau[\mathfrak{D}_R(\rho)^* K_{\text{rec}}(\nabla, g, R) \mathfrak{D}_R(\rho)].$$

Here $K_{\text{rec}}(\nabla, g, R)$ is the regional restriction of the universal record stiffness operator, τ is the canonical trace in the finite-factor approximation, and the paired susceptibility $M_{\text{rec}}(R)$ enters through the recovery-loss / record-susceptibility functional of the same defect (made explicit in Section 11). In the purely algebraic setting the same formula denotes the trace-quadratic normal form of the defect geometry, and the paired operator G_{acc} becomes the controlling object for record propagation.

The named strain quantities are defined only as readout channels of \mathcal{K}_R . *Factorization readout* is the readout measuring failure of independent regional product structure. *Coarse-graining readout* is the readout measuring irrecoverability under E_R . *Regional readout* is the combined localization and access-defect readout over a partition. They are not independent primitive actions; they are projections, restrictions, or approximations of one regional disturbance bookkeeping.

None of these refers to a parameter that flows. Each is a static feature of the structure.

3.6 Substrate-level identifications

The substrate layer reads the following operational signatures as static features of the totality.

Space. The stable separability pattern of the substrate. When the regional disturbance has a deep, well-separated minimum at a partition Π_* , the regions of Π_* act as neighborhoods. Spatial proximity is small information distance under the disturbance geometry G_{acc} ; spatial extension is the number and relation of stable independent regions.

Mass. Persistent localized disturbance. A stable substrate excitation at a region R is an algebraic perturbation that persists under restricted automorphisms and increases \mathcal{K}_R relative to a reference state. Its substrate-level mass readout is the local disturbance increment, weighted through K_{rec} , not an independent primitive charge.

Inertia. Resistance to algebraic support deformation. A persistent excitation has a finite space of identity-preserving deformations of its relational support. These are not translations through pre-existing space; they are changes in which subalgebraic factors, joins, and support relations are included in the observer-access cut. The Hessian of \mathcal{K}_R along this support-deformation manifold is the inertial readout, controlled by K_{rec} .

Gravity-like response. Connection variation of the same disturbance. A substrate with a non-integrable connection sector exhibits loop holonomy whose response to localized disturbance is the substrate-level analog of gravitational sourcing. Variation of the paired stiffness-susceptibility geometry under connection deformations supplies the response channel.

Gauge phase. The holonomy of an internal automorphism sector. Local choices of phase under an internal symmetry can disagree across overlapping regions; the closed-loop mismatch is the gauge readout of the same connection-dependent disturbance geometry.

These are organizational identifications, not completed derivations. The substrate layer's proposal is that these operational signatures are read more cleanly as different faces of one tenseless paired regional disturbance $(K_{\text{rec}}, M_{\text{rec}})$ than as separately constructed phenomena.

3.7 Non-Abelian selection and the Wilson-defect criterion

The substrate-level criterion for when an internal connection sector must be non-Abelian, rather than reducible to a $U(1)$ phase ledger, is supplied by an algebraic compatibility condition on the substrate's admissible coarsening transitions. The criterion applies to every internal automorphism sector of the substrate uniformly — gauge sectors, frame-bundle sectors of an internal

geometry, and any other compact internal automorphism group — and is invoked at the gauge-sector recovery of Section 11.5 and at any substrate sector whose holonomy is read through closed-loop transport.

Non-Abelian selection. The same recovery pattern also explains when a non-Abelian sector is needed. Let edge transports take values in a compact group G , with

$$U_{ij} \in G, \quad U_{ji} = U_{ij}^{-1}.$$

For a loop $\gamma = (i_0, i_1, \dots, i_n, i_0)$, define

$$H_\gamma = U_{i_0 i_1} U_{i_1 i_2} \cdots U_{i_n i_0}.$$

If all admissible coarsening transitions commute, an Abelian phase sector is sufficient. If there are admissible low-disturbance coarsenings with

$$[T_{RS}, T_{ST}] \neq 0,$$

then Abelian holonomy cannot represent the order-dependence of the regional join. The low-disturbance stationary connection must then live in a noncommutative compact sector, and closed-loop mismatch is measured by the conjugacy-invariant Wilson defect

$$W_\gamma = \Re \operatorname{Tr}(I - H_\gamma).$$

This does not derive the Standard Model. It gives a selection rule: non-Abelian holonomy is recovered when algebraic coarsening itself contains noncommuting transport data that cannot be encoded by $U(1)$ phases.

4 Type II₁ trace-quadratic normal form and cone recovery

The previous section closed the substrate layer with the paired bookkeeping $(K_{\text{rec}}, M_{\text{rec}})$, the master disturbance functional \mathcal{K}_R , and a list of substrate-level identifications. Before the observer layer is introduced, two operator-algebraic results that are used repeatedly downstream are recorded here. Both rely on tracial Type II₁ geometry, both anchor multiple later derivations to the same canonical trace, and neither requires any observer-side structure. Placing them here removes the forward references on which Sections 8–9 and Section 15 silently relied in earlier versions of this manuscript.

The first result, the Type II₁ isotropic quadratic disturbance normal-form theorem, is the structural reason the shared-generator argument of Section 9 produces a single canonical coefficient rather than independent inertial, gravity-like, and gauge coefficients. The second result, the cone-recovery trace lemma, is the structural reason the three primitive spatial retrieval projections of the observer cut carry equal trace, and is the load-bearing hypothesis of the Drell-Yan projection theorem of Section 15.10.

4.1 Conditional-expectation defects and the Type II₁ isotropic quadratic disturbance theorem

The quadratic paired-disturbance geometry adopted in the substrate layer should not remain a bare physical postulate. A rigorous foundation result of the framework is that, in a real finite von Neumann factor, the proposed quadratic paired disturbance geometry is not merely a matrix-proxy choice once an isotropic local defect sector is assumed. It is the local normal form of a positive, trace-covariant, isotropic functional measuring the failure of regional conditional expectations to compose over a join, and it controls both the stiffness K_{rec} and the susceptibility M_{rec} through their common canonical trace.

Let \mathcal{M} be a Type II₁ factor with faithful normal tracial state τ . Let $\mathcal{A}, \mathcal{B} \subset \mathcal{M}$ be von Neumann subalgebras and let

$$\mathcal{C} := \mathcal{A} \vee \mathcal{B}$$

be the von Neumann algebra they generate. Since \mathcal{M} is finite, each von Neumann subalgebra $\mathcal{N} \subset \mathcal{M}$ admits a unique τ -preserving conditional expectation [12, 13]

$$E_{\mathcal{N}} : \mathcal{M} \rightarrow \mathcal{N}.$$

The trace defines the noncommutative Hilbert space

$$L^2(\mathcal{M}, \tau), \quad \|x\|_{2,\tau}^2 := \tau(x^*x),$$

and each $E_{\mathcal{N}}$ extends to the orthogonal projection

$$P_{\mathcal{N}} : L^2(\mathcal{M}, \tau) \rightarrow L^2(\mathcal{N}, \tau).$$

Define the regional join defect on the joined support by

$$D_{\mathcal{A},\mathcal{B}} := (P_{\mathcal{C}} - P_{\mathcal{A}}P_{\mathcal{B}})|_{L^2(\mathcal{C},\tau)}.$$

Equivalently, for $x \in L^2(\mathcal{C}, \tau)$,

$$D_{\mathcal{A},\mathcal{B}}x = x - E_{\mathcal{A}}E_{\mathcal{B}}x.$$

This defect measures the part of a joined regional perturbation that is lost or distorted when the joined description is compressed sequentially through the two regional expectations.

Theorem 4.1 (Type II₁ isotropic quadratic disturbance normal form). *Let \mathcal{M} be a Type II₁ factor with trace τ , and let $\mathcal{A}, \mathcal{B} \subset \mathcal{M}$ be von Neumann subalgebras with $\mathcal{C} = \mathcal{A} \vee \mathcal{B}$. Then*

$$\sigma_{\mathcal{A},\mathcal{B}}(x) := \|D_{\mathcal{A},\mathcal{B}}x\|_{2,\tau}^2, \quad x \in L^2(\mathcal{C}, \tau),$$

is a positive, trace-intrinsic, internally unitarily covariant quadratic disturbance functional on local support perturbations. Moreover, if a local disturbance functional F of the defect variable satisfies

1. $F(0) = 0$;
2. $F(y) \geq 0$ near $y = 0$;
3. F is twice Fréchet differentiable at 0;
4. F is invariant under trace-preserving internal unitary conjugacies; and
5. the local defect sector is isotropic, so no internal defect direction is assigned a preferred stiffness,

then its leading nonzero term has the form

$$F(D_{\mathcal{A},\mathcal{B}}x) = c \|D_{\mathcal{A},\mathcal{B}}x\|_{2,\tau}^2 + o(\|D_{\mathcal{A},\mathcal{B}}x\|_{2,\tau}^2), \quad c \geq 0.$$

If the disturbance readout is nondegenerate on the defect sector, then $c > 0$.

Proof. Because \mathcal{M} is a Type II₁ factor, τ is finite, faithful, normal, and tracial. Hence

$$\langle x, y \rangle_{2,\tau} := \tau(y^*x)$$

defines a Hilbert-space inner product on $L^2(\mathcal{M}, \tau)$. For every von Neumann subalgebra $\mathcal{N} \subset \mathcal{M}$, the τ -preserving conditional expectation $E_{\mathcal{N}}$ is L^2 -contractive and extends to the orthogonal

projection $P_{\mathcal{N}}$ onto $L^2(\mathcal{N}, \tau)$. Therefore $P_{\mathcal{A}}, P_{\mathcal{B}}, P_{\mathcal{C}}$ are orthogonal projections on $L^2(\mathcal{M}, \tau)$, and $D_{\mathcal{A}, \mathcal{B}}$ is a bounded local defect operator on $L^2(\mathcal{C}, \tau)$. It follows immediately that

$$\sigma_{\mathcal{A}, \mathcal{B}}(x) = \langle D_{\mathcal{A}, \mathcal{B}}x, D_{\mathcal{A}, \mathcal{B}}x \rangle_{2, \tau} \geq 0,$$

with equality precisely when $D_{\mathcal{A}, \mathcal{B}}x = 0$.

Let $u \in \mathcal{M}$ be a unitary implementing an internal relabeling of the regional data:

$$\mathcal{A} \mapsto u\mathcal{A}u^*, \quad \mathcal{B} \mapsto u\mathcal{B}u^*, \quad \mathcal{C} \mapsto u\mathcal{C}u^*.$$

Since τ is tracial,

$$\|uxu^*\|_{2, \tau}^2 = \tau(ux^*xu^*) = \tau(x^*x) = \|x\|_{2, \tau}^2.$$

The expectations transform covariantly,

$$E_{u\mathcal{N}u^*}(uxu^*) = uE_{\mathcal{N}}(x)u^*,$$

so

$$D_{u\mathcal{A}u^*, u\mathcal{B}u^*}(uxu^*) = uD_{\mathcal{A}, \mathcal{B}}(x)u^*.$$

Thus

$$\sigma_{u\mathcal{A}u^*, u\mathcal{B}u^*}(uxu^*) = \sigma_{\mathcal{A}, \mathcal{B}}(x),$$

which proves internal unitary covariance.

Now let F be any positive twice differentiable local disturbance functional of a defect variable y with $F(0) = 0$. Since 0 is a local minimum, the first variation vanishes: $DF(0) = 0$. Taylor expansion gives

$$F(y) = \frac{1}{2}D^2F(0)[y, y] + o(\|y\|_{2, \tau}^2).$$

The Hessian is a positive semidefinite quadratic form. Internal unitary covariance and local isotropy rule out preferred directions in the defect sector; the only admissible leading quadratic form is therefore a scalar multiple of the L^2 inner product. Hence

$$D^2F(0)[y, y] = 2c\|y\|_{2, \tau}^2$$

for some $c \geq 0$. Setting $y = D_{\mathcal{A}, \mathcal{B}}x$ proves the claim. \square

The theorem provides the substrate-level normal form used by the physical reading of the paper [4, 5, 7]. The Type III₁ case relevant to local relativistic quantum field theory remains a modular version of this result: one must replace the trace geometry by the faithful-state GNS/modular geometry [9, 10, 6, 69, 70] and control the existence or replacement of ω -preserving expectations. But the finite factor theorem already removes the main matrix-proxy objection: quadratic paired regional disturbance is the local normal form of conditional-expectation defect under standard stability and covariance hypotheses, and the canonical trace anchors both K_{rec} and M_{rec} to the same Hilbert–Schmidt geometry.

Corollary 4.2 (Shared disturbance geometry and composition independence). *Assume the Type II₁ hypotheses above and let a persistent composite support be described by defect variables $D_{R, S}^{X_1, \dots, X_k}$ for all relevant regional joins. If its total substrate disturbance is the sum of local defect contributions,*

$$\Sigma_{\mathcal{A}}(X_1, \dots, X_k) = \sum_{R, S \subset \mathcal{A}} F\left(D_{R, S}^{X_1, \dots, X_k}\right),$$

with the same positive, covariant local functional F satisfying Theorem 4.1, then every contribution to the composite support enters through the same leading paired disturbance coefficient c :

$$\Sigma_{\mathcal{A}} = c \sum_{R, S \subset \mathcal{A}} \|D_{R, S}^{X_1, \dots, X_k}\|_{2, \tau}^2 + o\left(\sum_{R, S} \|D_{R, S}^{X_1, \dots, X_k}\|_{2, \tau}^2\right).$$

Therefore the inertial Hessian carried by K_{rec} and the gravity-like connection variation of the composite inherit one common trace-normalized disturbance geometry, and the record susceptibility carried by M_{rec} matches through the same canonical trace. Internal composition can change the shape and magnitude of the defect, but it cannot introduce a second canonical coupling for binding, kinetic, electromagnetic, or rest-like contributions.

Proof. Apply Theorem 4.1 to each local defect contribution. The same trace τ and the same local normal form determine each leading quadratic term. Summing defects changes the total quadratic form but not the source of its overall coefficient. The inertial stiffness and gravity-like connection response are variational readouts of this same total disturbance with respect to support and connection parameters, so they inherit the same local paired quadratic geometry. Distinct composition-dependent couplings would require distinct canonical quadratic geometries on the same defect sector, contradicting the single normal form under the stated covariance and isotropy hypotheses. \square

This corollary is the paper’s substrate-side version of composition independence. It says that if all physical contributions enter as defects of one persistent regional support, then all such contributions participate in the same paired disturbance geometry. The result does not yet produce spacetime free fall, but it prevents the substrate from assigning one gravity-like coefficient to rest-like persistence and another to internal binding contribution. The recovered free-fall statement is supplied downstream by combining this corollary with the observer-side recovery of trajectories (Section 11.1 and Section 11.6).

Tracial limitation. The same Type II₁ trace that makes Theorem 4.1 available also makes modular flow trivial:

$$\sigma_t^\tau = \text{id}.$$

Hence the Type II₁ theorem cannot by itself generate Connes–Rovelli thermal time [24], observer arrow, or Lorentzian kinematics. The allocation is explicit: Type II₁ supplies a clean substrate normal form and the shared-generator limit; nontrivial observer time and thermodynamics require state-dependent modular structure on the observer-access side, or a future Type III₁ extension of the theorem.

4.2 Cone-recovery trace lemma

Forward reference and load-bearing role. This subsection establishes a second tracial result that the observer layer and the charged-sector diagnostics rely on. The minimal observer-access dimension theorem of Section 6.5 (Theorem 6.3) fixes the number of primitive spatial access branches at $b = 3$ as a Čech-cohomological obstruction theorem on the projective cover. It does not fix the record-support trace weights of the corresponding primitive retrieval projections e_1, e_2, e_3 of $G_{\text{acc}}|_{\Sigma_{\text{obs}}^3}$. The present subsection supplies the weights. The argument proceeds in two steps: first, a substrate-stability selection that identifies the Lorentz-recovering sector as the unique ordinary-stable observer sector under finite, nondegenerate, refinement-stable, branch-neutral record recovery; second, the trace lemma proper, which uses the spatial subgroup of the Lorentz-recovery group to identify the primitive projections as a single orbit and the substrate trace as preserved on it. The two steps are independent: the substrate-stability selection identifies *which* ordinary observer sector is recovered, and the trace lemma extracts the weight equivalence *within* that sector.

Substrate-stability selection. An ordinary-stable observer sector is one in which the recoverable record geometry, read through the paired bookkeeping $(K_{\text{rec}}, M_{\text{rec}})$, satisfies four conditions: (1) *finiteness*, a finite generalized record-speed bound $c_*^2 = \lambda_{\min}(G_{\text{acc}}) < \infty$, excluding instantaneous/Galilean sectors; (2) *nondegeneracy*, a nonzero and well-defined spectrum of G_{acc}

on the access support, excluding collapsed or disordered sectors; (3) *refinement stability*, the property that observer readouts converge under admissible refinements of the access cut, where an admissible refinement is a trace-preserving conditional-expectation enlargement $E_\lambda \rightarrow E_{\lambda'}$ with $\mathcal{A}_\lambda \subset \mathcal{A}_{\lambda'}$ preserving the paired form $(K_{\text{rec}}, M_{\text{rec}})$; (4) *branch neutrality*, the property that no primitive spatial retrieval branch is invariantly distinguished by an admissible refinement, equivalently that for any pair e_i, e_j of primitive projections there exists an admissible refinement under which they are exchanged. Sectors failing (1) or (2) are excluded as non-recovering; sectors failing (3) are excluded as refinement artifacts; sectors failing (4) are excluded as branch-marked anisotropic sectors. The remaining sector is the ordinary-stable observer sector. The Lorentz-recovery program of Section 11.3 is the construction of the access cone for this sector, and the existence and irreducibility of the spatial subgroup $G_{\text{sp}} \subset G_{\text{Lor}}(\Sigma_{\text{obs}})$ used in Theorem 4.3 below is the operator-algebraic content of conditions (3) and (4) in the recovered sector.

The Type II₁ normal-form theorem fixes the local quadratic geometry of conditional-expectation defects but says nothing about how primitive spatial retrieval supports are distributed across the spatial-access layer. A second operator-algebraic result, conditional on the framework's Lorentz-recovery program (Section 11.3), forces the primitive trace weights to be equal. This subsection states and proves that result.

Setup. Let $\Sigma_{\text{obs}} = S_{\text{obs}} \oplus \tau_{\text{obs}}$ be the minimal $(b, d) = (3, 4)$ observer cut of Section 6.6. Let $G_{\text{acc}} = M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2}$ be the generalized record operator and let $G_{\text{acc}}|_{S_{\text{obs}}}$ denote its restriction to the spatial-access layer. The *primitive stable spatial-retrieval projections* are the minimal nonzero spectral support projections of $G_{\text{acc}}|_{S_{\text{obs}}}$, denoted

$$e_1, e_2, e_3, \quad e_i e_j = 0 \text{ for } i \neq j, \quad e_1 + e_2 + e_3 = \mathbf{1}_{R_3}.$$

The finite trace algebra they generate is

$$R_3 := \text{Alg}\{e_1, e_2, e_3\}.$$

This is the retrieval-support algebra induced by the spatial-access layer through the paired bookkeeping. It is distinct from S_{obs} itself (the geometric/projective access layer) and from the seven-element projective cover used in the projective access-cover theorem of Section 15.8. The primitive cyclic obstruction Ω_{ijk} of Definition 6.1 and the primitive retrieval projections e_i of the present setup are different objects on the same three-branch cover: Ω_{ijk} is the gauge invariant of the projective transition data and lives in the central phase group A , while e_i is a spectral support projection of the paired record operator and lives in R_3 . The branch count $b = 3$ is common to both and is fixed by Theorem 6.3; the weight $\tau(e_i)$ is the object of the present lemma.

The Lorentz-recovery group. Inside the framework, the Lorentz-recovery group is not the Lorentz group acting abstractly on an external spacetime. It is the group of automorphisms of the recovered observer support that preserve the structures defining the recovered record cone:

$$G_{\text{Lor}}(\Sigma_{\text{obs}}) := \{g \in \text{Aut}(\Sigma_{\text{obs}}) : g^* G_{\text{acc}} g = G_{\text{acc}}, \tau(g e g^{-1}) = \tau(e) \text{ for admissible retrieval projections } e\}.$$

Equivalently, $G_{\text{Lor}}(\Sigma_{\text{obs}})$ is the group of admissible observer-frame transformations that preserve the paired record ratio, the recovered cone, and the trace state on admissible record supports. The recovered Lorentz kinematics of Section 11.3 is the existence and identification of $G_{\text{Lor}}(\Sigma_{\text{obs}})$ as a nontrivial Lorentz-like group; the present lemma uses only its existence, not its full structure.

The spatial subgroup. The equal-trace argument needs only the spatial part of $G_{\text{Lor}}(\Sigma_{\text{obs}})$. Let

$$G_{\text{sp}} \subset G_{\text{Lor}}(\Sigma_{\text{obs}})$$

denote the subgroup that fixes the record-ordering coordinate τ_{obs} and acts on S_{obs} . In an ordinary Lorentz-recovered observer layer, G_{sp} is the recovered spatial-rotation group, effectively $SO(3)$ or its admissible representation on primitive retrieval supports.

Theorem 4.3 (Cone-recovery trace lemma). *Let $R_3 = \text{Alg}\{e_1, e_2, e_3\}$ be the primitive retrieval-support algebra defined above, with the trace τ inherited from the substrate Type II₁ factor. Suppose the observer layer recovers a Lorentz-like record cone in the sense that the spatial subgroup $G_{\text{sp}} \subset G_{\text{Lor}}(\Sigma_{\text{obs}})$ exists, acts on S_{obs} irreducibly, and preserves the record-support trace. Then the three primitive retrieval traces are equal:*

$$\tau(e_1) = \tau(e_2) = \tau(e_3).$$

Proof. The primitive projections e_i are generated as spectral support projections of $G_{\text{acc}}|_{S_{\text{obs}}}$, so any G_{sp} -invariant proper subset of $\{e_1, e_2, e_3\}$ generates a G_{sp} -invariant proper subalgebra of R_3 , which corresponds to a G_{sp} -invariant proper subspace of S_{obs} . By irreducibility of the G_{sp} action on S_{obs} , no such proper invariant subspace exists, so no proper G_{sp} -invariant subset of $\{e_1, e_2, e_3\}$ exists. Hence the primitive projections form a single G_{sp} -orbit. For any pair e_i, e_j there exists $g \in G_{\text{sp}}$ with $ge_i g^{-1} = e_j$. By trace preservation, $\tau(e_j) = \tau(ge_i g^{-1}) = \tau(e_i)$. Applying this for all pairs gives $\tau(e_1) = \tau(e_2) = \tau(e_3)$. \square

Remark 4.4 (Trace preservation as a Type II₁ consequence). Trace preservation $\tau(geg^{-1}) = \tau(e)$ for admissible record supports does not need to be imposed as a separate hypothesis on G_{sp} . The substrate is a Type II₁ factor on which the trace state τ is unique up to a positive scalar; every normal *-automorphism of a Type II₁ factor preserves the trace [4, 5]. Once G_{sp} is identified as a group of normal *-automorphisms of the retrieval-support algebra R_3 inherited from the substrate, trace preservation on R_3 follows automatically, and the trace-preservation clause in the definition of $G_{\text{Lor}}(\Sigma_{\text{obs}})$ is a consistency condition rather than an additional postulate.

Scope of the lemma and its dependency. Theorem 4.3 establishes equal-trace equivalence of the three primitive spatial retrieval projections *conditional on* the existence and irreducibility of the spatial Lorentz-recovery subgroup G_{sp} acting on S_{obs} . The recovered Lorentz kinematics of Section 11.3 is the framework's commitment that $G_{\text{Lor}}(\Sigma_{\text{obs}})$ is nontrivial and that its spatial subgroup is effectively $SO(3)$ -like with irreducible three-dimensional action. If the cone-recovery program of Section 11.3 closes, the hypotheses of Theorem 4.3 are satisfied and equal trace follows. If it does not close, both Section 11.3 and Theorem 4.3 share the same residual gap: they require the framework's full Lorentz-recovery derivation, listed as continuation work in Section 21. The lemma is therefore conditional in exactly the same sense as the existing Lorentz-recovery result; it does not introduce a new framework hypothesis.

Use of the lemma. Theorem 4.3 is used in Section 15.10 as one of three hypotheses in the conditional Drell-Yan projection theorem, where equal trace of the primitive spatial retrieval projections is the structural prerequisite for the channel-weight ratio to take a fixed value. Other future applications — to gravity-side channel structures, to Drell-Yan analogs at other observables, and to any precision diagnostic whose channel weight depends on primitive spatial-retrieval support equivalence — will use the same lemma.

Forward use as a predictive tool. The Type II₁ normal-form theorem and the Schur-complement reduction machinery developed in this section are used for the first time as a predictive tool in Section 5, which identifies the mass-shape sector of $(K_{\text{rec}}, M_{\text{rec}})$ and applies the visible-hidden Schur decomposition to derive closed-form predictions for the tau and baryonic mean access ratios under named selection rules. That section is the immediate downstream consumer of the substrate normal form; the observer-layer development of Section 6 and the

metrological completion of Section 7 carry the predictive content of the mass-shape sector into laboratory-comparable form.

5 Projective-access multiplicity rules and the mass-shape sector

The substrate-layer normal form of Section 4 establishes that the paired bookkeeping $(K_{\text{rec}}, M_{\text{rec}})$ admits a Type II₁ trace-quadratic structure, and that conditional-expectation defects in this geometry can be diagnosed by Schur-complement reductions on the defect Hilbert space. The present section uses that machinery for the first time as a predictive tool: it identifies a structural sector of $(K_{\text{rec}}, M_{\text{rec}})$ — the *mass-shape sector* — whose dominant readouts compress under a small set of selection rules into a closed-form prediction for the charged-leptonic and baryonic mean access ratios.

The construction is intentionally limited in scope. The section produces a core lemma (Theorem 5.8) covering the tau and baryonic mean access ratios as structural consequences of the framework’s substrate geometry, and a separate provisional extension (Section 5.7) covering the muon access ratio under an additional rule whose derivation remains open. The cleaning of the underlying mathematics — in particular, the projective-access ladder lemma (Lemma 5.3) — is stated in provisional form, with an explicitly named derivation target attached. The section closes with a falsifier statement and a list of named proof obligations.

This section is the first in which the substrate normal form is applied as a tool rather than developed as a target. It uses the substrate’s Type II₁ Schur-complement machinery (Section 4), the four-axis observer-access lattice $d = 4$ that will be established stably in Section 6 (anticipated here through the count $\dim(\Sigma_{\text{obs}}^3) + 1 = 4$ of Section 2.3), and the normalized access capacity $A = \alpha^{-1}$ produced by the charged-sector Schur fine-structure diagnostic of Section 15. Its predictive readouts r_τ , r_B , and (provisionally) r_μ are Level-4 dimensionless invariants in the sense of the observer metrological completion of Section 7; they become laboratory-comparable mass ratios under that completion.

5.1 Visible–hidden Schur architecture

The defect Hilbert space underlying $(K_{\text{rec}}, M_{\text{rec}})$ decomposes into a visible block on which observer cuts can read effective sector operators, and a hidden block that contributes to those operators only through Schur complementation.

Definition 5.1 (Visible–hidden decomposition of the substrate). A *visible–hidden decomposition* of the substrate bookkeeping is an orthogonal splitting

$$\mathcal{H}_{\text{phys}} = \mathcal{H}_{\text{vis}} \oplus \mathcal{H}_{\text{hid}},$$

under which the paired operators take the block form

$$K_{\text{rec}} = \begin{pmatrix} K_{\text{vis}} & B_K \\ B_K^* & K_{\text{hid}} \end{pmatrix}, \quad M_{\text{rec}} = \begin{pmatrix} M_{\text{vis}} & B_M \\ B_M^* & M_{\text{hid}} \end{pmatrix},$$

with $K_{\text{vis}}, M_{\text{vis}}$ acting on the visible block, $K_{\text{hid}}, M_{\text{hid}}$ on the hidden block, and B_K, B_M the off-block couplings. The *effective visible operator* is the Schur reduction

$$A_{\text{phys}} = K_{\text{vis}} - B_K K_{\text{hid}}^{-1} B_K^*,$$

under the gauge $M_{\text{vis}} = I_{\text{vis}}, M_{\text{hid}} = I_{\text{hid}}, B_M = 0$.

The visible–hidden decomposition is not a new substrate axiom; it is the application of the substrate normal form of Section 4 to a chosen visible block. Different choices of visible block produce different effective operators A_{phys} . The selection rules below restrict the choice of hidden block to those with minimum rank consistent with the framework’s predictive scope.

Remark 5.2 (Susceptibility gauge). The choice $M_{\text{vis}} = I$, $M_{\text{hid}} = I$, $B_M = 0$ places the entire nontrivial defect structure into K_{rec} . This is a working gauge, not a substrate axiom. In any other susceptibility gauge the same effective operator A_{phys} is produced; the gauge fixes the operator-theoretic representation, not the predictive content. We adopt this gauge throughout the section to avoid carrying redundant indices.

5.2 Projective-access ladder lemma (provisional form)

Let $A = \alpha^{-1}$ denote the normalized visible access capacity of the charged-sector observer (Section 15, where α is the framework’s predicted fine-structure constant). A is not an independent input; it is the Level-1 invariant of the charged-sector Schur diagnostic.

Lemma 5.3 (Projective-access ladder, provisional form). *Assume that an unresolved hidden access mode is not an independent trace contribution but a projective phase defect of the same access channel as A . Under the projective-access normalization used in this section, one unresolved projective phase contributes a suppression factor $1/\pi$. The first two hidden projective modes therefore contribute*

$$\Pi_1 = A/\pi, \quad \Pi_2 = A/\pi^2.$$

Lemma 5.3 is stated in provisional form: the factor $1/\pi$ is treated as a normalization lemma rather than a derived theorem. A derivation would require specifying the projective access space, its invariant measure, and the precise averaged Schur kernel whose normalized integral equals $1/\pi$.

Principle 5.4 (Derivation target for the projective ladder). There exists a projective access space \mathcal{P} , an observer-compatible invariant measure $d\mu_{\mathcal{P}}$ on \mathcal{P} , and a normalized Schur access kernel $\kappa_{\text{proj}} : \mathcal{P} \rightarrow \mathbb{R}_{\geq 0}$, such that

$$\int_{\mathcal{P}} \kappa_{\text{proj}} d\mu_{\mathcal{P}} = \frac{1}{\pi}.$$

If the same averaging operates independently on each unresolved projective mode, the k th-order unresolved projective access contributes A/π^k , and the ladder of Lemma 5.3 extends to $\Pi_k = A/\pi^k$ for all $k \geq 1$.

The shape of the framework’s commitment to Lemma 5.3 is therefore: the existence of a projective phase space with these properties is a structural prediction of the substrate normal form; the specific value $1/\pi$ per unresolved mode is a normalization lemma whose derivation is named here and is one of the central open problems of the program (cf. Section 5.9).

5.3 Second-order projective closure

The hidden block of the mass-shape sector is restricted to the first two projective modes by a closure condition on the visible–hidden decomposition.

Principle 5.5 (Second-order projective closure). The mass-shape sector is the visible–hidden decomposition in which the visible block carries the charged-sector mass-access readouts and the hidden block retains only the first two unresolved projective modes,

$$\mathcal{H}_{\text{hid}}^{\text{mass}} = \text{span}(h_1, h_2), \quad \dim \mathcal{H}_{\text{hid}}^{\text{mass}} = 2,$$

contributing Π_1 and Π_2 respectively. Higher unresolved projective modes h_k with $k \geq 3$ are assigned to separate boundary-obstruction sectors and do not enter the mass-shape effective operator.

This is a structural restriction on the visible block, not a numerical truncation: the magnitude $\Pi_3 \approx 4.4$ is not negligible by itself, but the mass-shape sector is defined as the minimal projectively-closed visible–hidden block. The assignment of higher projective modes to separate sectors is consistent with the framework’s general practice of decomposing into minimal closed blocks rather than carrying all available structure in one block.

5.4 Sector multiplicity and sign rules

The visible block of the mass-shape sector contains the access traces of the charged-leptonic excitations and the baryonic mean. The coefficients with which A , Π_1 , and Π_2 enter each access ratio are determined by support counts over the observer-access lattice.

The framework’s observer-access dimension is $d = 4$, comprising the three projective spatial directions of Σ_{obs}^3 and the one record-ordering coordinate τ_{obs} (Section 2.3). The natural integer support counts on a four-axis lattice are:

- the full ordered support $4! = 24$ (all permutations of the four axes),
- the nonempty subset support $2^4 - 1 = 15$ (all nonempty subsets of the four-axis support),
- the pure-axis support 4 (one boundary per axis),
- the pure-axis-plus-scalar support $4 + 1 = 5$ (four axes plus the common scalar closure).

These four counts are the dimensional invariants of the four-axis observer-access lattice. The selection rules below assign each to a specific sector readout.

Principle 5.6 (Sector multiplicity rules). On the mass-shape visible block, sector access trace multiplicities are assigned by observer-access support type:

- (M1) *Tau sector*. The fully excited charged-leptonic support carries the full four-axis ordered support trace ($4! = 24$ units of A), four first-projective boundary defects (4 units of Π_1), and one terminal second-projective closure defect (1 unit of Π_2).
- (M2) *Baryonic mean sector*. The symmetric baryonic support occupies the nonempty subset support of the four-axis lattice ($2^4 - 1 = 15$ units of A), and its pure-axis plus scalar closure modes relax through the first projective hidden boundary ($4 + 1 = 5$ units of Π_1).

The sign with which each hidden defect enters the effective visible operator is determined by whether the hidden block dresses or relaxes the visible support.

Principle 5.7 (Schur sign rule). Additive Schur defects ($+B_K K_{\text{hid}}^{-1} B_K^*$) describe *dressing*: the hidden block adds support to a visible sector that does not already saturate its observer-access support type. Subtractive Schur defects ($-B_K K_{\text{hid}}^{-1} B_K^*$) describe *relaxation*: the hidden block removes redundant support from a visible sector whose observer-access support type already over-counts. Charged-leptonic excitations are dressed; the baryonic mean is relaxed.

The structural distinction is that tau and muon are sequential charged-leptonic excitations (their observer-access support is built up by ordered traversal of the access lattice and admits further dressing), while the baryonic mean is a composite occupancy sector (its observer-access support is the full nonempty-subset lattice, which over-counts the pure-axis and scalar modes that must therefore be relaxed away).

5.5 Core lemma: tau and baryonic mean

The lemmas and rules of Sections 5.2 to 5.4 compose into a closed-form prediction for two of the three mass-shape access ratios.

Theorem 5.8 (Two-mode projective access core lemma). *Let $A = \alpha^{-1}$ be the normalized visible access capacity of the charged-sector observer, and let Π_1 and Π_2 be the first two projective hidden contributions of Lemma 5.3. Under Principle 5.5 (second-order projective closure), Principle 5.6 (sector multiplicity rules), and Principle 5.7 (Schur sign rule), the tau and baryonic mean access ratios are*

$$\begin{aligned} r_\tau &= 4!A + 4\Pi_1 + \Pi_2, \\ r_B &= (2^4 - 1)A - (4 + 1)\Pi_1, \end{aligned}$$

where $r_\tau = m_\tau/m_e$ and $r_B = (m_p + m_n)/(2m_e)$ are the access trace ratios of the tau lepton and the symmetric baryonic mean, respectively.

Proof. The visible block of the mass-shape sector carries the access traces $\{r_\tau, r_B\}$ in the dressing gauge of Remark 5.2. By Principle 5.5, the hidden block is two-dimensional, with modes h_1 and h_2 contributing Π_1 and Π_2 respectively. By Principle 5.6(M1), the tau visible support carries $4! = 24$ units of A , with hidden couplings of multiplicity 4 to h_1 and 1 to h_2 . By Principle 5.7, the tau hidden contributions are additive, so the effective tau access ratio is the sum $4!A + 4\Pi_1 + \Pi_2$. By Principle 5.6(M2), the baryonic mean visible support carries $2^4 - 1 = 15$ units of A , with hidden coupling of multiplicity $4 + 1 = 5$ to h_1 . By Principle 5.7, the baryonic mean hidden contribution is subtractive, so the effective baryonic mean access ratio is $(2^4 - 1)A - (4 + 1)\Pi_1$. \square

The proof composes the selection rules of Sections 5.2 to 5.4. Each rule is named, and each is open to derivation from deeper substrate axioms. Theorem 5.8 is the strongest form in which the framework currently makes the tau and baryon predictions; sharpening the rules into derived theorems is one of the central continuations of the program (Section 5.9).

5.6 Numerical diagnostics

The numerical content of Theorem 5.8 is assembled in Table 1, using the CODATA 2022 value $\alpha^{-1} = 137.035999177$ as input.

Sector	Predicted	Measured	Relative error
$r_\tau = 4!A + 4\Pi_1 + \Pi_2$	3477.228283	3477.23(24)	$\sim 5 \times 10^{-7}$
$r_B = (2^4 - 1)A - 5\Pi_1$	1837.440421	1837.418168	1.21×10^{-5}

Table 1: Numerical content of Theorem 5.8. The tau measured value uses $m_\tau = 1776.86(12)$ MeV, $m_e = 0.51099895$ MeV, giving $r_\tau = 3477.23 \pm 0.24$; the parenthesized uncertainty is on the last two digits. The baryon measured value uses the CODATA proton/electron and neutron/electron mass ratios.

The tau prediction sits within the experimental uncertainty on m_τ at the central-value level, with predicted–measured residual two orders of magnitude smaller than the measurement uncertainty itself. The baryon prediction has residual $\sim 10^{-5}$, comparable to the precision goals of typical phenomenological mass formulas in the charged sector.

5.7 Provisional extension: muon access ratio

The muon access ratio admits a closed-form expression in the same projective ladder, but the coefficient structure requires an additional rule that is not part of the core lemma.

Principle 5.9 (Branch-pair averaging rule (provisional)). The first nontrivial charged-leptonic excitation carries a three-branch support modulo projective orientation, contributing $(3/2)A$ to the visible block. Its first projective defect couples to the hidden mode h_1 averaged over branch-pair observer channels, producing the suppression coefficient

$$\frac{1}{3^2 \cdot 4} = \frac{1}{36}.$$

Under Principle 5.9, the muon access ratio is

$$r_\mu = \frac{3}{2}A + \frac{\Pi_1}{36},$$

with $r_\mu = m_\mu/m_e$ the muon-to-electron access ratio. Numerically, $r_{\mu,\text{pred}} = 206.765663$ against the measured $r_{\mu,\text{exp}} = 206.768283$, a relative residual of 1.27×10^{-5} .

Principle 5.9 is flagged as provisional and is not part of the core lemma. The combinatorial reading $36 = 3^2 \cdot 4$ is compatible with the framework's $b = 3$ branch structure (Section 15) and the four-axis observer-access lattice, but several alternative decompositions of 36 are possible, and the specific assignment “branch-pair observer channels” has not been derived from the substrate side. The muon line is therefore included in this section as a structural *candidate* rather than as a structural prediction, and it does not enter the core lemma's predictive content.

Remark 5.10 (Asymmetric reliance on the rules). The tau and baryon predictions of Theorem 5.8 rely only on the natural integer support counts of the four-axis observer-access lattice $(4!, 2^4 - 1, 4, 5, 1)$. The muon prediction of Section 5.7 additionally requires the branch-pair averaging coefficient $36 = 3^2 \cdot 4$. The asymmetry is structural, not stylistic: tau and baryon are pinned to the dimensional invariants of the access lattice; the muon coefficient requires a separate structural rule. Until Principle 5.9 is derived from a substrate-side argument, the muon line should be read as a constraint on the framework's open program rather than as a confirmed prediction.

5.8 Falsifier and predictive scope

The core lemma supplies a sharp falsifier on the tau mass.

Proposition 5.11 (Tau-sector falsifier). *Theorem 5.8 predicts*

$$r_\tau = 4!A + 4\Pi_1 + \Pi_2 = 3477.22828 \pm \delta,$$

where δ tracks the propagated uncertainty from α^{-1} (CODATA: $\delta \sim 10^{-8}$) and is therefore negligible relative to current experimental tau-mass uncertainty. If improved measurements of m_τ/m_e place its central value outside the interval $[3477.225, 3477.231]$, the core lemma in the form of Theorem 5.8 is falsified.

The baryon prediction is similarly constrained but with a wider current residual; its falsifier interval is approximately $[1837.40, 1837.48]$, with the prediction sitting near the upper edge.

A failure of either falsifier would not invalidate the substrate normal form of Section 4 or the metrological completion of Section 7; it would falsify the specific combination of selection rules in Theorem 5.8. The framework's commitment is to the combination, not to each rule independently.

5.9 Open problems and named proof obligations

The section's predictive content rests on five named open problems. Each is a derivation target from substrate-level axioms.

- (O1) *Projective-access ladder derivation.* Identify the projective access space \mathcal{P} , the invariant measure $d\mu_{\mathcal{P}}$, and the averaged Schur kernel κ_{proj} such that $\int_{\mathcal{P}} \kappa_{\text{proj}} d\mu_{\mathcal{P}} = 1/\pi$ as in Principle 5.4. This pins Lemma 5.3 as a theorem rather than a normalization lemma.
- (O2) *Closure condition.* Derive Principle 5.5 from a substrate-side minimality principle on the visible–hidden decomposition. The mass-shape sector should be identified as the unique closed visible–hidden block under a specified completeness criterion, rather than chosen by convention.
- (O3) *Support-type assignment.* Derive Principle 5.6: show that the tau visible support is ordered (4!) and the baryonic mean visible support is subset $(2^4 - 1)$ from the framework’s existing sector classification, rather than imposed by the multiplicity rule.
- (O4) *Schur sign rule.* Derive Principle 5.7 from a substrate-level positivity criterion. The framework should explain why charged-leptonic dressings are additive while baryonic mean defects are subtractive, in terms of the Schur complement structure on the visible block.
- (O5) *Branch-pair averaging rule.* Derive Principle 5.9 (the muon coefficient $1/36$) from the $b = 3$ branch structure and the four-axis observer-access lattice, or determine that no such derivation exists and treat the muon line as a separate sector requiring independent structural input.

Each open problem is a target of the framework’s continuation program. Pinning any one of them as a derived theorem strengthens the section’s overall standing; resolving (O1) elevates Lemma 5.3 to a derived ladder and locks the projective normalization that the whole section depends on. Pinning all five would convert Theorem 5.8 into a theorem with all premises derived from substrate axioms.

5.10 Section status

The mass-shape sector is the first section in which the framework’s substrate normal form (Section 4) and the observer-access lattice (Section 2.3, Section 6) compose into a concrete predictive readout. Its claim status, in the manuscript’s claim-status taxonomy (Section 10), is mixed:

- *Structural claims:* The visible–hidden Schur decomposition (Definition 5.1) is a direct application of the substrate normal form and is not itself open.
- *Provisional lemmas:* The projective-access ladder (Lemma 5.3) is stated in provisional form with derivation target (O1) named.
- *Selection rules:* The closure condition (Principle 5.5), multiplicity rule (Principle 5.6), and sign rule (Principle 5.7) are selection rules with derivation targets (O2)–(O4) named.
- *Core lemma:* The tau and baryonic mean predictions (Theorem 5.8) are theorems modulo the selection rules.
- *Provisional extension:* The muon prediction (Section 5.7) depends additionally on (O5) and is not part of the core lemma.
- *Falsifier:* Proposition 5.11 is a sharp test on the tau prediction with named interval.

This stratification is the section’s safeguard against numerological misreading: every coefficient has an assigned structural origin, every structural origin has a named open derivation target, and the predictive claim is the composition of the rules under those targets rather than an independent fit to the data.

6 The observer layer

The observer layer represents the experiential side of physics. It proposes that what observers call temporal order, the arrow of time, motion, and collapse are features of a memory-bearing access structure

$$\mathcal{O} = (\Lambda, \preceq, \{\mathcal{A}_\lambda\}, \{\mathcal{M}_\lambda\}),$$

where Λ is a directed set of access labels, $\mathcal{A}_\lambda \subseteq \mathcal{W}$ is the accessible subalgebra at label λ , and $\mathcal{M}_\lambda \subseteq \mathcal{A}_\lambda$ is the subalgebra of stable, decohered records at λ . The order \preceq is derived from record-recoverability: $\lambda \preceq \mu$ when records at μ permit recovery of records at λ via a reverse channel.

6.1 Two primitive operations and one stability criterion

An observer performs two primitive access operations: *cut* and *accumulate*. Record persistence is not a third free operation; it is an admissibility criterion imposed by the paired substrate disturbance barrier.

Cut. At each label λ , the observer selects a conditional expectation $E_\lambda : \mathcal{W} \rightarrow \mathcal{A}_\lambda$. The substrate state ω is not modified by the cut; what changes is the observer's accessible marginal. In the strengthened junction formulation, the cut is not only a restriction of attention but the first map in a modular embedding:

$$(\mathcal{W}, \tau, K_{\text{rec}}, M_{\text{rec}}) \xrightarrow{E_\lambda} (\mathcal{A}_\lambda, \phi_\lambda) \xrightarrow{\pi_{\phi_\lambda}} \mathcal{M}_\lambda^{\text{obs}} := \pi_{\phi_\lambda}(\mathcal{A}_\lambda)''.$$

The observer's effective time appears at this junction when ϕ_λ is faithful and nontracial, because the represented algebra then carries the modular automorphism group [9, 10] $\sigma_t^{\phi_\lambda}$. The substrate has not begun to evolve; rather, the observer has acquired a nontrivial modular flow on the algebra through which the substrate is represented.

Accumulate. The observer's record subalgebra grows along \preceq . Once acquired, records are not erased: $\lambda \preceq \mu$ implies $\mathcal{M}_\lambda \subseteq \mathcal{M}_\mu$ up to controlled error.

6.2 Local entropy-maximizing observer state

The observer cut is not an external agent placed on top of the substrate. For each access algebra \mathcal{A}_λ , the observer-effective state is the least-biased state compatible with the records already stabilized in \mathcal{M}_λ . In finite approximants one may write this as

$$\phi_\lambda = \arg \max_{\phi \in \mathcal{S}(\mathcal{A}_\lambda)} \{S(\phi) : \phi(m) = \omega(m) \text{ for all } m \in \mathcal{M}_\lambda\},$$

with additional conserved record constraints included when present. Thus nontracial observer data are induced by partial access and record constraints, not introduced as an independent physical substance. If ϕ_λ is faithful and nontracial, its modular automorphism group supplies an intrinsic scale on the accessible algebra. The operational order, however, remains the inclusion order of stable records.

6.3 Inclusion index

The label λ is not primitive substrate time. It is an index of inclusion for stable access:

$$\mathcal{A}_{\lambda_1} \subseteq \mathcal{A}_{\lambda_2} \subseteq \cdots, \quad \mathcal{M}_{\lambda_1} \subseteq \mathcal{M}_{\lambda_2} \subseteq \cdots,$$

up to the controlled recovery errors permitted by the observer budget. The word “later” therefore means: *represented at a larger or more informative stable record algebra*, not located at a later point in substrate time. When a faithful nonracial observer state is present, the associated modular automorphism parameter gives an intrinsic scale for this record order; the operational order itself is the inclusion chain.

6.4 Stability criterion

Record persistence is not left as a purely free observer postulate. A candidate record must be supported in a region whose erasure requires crossing a disturbance barrier generated by the paired $(K_{\text{rec}}, M_{\text{rec}})$ geometry. For a record support R , write

$$\text{Rec}(R) = 1 \iff \Delta\mathcal{K}^{\text{erase}}(R) \geq \Theta_{\text{rec}},$$

or, in a coarse threshold form,

$$\sigma_R^{\text{rec}} := \Phi_{\text{rec}}(\mathcal{K}_R) \geq \sigma_{\text{rec}}.$$

Here σ_R^{rec} is not an independent primitive functional; it is the record-stability readout of the paired master disturbance. The monotone record condition $\mathcal{M}_\lambda \subseteq \mathcal{M}_\mu$ is then interpreted as an observer-access rule restricted to high-barrier supports: records are stable because admissible low-disturbance updates cannot erase them without crossing the substrate stiffness threshold supplied by K_{rec} relative to the susceptibility cost supplied by M_{rec} .

6.5 Minimal stable observer-access dimension

Admissible observer readouts are locally organized by a $3+1$ -dimensional access geometry. This is not a spacetime postulate imposed on the substrate. Its status in the framework is fixed by an algebraic theorem on the observer cut: $3+1$ is the minimal stable form of a memory-bearing projective cut, where “minimal” is read as the smallest branch count carrying a non-removable projective obstruction. The geometric language (orientation, parallax, occlusion) is recovered downstream as an interpretation of this obstruction; it is not the argument.

Predictive use of the dimension count. The minimal observer-access dimension established here is the structural input to the sector multiplicity rules of Section 5: the integer support counts $4! = 24$, $2^4 - 1 = 15$, 4 , and 5 that fix the tau and baryonic mean mass-shape coefficients in the two-mode projective-access core lemma are direct combinatorial invariants of the four-axis lattice $(b, d) = (3, 4)$ of this subsection. The substrate’s commitment to a four-axis observer-access lattice and its commitment to those specific mass-shape coefficients are therefore not independent inputs.

The observable projective structure of an access cover. Let an observer-access cover of Σ_{obs}^3 have b primitive branches, represented by mutually orthogonal projections

$$p_1, \dots, p_b, \quad p_i p_j = 0 \text{ for } i \neq j, \quad \sum_{i=1}^b p_i = \mathbf{1}_{\Sigma_{\text{obs}}^3}.$$

The raw projective cover indexes the nonempty branch supports $p_I = \sum_{i \in I} p_i$ for $\emptyset \neq I \subseteq \{1, \dots, b\}$, of which there are $2^b - 1$. The observer-relevant content of the cover is not the raw support family but its quotient by branch-local relabeling and projective phase, formalized as follows.

Definition 6.1 (Projective transition data). Let A be an Abelian gauge group (the relabeling/phase group of the cover; in the framework’s central-phase reading $A = U(1)$). A *projective*

transition datum on the cover is a family $u_{ij} \in A$ defined for ordered pairs $i \neq j$ with $u_{ji} = u_{ij}^{-1}$. A branch-local gauge transformation is a family $g_i \in A$ acting by

$$u_{ij} \longmapsto g_i u_{ij} g_j^{-1}.$$

For each ordered triple i, j, k of distinct branches, the primitive cyclic obstruction on $\{i, j, k\}$ is

$$\Omega_{ijk} := u_{ij} u_{jk} u_{ki}.$$

A direct computation shows that Ω_{ijk} is invariant under branch-local gauge: for any choice of g_1, \dots, g_b ,

$$(g_i u_{ij} g_j^{-1})(g_j u_{jk} g_k^{-1})(g_k u_{ki} g_i^{-1}) = g_i u_{ij} u_{jk} u_{ki} g_i^{-1} = u_{ij} u_{jk} u_{ki},$$

the last equality by abelianness of A . Thus Ω_{ijk} descends to the projective quotient.

Definition 6.2 (Non-removable projective closure). An access cover with b primitive branches admits non-removable projective closure if the projective quotient supports a gauge-invariant cyclic obstruction Ω_{ijk} that is not forced to be trivial by the structure of the cover.

Theorem 6.3 (Minimal non-removable projective closure). Let b denote the number of primitive branches of an observer-access cover with projective transition data taking values in an Abelian gauge group A . Then:

- (i) $b = 1$ admits no relational projective datum; the projective quotient is trivial.
- (ii) $b = 2$ carries only a single pairwise transition u_{12} , which is removable: there exists a branch-local gauge g_i such that $g_1 u_{12} g_2^{-1} = \mathbf{1}_A$.
- (iii) $b = 3$ admits a non-removable gauge-invariant cyclic obstruction $\Omega_{123} \in A$.
- (iv) For $b \geq 4$, the gauge-invariant content of the cover is determined entirely by the family $\{\Omega_{ijk}\}_{1 \leq i < j < k \leq b}$ of triangle obstructions, subject to the tetrahedral relation

$$\Omega_{jkl} \Omega_{ikl}^{-1} \Omega_{ijl} \Omega_{ijk}^{-1} = \mathbf{1}_A \quad (1 \leq i < j < k < l \leq b).$$

No new independent obstruction arises at multi-overlap level ≥ 4 .

Hence the minimal b at which a non-removable projective obstruction first appears is $b = 3$. Covers with $b \geq 4$ are composite refinements built from triangle obstructions and are nonminimal as primitive observer supports.

Proof. For (i), with one branch the cover is $\{p_1\}$ and there is no pair (i, j) on which a transition datum is defined; the projective quotient is empty of relational content.

For (ii), the only transition is u_{12} . Choose $g_1 = u_{12}^{-1}$, $g_2 = \mathbf{1}_A$. Then $g_1 u_{12} g_2^{-1} = u_{12}^{-1} u_{12} = \mathbf{1}_A$. There are no triples on $\{1, 2\}$, so no Ω exists. Every two-branch transition is therefore a coboundary of a 0-cochain and the quotient is trivial.

For (iii), $\Omega_{123} = u_{12} u_{23} u_{31}$ is gauge-invariant by the computation above. It is not forced to take the value $\mathbf{1}_A$ by any constraint of the cover: any element of A can be realized as Ω_{123} by appropriate choice of u_{12}, u_{23}, u_{31} (e.g. $u_{12} = a$, $u_{23} = u_{31} = \mathbf{1}_A$ gives $\Omega_{123} = a$). Hence Ω_{123} is a genuine gauge-invariant degree of freedom.

For (iv), the cover is the nerve of a b -element index set, combinatorially a $(b - 1)$ -simplex. Projective transition data are A -valued 1-cochains on this simplex; branch-local gauge transformations are 0-cochains acting by Čech coboundary. The cyclic obstructions Ω_{ijk} are the 2-cochain δu . For any 4-element subset $\{i, j, k, l\}$, the four triangles $\{jkl\}, \{ikl\}, \{ijl\}, \{ijk\}$ bound the tetrahedron $\{ijkl\}$, and the identity $\delta(\delta u) = 0$ at the 3-simplex level gives exactly

the tetrahedral relation displayed. Since the $(b - 1)$ -simplex is contractible, $\check{H}^n(\Delta^{b-1}; A) = 0$ for $n \geq 1$, so no nontrivial gauge invariants exist beyond what is captured by the cocycle structure on triangles: every higher multi-overlap invariant is a combination of triangle obstructions modulo the tetrahedral relations. In particular, at $b = 4$ there are four triangle obstructions and one tetrahedral relation, giving three independent invariants, all of which are triangle obstructions. No invariant of multi-overlap order ≥ 4 exists. Hence $b = 3$ is minimal, and $b \geq 4$ is composite. \square

Lemma 6.4 (Four-branch decomposition by explicit cancellation). *Let four primitive branches 1, 2, 3, 4 carry pairwise transitions u_{ij} valued in an Abelian group A , with $u_{ji} = u_{ij}^{-1}$. Define the four triangular cyclic obstructions*

$$\Omega_{123} = u_{12}u_{23}u_{31}, \quad \Omega_{124} = u_{12}u_{24}u_{41}, \quad \Omega_{134} = u_{13}u_{34}u_{41}, \quad \Omega_{234} = u_{23}u_{34}u_{42},$$

corresponding to the four faces of the simplex on $\{1, 2, 3, 4\}$, each oriented as $i \rightarrow j \rightarrow k \rightarrow i$ with $i < j < k$. Then the tetrahedral coboundary is trivial:

$$\Omega_{234} \Omega_{134}^{-1} \Omega_{124} \Omega_{123}^{-1} = \mathbf{1}_A.$$

Proof. Substitute the definitions and expand each inverse using abelianness of A :

$$\Omega_{234} \Omega_{134}^{-1} \Omega_{124} \Omega_{123}^{-1} = (u_{23}u_{34}u_{42}) (u_{41}^{-1}u_{34}^{-1}u_{13}^{-1}) (u_{12}u_{24}u_{41}) (u_{31}^{-1}u_{23}^{-1}u_{12}^{-1}).$$

Tabulate the transition factors appearing in the product, using the antisymmetry $u_{ji} = u_{ij}^{-1}$ to identify each factor with one of the six ordered pairs $\{(1, 2), (1, 3), (1, 4), (2, 3), (2, 4), (3, 4)\}$:

Pair	Net contribution to the product
(1, 2)	u_{12} (from Ω_{124}) \cdot u_{12}^{-1} (from Ω_{123}^{-1}) = $\mathbf{1}_A$
(1, 3)	u_{13}^{-1} (from Ω_{134}^{-1}) \cdot u_{13} (from u_{31}^{-1} in Ω_{123}^{-1}) = $\mathbf{1}_A$
(1, 4)	u_{41}^{-1} (from Ω_{134}^{-1}) \cdot u_{41} (from Ω_{124}) = $\mathbf{1}_A$
(2, 3)	u_{23} (from Ω_{234}) \cdot u_{23}^{-1} (from Ω_{123}^{-1}) = $\mathbf{1}_A$
(2, 4)	u_{24}^{-1} (from u_{42} in Ω_{234}) \cdot u_{24} (from Ω_{124}) = $\mathbf{1}_A$
(3, 4)	u_{34} (from Ω_{234}) \cdot u_{34}^{-1} (from Ω_{134}^{-1}) = $\mathbf{1}_A$

Each of the six edges contributes one factor with positive orientation and one with negative orientation. The product is $\mathbf{1}_A$ by abelianness. \square

The Four-Branch Decomposition Lemma is the explicit verification, by pairwise cancellation, of the tetrahedral relation stated abstractly in part (iv) of Theorem 6.3. The general $b \geq 4$ case follows by applying Lemma 6.4 to each 4-element subset of branches.

Remark 6.5 (Cohomological framing). The argument above is the standard reading of projective transition data as a Čech 1-cochain with coefficients in A on the nerve of the cover, with branch-local gauge as the 0-coboundary action and cyclic obstructions as the resulting 2-cocycle δu . The theorem is the specialization of that machinery to a finite simplicial nerve. The abelianness of A is essential to the precise decomposition statement in part (iv); a non-Abelian gauge group would yield a non-Abelian \check{H}^1 classification with higher Postnikov data, and the “no new invariant beyond triangles” conclusion would require separate justification. The framework’s central-phase gauge group is Abelian throughout (the residual $U(1)$ footprint of Section 15.3), so the abelian restriction is the natural one here.

Remark 6.6 (Scope: central quotient and the universal observer cut). The restriction to Abelian coefficients A in Definition 6.1 is not a claim that the framework’s physics is Abelian. It is a claim about what the *primitive observer cut* reads. An observer compares record supports through quantities invariant under branch-local relabeling — phases, traces, determinants, central characters, scalar visibility weights — which are central elements of the projective transition data;

non-central transition data is, by construction, not directly observer-readable, because it rotates inside an internal fiber without changing external record localization. Non-Abelian transition data, when present (color, weak isospin, the full $SU(3) \times SU(2) \times U(1)$ structure of Section 15.3), therefore describes internal sector structure built *over* the observer-access layer rather than the universal access layer itself. The branch-count theorem accordingly fixes the number of primitive spatial access branches as a property of the central projective quotient, independent of which internal gauge representation a recorded excitation carries. This universality is the property required of an ordinary observer cut: spatial localization must not depend on the gauge content of the recorded excitation. Higher non-Abelian obstructions, if present, belong to refined or sector-specific structure and do not modify the primitive observer branch count.

Remark 6.7 (Disambiguation from connection holonomy). The cyclic obstruction Ω_{ijk} is a projective invariant of the observer-access cover and should not be confused with the connection holonomies of the charged-sector chapter (Section 15). The latter are Wilson-loop transports of an internal connection around closed paths; the former is a combinatorial invariant of branch-local relabeling on a finite projective cover.

Remark 6.8 (Refinement stability and the bisection objection). Theorem 6.3 is a statement about a single primitive observer cut. Iterated binary refinement of a two-branch cover can classify accessible records arbitrarily finely, but each individual primitive cut along the refinement history is a $b = 2$ cut and carries only removable relative phase data by part (ii). A non-removable obstruction could be generated only by promoting a refinement history to a higher composite cover, at which point the recovered projective content depends on the refinement history rather than on the primitive cut. A primitive observer-access cover that supports refinement-stable record recovery must therefore already carry the non-removable projective obstruction at the level of the cut itself, which by parts (i)–(iii) forces $b \geq 3$. Minimality then forces $b = 3$.

Corollary 6.9 (Minimal observer-access dimension). *Combining Theorem 6.3 with the record-ordering requirement, the minimal stable observer-access layer is*

$$O_{\text{acc}} = \Sigma_{\text{obs}}^3 \oplus \tau_{\text{obs}}, \quad \dim \Sigma_{\text{obs}}^3 = b = 3, \quad \dim O_{\text{acc}} = d = 4.$$

The value $d = 4$ is supplied directly by the addition of one record-ordering coordinate to the three-branch spatial-access layer, and is independently selected within the framework by the substrate-observer duality closure conditions developed in Section 15.12, namely the algebraic identities $21d/2 = 7d(d-1)/2$ and $4d = 2^d$, each of which fixes $d = 4$ uniquely. The duality identities are internal-consistency conditions on the framework's chosen primitives ($T_{\text{bare}} = 21/2$, $b = 3$, projective trace counting) and not external selection rules, but their closure at the same $d = 4$ supplies an algebraic check independent of the record-ordering interpretation.

Proof. Theorem 6.3 fixes $b = 3$. An observer cut is memory-bearing only if localized records are sequenced under accumulation, which adds one record-ordering coordinate τ_{obs} , giving $d = b + 1 = 4$. The independent algebraic closure at $d = 4$ is the result of Section 15.12. \square

The asymmetry between the three spatial access directions and the single record-ordering direction is therefore structural. Projective closure lives on Σ_{obs}^3 because non-removable cyclic obstructions live on its branch cover. The record-ordering direction τ_{obs} supplies sequencing along the inclusion-index \preceq and contributes a one-sided boundary trace to the access bookkeeping (Section 6.6), but it does not generate independent spatial projective sectors. The projective-cover diagnostics of the charged sector therefore use $b := \dim \Sigma_{\text{obs}}^3 = 3$ and $d := \dim O_{\text{acc}} = 4$, both derived from the present subsection.

Remark 6.10 (Branch count versus branch weight). Theorem 6.3 fixes the *number* of primitive branches; it does not fix their *trace weights*. Given the $(b, d) = (3, 4)$ observer cut, the paired record bookkeeping induces three primitive stable spectral support projections e_1, e_2, e_3

of $G_{\text{acc}}|_{\Sigma_{\text{obs}}^3}$ generating the retrieval-support algebra $R_3 = \text{Alg}\{e_1, e_2, e_3\}$. The equal-trace statement $\tau(e_1) = \tau(e_2) = \tau(e_3)$ is supplied separately by the cone-recovery trace lemma (Theorem 4.3 of Section 4.2), conditional on the Lorentz-recovery program of Section 11.3. The branch-count theorem and the branch-weight lemma are independent results and rest on independent hypotheses; in particular, the branch-count result does not require the Lorentz-recovery program to close.

6.6 One-sided record-boundary trace

The preceding subsection fixes the minimal observer-access split, but the record-order coordinate still requires a scalar normalization. This normalization is obtained from the inclusion structure of the observer cut rather than from an added time coordinate. Consider a normalized inclusion pair

$$\mathcal{A}_- \subset \mathcal{A}_+,$$

where \mathcal{A}_- is the already stabilized record algebra and \mathcal{A}_+ is the next admissible record algebra. Let the associated trace-preserving access maps be

$$E_- : \mathcal{W} \rightarrow \mathcal{A}_-, \quad E_+ : \mathcal{W} \rightarrow \mathcal{A}_+.$$

The algebraic increment of the observer cut is

$$B_{\text{rec}} := E_+ - E_-.$$

It extracts the newly record-accessible part of the substrate relative to the previous stable cut. Thus record order is represented by an inclusion increment, not by a primitive temporal displacement.

The elementary inclusion boundary is normalized to one scalar trace unit under the same observer-cut normalization used for projective-cover weights. Before memory orientation, Section 3.4 gives two trace-equivalent formal sides of that boundary. Type II₁ trace uniqueness assigns equal scalar weight to trace-equivalent sides. Since the full elementary boundary carries one scalar unit, each pre-oriented side carries one half of that unit. A memory-bearing observer cut then selects only the stabilizing, recorded side of the boundary. Therefore the record-order coordinate contributes one half-boundary trace unit,

$$\tau_{\partial \text{rec}} = \frac{1}{2},$$

while contributing no projective-cover branch.

This is the algebraic distinction between b and d . The three spatial directions produce projective localization branches and are counted by $b = 3$. The record-order coordinate is part of the full observer-access dimension $d = 4$, but it is a one-sided inclusion boundary rather than a fourth spatial projection. It can therefore contribute a scalar boundary weight without changing the projective branch count.

6.7 Access resolution and the observer-layer conjugate structure

The substrate is tenseless and carries no clock; the observer layer is where dynamical conjugate structure enters. The natural conjugates are dictated by the observer-cut geometry of Section 6.6 and the preceding $(b,d) = (3,4)$ decomposition.

Momentum and energy as observer-layer conjugates. Conjugate to the three spatial-access directions S_{obs} is a momentum-like variable that scales inversely with the spatial-access resolution of the cut. Conjugate to the record-ordering coordinate τ_{obs} is an energy-like variable

that scales inversely with the recovery-recoverable time-ordering resolution. These are not new substrate primitives; they are derived observer-layer quantities, in the same sense that time itself appears at the observer junction through the modular automorphism group of a nontracial cut. Momentum and energy live on the observer layer, not in the substrate.

Access depth as the framework’s native scale variable. Each observer cut $E_\lambda : \mathcal{W} \rightarrow \mathcal{A}_\lambda$ has a characteristic resolution depth, the access label λ already used in the modular-embedding language of Section 13. Internal holonomy sectors of \mathcal{W} are resolvable at depth λ only when the cut \mathcal{A}_λ separates them into distinct factors. At shallow access depths, gauge sectors nominally distinct in the deep algebra may not be separately resolvable by the cut; at deep access depths, additional internal sectors become resolvable. Access depth is therefore the framework’s native scale variable, conjugate to a probe energy through the same observer-layer geometry.

Which sectors enter a Schur reduction depends on λ . The Schur-complement book-keeping introduced in Section 6.8 reduces the joint quadratic form over the observer’s accessible algebra. The Schur sum runs over internal sectors that are simultaneously resolvable at the access depth in question. The resulting observer-effective inverse coupling therefore depends on which sectors are visible at λ : at the electroweak observer cut, $SU(3)$, $SU(2)$, and $U(1)$ are all separately resolvable, and the Schur expansion includes pairwise overlaps among them; at deep-QCD access depths where $SU(2)$ is broken (the W and Z are no longer resolvable propagating modes), the $SU(2)$ -overlap terms drop out and the Schur trace recomputes. This is the framework-native statement of what perturbative QFT calls the running of couplings. The variation across observer cuts is computable through the same machinery used at a single cut; the present text computes a single cut and treats access-depth variation as continuation work (Section 21).

6.8 Schur-complement reduction of observer-junction quadratic forms

Observer-effective couplings frequently arise from coupled quadratic forms on the observer junction, with a primary scalar readout coupled to a recoverable access sector. The sign and form of the resulting access dressings are not independent assumptions: they follow from positivity of the coupled observer-junction quadratic form together with the standard Schur-complement identity.

Coupled observer-junction quadratic form. At the observer junction, let x denote a scalar charged-readout amplitude and y denote the recoverable access-sector amplitudes. The framework’s paired $(K_{\text{rec}}, M_{\text{rec}})$ structure organizes the joint stiffness of these two components into a positive-definite quadratic form

$$Q(x, y) = K_{\text{bare}} x^2 + 2xBy + \langle y, Cy \rangle,$$

with $K_{\text{bare}} > 0$ the undressed charged stiffness, $C > 0$ the access-sector stiffness (a restriction of K_{rec} to the recoverable access channel), and B the coupling between charged readout and access modes (a restriction of the off-diagonal block of K_{rec}). Positivity of the joint form is the requirement that the substrate disturbance functional be bounded below on the coupled (x, y) space, and follows from positivity of K_{rec} .

Effective scalar stiffness by Schur complement. Minimizing $Q(x, y)$ over y at fixed x ,

$$\frac{\partial Q}{\partial y} = 2B^\dagger x + 2Cy = 0 \implies y_*(x) = -C^{-1}B^\dagger x,$$

and substituting back,

$$Q_{\text{eff}}(x) = \left(K_{\text{bare}} - BC^{-1}B^\dagger \right) x^2 =: K_{\text{eff}} x^2.$$

This is the standard Schur-complement / Feshbach partitioning of a coupled quadratic form. It is exact for the joint (x, y) system, not perturbative.

Schur-complement theorem for access dressing. The identity above immediately gives the following structural result.

Theorem 6.11 (Sign-forced access dressing). *Let $Q(x, y)$ be a positive-definite coupled observer-junction quadratic form with diagonal blocks (K_{bare}, C) and off-diagonal coupling B . Then the effective scalar stiffness at the observer junction is*

$$K_{\text{eff}} = K_{\text{bare}} - BC^{-1}B^\dagger,$$

with $BC^{-1}B^\dagger \succeq 0$ a positive semidefinite access-dressing operator. The observer-effective inverse coupling $\alpha_{\text{eff}}^{-1} = K_{\text{eff}}$ therefore satisfies $\alpha_{\text{eff}}^{-1} \leq \alpha_{\text{junction}}^{-1}$, with equality if and only if $B = 0$.

Minimal inter-sector Schur overlap. The same Schur-complement mechanism applies when two internal holonomy sectors are simultaneously unresolved by one observer cut. Let the primary trace weights of two such sectors be T_i and T_j . Under the trace-quadratic metric, a unit boundary amplitude in sector i is normalized by $T_i^{-1/2}$, and a unit boundary amplitude in sector j is normalized by $T_j^{-1/2}$. The minimal admissible overlap is the rank-one trace-norm preserving boundary coupling between these two normalized modes. If the elementary overlap block has unit scalar stiffness, the off-diagonal coupling has amplitude

$$T_i^{-1/2} T_j^{-1/2} = \frac{1}{\sqrt{T_i T_j}}.$$

Schur elimination of this unresolved overlap block subtracts the square of this amplitude. The minimal inter-sector correction is therefore

$$\Delta K_{ij}^{\text{overlap}} = -\frac{1}{T_i T_j}. \quad (1)$$

This reciprocal-product rule is the inter-sector analogue of access dressing: unresolved overlap lowers the observer-effective inverse stiffness because it enters through the negative Schur term $-BC^{-1}B^\dagger$. The new ingredient is not the sign but the minimal overlap normalization: rank one, trace-normalized sector amplitudes, and unit elementary overlap stiffness.

Three structural consequences. The theorem fixes three structural features of any observer-effective coupling that arises through Schur reduction, independently of the specific operator that produces the access dressing.

First, the sign of every access-dressing correction is negative: charged record modes coupled to recoverable access modes always lower the observer-effective inverse stiffness, never raise it. For example, the minus signs in the fine-structure access-dressing contributions $-3/(8\pi^4)$ and $-21/(32\pi^{12})$ of Section 15.7 are forced by this sign rule, not chosen.

Second, the access dressing is a single operator $BC^{-1}B^\dagger$, not a sum of independently-postulated corrections. The decomposition of this operator into a single-channel piece and a projective-cover piece is a spectral decomposition of one access-dressing object, not a postulate that several independent corrections exist.

Third, the mechanism is exact at the level of the coupled quadratic form. There is no perturbative truncation in the Schur step itself; truncation enters only when $BC^{-1}B^\dagger$ is expanded over

the access-mode spectrum. Any observer-effective coupling so obtained is therefore a calculation of the spectrum of one operator, not a perturbation series at arbitrary order.

The Schur-complement theorem does not by itself determine the coefficients of the corrections. It fixes the form $K_{\text{eff}} = K_{\text{bare}} - BC^{-1}B^\dagger$ and the sign of the dressing; the access-mode spectrum that fills in this operator is supplied by the access–recovery composition theorem of the next subsection.

6.9 Access–recovery composition for observer-effective couplings

The Schur-complement reduction of Section 6.8 reduces the access-dressing problem to determining the access-sector stiffness C that fills in $BC^{-1}B^\dagger$. This subsection identifies the primitive scalar access-dressing kernel as L_{acc}^{-2} , the squared inverse access Laplacian, from the requirement that a record disturbance be both accessed and recovered. The result is observer-layer machinery, applicable to any sector whose observer-effective coupling is obtained by Schur reduction over a recoverable access channel.

Access–recovery as a composed operation. A record disturbance at the observer junction enters the recoverable access channel through two distinct steps. First, the disturbance is accessed: it is registered against the observer’s accessible subalgebra through the conditional expectation E_λ , picking up an inverse-access propagator. Second, the registered disturbance must be recovered: it must propagate back to the primary readout through the same access channel, picking up a second inverse-access propagator. The primitive scalar access dressing is therefore not a single inverse access operator, but the composition of access and recovery,

$$\text{primitive access dressing} = \text{recovery} \circ \text{access}.$$

Both steps are governed by the access Laplacian L_{acc} , the substrate operator whose modes propagate record disturbances through the recoverable access channel. The composition is then

$$L_{\text{acc}}^{-1} \cdot L_{\text{acc}}^{-1} = L_{\text{acc}}^{-2}.$$

Theorem 6.12 (Access–recovery composition). *Under the requirements that (i) a record disturbance enter the recoverable access channel through a single primitive access step, and (ii) the disturbance return to the primary readout through a single primitive recovery step, the primitive scalar access-dressing kernel is the squared inverse access Laplacian,*

$$C_{\text{primitive}}^{-1} = L_{\text{acc}}^{-2}.$$

On the normalized Dirichlet access interval $u \in [0, 1]$, the modes $\psi_n(u) = \sqrt{2} \sin(n\pi u)$ are eigenstates with $L_{\text{acc}}\psi_n = (n\pi)^2\psi_n$ and $L_{\text{acc}}^{-2}\psi_n = (n\pi)^{-4}\psi_n$; the first inverse eigenvalue is π^{-4} .

Exclusion of lower-order terms. The composition argument simultaneously rules out a π^{-2} primitive contribution. A single inverse access operator L_{acc}^{-1} would correspond to an access step without recovery: a charged disturbance that enters the access channel and remains there. Such a configuration is not closed under record recovery, and therefore is not an admissible access dressing of a scalar primary readout. The first primitive contribution is therefore at π^{-4} , with no π^{-2} term, and no π^{-1} or π^0 term either.

Intermediate orders $\pi^{-3}, \pi^{-5}, \pi^{-6}, \pi^{-7}$ are excluded for the same structural reason: any odd power of π^{-1} corresponds to an unmatched access/recovery composition (one access without recovery, or one recovery without access), and any even power not of the form π^{-4k} corresponds to a partial composition that does not return to the charged readout. The admissible primitive scalar access dressings on a single channel are powers $(L_{\text{acc}}^{-2})^k$ for integer $k \geq 1$.

6.10 Bare trace from observer-side duality: $T_{\text{bare}} = 21/2$ derived from $d = 4$ and $b = 3$

The bare-trace value $T_{\text{bare}} = 21/2$ that governs the 42π term has so far been read from substrate-side trace counting over minimal chiral content. The observer-side duality established in the previous subsection inverts this logic: given the minimal stable observer-access dimension $d = 4$ and the minimal projective-closure multiplicity $b = 3$, the value $T_{\text{bare}} = 21/2$ is fixed by the duality and not by substrate content counting.

Duality equation as a constraint on T_{bare} . The substrate–observer duality at the bare-trace level is

$$d \cdot T_{\text{bare}} = \binom{d}{2}(2^b - 1),$$

with the left expression the substrate trace reading and the right expression the observer access-combinatorics reading. Solving for T_{bare} ,

$$T_{\text{bare}} = \frac{\binom{d}{2}(2^b - 1)}{d} = \frac{d-1}{2}(2^b - 1).$$

At the minimal stable observer-access dimension $d = 4$ and the minimal projective-closure multiplicity $b = 3$ (both established independently in the chapter, with $b = \dim \Sigma_{\text{obs}}^3 = 3$ identifying the projective multiplicity with the spatial-access dimension),

$$T_{\text{bare}} = \frac{4-1}{2}(2^3 - 1) = \frac{3}{2} \cdot 7 = \frac{21}{2}.$$

The bare-trace value is therefore not an input from substrate content counting. It is an output of the observer-side duality at the framework’s selected access dimension and minimal closure multiplicity.

Theorem 6.13 (Observer-side bare trace). *Given the minimal stable observer-access dimension $d = \dim O_{\text{acc}} = 4$ and the minimal non-removable projective-closure multiplicity $b = 3$, the substrate–observer duality determines the bare-trace value*

$$T_{\text{bare}} = \frac{d-1}{2}(2^b - 1) = \frac{21}{2}.$$

The bare junction stiffness $4\pi T_{\text{bare}} = 42\pi$ is therefore an observer-side prediction, independent of any specific substrate content realization.

Reversal of the epistemic dependence. This result reverses the logical dependence between the substrate’s minimal chiral content and the bare-trace value of the fine-structure formula. In the substrate-side reading, three-generation chiral content with one access-stabilizing doublet at $|Y_H| = 1/2$ is the input that produces $T_{\text{bare}} = 21/2$ as a trace sum. In the observer-side reading developed here, the duality fixes $T_{\text{bare}} = 21/2$ from $d = 4$ and $b = 3$, and the substrate’s chiral content must realize this trace value through admissible representations. The minimal substrate realization is three-generation standard chiral content with one minimal access-stabilizing doublet, but this is now a consequence of the trace constraint rather than its origin.

Consequence for the framework’s claim structure. The framework no longer claims that the fine-structure diagnostic depends on the Standard Model chiral content as an input. The bare-trace value is fixed by the observer-side duality; the substrate’s chiral content is then constrained to realize this trace value. That the minimal realization coincides with three-generation Standard Model charged content is an internal consistency between observer-layer geometry and admissible substrate content, not an imported assumption.

6.11 Trace weights at the minimal stable observer-access cut

The numerical agreements of Sections 15.9, 15.10, and 15.11 rest on specific trace-weight values T_1 , T_2 , $T_{b,\text{col}}$, T_s and the charged trace weight w_Y . This subsection collects the framework-internal derivations of these weights into a single place, distinguishing the two distinct structural theorems they invoke, supplying the algebraic definitions on which each rests, and making explicit the route from substrate inputs to numerical value. The fine-structure trace $T_1 = 137.035\,999\,176\dots$ is established in Section 15.9 through a single Schur-trace identity and is not re-derived here. The charged trace weight $w_Y = 3/8$ admits both a derivation from the minimal-chiral-content recoverability-obstruction equations of Section 15.4 and the Abelian ledger embedding normalization theorem stated in this appendix (Theorem 6.18).

Two classes of trace-weight object

The framework reads trace weights through two structurally distinct classes of object.

Observation supports. An observation support is a projector P_{min} on the resolved retrieval algebra whose image carries the branch coordinates the neutral-current observable resolves. Coordinates resolved by the observable count as branches; coordinates the observable integrates over or thresholds on do not. For a quark-flavor support read through the neutral-current vertex, the branches are the internal $SU(3)$ color components of the fundamental representation. For a detector-level charged-lepton record at the angular-asymmetry vertex, the branches are charge orientation, projective direction, and kinematic-closure magnitude. The branch count d_{supp} is determined by the observable’s resolution structure, not by the underlying particle’s gauge representation.

Adjoint connection blocks. An adjoint connection block is the holonomy support of a gauge symmetry, with internal dimension equal to the adjoint representation dimension of the underlying Lie algebra. The $SU(2)$ neutral block read as an adjoint holonomy in Section 15.10 and the $SU(3)$ adjoint block of the strong-coupling readout in Section 15.11 are both adjoint connection blocks. The branches in this case are the generators of the underlying gauge group, counted with adjoint Casimir normalization.

The two classes use different trace-weight theorems. The framework has both because the substrate–observer architecture admits both observation supports (resolved by readout structure) and connection blocks (resolved by gauge structure).

Branch coordinates: algebraic definition

Definition 6.14 (Branch coordinate). Let \mathcal{A} be a resolved retrieval algebra with faithful trace τ , and let $O \in \mathcal{A}$ be an observable. A coordinate c of \mathcal{A} is a *branch coordinate of O* if the conditional expectation $\mathbb{E}_\tau[O \mid c]$ is non-constant in c ; equivalently, O lies in a subalgebra of \mathcal{A} that resolves c . A coordinate that has been integrated, marginalized, thresholded, or otherwise reduced to a $\{0, 1\}$ indicator before O is constructed is not a branch coordinate of O .

For the leptonic neutral-current angular-asymmetry observable A_{FB} through a reconstructed charged-lepton final state, the branch coordinates of the charged-lepton record are charge orientation, projective direction, and kinematic-closure magnitude: A_{FB} depends non-trivially on each of the three. Detector-level coordinates used only for event selection or calibration (isolation, shower shape, track quality, fiducial geometry) are thresholded into $\{0, 1\}$ indicators before A_{FB} is defined and do not count as branch coordinates. They belong to the experiment-and-channel projection operator $P_{\text{exp},\ell}$, not to the minimal observation support P_{min} .

Definition 6.15 (Minimal observation support). For an observable O on the resolved retrieval algebra \mathcal{A} , the *minimal observation support* $P_{\min}(O)$ is the smallest projector in \mathcal{A} whose image carries every branch coordinate of O . The dimension of the image, denoted $d_{\text{supp}}(O)$, is the branch count of O .

For the leptonic neutral-current A_{FB} observable through a charged-lepton final state, $d_{\text{supp}} = 3$. For a quark-flavor color-carrying support read through the neutral-current vertex, $d_{\text{supp}} = 3$ (the three internal color components of the fundamental representation). For the $SU(2)$ weak-isospin doublet observation, $d_{\text{supp}} = 2$ (the two doublet components).

Observation-support trace weight: closure-normalized theorem

Theorem 6.16 (Observation-support trace weight). *On the minimal stable observer-access dimension $(b, d) = (3, 4)$ of Section 6, every fundamental observation support inherits a substrate closure-unit normalization $C_{\text{supp}}(b, d) = d/b = 4/3$. The trace weight of a fundamental observation support with branch count d_{supp} is*

$$T_{\text{supp}} = \frac{b \cdot d_{\text{supp}}}{C_{\text{supp}}(b, d)} = \frac{b^2 d_{\text{supp}}}{d}. \quad (2)$$

Sketch. The branch count d_{supp} is the cardinality of resolved branch coordinates by Definition 6.15. The substrate factor b enters through the projective multiplicity of Section 6 (the number of primitive projective branches of the spatial observer cut). The closure-unit normalization $C_{\text{supp}}(b, d) = d/b$ is the substrate-imposed ratio between branch and record dimensions on the minimal observer-access dimension $(b, d) = (3, 4)$; it is the substrate's per-branch share of the record-ordering coordinate. The product $b \cdot d_{\text{supp}}$ is the unnormalized support trace summed over projective branches and observation branches; dividing by the closure unit per branch gives the trace weight $T_{\text{supp}} = b^2 d_{\text{supp}}/d$. \square

Corollary: weak-isospin doublet, $d_{\text{supp}} = 2$. The neutral $SU(2)$ observation support resolves the two doublet states at the neutral-current vertex. With $d_{\text{supp}} = 2$ and $(b, d) = (3, 4)$,

$$T_2^{(\text{supp})} = \frac{9 \cdot 2}{4} = \frac{9}{2}.$$

This coincides with the value $T_2 = 9/2$ used in the neutral-boundary visibility lemma (Equation (27)) and produces the leptonic effective angle of Equation (28). The same value is obtained through the adjoint $SU(2)$ connection-block route; the framework reads T_2 through either route consistently because the two theorems coincide numerically on $(b, d) = (3, 4)$ for $SU(2)$.

Corollary: fundamental color support, $d_{\text{supp}} = 3$. A fundamental $SU(3)$ color-carrying support resolves three internal color branches. With $d_{\text{supp}} = 3$,

$$T_{b,\text{col}} = \frac{9 \cdot 3}{4} = \frac{27}{4}.$$

This recovers Equation (31). The substrate closure unit $C_{\text{supp}}(3, 4) = 4/3$ coincides numerically with the $SU(3)$ fundamental quadratic Casimir $C_2(\mathbf{3}) = 4/3$; the framework reads the substrate origin as primary and the gauge-theoretic re-expression as a secondary identification, valid for gauge-charged supports but not generally.

Corollary: minimal charged-lepton detector record, $d_{\text{supp}} = 3$. A reconstructed detector-level charged-lepton observation at the neutral-current angular-asymmetry vertex resolves three branch coordinates by Definition 6.14: charge orientation (sign of electric charge, distinguishing ℓ^+ from ℓ^-), projective direction (track or shower axis, defining Collins–Soper geometry), and kinematic-closure magnitude (energy or momentum, closing the event kinematics). With $d_{\text{supp}} = 3$,

$$T_{\text{ch}} = \frac{9 \cdot 3}{4} = \frac{27}{4} = T_{b,\text{col}}.$$

The trace weight of a minimal charged-lepton observation equals that of a fundamental color-carrying support, not because charged leptons are color triplets in the Standard Model (they are not, and the framework does not claim they are) but because both supports realize three branch coordinates on the same $(b, d) = (3, 4)$ substrate. The unification is a structural feature of the framework’s substrate-derived trace bookkeeping at the observer-access cut, not a particle-physics identification or a metaphor between leptons and color triplets. The detector-level charged-lepton record is a fundamental observation support in the framework’s sense, with the three branches algebraically determined by Definition 6.14 from the angular-asymmetry observable’s resolution structure.

Adjoint connection block: generator-count theorem

Theorem 6.17 (Adjoint connection block trace weight). *For an internal gauge group G with adjoint representation of internal dimension $\dim G_{\text{adj}}$ and adjoint quadratic Casimir $C_2(G_{\text{adj}})$, the adjoint connection block trace weight on the $b = 3$ observer projective cut is*

$$T_{\text{adj}}(G) = \frac{b \cdot \dim G_{\text{adj}}}{C_2(G_{\text{adj}})}. \quad (3)$$

Corollary: $SU(3)$ strong adjoint, $T_s = 8$. For $SU(3)$, $\dim G_{\text{adj}} = 8$ and $C_2(G_{\text{adj}}) = N = 3$, giving

$$T_s = \frac{3 \cdot 8}{3} = 8.$$

This recovers Equation (50) and supplies the strong-coupling readout of Section 15.11.

Status of T_2 under both theorems. For $SU(2)$, $\dim G_{\text{adj}} = 3$ and $C_2(G_{\text{adj}}) = N = 2$, giving $T_2^{(\text{adj})} = 3 \cdot 3/2 = 9/2$. This coincides numerically with $T_2^{(\text{supp})} = 9/2$ from Theorem 6.16. The framework reads T_2 through either route consistently. This double coverage does not extend to $T_{b,\text{col}}$ (which has only an observation-support route; no adjoint-block interpretation exists for a fundamental color-carrying support) or to T_s (which has only an adjoint-connection-block route; the substrate observation-support theorem with $d_{\text{supp}} = 32/9$ would be required and is not integer-valued).

Abelian ledger embedding normalization theorem

Theorem 6.18 (Abelian ledger embedding normalization). *A single $U(1)$ Abelian phase carries no intrinsic non-trivial trace normalization on the substrate observer cut: an Abelian phase operator at one point of the substrate is a complex number times the identity on its support, and a trace over the support returns only the support dimension times the phase. The non-trivial central Abelian visibility weight w_Y is therefore fixed only after specifying the chiral support into which the phase ledger is embedded.*

For the framework’s minimal electroweak chiral content of Section 15.4, the chiral ledger has dimension eight: the left-handed $SU(2)$ -doublet states of one generation, comprising the lepton doublet L_L (two components) and the quark doublet Q_L in three color states (six components),

with $2 + 6 = 8$. Right-handed singlets (e_R, u_R, d_R) carry no weak isospin and do not enter the chiral ledger. The observer cut resolves three projective weak branches inside this ledger — the three projective branches of the spatial observer cut of Section 6. The central Abelian visibility weight is the ratio of projective weak branches to chiral ledger dimension,

$$w_Y = \frac{b}{\dim(\text{chiral ledger})} = \frac{3}{8}. \quad (4)$$

Sketch. The numerator is the projective multiplicity of the observer cut from Section 6, contributing one resolved weak-active branch per spatial projective direction. The denominator is the dimension of the chiral support into which the Abelian phase is embedded: an Abelian phase operator on the substrate observer cut carries trace only through its embedded support, and on the minimal electroweak chiral content of Section 15.4 this support is the eight-component left-handed doublet sector. The ratio is the framework’s embedded visibility weight: the fraction of chiral ledger components visible to the observer cut at the projective multiplicity b . Right-handed singlets contribute to the Abelian phase computation in other diagnostic contexts but do not normalize the visibility weight, because they decouple from the weak-active embedding that defines the support into which the phase ledger enters. \square

Relation to the trace-ratio derivation of Section 15.5. Section 15.5 establishes $w_Y = 3/8$ through the $5/3$ hypercharge normalization $\text{Tr}(Y^2)/\text{Tr}(T_3^2) = 10/3 \div 2 = 5/3$ obtained from the anomaly-cancellation system of Section 15.4, with $\sin^2 \theta_W = (3/5)/(1 + 3/5) = 3/8$ at the substrate normalization scale. That route requires the full anomaly-cancellation system to fix the hypercharge values Y_Q, Y_u, Y_d, Y_L, Y_e first. Theorem 6.18 requires only the chiral ledger dimension and the observer-cut projective multiplicity; no hypercharge values are used as inputs. The two derivations agree numerically because the framework’s anomaly-cancelled hypercharge assignments and chiral ledger structure are mutually compatible features of the same minimal-chiral-content construction. Theorem 6.18 identifies w_Y as an algebraic ratio of branch counts, while the trace-ratio derivation extracts the same value from the Abelian and weak stiffness traces. The framework reads w_Y through either route consistently; this is the same kind of double coverage that T_2 admits under the observation-support theorem (6.16) and the adjoint connection block theorem (6.17).

Summary of trace-weight bookkeeping

The trace weights used numerically in the present paper are collected below, with theorem provenance and substrate inputs:

Weight	Theorem	Substrate inputs	Value
T_1	Schur-trace identity (§15.9)	$b = 3, d = 4, w_Y = 3/8$	137.035 999 176 ...
T_2	6.16 or 6.17	$(b, d) = (3, 4), d_{\text{supp}} = 2$ (or $SU(2)$ adj.)	9/2
$T_{b,\text{col}} = T_{\text{ch}}$	6.16	$(b, d) = (3, 4), d_{\text{supp}} = 3$	27/4
T_s	6.17	$b = 3, SU(3)$ adjoint	8
w_Y	6.18 (Abelian ledger embedding)	$b = 3, \text{eight-component chiral ledger}$	3/8

The trace-weight bookkeeping of the present paper rests on three structural theorems. Theorem 6.16 unifies $T_2, T_{b,\text{col}}$, and T_{ch} under one substrate-derived structural object: the closure-normalized observation-support trace weight, with d_{supp} as the only observation-type-dependent input. Theorem 6.17 handles the structurally distinct case of adjoint connection blocks (T_s , and equivalently T_2 when read as an adjoint block). Theorem 6.18 handles the Abelian ledger embedding normalization, fixing w_Y as a branch-count ratio between the observer-cut projective multiplicity and the chiral ledger dimension. The substrate input $(b, d) = (3, 4)$ is the same minimal observer-access dimension throughout: $b = 3$ enters all three theorems through the

observer-cut projective multiplicity, and $d = 4$ enters the observation-support theorem through the closure-unit normalization $C_{\text{supp}} = d/b$. No separate “weak,” “strong,” or “charged” substrate parameter is introduced. The three theorems together produce T_1 (via the Schur-trace identity of Section 15.9 into which w_Y feeds), T_2 , $T_{b,\text{col}}$, T_{ch} , T_s , and w_Y from the substrate-observer architecture without external numerical inputs.

6.12 Wavefunction and collapse from accessible-algebra restriction

The observer statement also hosts the standard mathematical structure of a wavefunction and measurement collapse. The interpretive restriction is essential: the wavefunction is not the state of the entire substrate, and collapse is not a change of the substrate state. Both are observer-effective objects associated with an accessible algebra and a recorded outcome.

Let $\mathcal{O} = (\Lambda, \preceq, \{\mathcal{A}_\lambda\}, \{\mathcal{M}_\lambda\})$ be the observer. At label λ , the observer has access only to $\mathcal{A}_\lambda \subseteq \mathcal{W}$. The observer-effective state is the restriction

$$\omega_\lambda = \omega|_{\mathcal{A}_\lambda}.$$

Represented on a Hilbert space, it may be written as a density operator ρ_λ . If the restricted state is pure, one may write $\rho_\lambda = |\psi_\lambda\rangle\langle\psi_\lambda|$. The wavefunction is therefore read as the observer-effective state on the accessible algebra. It is not required to be a fundamental substrate object.

Let a record-compatible measurement inside \mathcal{A}_λ be represented by projectors $\{P_k\} \subset \mathcal{A}_\lambda$. The observer assigns Born probabilities

$$p(k) = \omega_\lambda(P_k) = \text{Tr}(P_k\rho_\lambda).$$

If outcome k is written into the later record algebra $\mathcal{M}_{\lambda'}$, the observer-effective state updates by Lüders conditionalization,

$$\rho_{\lambda'} = \frac{P_k\rho_\lambda P_k}{\text{Tr}(P_k\rho_\lambda)}.$$

The substrate state ω on \mathcal{W} is not modified by this update. What changes is the observer’s conditional state on the accessible algebra after the record has been written. Collapse is record-conditioned access, not substrate-level dynamics.

6.13 Observer covariance through invariant record relations

Different observer cuts $\mathcal{O}_1, \mathcal{O}_2, \mathcal{O}_3$ need not agree on the same observer-layer description, on a single simultaneity foliation, or on a single coordinate decomposition of the access geometry. What they must agree on is the invariant relational content of the records they recover. Writing $\mathcal{R}_\mathcal{O}$ for the record content recovered through a cut \mathcal{O} , the operative relation is

$$\mathcal{R}_{\mathcal{O}_1} \sim \mathcal{R}_{\mathcal{O}_2},$$

not strict equality. The equivalence holds at the level of substrate-invariant relational content: intersection of accessible record algebras, causal accessibility through record-recoverability order, holonomy class of charged sectors, spectra of paired stiffness and susceptibility operators, charge-sector identification, and mass-readout ratios obtained from singular-value structure. Different observers carry different cuts and may experience different modular parameters, but they read the same physics in the sense that they identify the same equivalence class of record relations. The substrate is common; observer cuts differ; physical content is the equivalence class invariant under change of cut.

6.14 Projective observer readout

The readout that recovers physically meaningful invariants from the observer-access structure is projective. An overall record-amplitude rescaling of an accessible support is not an observer-distinguishable operation: only relative phases, relative amplitudes, and overlap ratios between accessible record states carry physical readout content. The operational object at a label λ is therefore the projective record space

$$\mathbb{P}(\mathcal{H}_\lambda^{\text{rec}}) = \mathcal{H}_\lambda^{\text{rec}} / \sim, \quad |\psi\rangle \sim e^{i\theta} \lambda |\psi\rangle \quad (\lambda \in \mathbb{R}_{>0}, \theta \in \mathbb{R}),$$

with $\mathcal{H}_\lambda^{\text{rec}}$ the Hilbert-space representative of stable records at the observer-access cut. The implication is that any framework-native diagnostic of observer-level content is a function of projective record invariants only. In particular, fine-structure-type ratios that compare a minimal charged holonomy readout to an isotropic record-cone access capacity are projective ratios by construction, so their values depend on the invariant projective record geometry and not on conventionally chosen representatives.

6.15 Observer-level identifications

The observer layer reads the following phenomena as features of record-ordered access.

Time. The observer-effective parameter reconstructed from the inclusion order \preceq of stable accessible records. Time is not a substrate coordinate in this framework; it is record order along an inclusion-indexed chain of cuts. When the local entropy-maximizing state ϕ_λ is faithful and nonracial, modular flow gives a canonical automorphism parameter for the accessible algebra; this modular parameter is a representation of observer-internal scale, while the operational arrow is the monotone inclusion of records. Different observers may therefore carry different access orders and different modular time scales.

Proper record time. When the observer-access structure is enriched by an invariant maximum access speed c_* derived from the generalized record operator G_{acc} , a stable excitation's support history also carries an observer-effective proper record time τ . This is not a substrate clock. It is a kinematic reconstruction rule on record-ordered support histories.

Arrow of time. The monotone growth of \mathcal{M}_λ along \preceq . The arrow is not a substrate feature. It is the fact that records, once acquired, cannot be un-lived from inside the observer.

Motion. The variation of relational support of stable substrate excitations along \preceq .

Light cone. The boundary of \mathcal{A}_λ at each λ . Operators outside \mathcal{A}_λ are inaccessible at λ .

Collapse. Conditionalization inside \mathcal{A}_λ . When a measurement returns a value, the observer's effective state on \mathcal{A}_λ is the conditional. The substrate state ω on \mathcal{W} is untouched.

6.16 Order and arrow

The substrate carries tenseless orderings on its admissible cuts and on its disturbance landscape. It carries no oriented arrow. A partial order is substrate-level and tenseless; an arrow is experiential and oriented. The observer does not create the substrate's orderings; the observer creates the experienced arrow by traversing those orderings while preserving records.

This is a locational claim, not by itself a derivation of thermodynamic irreversibility. The disturbance-barrier condition links observer memory to substrate stiffness, but thermodynamics

only appears after the substrate and observer layers are combined. The next section makes that recovery explicit.

7 Observer metrological completion

The substrate–observer architecture introduced in Sections 3 and 6 closes the framework in dimensionless form. The substrate supplies the paired bookkeeping $(K_{\text{rec}}, M_{\text{rec}})$ on the defect Hilbert space and the generalized record operator

$$G_{\text{acc}} = M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2};$$

the observer junction supplies a stable record cone, modular record-time, sector projections, charged-holonomy, thermal information, and countable records. Every invariant produced by this joint architecture is dimensionless: generalized eigenvalues $\lambda_X = \text{spec}(P_X G_{\text{acc}} P_X)$, sector traces Tr_X , Schur defects Δ_X , sector ratios ρ_{XY} , and dimensionless couplings α_X .

A dimensional laboratory constant cannot be a substrate output. Newton’s gravitational constant has SI dimensions $\text{m}^3 \text{kg}^{-1} \text{s}^{-2}$; asking the substrate to produce its numerical value in those units would require the substrate to already carry the observer’s conventions for length, mass, and time. The same circularity afflicts every dimensional constant of physics, including h , e , k_B , N_A , and the unit-fixed Planck mass and Planck time. The framework therefore requires a structural layer distinct from both substrate dynamics and observer access: a gauge-fixing of the observer’s measurement language. We call this layer the *observer metrological completion*.

The placement of the layer is forced. The objects to be gauged — record cone, modular record-time, sector projections, thermal information, countable records — do not exist before Section 6; placing the completion earlier would be circular. They must exist before any recovered effective law is read off in dimensional form; placing the completion later would leave the recovered-thermodynamics layer (Section 8) and below in dimensional ambiguity. The completion therefore sits between the observer layer and the recovered-thermodynamics layer.

The structural content of the completion is operator-theoretic rather than scalar. The seven SI defining constants do not enter as conversion multipliers attached to G_{acc} ; they enter as a coordinate change on the metrological block of $(K_{\text{rec}}, M_{\text{rec}})$ under which G_{acc} is left invariant. This is the central technical claim of the section, made precise in Sections 7.4 and 7.5.

We adopt the convention throughout that h is the cycle-action quantum and $\hbar = h/(2\pi)$ is the angular-action quantum. The action gauge below is fixed by h wherever an unambiguous primitive phase cycle is selected; wherever \hbar appears in derived dimensional constants, it denotes $h/(2\pi)$. Numerical values of $\Delta\nu_{\text{Cs}}$, c , h , e , k_B , N_A , and K_{cd} are treated as the fixed defining constants of the 2019 SI redefinition; in particular $\Delta\nu_{\text{Cs}}$ is treated as a pure integer, with the second a derived unit fixed by that integer rather than a primitive of the substrate layer.

7.1 Level hierarchy and the dimensionless baseline

The substrate–observer architecture and its metrological completion stratify cleanly into five levels.

- **Level 0.** The primitive substrate bookkeeping pair $(K_{\text{rec}}, M_{\text{rec}})$ as positive operators on the defect Hilbert space. Both operators are gauge-dependent representatives of the same generalized structure.
- **Level 1.** The invariant dimensionless generalized record operator $G_{\text{acc}} = M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2}$, together with the observer-junction sector projections $\{P_X\}$, the dimensionless time coordinate τ_{obs} , the three-direction spatial cover Σ_{obs}^3 , and the charged-sector projection P_{ch} . All Level-1 quantities are dimensionless.

- **Level 2.** The seven primary metrological gauges, fixed by the seven SI defining constants. This is the layer constructed in the present section.
- **Level 3.** The dimensional laboratory constants G , R , ε_0 , μ_0 , particle masses in kilograms, lifetimes in seconds, energies in joules, currents in amperes. Each is reconstructed by dimensional algebra from Levels 1 and 2.
- **Level 4.** The substrate-predictive targets: dimensionless invariants of G_{acc} and its sector projections. These are the only quantities the framework predicts in the strict sense. Examples include the fine-structure constant α , the lepton-mass ratios, mixing angles, dimensionless gravitational couplings $\alpha_G(m)$, the Schur defects, the access ratios r_τ, r_B of Section 5, and entropy ratios.

The relation $K_{\text{rec}} \leftrightarrow M_{\text{rec}} \leftrightarrow G_{\text{acc}}$ is gauge-redundant: distinct positive pairs can produce the same G_{acc} . What is physically invariant is the generalized record geometry, encoded equivalently as the spectrum of G_{acc} , the Rayleigh quotient $\langle v, K_{\text{rec}} v \rangle / \langle v, M_{\text{rec}} v \rangle$, or the family of sector-projected operators $A_X = P_X G_{\text{acc}} P_X$.

The dimensionless baseline produced jointly by the substrate and observer layers is the family

$$\mathcal{B}_{\text{rec}} = \{ G_{\text{acc}}, \{P_X\}_X, \tau_{\text{obs}}, \Sigma_{\text{obs}}^3, P_{\text{ch}}, \text{Tr}, \preceq \}.$$

A dimensional laboratory realization is a map

$$\Phi_{\text{SI}} : \mathcal{B}_{\text{rec}} \longrightarrow \text{SI}$$

that assigns positive conversion factors to a finite set of independent dimensional axes, such that every SI-side quantity computable from the framework is the product of a Level-1 invariant and a finite combination of those factors. The remainder of this section shows that this map factors through exactly seven independent axes and constructs the SI realization explicitly.

7.2 The seven metrological scales as a coordinate basis

The seven independent dimensional axes any laboratory observer must fix are time, length, action, charge, thermal information, amount of substance, and photometric response. The number seven is not chosen for SI compatibility; it is the rank of the dimensional space generated by the recoverable observer coordinates of Section 6, after the substrate-side null-cone identification of length with time is imposed. The seven defining constants of the 2019 SI supply one canonical bijection between these axes and a closed set of dimensional anchors.

Definition 7.1 (Metrological coordinate basis). The *metrological coordinate basis* is the seven-tuple of dimensionless record-side axes

$$\mathbf{e} = (e_t, e_x, e_S, e_q, e_\Theta, e_N, e_J),$$

labelling, respectively, record-time, record-distance, record-action, record-charge, thermal information, count, and photometric response. The corresponding dimensionless record-side coordinates are

$$(\tau_{\text{rec}}, x_{\text{rec}}, S_{\text{rec}}, q_{\text{rec}}, \Theta_{\text{rec}}, N_{\text{rec}}, J_{\text{rec}}).$$

A *metrological completion* is a positive diagonal map

$$\Gamma_{\text{SI}} = \text{diag}(\gamma_t, \gamma_x, \gamma_S, \gamma_q, \gamma_\Theta, \gamma_N, \gamma_J)$$

relating record-side coordinates to laboratory SI coordinates by $v_{\text{SI}} = \Gamma_{\text{SI}} v_{\text{rec}}$.

Concretely the SI-side coordinates obtained from the seven record-side axes are

$$\begin{aligned} t_{\text{SI}} &= \gamma_t \tau_{\text{rec}}, & x_{\text{SI}} &= \gamma_x x_{\text{rec}}, & S_{\text{SI}} &= \gamma_S S_{\text{rec}}, & q_{\text{SI}} &= \gamma_q q_{\text{rec}}, \\ T_{\text{SI}} &= \gamma_\Theta \Theta_{\text{rec}}, & n_{\text{SI}} &= \gamma_N N_{\text{rec}}, & I_{v,\text{SI}} &= \gamma_J J_{\text{rec}}. \end{aligned}$$

All record-side variables are dimensionless. Only after the seven γ 's are fixed do the SI-side quantities acquire their unit content.

7.3 Canonical metrology-block normalization

The metrological block of $(K_{\text{rec}}, M_{\text{rec}})$ admits a canonical normalization in which all seven primitive axes carry unit stiffness and unit susceptibility.

Definition 7.2 (Canonical metrology-block normalization). The *canonical metrology-block normalization* is the gauge choice

$$K_{\text{met}} = M_{\text{met}} = I_7,$$

in which the metrology block of the paired bookkeeping is the seven-dimensional identity on the metrological coordinate basis of Definition 7.1.

In this gauge the metrology block contributes

$$A_{\text{met}} = M_{\text{met}}^{-1/2} K_{\text{met}} M_{\text{met}}^{-1/2} = I_7$$

to the generalized record operator. The choice is conventional, not forced: any positive pair $(K_{\text{met}}, M_{\text{met}})$ producing $A_{\text{met}} = I_7$ would serve. Identity normalization is convenient because it makes the metrology block invisible at Level 1 — the seven primitive axes contribute trivially to G_{acc} until laboratory coordinates are imposed. The convention does not preclude anisotropic metrological gauges; those are reached by post-composition with a non-trivial similarity on the seven-dimensional block. The remainder of this section works in the canonical gauge.

7.4 Coordinate transformation of $(K_{\text{rec}}, M_{\text{rec}})$ and invariance of G_{acc}

The change of coordinates $v_{\text{SI}} = \Gamma_{\text{SI}} v_{\text{rec}}$ acts on the quadratic forms supporting K_{rec} and M_{rec} . The record-side quadratic form

$$Q_K^{\text{rec}}(v_{\text{rec}}) = \langle v_{\text{rec}}, K_{\text{rec}} v_{\text{rec}} \rangle$$

expressed in SI-side coordinates is

$$Q_K^{\text{SI}}(v_{\text{SI}}) = \langle \Gamma_{\text{SI}}^{-1} v_{\text{SI}}, K_{\text{rec}} \Gamma_{\text{SI}}^{-1} v_{\text{SI}} \rangle = \langle v_{\text{SI}}, \Gamma_{\text{SI}}^{-T} K_{\text{rec}} \Gamma_{\text{SI}}^{-1} v_{\text{SI}} \rangle.$$

The same holds for M_{rec} . We therefore have:

Proposition 7.3 (Metrological transformation rule). *Under the metrological coordinate change $v_{\text{SI}} = \Gamma_{\text{SI}} v_{\text{rec}}$ with Γ_{SI} positive diagonal, the substrate-observer pair transforms covariantly:*

$$K_{\text{rec}}^{\text{SI}} = \Gamma_{\text{SI}}^{-T} K_{\text{rec}} \Gamma_{\text{SI}}^{-1}, \quad M_{\text{rec}}^{\text{SI}} = \Gamma_{\text{SI}}^{-T} M_{\text{rec}} \Gamma_{\text{SI}}^{-1}.$$

The generalized record operator is invariant:

$$G_{\text{acc}}^{\text{SI}} = (M_{\text{rec}}^{\text{SI}})^{-1/2} K_{\text{rec}}^{\text{SI}} (M_{\text{rec}}^{\text{SI}})^{-1/2} = G_{\text{acc}}.$$

Proof. The transformation of K_{rec} and M_{rec} is the standard congruence transformation of quadratic forms under change of basis. For the invariance of G_{acc} , note that Γ_{SI} is positive diagonal, hence symmetric ($\Gamma_{\text{SI}}^T = \Gamma_{\text{SI}}$), and in the canonical metrology-block normalization commutes with K_{rec} and M_{rec} blockwise: on the seven-dimensional metrology block both operators are themselves diagonal (and equal to I_7 by Definition 7.2), and on the physics block Γ_{SI} acts as the identity by Principle 7.5. Therefore $\Gamma_{\text{SI}}^{-1} M_{\text{rec}} \Gamma_{\text{SI}}^{-1} = M_{\text{rec}} \Gamma_{\text{SI}}^{-2}$ blockwise, and

$$(M_{\text{rec}}^{\text{SI}})^{-1/2} = (M_{\text{rec}} \Gamma_{\text{SI}}^{-2})^{-1/2} = \Gamma_{\text{SI}} M_{\text{rec}}^{-1/2},$$

using commutativity of Γ_{SI} and M_{rec} in functional calculus. Substituting,

$$G_{\text{acc}}^{\text{SI}} = \Gamma_{\text{SI}} M_{\text{rec}}^{-1/2} (\Gamma_{\text{SI}}^{-1} K_{\text{rec}} \Gamma_{\text{SI}}^{-1}) M_{\text{rec}}^{-1/2} \Gamma_{\text{SI}} = M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2} = G_{\text{acc}},$$

where the Γ_{SI} factors cancel against Γ_{SI}^{-1} blockwise. Equivalently, the Rayleigh quotient $\langle v, K_{\text{rec}} v \rangle / \langle v, M_{\text{rec}} v \rangle$ is invariant under simultaneous congruence of numerator and denominator by the same positive diagonal map. \square

Remark 7.4. The invariance of G_{acc} under Proposition 7.3 is the operator-theoretic statement that the SI realization is a gauge, not a substrate input. The substrate's invariant content sits in the dimensionless generalized spectrum; the SI coordinates change only the representation of K_{rec} and M_{rec} along the seven metrological axes. Any falsification of the framework on the basis of laboratory measurements must therefore turn on the Level-1 invariants, never on the Level-3 dimensional constants.

In the canonical metrology-block normalization (Definition 7.2), the SI-side representation of the metrology block is

$$K_{\text{met}}^{\text{SI}} = M_{\text{met}}^{\text{SI}} = \Gamma_{\text{SI}}^{-2} = \text{diag}(\gamma_t^{-2}, \gamma_x^{-2}, \gamma_S^{-2}, \gamma_q^{-2}, \gamma_\Theta^{-2}, \gamma_N^{-2}, \gamma_J^{-2}).$$

The seven SI defining constants enter only through these diagonal entries. The block's contribution to $G_{\text{acc}}^{\text{SI}}$ remains I_7 .

7.5 Metrology–physics direct-sum decomposition

The metrological block and the physical content sit in orthogonal blocks of the substrate–observer pair.

Principle 7.5 (Metrology–physics decomposition). The defect Hilbert space underlying $(K_{\text{rec}}, M_{\text{rec}})$ admits a direct-sum decomposition

$$\mathcal{H} = \mathcal{H}_{\text{met}} \oplus \mathcal{H}_{\text{phys}},$$

into a seven-dimensional metrology block and the remaining physics block. Under this decomposition the paired bookkeeping splits as

$$K_{\text{rec}} = K_{\text{met}} \oplus K_{\text{phys}}, \quad M_{\text{rec}} = M_{\text{met}} \oplus M_{\text{phys}},$$

and the generalized record operator splits accordingly:

$$G_{\text{acc}} = A_{\text{met}} \oplus A_{\text{phys}} = I_7 \oplus M_{\text{phys}}^{-1/2} K_{\text{phys}} M_{\text{phys}}^{-1/2}.$$

The metrological coordinate change acts only on the first summand: $\Gamma_{\text{SI}}^{\text{total}} = \Gamma_{\text{SI}} \oplus I_{\text{phys}}$.

The decomposition has direct consequences. First, the SI shell touches only the metrology block; the physics block is invariant under metrological coordinate change in the strongest possible sense, namely the Γ_{SI} map acts as the identity on it. Second, all substrate predictions live in the spectrum of A_{phys} and its sector projections; the metrology block contributes the trivial spectrum $\{1\}^{\times 7}$ and adds no predictive content. Third, the framework's inverse problem is well-posed: laboratory measurements, pulled back through Γ_{SI}^{-1} , supply data on A_{phys} but cannot constrain the trivial metrology block. The metrological completion is therefore the unique extension of the dimensionless framework that produces laboratory-comparable predictions without contaminating the substrate's predictive content.

The mass-shape sector of Section 5 is the first nontrivial application of the physics block of Principle 7.5: its predictive readouts r_τ, r_B , and (provisionally) r_μ are Level-4 invariants of A_{phys} that become laboratory-comparable mass ratios only after the metrological completion of the present section.

7.6 Gauge conditions from the modern SI

The seven SI defining constants supply seven gauge conditions on Γ_{SI} , one per metrological axis.

Time gauge

The substrate contains recoverable ordering and modular record recurrence but not seconds. Let $\nu_{\text{Cs,rec}}$ be the dimensionless cycle-count per record-time unit of the chosen caesium-133 ground-state hyperfine transition, treated as a substrate-stable recurrence and normalized so that $\nu_{\text{Cs,rec}} = 1$. Physical frequency is $\nu_{\text{Cs,SI}} = \nu_{\text{Cs,rec}}/\gamma_t$, and the SI condition $\nu_{\text{Cs,SI}} = \Delta\nu_{\text{Cs}}$ fixes

$$\gamma_t = 1/\Delta\nu_{\text{Cs}}.$$

One record-time unit corresponds to $1/\Delta\nu_{\text{Cs}}$ seconds after the caesium clock has been selected.

Length gauge

The substrate-level observer cone supplies a normalized null-recoverability speed $c_* = 1$ in internal units (Sections 11.1 and 11.3). Physical speed is $c_{\text{SI}} = \gamma_x/\gamma_t$, and the SI condition $c_{\text{SI}} = c$ forces

$$\gamma_x = c\gamma_t = c/\Delta\nu_{\text{Cs}}.$$

The framework supplies the necessity of a null-cone conversion between record distance and record time; the SI supplies the numerical value.

Action gauge

Let S_{rec} be the dimensionless record-action count, normalized so that one primitive phase cycle has $S_{\text{cycle,rec}} = 1$. The SI condition $h_{\text{SI}} = h$ fixes

$$\gamma_S = h.$$

Mass is not an independent gauge. Since action has dimensions of mass \times length² \times time⁻¹, the mass scale is forced:

$$\gamma_m = \gamma_S \gamma_t / \gamma_x^2 = h \Delta\nu_{\text{Cs}} / c^2.$$

Mass is not introduced as a substrate substance; it is the inertial burden of a record support after the observer has fixed action, time, and length units.

Charge gauge

Let q_{rec} be the dimensionless record-charge coordinate, normalized so that the elementary record-charge support carries $q_{e,\text{rec}} = 1$. The SI condition $q_{e,\text{SI}} = e$ fixes

$$\gamma_q = e.$$

Current is charge transport per unit record-time, with scale

$$\gamma_I = \gamma_q / \gamma_t = e \Delta\nu_{\text{Cs}}.$$

The ampere is not a primitive metrological axis; it is the time-quotient of the charge gauge.

Thermal-information gauge

Temperature is not a primitive substrate dimension; it is the observer's conversion between record energy and record thermal information. Let Θ_{rec} be the dimensionless thermal record coordinate. The energy scale is forced by the action and time gauges:

$$\gamma_E = \gamma_S / \gamma_t = h \Delta\nu_{\text{Cs}}.$$

The Boltzmann bridge $\gamma_E = k_B \gamma_\Theta$ at one normalized thermal-information unit gives

$$\gamma_\Theta = \gamma_E / k_B = h \Delta\nu_{\text{Cs}} / k_B.$$

The kelvin is the observer-side conversion factor between record energy and thermal information.

Counting gauge

The mole is a coarse-graining convention. Let N_{rec} be the dimensionless microscopic record count. The mole anchors a fixed coarse-graining factor N_A :

$$n_{\text{SI}} [\text{mol}] = N_{\text{rec}} / N_A, \quad \gamma_N = 1 / N_A.$$

One mole equals N_A elementary record entities by definition.

Photometric gauge

The photometric axis is structurally asymmetric. The first six gauges are pure scale conversions on positive scalar axes; the seventh involves an observer-response weighting that is empirically calibrated to a fixed observer class (the CIE photopic standard observer) rather than substrate-derivable. The radiative-power scale is fixed by the action and time gauges,

$$\gamma_P = \gamma_E / \gamma_t = h \Delta\nu_{\text{Cs}}^2;$$

the photometric weighting is then imposed as a response function $V(\nu)$ on the radiative sector, with $V(540 \times 10^{12} \text{ Hz}) = 1$ by construction and $V(\nu) \leq 1$ elsewhere. The luminous intensity at frequency ν and solid angle Ω is

$$I_{\nu,\text{SI}}(\nu) = K_{\text{cd}} V(\nu) P_{\text{rad}}(\nu) \Omega^{-1},$$

with $K_{\text{cd}} = 683 \text{ lm/W}$ the SI-fixed luminous efficacy. In framework terms the photometric channel is a response-weighted projection of the radiative sector:

$$P_{\text{photo}}(\nu) = V(\nu) P_{\text{rad}}(\nu),$$

and γ_J is the composition of the radiative-power scale with the K_{cd} -anchored response weighting. The candela closes the SI base-unit list but does so by metrological convention rather than by substrate conditioning. We retain it for SI closure and flag its asymmetric status: it is the seventh gauge required for laboratory closure but the one most clearly outside the framework's substrate-observer dynamics, since $V(\nu)$ is an observer-class calibration rather than a substrate invariant.

7.7 Derived dimensional scales

Once the seven primary gauges are fixed, every other SI scale is forced by dimensional algebra. Using

$$\gamma_t = 1/\Delta\nu_{\text{Cs}}, \quad \gamma_x = c/\Delta\nu_{\text{Cs}}, \quad \gamma_S = h, \quad \gamma_q = e, \quad \gamma_\Theta = h \Delta\nu_{\text{Cs}}/k_B, \quad \gamma_N = 1/N_A,$$

the secondary scales follow:

$$\begin{aligned} \gamma_v &= \gamma_x/\gamma_t = c, \\ \gamma_E &= \gamma_S/\gamma_t = h \Delta\nu_{\text{Cs}}, \\ \gamma_m &= \gamma_E/c^2 = h \Delta\nu_{\text{Cs}}/c^2, \\ \gamma_p &= \gamma_S/\gamma_x = h \Delta\nu_{\text{Cs}}/c, \\ \gamma_F &= \gamma_E/\gamma_x = h \Delta\nu_{\text{Cs}}^2/c, \\ \gamma_P &= \gamma_E/\gamma_t = h \Delta\nu_{\text{Cs}}^2, \\ \gamma_I &= \gamma_q/\gamma_t = e \Delta\nu_{\text{Cs}}, \\ \gamma_V &= \gamma_E/\gamma_q = h \Delta\nu_{\text{Cs}}/e, \\ \gamma_{R\Omega} &= \gamma_V/\gamma_I = h/e^2. \end{aligned}$$

The resistance scale $\gamma_{R\Omega} = h/e^2$ is the von Klitzing constant; its appearance in the derived list is the dimensional signature of the action and charge gauges, not an additional input.

7.8 Gravitational coupling: G as metrological, α_G as physical

Newton's gravitational constant has SI dimensions $[G] = \text{length}^3 \text{mass}^{-1} \text{time}^{-2}$. Its scale factor is forced:

$$\gamma_G = \gamma_x^3 / (\gamma_m \gamma_t^2) = \frac{(c/\Delta\nu_{\text{Cs}})^3}{(h \Delta\nu_{\text{Cs}}/c^2) (1/\Delta\nu_{\text{Cs}})^2} = \frac{c^5}{h \Delta\nu_{\text{Cs}}^2}.$$

If the framework produces a dimensionless gravitational coupling G_{rec} from the paired regional disturbance and the connection-variation readout, then its SI value is

$$G_{\text{SI}} = G_{\text{rec}} \frac{c^5}{h \Delta\nu_{\text{Cs}}^2}, \quad G_{\text{rec}} = G_{\text{SI}} \frac{h \Delta\nu_{\text{Cs}}^2}{c^5}.$$

With the CODATA-2022 value $G_{\text{SI}} \approx 6.674 \times 10^{-11} \text{m}^3 \text{kg}^{-1} \text{s}^{-2}$, the dimensionless record-side number is $G_{\text{rec}} \approx 1.5 \times 10^{-66}$. This is not a substrate prediction. Up to a factor of 2π ,

$$G_{\text{rec}} = 2\pi (\ell_P/\lambda_{\text{Cs}})^2,$$

where $\ell_P = \sqrt{\hbar G/c^3}$ is the Planck length and $\lambda_{\text{Cs}} = c/\Delta\nu_{\text{Cs}}$ is the caesium-clock wavelength. The small value of G_{rec} encodes the enormous gap between the laboratory clock and the substrate's intrinsic action scale; it carries no substrate content beyond what is already in the metrological choices $\{\Delta\nu_{\text{Cs}}, h, c\}$.

The substrate target is therefore neither G_{SI} nor G_{rec} in the form above. The substrate target is a Level-4 dimensionless gravitational coupling intrinsic to the matter sector, such as the gravitational fine-structure constant

$$\alpha_G(m) = \frac{G m^2}{\hbar c}.$$

Writing $m_{\text{SI}} = \mu_{\text{rec}} \gamma_m$ and $G_{\text{SI}} = G_{\text{rec}} \gamma_G$, all SI conversion factors cancel:

$$\alpha_G(m) = G_{\text{rec}} \mu_{\text{rec}}^2 \frac{\gamma_G \gamma_m^2}{\hbar c} = 2\pi G_{\text{rec}} \mu_{\text{rec}}^2,$$

using $\gamma_G \gamma_m^2 / (\hbar c) = (c^5 / h \Delta \nu_{Cs}^2) \cdot (h \Delta \nu_{Cs} / c^2)^2 / (\hbar c) = 2\pi$. The right-hand side is purely dimensionless and lies in Level 4. For the electron, $\alpha_G(m_e) \approx 1.75 \times 10^{-45}$; for the proton, $\alpha_G(m_p) \approx 5.9 \times 10^{-39}$. These ratios are observer-invariant and substrate-meaningful in the strict Level-4 sense; G_{SI} is not.

The clean statement is:

G is metrological. α_G is physical.

7.9 Universal gas constant as a derived bridge

The universal gas constant is fixed once the thermal-information gauge and the counting gauge are in place. From $\gamma_\Theta = h \Delta \nu_{Cs} / k_B$ and $\gamma_N = 1 / N_A$,

$$R = N_A k_B.$$

This is not an independent metrological constant but the product of the counting gauge inverse and the thermal-information bridge. In framework terms: k_B converts thermal record information to record energy; N_A converts elementary record count to molar count; R is the molarized thermal-information bridge. R sits at Level 3, not Level 2.

7.10 Vacuum electromagnetic constants as α -dependent reconstructions

The vacuum permittivity ϵ_0 and permeability μ_0 are no longer exact defining constants under the 2019 SI. They are Level-3 reconstructions from the fine-structure constant α , which is Level 4:

$$\epsilon_0 = \frac{e^2}{4\pi \alpha \hbar c}, \quad \mu_0 = \frac{1}{\epsilon_0 c^2}.$$

Within the present framework, α is produced by the projective access-cover diagnostic of Section 15, which compresses to a Schur-complement trace on one substrate-cut block operator and is a Level-1 invariant of G_{acc} . Once α is fixed at Level 4, ϵ_0 and μ_0 follow as Level-3 reconstructions through the action, time, length, and charge gauges. The electromagnetic vacuum constants sit at the same hierarchical level as G : observer-side metrological constructs whose dimensional values carry no substrate content beyond the dimensionless α that drives them.

7.11 Dimensional-weight predictive criterion

Every laboratory quantity admits a dimensional decomposition along the seven metrological axes. For any SI-side quantity Q_{SI} ,

$$Q_{SI} = Q_{rec} \prod_{X \in \{t, x, S, q, \Theta, N, J\}} \gamma_X^{a_X(Q)},$$

where the exponents $a_X(Q) \in \mathbb{Z}$ (or \mathbb{Q} in rare cases) are determined by the SI dimension of Q along each axis and Q_{rec} is the Level-1 record-side content. This decomposition gives a sharp predictive criterion:

Principle 7.6 (Dimensional-weight predictive criterion). A quantity Q is substrate-predictive in the strict Level-4 sense if and only if its metrological weights vanish or cancel in a ratio:

$$\sum_X a_X(Q) \cdot e_X = 0 \quad (\text{direct}), \quad \text{or} \quad \frac{a_X(Q_1)}{a_X(Q_2)} = \text{const for all } X \quad (\text{ratio}).$$

Quantities satisfying this criterion lie in the kernel of the metrological rescaling action on Γ_{SI} and are invariants of the substrate-observer reading.

This criterion partitions laboratory observables cleanly. Substrate-predictive (Level 4): the fine-structure constant α , the effective weak mixing angle $\sin^2 \theta_{\text{eff}}$, all mass ratios m_i/m_j including the projective-access readouts r_τ, r_B, r_μ of Section 5, lifetime ratios, branching ratios, the gravitational fine-structure constants $\alpha_G(m)$, the Schur defects, trace-normalized sector spectra, and dimensionless entropy ratios. Non-predictive in the substrate sense (Level 3): G in $\text{m}^3\text{kg}^{-1}\text{s}^{-2}$, R in $\text{J mol}^{-1}\text{K}^{-1}$, ε_0 in F/m , particle masses in kilograms, lifetimes in seconds, energies in joules. These are not meaningless — they are the laboratory readouts of Level-4 invariants through Γ_{SI} — but they are not what the substrate is competent to predict.

7.12 Full Level-2 metrological map

The Level-2 map is summarized in Table 2, with two Level-3 derivations included for orientation.

Quantity	Record-side object	SI gauge
<i>Level 2 (primary):</i>		
time	τ_{rec}	$\gamma_t = 1/\Delta\nu_{\text{Cs}}$
length	x_{rec}	$\gamma_x = c/\Delta\nu_{\text{Cs}}$
action	S_{rec}	$\gamma_S = h$
charge	q_{rec}	$\gamma_q = e$
temperature	Θ_{rec}	$\gamma_\Theta = h \Delta\nu_{\text{Cs}}/k_B$
amount	N_{rec}	$\gamma_N = 1/N_A$
photometry	J_{rec}	fixed by K_{cd} and $V(\nu)$
<i>Level 3 (derived from Levels 1 and 2):</i>		
mass	derived	$\gamma_m = h \Delta\nu_{\text{Cs}}/c^2$
energy	derived	$\gamma_E = h \Delta\nu_{\text{Cs}}$
current	derived	$\gamma_I = e \Delta\nu_{\text{Cs}}$
resistance	derived	$\gamma_{R\Omega} = h/e^2$
gravitation	G_{rec} (Level 1)	$\gamma_G = c^5/(h \Delta\nu_{\text{Cs}}^2)$
gas constant	—	$R = N_A k_B$
vacuum permittivity	α (Level 4)	$\varepsilon_0 = e^2/(4\pi \alpha \hbar c)$

Table 2: Level-2 metrological map. The upper block fixes the seven primary gauges; the lower block illustrates Level-3 derivation. Predictive substrate targets (Level 4) are not listed here; they are by construction outside the metrological map.

7.13 Observer metrological completion theorem

The structural content of this section is captured by the following theorem.

Theorem 7.7 (Observer metrological completion). *Let $(K_{\text{rec}}, M_{\text{rec}})$ be the paired bookkeeping of the substrate–observer architecture, with generalized record operator $G_{\text{acc}} = M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2}$, and assume the observer junction recovers dimensionless record coordinates for the seven metrological axes (Definition 7.1). The following hold.*

- (i) Existence and uniqueness of the SI completion. *The metrological coordinate change $\Gamma_{\text{SI}} = \text{diag}(\gamma_t, \gamma_x, \gamma_S, \gamma_q, \gamma_\Theta, \gamma_N, \gamma_J)$ is the unique positive diagonal map such that the seven SI*

defining conditions hold:

$$\begin{aligned}\gamma_t &= 1/\Delta\nu_{\text{Cs}}, & \gamma_x &= c/\Delta\nu_{\text{Cs}}, & \gamma_S &= h, & \gamma_q &= e, \\ \gamma_\Theta &= h \Delta\nu_{\text{Cs}}/k_B, & \gamma_N &= 1/N_A,\end{aligned}$$

with γ_J fixed by K_{cd} together with the photometric response function $V(\nu)$.

- (ii) Operator-theoretic covariance. *The paired bookkeeping transforms covariantly under Γ_{SI} via $K_{\text{rec}}^{\text{SI}} = \Gamma_{\text{SI}}^{-T} K_{\text{rec}} \Gamma_{\text{SI}}^{-1}$ and $M_{\text{rec}}^{\text{SI}} = \Gamma_{\text{SI}}^{-T} M_{\text{rec}} \Gamma_{\text{SI}}^{-1}$, and the generalized record operator G_{acc} is invariant.*
- (iii) Direct-sum compatibility. *Under the metrology–physics decomposition $\mathcal{H} = \mathcal{H}_{\text{met}} \oplus \mathcal{H}_{\text{phys}}$, Γ_{SI} acts as the identity on $\mathcal{H}_{\text{phys}}$.*
- (iv) Predictive kernel. *Under the completion, every Level-3 dimensional constant is reconstructed by dimensional algebra from Levels 1 and 2, while substrate-predictive content is confined to Level-4 invariants: quantities satisfying the dimensional-weight cancellation criterion of Principle 7.6.*

Corollary 7.8 (Dimensionless prediction kernel). *The substrate’s empirical content under any metrological completion Γ_{SI} is exactly the kernel of the rescaling action $\gamma_X \mapsto \lambda_X \gamma_X$ on Γ_{SI} . The fine-structure constant α , the lepton-mass ratios, the projective-access readouts r_τ, r_B, r_μ of Section 5, the mixing angles, the dimensionless sector traces, the Schur defects, the gravitational fine-structure constants $\alpha_G(m)$, and the access-cone compatibility ratios all lie in this kernel. The numerical SI values of G , R , ε_0 , μ_0 , particle masses in kilograms, lifetimes in seconds, and analogous dimensional constants do not.*

Corollary 7.9 (Falsifier restriction). *Any falsifying observation against the framework must turn on a Level-4 invariant. A measurement reporting an anomalous laboratory value of G , R , ε_0 , a particle mass in kilograms, or any other Level-3 quantity does not by itself falsify the framework; the framework is empty of predictions about Level-3 quantities except through Level-4 invariants and the SI completion. The falsification criteria of Section 21 should be read under this restriction; in particular the tau-mass and baryon-mean falsifiers of Proposition 5.11 are Level-4 falsifiers (they constrain r_τ and r_B as dimensionless invariants), not laboratory-mass falsifiers.*

Operational reading. Two slogans capture the content of Theorem 7.7:

The SI does not derive K_{rec} and M_{rec} . The SI lets laboratory physics be pulled back into the dimensionless geometry of G_{acc} .

Normalize the metrology axes to identity; place the seven SI constants in the external metrological map Γ_{SI} ; keep the physics in dimensionless sector readouts of G_{acc} .

This completes the metrological layer. The recovered effective laws of the subsequent sections (Sections 8, 11, 14 and 15) are now read off the dimensionless geometry of G_{acc} and projected into SI form through Theorem 7.7; the framework’s prediction targets are confined to Corollary 7.8; and the falsification criteria of Section 21 are correspondingly restricted by Corollary 7.9.

8 Recovered thermodynamics from substrate–observer coupling

Thermodynamics is not placed inside the substrate alone and is not assigned to an observer alone. The substrate supplies static regional disturbance, conditional-expectation defect, and

the paired connection-dependent stiffness/susceptibility $(K_{\text{rec}}, M_{\text{rec}})$; the observer supplies cuts, accessible records, and an entropy gradient over those records. Only the pair

$$(\mathcal{W}, \omega, S; \mathcal{O})$$

contains enough structure to define a thermodynamic exchange.

This separation is important. A bare substrate has no experienced heat flow, because it has no oriented record access. A bare observer cut has no physical content, because without substrate disturbance there is no cost whose variation can be read as flux. Thermodynamics is therefore recovered at the interface:

$$\text{regional disturbance flux} + \text{observer record entropy} \implies \text{thermodynamic balance.}$$

8.1 Cuts, entropy, and disturbance flux

Let $E_R : \mathcal{W} \rightarrow \mathcal{A}_R$ be a modularly stable observer cut, and let $\omega_R = \omega \circ E_R$ be the accessible state. The observer-side entropy of the cut is the entropy of accessible records. In finite approximants this is

$$S_R(\omega) = -\text{Tr}(\rho_R \log \rho_R),$$

while in the von Neumann algebraic setting the intended replacement is the relative modular entropy associated with the restricted state and its reference cut. This entropy is not a substrate scalar by itself; it is the entropy of a state as viewed through an admissible access map, and its second variation along admissible deformations is the record susceptibility carried by M_{rec} .

The substrate-side quantity is the variation of \mathcal{K}_R across the same cut. For a cut deformation, boundary support deformation, or connection variation ξ , define

$$\delta Q_R(\xi) := \delta_R \mathcal{K}_R(\rho, \nabla, g; \xi),$$

where δ_R denotes variation measured across the boundary between the accessible algebra \mathcal{A}_R and its complement in the cut decomposition. In a boundary form this may be read as a disturbance flux

$$J_{\partial R} \sim \langle \mathfrak{D}_R, K_{\text{rec}} \nabla_n \mathfrak{D}_R \rangle_{\partial R},$$

whenever an appropriate normal derivative or shell-gradient is available. The key point is that the heat-like quantity is not an additional source; it is the observer-accessible flux of the same paired regional disturbance.

8.2 Recovered Clausius relation

The modular data of (\mathcal{W}, ω) supply the scale relating record-entropy change to access-flow change. Denote this scale by Θ_R . A cut is in local substrate–observer thermodynamic equilibrium when

$$\boxed{\delta Q_R = \Theta_R \delta S_R.}$$

This equation is not imposed at the substrate level. It is recovered only when disturbance flux through K_{rec} is paired with observer-access entropy whose Hessian is carried by M_{rec} . In this sense, thermodynamics is neither fundamental time nor merely psychological memory; it is the equilibrium bookkeeping of regional disturbance exchange across record-bearing cuts.

8.3 Relation to Jacobson

Jacobson’s 1995 argument [25] treats the Einstein equation as an equation of state by imposing $\delta Q = T dS$ on all local Rindler horizons, with horizon entropy proportional to area and Unruh temperature fixing the local temperature scale. The present framework keeps the structural

moral but changes the primitive data. Jacobson begins with spacetime causal horizons; here the primitive object is a modularly stable algebraic cut. Jacobson's heat is energy flux through a horizon; here the corresponding flux is variation of the paired master disturbance across an observer cut. Jacobson's entropy is horizon entropy; here the entropy is accessible record entropy or its modular relative-entropy replacement.

The claim is therefore narrower than Jacobson's. This paper does not derive the Einstein equation. It proposes that gravity-like response becomes equation-of-state-like when the connection variation of one paired regional disturbance is read through record-bearing cuts. The inertia/gravity alignment of the next section is not itself thermodynamics; it supplies the shared generator whose variation becomes δQ_R once an observer cut is introduced.

8.4 Why the placement matters

The Shared-Connection principle can be stated as a substrate bookkeeping result: inertia and gravity-like response are not independent costs for the same paired regional disturbance. The thermodynamic relation needs more. It needs an observer cut, an entropy of accessible records, and modular stability of that cut. Therefore the logical order is

$$\begin{aligned} \text{paired regional disturbance } (K_{\text{rec}}, M_{\text{rec}}) &\Rightarrow \text{inertial/gravity-like readouts} \\ &\Rightarrow \text{observer cut} \Rightarrow \text{thermodynamic balance.} \end{aligned}$$

9 Shared-Connection principle from one paired regional disturbance and Type II₁ trace uniqueness

The primary substrate-side route to the alignment of inertial response and gravity-like response is not a literal equality between two separately postulated forces. It is the compression of both readouts into one paired regional disturbance bookkeeping $(K_{\text{rec}}, M_{\text{rec}})$. The structural argument is the following.

The Type II₁ isotropic quadratic normal-form theorem of Section 4.1 (Theorem 4.1) says that, for Hilbert–Schmidt conditional-expectation defects, any positive, continuous, twice-differentiable, fully unitary-covariant disturbance functional has leading form

$$F(y) = c \tau_H(y^* y) + o(\|y\|_{HS}^2), \quad c > 0,$$

where τ_H is the canonical trace on the Hilbert–Schmidt geometry of $L^2(\mathcal{M}, \tau)$. The Murray–von Neumann uniqueness of the trace on a Type II₁ factor says that the substrate has no second canonical tracial weight from which to build a second independent leading defect geometry. Therefore any admissible readout that is a variation of the same paired regional disturbance inherits the same canonical quadratic bookkeeping, weighted through K_{rec} and measured by M_{rec} .

Proposition 9.1 (Shared generator for inertial and gravity-like readouts). *Let (\mathcal{M}, τ) be a Type II₁ factor and let X be a persistent excitation whose admissible support and connection variations are encoded by Hilbert–Schmidt defects satisfying the hypotheses of the Type II₁ normal-form theorem. Let*

$$\mathcal{K}_R(\rho, \nabla, g) = \tau[\mathfrak{D}_R(\rho)^* K_{\text{rec}}(\nabla, g, R) \mathfrak{D}_R(\rho)]$$

be the local regional disturbance. Define the inertial readout along an algebraic support-deformation direction v by

$$\mathcal{I}_R(v) = \left. \frac{d^2}{d\epsilon^2} \mathcal{K}_{R+\epsilon v}(\rho, \nabla, g) \right|_{\epsilon=0},$$

and the gravity-like readout along an accessible connection or metric-like variation h by

$$\mathcal{G}_R(h) = \left. \frac{d}{d\epsilon} \mathcal{K}_R(\rho, \nabla, g + \epsilon h) \right|_{\epsilon=0}.$$

Then \mathcal{I}_R and \mathcal{G}_R are not independent primitive couplings. They are variational projections of the same paired trace-quadratic disturbance geometry. In weak, isotropic, low-disturbance normalizations where the connection variation induced by an algebraic support deformation is represented by the same positive operator Δ_∇ , both reduce to the same leading coefficient,

$$\mathcal{I}_R(v) \simeq q_R^2 v^T \Delta_\nabla v, \quad \mathcal{G}_R(v) \simeq q_R^2 v^T \Delta_\nabla v,$$

with equality understood as a limiting shared-generator result, not as a primitive identity between separate forces.

Proof sketch. Both readouts are derivatives of one admissible local disturbance functional applied to defect variables generated by the same persistent support. By the normal-form theorem, the leading local geometry of those defects is anchored to the same canonical trace and hence to the same scalar coefficient c . Independent leading couplings would require independent canonical quadratic geometries on the same defect sector. A Type II₁ factor supplies only one faithful normal tracial state up to scale, and the full-covariance normal form supplies only one leading Hilbert–Schmidt defect norm. Thus the two readouts share a generator. In the weak isotropic limit the support and connection variations are represented through the same connection stiffness, giving the displayed common quadratic form. The paired susceptibility M_{rec} controls the same defect through entropy/distinguishability, so that the common eigenvalue of $G_{\text{acc}} = M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2}$ governs both readouts. \square

This is the algebraic-substrate analog of the equivalence principle, not the full Lorentzian general-relativistic equivalence principle. In ordinary language, the equivalence principle says that inertial response and gravitational response cannot be operationally separated by the free fall of a test body. In the present framework that alignment is not introduced as a spacetime postulate. It is the observer-effective expression of a substrate bookkeeping fact: a Type II₁-like disturbance geometry supplies one canonical tracial quadratic coefficient for persistent regional defects, and the paired $(K_{\text{rec}}, M_{\text{rec}})$ structure carries it consistently through the generalized record operator. To obtain the empirical free-fall statement, the observer-side recovery of support motion is still needed. That recovery is supplied in Section 11.1 as least-disturbance record continuation, and its weak-field version is supplied in Section 11.6.

9.1 Alternative motivation: support re-embedding as paired disturbance variation

A persistent excitation has one substrate cost for admissible re-embedding of its support,

$$\mathcal{C}_X(v; \nabla) = \mathcal{K}_{R+v}(\rho, \nabla, g) - \mathcal{K}_R(\rho, \nabla, g).$$

If inertia and gravity-like response used different primitive functionals for the same support deformation, the identity of X would depend on which observer protocol was used. If the two operations merely commuted but remained distinct, the joined disturbance would still contain duplicate bookkeeping. Since the substrate has no primitive temporal order deciding which operation happens first, regional join disturbance must be path-independent:

$$\delta_I \delta_G \mathcal{K}_{RVS} = \delta_G \delta_I \mathcal{K}_{RVS}.$$

Minimal bookkeeping then collapses the two generators to one paired connection-dependent disturbance. The trace-uniqueness derivation above is the sharper mathematical form of this intuition.

9.2 Composition independence: the binding-energy question

The Shared-Connection principle is not strong enough if it applies only to a single elementary support. The empirical content of the equivalence principle is composition-insensitive: different bodies can contain different fractions of rest-like persistence, binding contribution, kinetic contribution, electromagnetic contribution, or other internal energy. A substrate account must therefore explain why internal contributions to a persistent composite do not carry independent gravity-like charges.

The framework's answer is that a composite excitation is still measured by the same paired master disturbance functional. Let a compound support in a region A be represented by persistent sub-supports X_1, \dots, X_k . Its total substrate disturbance is one defect geometry,

$$\mathcal{K}_A(X_1, \dots, X_k) = \tau[\mathfrak{D}_A^* K_{\text{rec}} \mathfrak{D}_A],$$

where the composite defect may decompose into physically labeled components. The observer's physical labels — rest-like, binding, kinetic, electromagnetic, gauge, or internal — describe different components of the defect, not different disturbance functionals. The Type II₁ normal form gives

$$\mathcal{K}_A = c \|\mathfrak{D}_A\|_{2,\tau,K_{\text{rec}}}^2 + o(\|\mathfrak{D}_A\|_{2,\tau}^2),$$

with the same c for all local defect contributions. Internal composition can change the magnitude and shape of the defect, but it cannot introduce separate canonical couplings c_{rest} , c_{bind} , or c_{em} . Such couplings would require separate canonical trace-quadratic geometries, which the finite-factor substrate does not supply.

A simple worked bookkeeping example makes the point explicit. Suppose the local defect associated with a composite support decomposes, relative to an observer's physical labels, as

$$\mathfrak{D}_X = \mathfrak{D}_{\text{rest}} + \mathfrak{D}_{\text{em}} + \mathfrak{D}_{\text{bind}} + \mathfrak{D}_{\text{kin}}.$$

Then the leading substrate disturbance is not a sum with independent gravitational weights. It is one paired Hilbert–Schmidt norm:

$$\mathcal{K}_X = c \|\mathfrak{D}_X\|_{2,\tau,K_{\text{rec}}}^2 + o(\|\mathfrak{D}_X\|_{2,\tau}^2),$$

so that, when the explicit stiffness weighting is locally suppressed from the notation,

$$\mathcal{K}_X/c = \|\mathfrak{D}_{\text{rest}}\|_2^2 + \|\mathfrak{D}_{\text{em}}\|_2^2 + \|\mathfrak{D}_{\text{bind}}\|_2^2 + \|\mathfrak{D}_{\text{kin}}\|_2^2 + 2 \sum_{i<j} \Re\langle \mathfrak{D}_i, \mathfrak{D}_j \rangle_{2,\tau} + o(\|\mathfrak{D}_X\|_2^2).$$

The cross terms matter because composition changes the defect shape. What is forbidden, under the one-trace normal form, is a replacement such as

$$c_{\text{rest}} \|\mathfrak{D}_{\text{rest}}\|_2^2 + c_{\text{em}} \|\mathfrak{D}_{\text{em}}\|_2^2 + c_{\text{bind}} \|\mathfrak{D}_{\text{bind}}\|_2^2,$$

with independently canonical coefficients. That expression would require separate preferred tracial quadratic geometries for different internal sources. The Type II₁ substrate supplies only the single Hilbert–Schmidt geometry induced by τ and the paired $(K_{\text{rec}}, M_{\text{rec}})$ structure.

Thus the substrate-side equivalence statement is stronger than mass cancellation for an elementary support:

all contributions to persistent regional disturbance couple through one leading paired defect geometry.

The observer-side reconstruction of trajectories is still needed to translate this into composition-independent recorded free fall.

10 Scope, modularity, and claim status

The framework’s claims sit in three classes.

10.1 Consistency claims

The substrate-side features and observer-side features named above can be implemented together on one unchanged finite substrate. Finite validations verify coexistence and bookkeeping consistency. They do not establish naturalness or uniqueness.

10.2 Reconstruction claims

Given the substrate-side structure $(K_{\text{rec}}, M_{\text{rec}})$ and the observer’s access rules, the observer can reconstruct Newton-like mechanics, wavefunction-and-collapse, Maxwell-like equations, and a weak-field gravity response as effective descriptions of recorded experience.

10.3 Locational claims

Stable separability, persistence, stiffness, gravity-like response, and gauge connection are placed on the substrate side as readouts of paired regional disturbance. Time, the arrow, motion, collapse, and field evolution are placed on the observer side.

10.4 Scope of the claims

The framework should be read as a reconstruction program, not as a completed derivation of established physics. It studies the hypothesis that physical time, motion, thermodynamics, and field-like laws can be represented as observer-effective structures arising from access to a tenseless algebraic substrate. The paper does not assume as established fact that nature is such a substrate.

Several ingredients are specified constructively rather than derived uniquely from primitive axioms: the admissible regions, conditional expectations, connection sectors, source readouts, paired stiffness and susceptibility operators $(K_{\text{rec}}, M_{\text{rec}})$ and their structural decompositions, and observer-access rules. The compression introduced here reduces the number of independent disturbance readout sectors by treating them as readouts of one paired regional disturbance functional, but it does not yet derive that functional uniquely. The finite validations therefore demonstrate internal compatibility and feasibility, not naturalness, uniqueness, or empirical confirmation.

The recovered structures are structural analogues unless the additional continuum hypotheses stated later are satisfied. In particular, the paper does not claim an unconditional derivation of the Einstein field equations, the Standard Model, measured coupling constants, a physical continuum spacetime, or full diffeomorphism invariance. Section 12.2 gives the conditional continuum calculation required for the gravity claim: if the observer-limit action has an Einstein–Hilbert leading term and if the defect sector supplies a conserved stress tensor through metric variation, then the junction stationarity condition gives the Einstein equation. The coordinate-graph models used in the validation appendix are benchmarks, not substrate-intrinsic definitions. The dynamic metric-like field used there is a finite symmetric positive matrix field, not a Lorentzian metric.

Finally, “observer” does not mean a conscious agent. It denotes a memory-bearing access structure, in the operational sense used in quantum information and algebraic reconstruction.

10.5 Modularity

The two layers are modular. A reader interested only in the structural organization of physics can use the substrate layer without committing to the observer layer. A reader interested only in the foundations of temporal experience can use the observer layer without committing to the substrate layer. Each layer stands on its own. Together they yield a unified picture in which substrate phenomena and observer phenomena are explicitly separated and explicitly related, the bridge being the observer's cut and record accumulation.

11 Junction recovery principle for effective laws

The thermodynamic relation of Section 8 is recovered only after both layers are present. The same rule is used for all the other physical sectors in this paper. An effective law is counted as recovered, rather than merely stipulated as an observer readout, when it follows from three ingredients:

paired regional disturbance + observer access + junction stationarity/compatibility.

The substrate supplies tenseless objects: conditional-expectation defect, paired regional disturbance ($K_{\text{rec}}, M_{\text{rec}}$), support stiffness through K_{rec} , record susceptibility through M_{rec} , connection response, holonomy, and algebraic shell growth. The observer supplies record order, accessible algebras, stable cuts, and finite or modular bandwidth. The junction principle is a balance, least-action, stationarity, conservation, or invariant-access condition that is not meaningful on either layer alone.

Thus the target is not to say that Newtonian mechanics, Maxwell equations, Lorentz kinematics, and weak gravity are hand-selected observer readouts. The target is to recover each as the effective equation forced when the observer is allowed to read one paired substrate disturbance through a stable access cut. The recoveries below are still limited: they do not establish standard physics in nature, do not determine numerical constants, and do not solve the Type III extension. But their logical status is stronger than a witness language. They are junction recoveries of the framework.

11.1 Recovered Newtonian mechanics as least-disturbance record continuation

Newtonian mechanics is recovered at the junction of support disturbance and record order. The substrate supplies a stable support manifold \mathcal{S}_X for a persistent excitation X and a local disturbance functional $\mathcal{K}_X(q)$, where q^a are coordinates on the support-deformation manifold. These coordinates are not background space; they are parameters on the family of support embeddings accessible to the observer. The observer supplies a record chain $\lambda \mapsto \mathcal{A}_\lambda$ and an internal label parameter $t = t(\lambda)$.

The junction principle is least-disturbance record continuation. Among support histories compatible with the observer's record chain, the realized effective history is stationary for

$$\mathcal{I}_X[q] = \int \left(\frac{1}{2} \dot{q}^a K_{ab}(q) \dot{q}^b - \mathcal{K}_X(q) \right) dt,$$

where

$$K_{ab}(q) = \frac{\partial^2 \mathcal{K}_X}{\partial q^a \partial q^b}$$

is the support-stiffness Hessian inherited from the paired master disturbance, weighted by K_{rec} . A record is unstable when the conditional expectation representing the observer's accessible memory fails to approximately preserve the previous support state under continuation; equivalently,

the recovery error between successive observer cuts grows rather than remaining bounded. In this sense, a non-stationary path is not merely dynamically disfavored. It is a path along which the observer's algebra cannot maintain a stable conditional memory of the previous record.

The first variation gives

$$K_{ab} \ddot{q}^b + \Gamma_{abc} \dot{q}^b \dot{q}^c = -\partial_a \mathcal{K}_X,$$

with Γ_{abc} the connection induced by the support-stiffness metric when K_{ab} varies over the support manifold. In the local constant-stiffness limit this reduces to

$$K_{ab} \ddot{q}^b = -\partial_a \mathcal{K}_X.$$

Thus Newton-like mechanics is not substrate evolution. It is the observer's stable-record representation of stationary support continuation through a timeless paired disturbance landscape.

11.2 Finite record recoverability and the observer-junction light limit

Before introducing Lorentz kinematics, the speed limit must be located correctly. A tenseless substrate should not be assigned a primitive velocity bound by itself: velocity already presupposes both a distance and an oriented time parameter. The substrate supplies algebraic distance, disturbance barriers, modular stability, and conditional-expectation structure; the observer supplies record order, cuts, and an internal proper record increment. The finite light-like bound is therefore a substrate-observer junction quantity, not a substrate scalar alone.

Principle 11.1 (Finite Record Recoverability). A memory-bearing observer can recover new stable records only through admissible cut updates that preserve prior records, remain modularly stable, and do not cross forbidden disturbance-erasure barriers. The maximum rate at which newly recoverable stable records can enter the observer's accessible algebra defines the observer-junction speed.

It is useful to distinguish this speed from a substrate propagation rate. If a chosen substrate flow or connection update admits a rate of disturbance spread, write it schematically as v_{sub} . This is not yet the speed experienced by an observer, because propagation is weaker than record formation. A substrate disturbance becomes an event for an observer only if it passes three filters:

1. it is acquired into the accessible algebra \mathcal{A}_λ ;
2. it decoheres or stabilizes into the record algebra \mathcal{M}_λ ;
3. it remains monotonically recoverable along \preceq .

The observer-recoverability speed c_* is the maximal rate satisfying all three requirements. Consequently

$$c_* \leq v_{\text{sub}},$$

and the inequality can be strict.

Let $\mathcal{O} = (\Lambda, \preceq, \{\mathcal{A}_\lambda\}, \{\mathcal{M}_\lambda\})$ be an observer. The substrate supplies an intrinsic algebraic distance $d_\omega(R, S)$ between record supports, defined by the state and regional conditional-expectation structure rather than by coordinates. The observer supplies a proper record increment $\Delta\tau_\lambda > 0$ between adjacent labels. For a transition $\lambda \rightarrow \lambda^+$, define the admissible newly recoverable transitions by

$$\mathcal{R}_\lambda := \{(R, S) : R \subset \mathcal{A}_\lambda, S \subset \mathcal{A}_{\lambda^+}, S \text{ contains a newly recoverable stable record caused by } R\},$$

subject to the three record-recoverability constraints

$$S \subset \mathcal{A}_{\lambda^+}, \quad S \subset \mathcal{M}_{\lambda^+}, \quad \mathcal{M}_\lambda \subseteq \mathcal{M}_{\lambda^+},$$

together with modular stability and the disturbance-barrier condition

$$\|E_\lambda \sigma_t^\omega - \sigma_t^\omega E_\lambda\|_{2,\omega} \leq \epsilon, \quad \Delta \mathcal{K}^{\text{erase}}(S) \geq \Theta_{\text{rec}}.$$

The one-step record expansion is

$$\Delta_\lambda^{\text{rec}} := \sup_{(R,S) \in \mathcal{R}_\lambda} d_\omega(R, S),$$

and the observer-recoverability speed is

$$c_* := c_{\mathcal{O}} := \sup_\lambda \frac{\Delta_\lambda^{\text{rec}}}{\Delta \tau_\lambda}.$$

The corresponding access cone is

$$d_\omega(R_\lambda, R_\mu) \leq c_* |\tau_\mu - \tau_\lambda|.$$

This cone is not a statement that the substrate cannot contain nonlocal relations, nor that v_{sub} is the physical speed of light, nor that c_* has a universal numerical value in the present work. It is a statement that a record-bearing observer cannot make stable, ordered records recoverable faster than the admissible cut-update rate.

11.2.1 Paired record bandwidth from $(K_{\text{rec}}, M_{\text{rec}})$

The expected reduction form is the generalized eigenvalue of the paired stiffness–susceptibility operator. Let α_s be admissible algebraic deformations of the observer-access inclusion. These are not spatial translations in a pre-existing manifold; they are deformations of which regional factors, conditional expectations, and stable records are included in the cut. For a deformation direction v , the support-stiffness form is

$$\kappa_R(v, v) := \left. \frac{d^2}{ds^2} \mathcal{K}_{\alpha_s v}(R) \right|_{s=0} = \langle v, K_{\text{rec}} v \rangle,$$

the mass-like side of the bandwidth law. It measures the quadratic disturbance cost required to deform the accessible support while preserving record identity. The record-susceptibility form is

$$\chi_R(v, v) := \left. \frac{d^2}{ds^2} D(\phi_{\alpha_s v}(R) \parallel \phi_R) \right|_{s=0} = \langle v, M_{\text{rec}} v \rangle,$$

the entropy-like side. It measures how rapidly recoverable record information changes under the same algebraic deformation. In the sharper recovery version, χ_R is not merely an abstract entropy curvature. If $\mathcal{R}_R : \mathcal{A}_R \rightarrow \mathcal{W}$ is an admissible recovery channel associated with the cut, the junction recovery loss is

$$\mathcal{L}_R(\rho) := S(\rho \parallel \mathcal{R}_R E_R(\rho)),$$

whenever the relative entropy is finite. For a deformation $\rho_s = \alpha_s(\rho)$ generated by an admissible derivation, one may take

$$\chi_R(v, v) = \left. \frac{d^2}{ds^2} \mathcal{L}_R(\rho_s) \right|_{s=0}$$

up to the normalization used for modular record time. Thus record susceptibility is the Hessian of observer reconstruction loss. The zero-loss case is a cut whose accessible algebra recovers the relevant substrate state; positive loss is scalar unrecoverability. The record-speed scale is therefore not introduced independently; it is determined by the paired $(K_{\text{rec}}, M_{\text{rec}})$ structure.

Proposition 11.2 (Record-bandwidth bound from the paired generator). *Let R be an observer-access region with paired master disturbance operators $(K_{\text{rec}}, M_{\text{rec}})$, both positive. A deformation mode v propagated at inclusion rate r is stably recordable only if*

$$r^2 \chi_R(v, v) \leq \kappa_R(v, v), \quad \text{i.e.,} \quad r^2 \langle v, M_{\text{rec}} v \rangle \leq \langle v, K_{\text{rec}} v \rangle.$$

Consequently the local observer record-bandwidth bound is the generalized eigenvalue threshold

$$c_*^2(R) := \inf_{v \neq 0} \frac{\langle v, K_{\text{rec}} v \rangle}{\langle v, M_{\text{rec}} v \rangle} = \lambda_{\min} \left(M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2} \right) = \lambda_{\min}(G_{\text{acc}}),$$

with the infimum taken over modes that remain in the stable record sector. Equivalently, stable modes solve the generalized eigenvalue problem

$$K_{\text{rec}} v = c^2 M_{\text{rec}} v,$$

and the observer cone is the recoverable part of this generalized spectrum.

Proof. A deformation at inclusion rate r changes the observer state by an amount whose second-order distinguishability cost is $r^2 \chi_R(v, v)$ per unit record increment. The same mode is stabilized by the disturbance stiffness $\kappa_R(v, v)$ supplied by the substrate. If the distinguishability growth exceeds the available stiffness, the reverse channel cannot keep previous records approximately recoverable and the deformation is not a stable element of the record algebra. Thus stable recordability requires $r^2 \chi_R(v, v) \leq \kappa_R(v, v)$ for every admissible mode. Solving for the largest uniformly admissible rate gives the Rayleigh-quotient bound above. The Rayleigh quotient $\langle v, K_{\text{rec}} v \rangle / \langle v, M_{\text{rec}} v \rangle$ equals the spectrum of $M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2} = G_{\text{acc}}$ since positive M_{rec} admits a symmetric square root, and the minimum of the quotient is the smallest eigenvalue of G_{acc} . \square

This proposition is the framework-native version of the relation between mass and entropy in the access speed. Mass-like inertia is not inserted as a separate constant; it is the stiffness Hessian carried by K_{rec} of the same paired regional disturbance. Entropy is not an independent thermodynamic add-on; it is the susceptibility of record distinguishability carried by M_{rec} under the same deformation. The speed bound is the lowest generalized eigenvalue of their pair through G_{acc} .

In a finite or continuum model, a universal value of c_* would require the generalized spectrum of $(K_{\text{rec}}, M_{\text{rec}})$ to converge under refinement to the substrate access spectrum G_{acc} , and to be independent of the admissible observer cut. That universality is not assumed here.

11.2.2 Compatibility without metric coincidence

The substrate framework requires only

$$\text{spec} \left(M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2} \right) \rightarrow \text{spec}(G_{\text{acc}}),$$

where G_{acc} is the substrate's intrinsic access-spectrum operator. It does not impose

$$K_{\text{rec}} = M_{\text{rec}}^{1/2} L_{\text{acc}} M_{\text{rec}}^{1/2}$$

as a definition. The metric-compatible form is a possible realization of the compatibility condition — one in which G_{acc} literally equals the access Laplacian L_{acc} — but the framework only requires that the generalized ratio of the two independently constructed operators stabilizes under refinement to the same access spectrum. Both K_{rec} and M_{rec} are then nontrivial substrate-built operators, and neither is defined from the other.

11.2.3 Extreme regimes

The same paired bandwidth formula gives precise framework-native placements for several limiting cases. In a black-hole-like region \mathcal{B} , exterior recovery fails when the entropy susceptibility carried by $M_{\text{rec},\mathcal{B}}^{\text{ext}}$ overwhelms the support stiffness carried by $K_{\text{rec},\mathcal{B}}^{\text{ext}}$,

$$\lambda_{\min}\left((M_{\text{rec},\mathcal{B}}^{\text{ext}})^{-1/2} K_{\text{rec},\mathcal{B}}^{\text{ext}} (M_{\text{rec},\mathcal{B}}^{\text{ext}})^{-1/2}\right) \rightarrow 0.$$

Then

$$c_{*,\text{ext}}^2(\mathcal{B}) \rightarrow 0,$$

so the horizon is an exterior record-bandwidth collapse rather than a primitive substrate wall. Interior information need not be destroyed in \mathcal{W} ; it becomes nonrecoverable to the exterior algebra because the reverse channel is dominated by record susceptibility at the boundary cut.

In quantum regimes, the defect covariance $P_{\mathfrak{D}}$ and the record susceptibility M_{rec} dominate the observer readout. Large M_{rec} -eigenvalue ratios relative to K_{rec} produce probabilistic, noisy, or nonseparable records; small ratios give stiff, low-noise records and therefore a classical-looking stable history. The bandwidth law therefore places the quantum/classical transition in the generalized spectrum of $(K_{\text{rec}}, M_{\text{rec}})$.

In cosmological regimes, the relevant region is a large accessible domain Ω . The global disturbance baseline contributes to the observer-effective vacuum offset, while the large-scale bandwidth is

$$c_{*,\Omega}^2 = \inf_v \frac{\langle v, K_{\text{rec},\Omega} v \rangle}{\langle v, M_{\text{rec},\Omega} v \rangle}.$$

A cosmological-constant-like term is then not a localized matter disturbance but the large-region baseline part of the paired disturbance-action dictionary. Its effect on large-scale causal appearance is mediated by the same pair of quantities: global stiffness through K_{rec} and global record susceptibility through M_{rec} .

11.2.4 Finite generalized-eigenvalue benchmark

The generalized-eigenvalue statement can be tested in a deliberately minimal finite model without assigning any physical value to c_* . Let the admissible support-deformation space be $V_R \cong \mathbb{R}^6$. Choose two positive-definite matrices realizing the structural decompositions, with each ingredient SPD on V_R and the coefficients chosen so that $K_{\text{rec}}, M_{\text{rec}} > 0$. For a direct record-stability simulation, a deformation mode v updated at inclusion rate r is declared stable exactly when

$$r^2 \langle v, M_{\text{rec}} v \rangle \leq \langle v, K_{\text{rec}} v \rangle.$$

The eigenvalue prediction is not separately chosen; it is the Rayleigh-quotient threshold

$$c_*^2 = \lambda_{\min}\left(M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2}\right) = \inf_{v \neq 0} \frac{\langle v, K_{\text{rec}} v \rangle}{\langle v, M_{\text{rec}} v \rangle}.$$

In a representative run using the structural decomposition with simple admissible choices for each ingredient and unit normalization weights, the generalized eigenvalue calculation reproduces the same critical threshold to which Monte Carlo random-direction sampling converges, and the critical-direction test makes the transition explicit: rates below the threshold leave the inequality positive, rates equal to the threshold give equality on the critical eigenvector, and rates above the threshold violate the inequality. This is the logical necessity test of the paired generalized eigenvalue construction: when the same $(K_{\text{rec}}, M_{\text{rec}})$ define direct record stability, the first unstable mode is exactly the generalized eigenmode, and the record-bandwidth threshold is the corresponding eigenvalue bound on G_{acc} .

11.2.5 Stability envelope and null record modes

The record-bandwidth bound can be phrased geometrically without assigning a numerical value to c_* . For a region R , define the local record-stability envelope

$$\mathcal{C}_R := \{(r, v) : r^2 \langle v, M_{\text{rec}} v \rangle \leq \langle v, K_{\text{rec}} v \rangle\}.$$

Its boundary,

$$r^2 \langle v, M_{\text{rec}} v \rangle = \langle v, K_{\text{rec}} v \rangle,$$

is the null recoverability boundary of the observer cut. The critical modes that first saturate this boundary satisfy the generalized eigenvalue problem

$$K_{\text{rec}} v_* = c_*^2 M_{\text{rec}} v_*.$$

Thus the generalized eigenvector is not, by itself, a photon or a physical light ray. It is the first saturating algebraic deformation mode of the paired record-stability problem. A photon-like interpretation would require an additional identification: the saturating eigenspace must lie in the accessible holonomy/gauge sector, carry the appropriate transverse degrees of freedom, and obey the Maxwell-type stationarity condition of the low-disturbance readout. The framework therefore treats critical eigenvectors as null record modes; identifying a null record mode with light is a reduction target, not an assumption.

The same language clarifies the failure modes. If $c_* = \infty$, there is no finite recoverability cone and the observer layer becomes Galilean-like rather than Lorentzian. If $c_* = 0$ in an ordinary region, distinct stable records cannot propagate and extended dynamics fails. If $c_*(R, v)$ depends on direction, the local geometry is anisotropic, closer to a Finsler or medium-like record geometry than to Lorentz spacetime. If different admissible observers disagree about the boundary $r^2 \langle v, M_{\text{rec}} v \rangle = \langle v, K_{\text{rec}} v \rangle$, then there is no observer-independent causal structure. Consequently the eigenvalue bound supplies the speed scale, but Lorentz kinematics requires the additional invariant-cone hypotheses stated in the next subsection.

For black-hole-like cuts, the same envelope gives a precise threshold formulation. A horizon is approached when the exterior envelope collapses in the interior-to-exterior directions,

$$\lambda_{\min} \left((M_{\text{rec}, \partial \mathcal{B}}^{\text{ext}})^{-1/2} K_{\text{rec}, \partial \mathcal{B}}^{\text{ext}} (M_{\text{rec}, \partial \mathcal{B}}^{\text{ext}})^{-1/2} \right) \rightarrow 0,$$

or falls below the recovery threshold of the exterior observer. This resembles an entropy-area constraint only after one proves that the boundary stiffness scales with area and that the susceptibility readout coincides with physical entropy. Recovering the Bekenstein–Hawking [33, 34] coefficient is therefore a separate reduction target, not a result of the present finite diagnostic.

11.2.6 Light as the saturating record sector

The electromagnetic or gauge/light sector is identified with the finite speed only when its records saturate the recoverability bound. More precisely, if the massless gauge-record sector is the minimal-disturbance stable sector for which

$$d_\omega(R_\lambda, R_{\lambda+}) = c_{\mathcal{O}} \Delta \tau_\lambda$$

on admissible transitions, then

$$c_{\text{light}} = c_{\mathcal{O}}.$$

Thus the framework does not say that light measures how fast the substrate moves, and it does not require $c_{\text{light}} = v_{\text{sub}}$. It says that light is the fastest stable way an observer can receive decohered, monotonically recoverable records through the substrate–observer junction. The electromagnetic sector inherits the paired bookkeeping with a matching coupling renormalization:

$$K_{\text{rec}}^{\text{em}} = g_{\text{eff}}^{-2} K_{\text{rec}}, \quad M_{\text{rec}}^{\text{em}} = g_{\text{eff}}^{-2} M_{\text{rec}}.$$

The matched pair ensures that the electromagnetic record cone

$$c_{\text{em}}^2 = \lambda_{\min} \left((M_{\text{rec}}^{\text{em}})^{-1/2} K_{\text{rec}}^{\text{em}} (M_{\text{rec}}^{\text{em}})^{-1/2} \right) = \lambda_{\min} \left(M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2} \right) = c_*^2$$

coincides with the universal record cone: the g_{eff}^{-2} factor cancels exactly between numerator and denominator, so the electromagnetic coupling g_{eff} controls the Maxwell amplitude scale without altering the universal record cone. This is the speed that enters the Lorentz recovery below. We write $c_* := c_{\mathcal{O}}$ throughout.

11.3 Recovered Lorentz kinematics from invariant modular-access cones

Lorentz kinematics is the hardest sector. The preceding section gives a finite observer-junction speed c_* as a maximum stable record-access rate determined by the lowest generalized eigenvalue of G_{acc} . A finite speed alone, however, is still not enough to force Lorentz transformations. The stronger substrate–observer recovery requires an invariant access cone: all admissible observers must agree on which record events can be joined by finite record recoverability.

The substrate supplies an algebraic distance, not a coordinate distance. For regions or record supports R, S , write $d_{\omega}(R, S)$ for the state-dependent algebraic distance defined by conditional-expectation attenuation, modular distinguishability, or shell depth in the algebraic coarsening category. The observer supplies stable record increments and cut updates constrained by the Finite Record Recoverability Principle. The junction principle is invariant accessibility:

$$d_{\omega}(R_{\lambda}, R_{\mu}) \leq c_* |\tau_{\lambda} - \tau_{\mu}|$$

must have the same truth value for every observer using the same modular state class and stable cut category.

When the accessible record histories admit a two-dimensional normal form spanned by record duration and one support-separation variable, the invariant of this access cone is

$$ds_{\text{acc}}^2 = c_*^2 d\tau^2 - d_{\omega}^2.$$

Transformations preserving the access cone preserve ds_{acc}^2 , and the usual Lorentz factor follows in an observer chart:

$$d\tau_{\text{prop}} = dt \sqrt{1 - v^2/c_*^2}, \quad \gamma = (1 - v^2/c_*^2)^{-1/2}.$$

A stable excitation with support stiffness m then has effective action

$$S_{\text{obs}} = -mc_*^2 \int d\tau_{\text{prop}},$$

which gives

$$p = \gamma m v, \quad E = \gamma m c_*^2, \quad E^2 - p^2 c_*^2 = m^2 c_*^4.$$

At rest, $E_0 = m c_*^2$.

This is a recovery only under the invariant modular-access-cone hypothesis. It is stronger than a bandwidth compatibility statement, but it still marks the open technical point honestly: on a purely tracial Type II₁ substrate the modular flow is trivial, so nontrivial Lorentzian access geometry requires either a Type III extension or an observer cut whose accessible algebra carries nontrivial modular data. The law is recovered at the junction; it is not proved from the Type II₁ substrate alone.

Operator-algebraic consequence for primitive retrieval supports. The recovery of a Lorentz-like access cone has a direct operator-algebraic consequence for the primitive spatial retrieval projections of $G_{\text{acc}}|_{S_{\text{obs}}}$: under the irreducibility of the spatial subgroup $G_{\text{sp}} \subset G_{\text{Lor}}(\Sigma_{\text{obs}})$ that the recovered cone provides, these primitive projections form a single G_{sp} -orbit and share a common trace. This result is stated and proved as the cone-recovery trace lemma (Theorem 4.3) in Section 4.2, where it is used as a structural ingredient for sector readouts that depend on equal-trace primitive support equivalence. The lemma is conditional on the same Lorentz-recovery program described in this subsection: if the recovered $G_{\text{Lor}}(\Sigma_{\text{obs}})$ has the spatial irreducibility property, equal trace of primitive retrieval projections follows; if not, both this subsection and Theorem 4.3 share the same residual gap.

11.4 Wavefunction and collapse: observer-side reading

The restriction-to-accessible reading of measurement collapse and the wavefunction is observer-layer machinery, developed in Section 6.12 of the observer-layer chapter. There the wavefunction ρ_λ is identified as the observer-effective state on \mathcal{A}_λ , Born probabilities arise from $\omega_\lambda(P_k)$ on record-compatible projectors, and Lüders conditionalization $\rho_{\lambda'} = P_k \rho_\lambda P_k / \text{Tr}(P_k \rho_\lambda)$ updates the observer-effective state without touching the substrate state ω . The junction-recovery reading is therefore: ordinary measurement is the special case of accessible-algebra restriction in which the cut \mathcal{A}_λ contains the record-compatible projectors, and collapse is record-conditioned access rather than substrate-level dynamics.

11.5 Recovered Maxwell-type equations from a paired holonomy readout

The Maxwell sector is recovered by applying the same junction rule to the substrate's internal connection sector through the paired $(K_{\text{rec}}, M_{\text{rec}})$ bookkeeping. The substrate supplies a compact internal automorphism sector with Abelian edge transport in the weak-field sector. The observer supplies accessible gauge records and conserved source records. The junction principle is stationary low-disturbance holonomy response, understood as the Abelian connection readout of the paired master regional disturbance.

Let A be the observer-accessible Abelian connection 1-cochain and let

$$F = dA$$

be its curvature. The identity

$$dF = d^2A = 0$$

is substrate-side algebra: it follows from nilpotency of the coboundary and requires no source equation. The sourced equation is recovered only after an observer cut supplies a conserved accessible current J . Conservation is not a separate dynamical law here; it is the compatibility condition that the current record can be written through the same access cut without creating or destroying charge inside the inaccessible complement:

$$\delta J = 0.$$

Define the holonomy-disturbance action on the accessible connection sector through the matched pair $(K_{\text{rec}}^{\text{em}}, M_{\text{rec}}^{\text{em}}) = (g_{\text{eff}}^{-2} K_{\text{rec}}, g_{\text{eff}}^{-2} M_{\text{rec}})$ by

$$\mathcal{A}_{\text{em}}[A; J] = \frac{1}{2g_{\text{eff}}^2} \langle dA, dA \rangle_\omega - \langle J, A \rangle_\omega,$$

where the inner product is the observer-accessible quadratic form induced by the cut, the state, and the substrate's connection data. This action is not an independent primitive sector; it is the Abelian holonomy-sector quadratic expansion of \mathcal{K}_R under accessible connection variations,

with the matched paired bookkeeping in the electromagnetic sector. Stationarity with respect to accessible gauge variations $A \mapsto A + \eta$, modulo pure-gauge directions, gives

$$\delta dA = g_{\text{eff}}^2 J.$$

Together with $F = dA$, this yields the Maxwell-type system

$$\boxed{dF = 0, \quad \delta F = g_{\text{eff}}^2 J.}$$

Thus the field equations are not obtained by defining J from F . The current is an observer-accessible conserved record, while F is the stationary low-disturbance holonomy response to that record. The coupling g_{eff} controls the amplitude through the matched pair, but does not alter the universal record cone $c_*^2 = \lambda_{\min}(G_{\text{acc}})$, because the g_{eff}^{-2} factor cancels between the matched $K_{\text{rec}}^{\text{em}}$ and $M_{\text{rec}}^{\text{em}}$.

Finite U(1) stationarity benchmark. The validation companion includes a finite sourced U(1) solve in which the accessible connection is varied while the same paired master disturbance supplies the quadratic holonomy cost. The important point of that computation is not that a lattice script proves Maxwell theory. It is that relaxing the accessible connection lowers the single paired disturbance functional and satisfies the discrete stationarity equation in the presence of a conserved current. This is the finite counterpart of the claim made here: the Abelian gauge law is an accessible-connection stationarity readout of \mathcal{K}_R , not a separately inserted strain sector.

Gauge invariance is also a junction property. Under $A \mapsto A + d\chi$,

$$F \mapsto dA + d^2\chi = F,$$

and the source term changes by

$$\langle J, d\chi \rangle_\omega = \langle \delta J, \chi \rangle_\omega = 0$$

for conserved accessible current. Therefore gauge invariance and charge conservation are two sides of the same access-compatible stationarity condition.

The framework's substrate-side selection rule for when a non-Abelian internal connection is required is developed in Section 3.7: low-disturbance algebraic coarsenings whose generators do not commute force the stationary connection into a noncommutative compact sector, with closed-loop mismatch measured by the conjugacy-invariant Wilson defect $W_\gamma = \Re \text{Tr}(I - H_\gamma)$. The Maxwell sector recovered above is the Abelian special case in which all admissible coarsening transitions commute.

11.6 Recovered weak-field gravity from stationary accessible connection response

The weak-gravity sector is recovered from the same substrate-observer pattern. The substrate supplies localized regional disturbance and the paired connection-derived geometry $(K_{\text{rec}}, M_{\text{rec}})$. The observer supplies an accessible scalar or metric-like connection potential. The junction principle is stationary accessible connection response to paired disturbance.

Let X be a persistent excitation. Its source is not independently declared; it is the accessible density associated with the local disturbance increment,

$$T[X] = \Pi_0(\delta\mathcal{K}_R(X) p_X),$$

where p_X is the support profile and Π_0 removes the null mode of the accessible connection operator. Let Δ_∇ be the positive connection-stiffness operator induced by K_{rec} on the accessible cut. The observer-accessible potential ϕ is recovered by stationarity of

$$\mathcal{A}_{\text{grav}}[\phi; X] = \frac{1}{2}\langle \phi, \Delta_\nabla \phi \rangle_\omega - \kappa \langle T[X], \phi \rangle_\omega.$$

For all admissible variations η in the gauge-fixed accessible sector,

$$\delta\mathcal{A}_{\text{grav}}[\phi; X](\eta) = \langle \eta, \Delta_{\nabla}\phi - \kappa T[X] \rangle_{\omega}.$$

Stationarity gives the recovered weak-field equation

$$\boxed{\Delta_{\nabla}\phi = \kappa T[X].}$$

This has the same logical form as the recovered thermodynamic balance: a substrate quantity, here $\delta\mathcal{K}_R(X)$, is paired with an observer-access quantity, here ϕ , and the effective equation is the stationarity condition at their junction. The equation is not a postulated Poisson law; it is the low-disturbance accessible connection response to a paired regional disturbance source.

Mass-independent test acceleration is recovered as a weak-limit consequence of the shared-generator principle. Let Y be a test excitation with support-stiffness scalar $m_Y > 0$ carried by K_{rec} . Its inertial readout is the support Hessian of \mathcal{K}_R , while its gravity-like coupling is the connection variation of the same \mathcal{K}_R . In the weak, isotropic, low-disturbance limit both readouts reduce to the same trace-normalized coefficient. The force read by the observer is therefore

$$F_{\text{grav}}(Y) = m_Y g_{\nabla}, \quad g_{\nabla} = -\nabla_{\text{acc}}\phi,$$

while the inertial readout is

$$m_Y a_Y = F_{\text{grav}}(Y).$$

Therefore

$$a_Y = g_{\nabla},$$

independent of the test support's own disturbance magnitude. If Y is composite, Corollary 4.2 of Section 12 implies that binding, kinetic, electromagnetic, and rest-like internal contributions all enter through the same leading paired regional disturbance geometry; the cancellation is therefore composition-insensitive in the weak-limit observer readout.

This recovery is still weak-field and observer-effective. It is not a derivation of Einstein's equation, diffeomorphism invariance, or Lorentzian geodesic motion. Those require the Type III/modular extension and the invariant access-cone construction. What is recovered here is the Newtonian weak-field structure: paired regional disturbance sources an accessible connection, and test support acceleration is independent of the test body's own disturbance coefficient because the same paired master disturbance defines both inertial and gravity-like readouts in the weak limit.

12 Algebraic necessity program: paired disturbance, access, shell growth, non-Abelian selection, and observer cuts

The preceding sections define the observer-effective readouts on top of the algebraic substrate. The necessity program records the stronger claim that the framework must pursue: the familiar equations should not merely be compatible with engineered readouts, but should arise as stability or stationarity conditions of the substrate itself. This section recasts four remaining choices as substrate-facing derivation targets. The claims below remain programmatic in part, but they remove coordinate assumptions from the substrate core and are formulated so that future numerics and operator-algebraic models can test them without appealing to an ambient coordinate system.

The Type II₁ isotropic quadratic disturbance normal-form theorem (Theorem 4.1) and the cone-recovery trace lemma (Theorem 4.3) are now collected in Section 4, immediately after the substrate layer. Both are tracial results that anchor multiple later derivations to the same canonical trace; placing them next to the substrate definitions removes a forward-reference chain

that ran from Section 9 and Section 15 through to the necessity program. The present section continues with the remaining items of the necessity program: spectral covariance of defect modes from observer coarse-graining (Section 12.1), the continuum Einstein-translation schema (Section 12.2), the proof obligations for an airtight Einstein reduction (Section 12.3), and the remaining structural items on disturbance stationarity, algebraic shell growth, modular observer access, and non-Abelian selection.

12.1 Spectral covariance of defect modes from observer coarse-graining

The Type II₁ normal-form theorem of Section 4.1 fixes the local quadratic geometry of conditional-expectation defect, but it does not yet explain why a paired trace-quadratic disturbance should have the kind of continuum asymptotics that can produce an Einstein–Hilbert term. The missing bridge is spectral. If the generalized record operator $G_{\text{acc}} = M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2}$ has a Laplace-type continuum limit, and if the observer-access defect covariance is controlled only by the spectrum of G_{acc} , then the master disturbance becomes a spectral weighted trace. Heat-kernel methods can then generate a local expansion in volume, scalar curvature, and higher-curvature invariants.

This subsection records the precise justification for the spectral assumption. It is not an additional physical law. It is the least-biased covariance assignment available to an observer whose access algebra cannot distinguish microscopic defect labels beyond their disturbance cost under the paired geometry.

Exact finite spectral rewriting. Let $\mathcal{H}_{\mathfrak{D}}(R)$ be the finite defect Hilbert space associated with the regional defect \mathfrak{D}_R . Let

$$K := K_{\text{rec}}(\nabla, g, R) \geq 0$$

be the stiffness operator on that defect space and $M := M_{\text{rec}}(\nabla, g, R) > 0$ the matching susceptibility. The generalized record operator is

$$G_{\text{acc}} = M^{-1/2} K M^{-1/2}.$$

If $G_{\text{acc}}\phi_\alpha = \lambda_\alpha\phi_\alpha$, and

$$\mathfrak{D}_R = \sum_{\alpha} d_{\alpha} M^{1/2} \phi_{\alpha},$$

then the finite disturbance is exactly

$$\mathcal{K}_R = \langle \mathfrak{D}_R, K \mathfrak{D}_R \rangle = \sum_{\alpha} \lambda_{\alpha} |d_{\alpha}|^2 \cdot \langle \phi_{\alpha}, \phi_{\alpha} \rangle.$$

Equivalently, if $P_{\mathfrak{D}}$ denotes the defect covariance or density on $\mathcal{H}_{\mathfrak{D}}(R)$, then

$$\mathcal{K}_R = \text{Tr}(K P_{\mathfrak{D}})$$

and, in the G_{acc} -normalized basis,

$$\mathcal{K}_R = \text{Tr}\left(G_{\text{acc}} M^{1/2} P_{\mathfrak{D}} M^{1/2}\right).$$

Thus the master disturbance is always a paired spectral energy once the defect space and the operators $(K_{\text{rec}}, M_{\text{rec}})$ are fixed. What is not automatic is that this weighted trace becomes a universal spectral action rather than a mode-by-mode memory of microscopic defect preparation.

Spectral covariance as the universal regime. A universal spectral action follows if the defect covariance is approximately a function of the generalized record operator,

$$P_{\mathfrak{D}} \approx f(G_{\text{acc}}/\Lambda^2).$$

Then

$$\mathcal{K}_R = \text{Tr}(K f(G_{\text{acc}}/\Lambda^2)) = \Lambda^2 \text{Tr} F(G_{\text{acc}}/\Lambda^2), \quad F(x) := xf(x),$$

which is a spectral-action-type expression. The question is therefore not whether \mathcal{K}_R can be diagonalized. It can. The question is why observer-access covariance should commute with G_{acc} , so that only spectral shells matter.

Coarse-graining justification. The observer does not access the substrate's full microscopic labels. The observer has stable access only to record-preserving quantities. If G_{acc} is the generator of recoverability cost (the lowest eigenvalue of G_{acc} being c_*^2), then two defect modes in the same spectral shell of G_{acc} have the same observable disturbance cost. An observer whose algebra contains no additional stable labels cannot assign a preferred basis inside such a shell. Therefore the covariance must be invariant under all unitaries that preserve G_{acc} . This is the algebraic meaning of coarse-graining over unobserved defect labels.

Theorem 12.1 (Spectral covariance from observer coarse-graining). *Let G_{acc} be the positive self-adjoint generalized record operator on a finite regional defect Hilbert space $\mathcal{H}_{\mathfrak{D}}$, and let $P_{\mathfrak{D}} \geq 0$ be the observer-access defect covariance. Suppose:*

1. *the observer fixes the total accessible defect weight $\text{Tr} P_{\mathfrak{D}} = N_{\mathfrak{D}}$;*
2. *the observer fixes the mean disturbance $\text{Tr}(G_{\text{acc}}P_{\mathfrak{D}}) = E_{\mathfrak{D}}$;*
3. *the observer has no stable labels inside a spectral shell of G_{acc} , so $UG_{\text{acc}}U^* = G_{\text{acc}} \implies UP_{\mathfrak{D}}U^* = P_{\mathfrak{D}}$ for every unitary U on $\mathcal{H}_{\mathfrak{D}}$.*

Then $P_{\mathfrak{D}}$ commutes with G_{acc} and is block-spectral:

$$P_{\mathfrak{D}} = \sum_{\lambda} p_{\lambda} \Pi_{\lambda},$$

where Π_{λ} is the spectral projection of G_{acc} . In the nondegenerate or smooth-shell limit this is equivalently

$$P_{\mathfrak{D}} = f(G_{\text{acc}}/\Lambda^2)$$

for some spectral filter f and access scale Λ . If, in addition, $P_{\mathfrak{D}}$ maximizes von Neumann entropy subject to the two constraints above, then

$$P_{\mathfrak{D}} = Z^{-1} \exp[-\beta G_{\text{acc}}/\Lambda^2].$$

Proof. Let $G_{\text{acc}} = \sum_{\lambda} \lambda \Pi_{\lambda}$ be the spectral decomposition. If $UG_{\text{acc}}U^* = G_{\text{acc}}$, then U is arbitrary inside each eigenspace $\Pi_{\lambda}\mathcal{H}_{\mathfrak{D}}$ and cannot mix eigenspaces with distinct eigenvalues. The invariance condition $UP_{\mathfrak{D}}U^* = P_{\mathfrak{D}}$ for all such U implies, by Schur's lemma on each eigenspace, that $P_{\mathfrak{D}}$ is scalar on every eigenspace and has no off-diagonal blocks between distinct eigenspaces. Hence $P_{\mathfrak{D}} = \sum_{\lambda} p_{\lambda} \Pi_{\lambda}$, so $[P_{\mathfrak{D}}, G_{\text{acc}}] = 0$. If the spectrum is nondegenerate, set $f(\lambda/\Lambda^2) = p_{\lambda}$. If the eigenspaces are degenerate but the shell weights vary smoothly with λ , the same relation holds in the coarse spectral calculus.

For the maximum-entropy statement, maximize $S(P_{\mathfrak{D}}) = -\text{Tr}(P_{\mathfrak{D}} \log P_{\mathfrak{D}})$ subject to $\text{Tr} P_{\mathfrak{D}} = N_{\mathfrak{D}}$ and $\text{Tr}(G_{\text{acc}}P_{\mathfrak{D}}) = E_{\mathfrak{D}}$. Variation gives $-\log P_{\mathfrak{D}} - I - \alpha I - \beta G_{\text{acc}} = 0$, so $P_{\mathfrak{D}} = Ce^{-\beta G_{\text{acc}}}$, with $C = Z^{-1}$ fixed by $\text{Tr} P_{\mathfrak{D}} = N_{\mathfrak{D}}$. Restoring the observer access scale gives the displayed form. \square

Maximum entropy and modular equilibrium. The theorem gives two compatible readings. In the coarse-graining reading, $P_{\mathfrak{D}} = f(G_{\text{acc}}/\Lambda^2)$ follows from ignorance of microscopic labels inside spectral shells. In the maximum-entropy reading, the exponential filter follows because the observer fixes only total defect weight and average disturbance. If the observer representation is nonracial, the same structure can be read modularly. Let

$$H_{\text{mod}} = -\log \rho_{\mathcal{O}}$$

be the modular Hamiltonian of the observer-access state. In local equilibrium of the defect sector one expects a KMS/Gibbs form $P_{\mathfrak{D}} \propto e^{-\beta H_{\text{mod}}}$. If the modular generator is controlled in the defect sector by the generalized record operator, $H_{\text{mod}} \approx G_{\text{acc}}/\Lambda^2$, then

$$P_{\mathfrak{D}} \approx Z^{-1} e^{-\beta G_{\text{acc}}/\Lambda^2}.$$

Laplace-type continuum bridge. If, under refinement of the regional substrate,

$$G_{\text{acc}}(\nabla_n, g_n, R) \longrightarrow -\nabla_g^2 + E$$

as a Laplace-type operator on an observer-effective continuum geometry, then the spectral weighted trace becomes amenable to heat-kernel asymptotics. In schematic form,

$$\text{Tr } F(G_{\text{acc}}/\Lambda^2) \sim c_0 \Lambda^d \int_M \sqrt{|g|} d^d x + c_1 \Lambda^{d-2} \int_M R \sqrt{|g|} d^d x + c_2 \Lambda^{d-4} \int_M R^2 \sqrt{|g|} d^d x + \dots,$$

where R^2 denotes quadratic curvature invariants and bundle-curvature terms. Thus the regional disturbance has the possible low-curvature form

$$\mathcal{K}_R^{\text{geom}} \sim \int_M (a_0 + a_1 R + a_2 R^2 + a_3 R_{\mu\nu} R^{\mu\nu} + \dots) \sqrt{|g|} d^d x.$$

In the long-distance observer regime, if $a_1 \neq 0$ and higher-curvature terms are suppressed, this becomes the Einstein–Hilbert reduction target

$$\mathcal{K}_R^{\text{geom}} \approx \frac{1}{16\pi G_{\text{eff}}} \int_M (R - 2\Lambda_{\text{eff}}) \sqrt{|g|} d^d x.$$

The present theorem does not prove this final limit. It reduces the problem to showing that the generalized record operator is Laplace-type, local, and observer-diffeomorphism invariant in the continuum limit, and that the scalar-curvature coefficient does not vanish.

Where the argument can fail. The spectral reduction fails if the observer has stable access to microscopic labels not generated by G_{acc} , because then $P_{\mathfrak{D}}$ need not commute with G_{acc} . It also fails if the defect sector is far from maximum-entropy or modular equilibrium, if G_{acc} has no local Laplace-type continuum limit, if nonlocal tails survive refinement, or if the spectral filter makes $a_1 = 0$. These failure modes are useful: they make the reduction checkable inside the framework. A substrate that cannot satisfy them may still support finite disturbance bookkeeping, but it would not supply a route to Einstein–Hilbert dynamics.

Reduction status. The result of this subsection is therefore a conditional bridge, not a completed derivation of general relativity. It proves that observer coarse-graining and maximum-entropy assignment force the defect covariance to be spectral whenever the generalized record operator is the only stable label. It also shows why this matters: the master disturbance then becomes a spectral weighted trace, and spectral traces of Laplace-type operators are the known

mathematical setting in which volume, scalar curvature, and higher-curvature terms appear. The remaining derivation problem is sharp:

$$\tau(\mathfrak{D}_R^* K_{\text{rec}} \mathfrak{D}_R) \implies \text{Tr} F(G_{\text{acc}}/\Lambda^2) \implies \int_M (R - 2\Lambda) \sqrt{|g|} d^d x,$$

after a controlled observer-continuum limit. The first arrow is supplied by spectral covariance through G_{acc} ; the second still requires a Laplace-type continuum theorem and a nonzero scalar-curvature coefficient.

12.2 Continuum Einstein-translation of the junction variation

The preceding spectral bridge gives the correct mathematical place where an Einstein–Hilbert term can appear, but it does not by itself execute the macroscopic variation. This subsection records the exact continuum translation required by the gravity claim. It is deliberately stated as a conditional theorem schema, because the conclusion is only as strong as the continuum hypotheses: locality of the generalized record operator limit, spectral covariance of the defect sector, emergence of an observer-effective Lorentzian metric, control of boundary terms, and a conserved defect-source tensor.

The purpose is to separate three statements that are easy to conflate:

1. *Architecture*: gravity-like response is the connection or metric variation of the paired master regional disturbance.
2. *Continuum dictionary*: the long-distance observer representation rewrites that disturbance as a local action of an effective metric and defect sources.
3. *Einstein execution*: variation of that local action with respect to the effective metric gives the Einstein tensor on the geometric side and stress-energy on the defect side.

Only the third item answers the question: what does the variation actually produce?

Macroscopic dictionary. Let a stable observer scaling limit of access labels produce an effective continuum region Ω with metric-like field $g_{\mu\nu}$, volume density $\sqrt{|g|} d^d x$, and admissible metric variations $\delta g_{\mu\nu}$ compactly supported in Ω or accompanied by the appropriate boundary term. The dictionary is

paired spectral trace \longrightarrow geometric effective action,

$$\text{Tr} F(G_{\text{acc}}/\Lambda^2) \longrightarrow \int_{\Omega} (a_0 \Lambda^d + a_1 \Lambda^{d-2} R[g] + a_2 \Lambda^{d-4} I_2[g] + \dots) \sqrt{|g|} d^d x,$$

localized defect sector \longrightarrow observer-effective matter/source action,

$$\mathfrak{D}_R, P_{\mathfrak{D}}, \mathcal{L}_R \longrightarrow S_{\text{def}}[g, \Psi],$$

junction stationarity $\longrightarrow \delta_g S_{\text{obs}} = 0$.

Here Ψ denotes the continuum fields or collective defect variables that summarize persistent regional disturbance at the observer scale. The term $I_2[g]$ denotes curvature-squared and higher derivative invariants.

The observer-effective action therefore has the form

$$S_{\text{obs}}[g, \Psi] = S_{\text{geo}}[g] + S_{\text{def}}[g, \Psi] + S_{\text{bdry}}[g],$$

with

$$S_{\text{geo}}[g] = \int_{\Omega} (A_0 + A_1 R[g] + A_2 I_2[g] + \dots) \sqrt{|g|} d^d x.$$

The constants A_i are not independently tuned gravitational couplings at this stage. They are continuum coefficients determined, if the reduction succeeds, by the spectral filter, cutoff scale, trace normalization, and the spectrum of G_{acc} .

Metric variation of the geometric term. Assume first that the long-distance regime is dominated by the volume and scalar-curvature terms,

$$S_{\text{geo}}[g] = \int_{\Omega} (A_0 + A_1 R[g]) \sqrt{|g|} d^d x + S_{\text{bdry}}[g].$$

The boundary functional is the observer-effective analogue of the Gibbons–Hawking–York correction [40, 41]: it is included, or variations are restricted, so that no uncontrolled normal-derivative variation remains on $\partial\Omega$. The standard metric variations are

$$\delta\sqrt{|g|} = -\frac{1}{2}\sqrt{|g|} g^{\mu\nu} \delta g_{\mu\nu},$$

and, after the boundary contribution is cancelled or fixed,

$$\delta\left(\int_{\Omega} R\sqrt{|g|} d^d x\right) = \int_{\Omega} G^{\mu\nu} \delta g_{\mu\nu} \sqrt{|g|} d^d x,$$

where $G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu}$. Therefore

$$\delta S_{\text{geo}} = \int_{\Omega} (A_1 G^{\mu\nu} - \frac{1}{2}A_0 g^{\mu\nu}) \delta g_{\mu\nu} \sqrt{|g|} d^d x + \delta S_{\text{higher}},$$

where δS_{higher} is the variation of the curvature-squared and higher-order terms. In the low-curvature regime where those terms are negligible, the geometric variation produces precisely the Einstein tensor plus a cosmological term.

Defect-source variation and stress-energy. The matter/source side is not introduced as a separate gravitational charge. It is the continuum readout of persistent regional paired defect structure. Define the observer-effective stress tensor by the metric variation of the defect action:

$$T_{\mu\nu}^{\text{def}} := -\frac{2}{\sqrt{|g|}} \frac{\delta S_{\text{def}}}{\delta g^{\mu\nu}}.$$

Equivalently,

$$\delta S_{\text{def}} = -\frac{1}{2} \int_{\Omega} T_{\mu\nu}^{\text{def}} \delta g^{\mu\nu} \sqrt{|g|} d^d x.$$

This is the continuum form of the framework’s source statement: the source is the metric/access variation of the same paired regional disturbance whose support Hessian through K_{rec} gives inertia. Rest-like, binding, electromagnetic, kinetic, gauge, and internal contributions may change the defect action S_{def} , but they enter a single tensor $T_{\mu\nu}^{\text{def}}$ because they are continuum components of one persistent defect sector rather than independent gravity-like charges.

Junction stationarity. The macroscopic junction condition is

$$\delta_g S_{\text{obs}}[g, \Psi] = 0$$

for all admissible metric-like observer variations. Combining the previous two variations gives

$$\int_{\Omega} \left(A_1 G^{\mu\nu} - \frac{1}{2}A_0 g^{\mu\nu} - \frac{1}{2}T_{\text{def}}^{\mu\nu} + H_{\text{higher}}^{\mu\nu} \right) \delta g_{\mu\nu} \sqrt{|g|} d^d x = 0,$$

where $H_{\mu\nu}^{\text{higher}}$ denotes the contribution from higher-curvature spectral terms. Since the variation is arbitrary inside the admissible observer-continuum sector,

$$A_1 G^{\mu\nu} - \frac{1}{2}A_0 g^{\mu\nu} + H_{\text{higher}}^{\mu\nu} = \frac{1}{2}T_{\text{def}}^{\mu\nu}.$$

If the higher-curvature corrections are negligible and $A_1 \neq 0$, this becomes

$$\boxed{G_{\mu\nu} + \Lambda_{\text{eff}} g_{\mu\nu} = \kappa_{\text{eff}} T_{\mu\nu}^{\text{def}}},$$

with $\Lambda_{\text{eff}} = -A_0/(2A_1)$ and $\kappa_{\text{eff}} = 1/(2A_1)$.

Theorem 12.2 (Conditional Einstein-translation theorem schema). *Suppose an observer-junction scaling limit satisfies the following conditions.*

1. *The generalized record operator $G_{\text{acc}} = M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2}$ has a local Laplace-type continuum limit on an observer-effective metric geometry (Ω, g) .*
2. *The defect covariance is spectral, $P_{\mathfrak{D}} = f(G_{\text{acc}}/\Lambda^2)$, so that the master disturbance reduces to a spectral weighted trace.*
3. *The heat-kernel expansion of that trace has a nonzero scalar-curvature coefficient A_1 and a controlled long-distance regime in which higher-curvature variations are negligible or explicitly retained as corrections.*
4. *The persistent defect sector admits a continuum source action $S_{\text{def}}[g, \Psi]$, and its metric variation defines a conserved stress tensor $T_{\mu\nu}^{\text{def}}$.*
5. *Boundary terms are fixed or cancelled so that the metric variation is well posed.*

Then the observer-junction stationarity condition $\delta_g S_{\text{obs}} = 0$ gives

$$G_{\mu\nu} + \Lambda_{\text{eff}} g_{\mu\nu} = \kappa_{\text{eff}} T_{\mu\nu}^{\text{def}}$$

in the low-curvature sector, with higher-curvature corrections when the suppressed spectral terms are retained.

Conservation and consistency. The equation is consistent only if the source is conserved relative to the observer-effective connection. The geometric side satisfies the contracted Bianchi identity, $\nabla_{\mu} G^{\mu\nu} = 0$. Therefore the low-curvature equation requires

$$\nabla_{\mu} T_{\text{def}}^{\mu\nu} = 0$$

up to exchange with higher-curvature or boundary sectors. In this framework that conservation law is not an independent matter postulate. It is the continuum expression of record-compatible paired-defect persistence under admissible access variations: a source that cannot be stably transported through the observer cut is not a conserved macroscopic source.

12.3 Proof obligations for an airtight Einstein reduction

The conditional theorem above identifies the correct variational target, but it is not yet a proof that the framework derives general relativity. The remaining problem is not rhetorical. Four mathematical obligations must be discharged before the Einstein-translation can be promoted from theorem schema to theorem.

Obligation 1: the exact action and the higher-curvature problem. The first task is to prove that the paired generator's spectral trace has an Einstein–Hilbert leading term and that the remaining terms are either suppressed in the observer regime or retained as controlled corrections:

$$\tau[\mathfrak{D}_R^* K_{\text{rec}} \mathfrak{D}_R] \simeq \text{Tr } F(G_{\text{acc}}/\Lambda^2) \sim A_0 \int_{\Omega} \sqrt{|g|} d^d x + A_1 \int_{\Omega} R \sqrt{|g|} d^d x + \sum_{j \geq 2} A_j \int_{\Omega} I_j[g] \sqrt{|g|} d^d x.$$

The Einstein reduction requires $A_1 \neq 0$ and the higher-curvature Euler–Lagrange tensors $H_{\mu\nu}^{(j)}$ to satisfy

$$\left\| \sum_{j \geq 2} A_j H_{\mu\nu}^{(j)} \right\| \ll |A_1| \|G_{\mu\nu}\|$$

in the low-curvature, long-distance sector, or else the framework predicts a higher-derivative gravity theory rather than Einstein gravity. The appropriate proof machinery is the heat-kernel or spectral-action expansion [81, 63, 64] for a Laplace-type positive operator on the paired generator G_{acc} . The structural reading of gravity as an induced-by-coarse-graining response is a substrate-level analog of the induced-gravity program [71, 72] and of the entropic gravity proposals [73, 74], with the present paper differing in primitive data: the substrate is a paired regional-disturbance bookkeeping rather than a vacuum-fluctuation or holographic-screen primitive. Thus G_{acc} cannot remain an arbitrary positive operator. It must be shown to have a continuum principal symbol

$$\sigma_2(G_{\text{acc}})(x, \xi) = g^{\mu\nu}(x)\xi_\mu\xi_\nu \mathbf{1}$$

up to bundle endomorphism terms, with locality, ellipticity or hyperbolic Wick-rotated control, and trace-class spectral filtering.

Obligation 2: the variation class and diffeomorphism-like invariance. The metric variation used in Theorem 12.2 is legitimate only if continuum metric changes correspond to admissible algebraic variations of the substrate–observer system. For every compactly supported infinitesimal observer diffeomorphism generated by a vector field X , there should be a corresponding family of admissible algebraic automorphisms α_s , preferably implemented in finite approximants by unitaries u_s , such that

$$\alpha_s(a) = u_s a u_s^*, \quad \left. \frac{d}{ds} \alpha_s(a) \right|_{s=0} = \delta_X(a),$$

and whose macroscopic action on the observer representation is the Lie derivative of the effective fields: $\delta_X g_{\mu\nu} = \mathcal{L}_X g_{\mu\nu}$, $\delta_X \Psi = \mathcal{L}_X \Psi$. The substrate-side invariance requirement is

$$\mathcal{K}_R(\alpha_s(\rho), \alpha_s \nabla, \alpha_s g) = \mathcal{K}_R(\rho, \nabla, g)$$

up to the controlled change of observer cut. This obligation also fixes the conservation law. If S_{obs} is invariant under the lifted algebraic diffeomorphism class, then variation along X gives, after integration by parts and arbitrariness of X , the covariant conservation constraint on the total source sector. In this reading, $\nabla_\mu T_{\text{def}}^{\mu\nu} = 0$ is the continuum shadow of record-compatible algebraic automorphism invariance.

Obligation 3: boundary terms and the access-cut contribution. The metric variation of $\int_\Omega R\sqrt{|g|}$ is not well posed on a region with boundary unless the normal-derivative boundary variation is cancelled or the variation class is restricted. In the present framework the boundary is not an auxiliary surface. It is the observer access-cut. Therefore the stronger target is not to add S_{GHY} by hand, but to derive the necessary boundary functional from access-cut defect and recovery loss:

$$S_{\text{cut}}[E_R, \rho] := \mathcal{L}_R(\rho) + B_R(\rho, \nabla, g) \longrightarrow 2A_1 \int_{\partial\Omega} K_{\partial\Omega} \sqrt{|h|} d^{d-1}x + S_{\text{edge}}[h, \Psi]$$

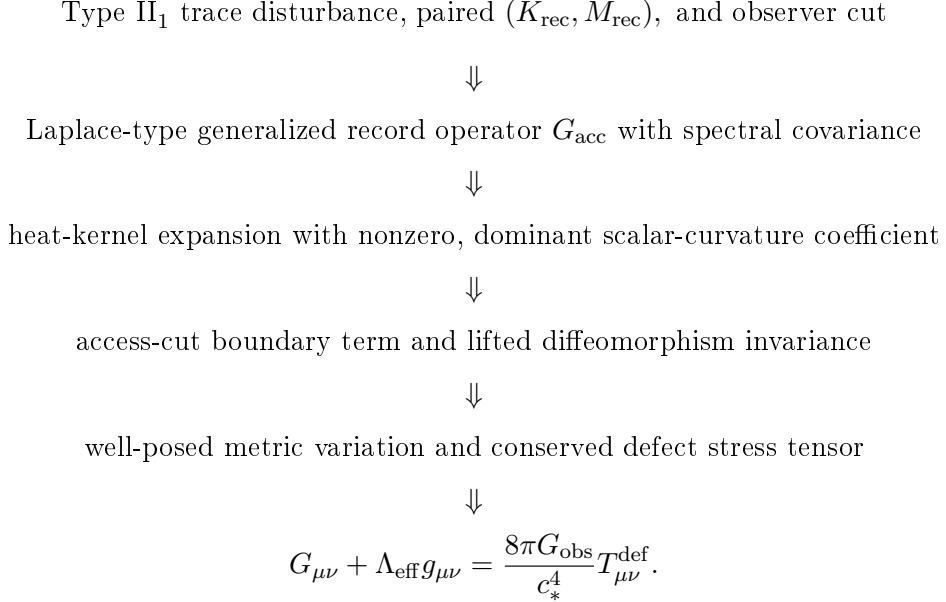
in the observer continuum limit. The proof obligation is that the first variation of this cut functional cancels exactly the normal-derivative part of the scalar-curvature variation.

Obligation 4: normalization and the emergence of Newton’s constant. The variational calculation yields $\kappa_{\text{eff}} = 1/(2A_1)$. To identify this with the physical value $8\pi G/c_*^4$, the framework must derive A_1 from the substrate trace normalization, the spectral filter, and the observer modular scale. Because c_* is already defined by the lowest generalized eigenvalue of G_{acc} , the desired substrate reading of Newton’s constant is

$$G_{\text{obs}} = \frac{1}{16\pi} \left(\inf_{v \neq 0} \frac{\langle v, K_{\text{rec}} v \rangle}{\langle v, M_{\text{rec}} v \rangle} \right)^2 \frac{1}{A_1}.$$

Thus small observed gravitational coupling corresponds, in this framework, to a large scalar-curvature stiffness coefficient A_1 relative to the observer's record-bandwidth scale.

Summary of the proof program. The completed Einstein proof therefore requires the chain



The framework already supplies the first line and the variational architecture. The remaining work is to prove the spectral, boundary, invariance, and normalization links without inserting the Einstein–Hilbert action or Newton coupling as independent assumptions.

12.4 Regional disturbance stationarity as the source of field actions

The weak-field and metric-like actions used above can be read in two ways. The weaker reading is engineering: one chooses a quadratic action whose Euler–Lagrange equation has the desired form. The stronger reading is substrate necessity: a persistent partition must settle into a low-disturbance configuration, and the familiar action is the second-order expansion of the paired master regional disturbance around a stable support.

Let Φ denote a collective field variable on a region R . It may be a scalar readout ϕ , a symmetric response tensor h , or a gauge connection perturbation. The primitive object is not an action but a paired regional disturbance readout

$$\mathcal{K}_R[\Phi] = \mathcal{K}_R^{(0)} + \langle J_R, \Phi \rangle + \frac{1}{2} \langle \mathfrak{D}_R \Phi, K_{\text{rec}} \mathfrak{D}_R \Phi \rangle + O(\|\Phi\|^3),$$

where \mathfrak{D}_R is the substrate comparison operator on admissible regional re-embeddings, K_{rec} is the positive stiffness form induced by the connection sector, and J_R is the first-order disturbance source created by localized support asymmetry. A stable regional readout is defined by $\delta\mathcal{K}_R[\Phi] = 0$. To quadratic order this gives

$$\mathfrak{D}_R^* K_{\text{rec}} \mathfrak{D}_R \Phi = -J_R.$$

Thus the finite actions used earlier are not fundamental postulates. They are normal forms for the quadratic approximation to low-disturbance regional stationarity through the paired $(K_{\text{rec}}, M_{\text{rec}})$ structure.

Diagnostic. A field equation counts as substrate-derived only to the extent that its operator and source arise from derivatives of \mathcal{K}_R :

$$K_{\text{phys}} = \nabla^2 \mathcal{K}_R|_{\Phi=0}, \quad J_{\text{phys}} = \nabla \mathcal{K}_R|_{\Phi=0}.$$

12.5 Algebraic shell growth instead of coordinate distance

Coordinate-graph benchmarks are retained only in the validation appendix. They are not part of the substrate definition or the fundamental refinement language. A region is represented instead by a subalgebra $\mathcal{A}_R \subset \mathcal{W}$, and coarsening is represented by the join $\mathcal{A}_{R \vee S} = \mathcal{A}_R \vee \mathcal{A}_S$. Define an intrinsic adjacency relation by nontrivial conditional-expectation coupling or by nonzero boundary disturbance readout:

$$R \sim S \iff \|E_R E_S - E_{R \vee S}\|_\omega > \varepsilon \quad \text{or} \quad \Phi_{\text{bdry}}(\mathcal{K}_{R,S}) > \varepsilon.$$

The algebraic distance is the minimal join-chain length

$$d_{\mathcal{A}}(R, S) = \min\{n : R = R_0 \sim R_1 \sim \dots \sim R_n = S\}.$$

The intrinsic shell around R is

$$\mathcal{S}_R(n) = \{S : d_{\mathcal{A}}(R, S) = n\}, \quad N_R(n) = |\mathcal{S}_R(n)|.$$

A conserved disturbance flux diluted uniformly across algebraic shells gives $F_R(n) \propto 1/N_R(n)$. Therefore an inverse-square law is obtained when the substrate has algebraic shell-growth dimension three: $N_R(n) \sim Cn^2 \implies F_R(n) \sim 1/n^2$.

12.6 Observer access from modular spectral data

Given a faithful state ω on a von Neumann algebra \mathcal{W} , Tomita–Takesaki theory supplies a modular automorphism group $\sigma_t^\omega : \mathcal{W} \rightarrow \mathcal{W}$. In the present framework, modular flow is not identified with external time. It is used to measure how rapidly accessible regional distinctions change under the state-algebra pair (\mathcal{W}, ω) .

Let \mathcal{A}_R be an accessible regional algebra and let $P_{\leq \Omega}$ denote the spectral projector onto modular frequencies resolvable by the observer’s record algebra. Define the modularly accessible neighborhood by

$$\mathcal{N}_\Omega(R) = \{S : \|P_{\leq \Omega}(\sigma_{\delta\tau}^\omega(\mathcal{A}_R)) - P_{\leq \Omega}(\mathcal{A}_S)\|_\omega < \epsilon\}.$$

The bandwidth radius is then an induced quantity,

$$r_*(\omega, \Omega) = \sup_{S \in \mathcal{N}_\Omega(R)} d_{\mathcal{A}}(R, S),$$

and the observer access speed is

$$c_*(\mathcal{W}, \omega, \Omega) = \frac{r_*(\omega, \Omega)}{\delta\tau_\omega}.$$

The remaining derivation target is to show that the gauge-action dispersion coefficient equals this modular access speed when record stability and low-disturbance propagation are both imposed, and that it coincides with the lowest generalized eigenvalue of G_{acc} .

12.7 Modularly stable observer cuts

A candidate accessible subalgebra $\mathcal{A} \subset \mathcal{W}$ with conditional expectation $E_{\mathcal{A}}$ is $(\omega, \delta\tau, \epsilon)$ -stable when

$$\|E_{\mathcal{A}}\sigma_t^\omega - \sigma_t^\omega E_{\mathcal{A}}\|_{2,\omega} \leq \epsilon, \quad |t| \leq \delta\tau,$$

and record-stable when $\Delta\mathcal{K}^{\text{erase}}(\mathcal{A}) \geq \Theta_{\text{rec}}$. The observer’s admissible cuts are therefore not arbitrary projections of the substrate. They are the subalgebras that are simultaneously approximately invariant under the modular flow and protected by a paired disturbance barrier through $(K_{\text{rec}}, M_{\text{rec}})$.

12.8 Hierarchy of necessity statements

The upgraded hierarchy is:

- equations of motion = low-disturbance paired stationarity;
- $1/r^2$ = inverse algebraic shell growth;
- c_* = lowest generalized eigenvalue of $G_{\text{acc}} = M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2}$ on stable cuts;
- non-Abelian response = selected by noncommuting coarsening holonomy.

These statements do not yet prove that nature is this substrate. They do make the tests non-circular: operators, speeds, flux laws, admissible observer cuts, and gauge actions must be recovered from conditional-expectation defect, paired generalized record geometry, algebraic shell growth, and coarsening holonomy rather than inserted as coordinate formulas.

13 Type II₁ to Type III₁ observer junction

The framework deliberately keeps two facts apart. The substrate side uses a Type II₁-like trace geometry because the trace gives a clean quadratic normal form for the paired $(K_{\text{rec}}, M_{\text{rec}})$ defect disturbance and a single coefficient for the shared-generator principle. The observer side, however, needs nontrivial modular structure if it is to support thermal time, modular access cones, and local-QFT-like behavior. The conclusion is: Type III₁ [8, 7] is not a new substrate postulate. It is a possible state-dependent observer-junction representation of accessible records, where time appears as modular flow in the nontracial representation selected by the observer access cut. The general structural pattern — a tracial substrate together with a non-tracial observer state generating modular flow and a crossed-product core — is closely related to recent work on observer algebras and the crossed product in semiclassical gravity [17, 18, 19].

13.1 Negative constraint: a normal cut of Type II₁ is not enough

Let (\mathcal{W}, τ) be a Type II₁ factor with faithful normal trace. If $\mathcal{A} \subset \mathcal{W}$ admits a τ -preserving conditional expectation, then \mathcal{A} is finite: $\tau|_{\mathcal{A}}$ is a faithful normal tracial state. An ordinary trace-preserving observer cut cannot turn a Type II₁ algebra into a Type III algebra. For a faithful normal state ω on \mathcal{W} , $\omega(x) = \tau(hx)$ for a positive density h , and the modular flow is inner: $\sigma_t^\omega(x) = h^{it} x h^{-it}$. This does not change the factor type. The Type III₁ structure required for modular time must enter through the observer representation or a scaling limit, not through a literal trace-preserving subfactor.

13.2 Positive bridge: non-tracial observer states and scaling limits

The proposed bridge has three ingredients:

tracial substrate geometry $(K_{\text{rec}}, M_{\text{rec}})$ + non-tracial observer state + scaling limit of record cuts.

At a finite level, let \mathcal{W}_n be finite trace approximants. An observer cut is represented by an accessible state $\omega_n(x) = \tau_n(\rho_n x)$ with modular Hamiltonian $H_n = -\log \rho_n$ and inner modular flow $\sigma_t^{\omega_n}(x) = \rho_n^{it} x \rho_n^{-it}$. The candidate Type III₁ observer algebra is

$$\mathcal{M}_\lambda^{\text{obs}} := \pi_{\omega_\lambda}(\mathcal{A}_\lambda)''.$$

A useful finite diagnostic is the modular ratio set Γ_n , generated by logarithmic ratios of eigenvalues of ρ_n . The Type III₁-like observer limit is suggested when $\bigcup_n \Gamma_n = \mathbb{R}$.

13.3 Access-cut as conditional expectation and modular embedding

The junction associated with λ is the modular embedding chain

$$(\mathcal{W}, \tau, K_{\text{rec}}, M_{\text{rec}}) \xrightarrow{E_\lambda} (\mathcal{A}_\lambda, \phi_\lambda) \xrightarrow{\pi_{\phi_\lambda}} \mathcal{M}_\lambda^{\text{obs}} \xrightarrow{\sigma_t^{\phi_\lambda}} \text{record flow}.$$

This is the formal statement that the junction is where time appears: a faithful nonracial observer state on the represented access algebra carries the Tomita–Takesaki modular automorphism group [9, 10, 6], while the substrate remains tenseless. The geometric content of modular flow in the local-QFT setting — that the modular flow of a wedge algebra is implemented by a Lorentz boost — is the classical Bisognano–Wichmann theorem [66, 67]; the present framework’s substrate-side reading is that boost-like access transformations are downstream consequences of the same modular structure that supplies observer-effective time.

Definition 13.1 (Modular observer embedding). A modular observer embedding of the paired tracial substrate presentation $(\mathcal{W}, \tau, K_{\text{rec}}, M_{\text{rec}})$ at access label λ is a quadruple $\mathcal{J}_\lambda = (E_\lambda, \mathcal{A}_\lambda, \phi_\lambda, \pi_{\phi_\lambda})$ such that $E_\lambda : \mathcal{W} \rightarrow \mathcal{A}_\lambda$ is an admissible access conditional expectation, ϕ_λ is a faithful normal record-compatible state, and $\mathcal{M}_\lambda^{\text{obs}} = \pi_{\phi_\lambda}(\mathcal{A}_\lambda)''$ is the represented observer algebra. It is *time-bearing* when $\sigma_t^{\phi_\lambda}$ is nontrivial, and *Type III₁-bearing* in the scaling-limit sense when the represented observer algebras have dense modular spectrum.

13.4 Crossed-product observer core

Given the represented observer algebra and its modular flow, define the modular crossed product [11]

$$\mathcal{C}_\lambda := \mathcal{M}_\lambda^{\text{obs}} \rtimes_{\sigma_{\phi_\lambda}} \mathbb{R}.$$

This algebra is the observer-junction core. The logical direction is: access cut \Rightarrow nonracial observer state \Rightarrow modular flow \Rightarrow crossed-product core, not crossed product \Rightarrow substrate evolution.

13.5 Support stiffness as a noncommutative derivation

Let δ be a closable $*$ -derivation, integrating locally to $\alpha_s = \exp(s\delta)$. For $\rho_s := \alpha_s(\rho)$, the inertial stiffness readout is

$$\mathcal{I}_\lambda(\delta) := \left. \frac{d^2}{ds^2} \mathcal{K}_\lambda(\rho_s, \nabla, g) \right|_{s=0},$$

where $\mathcal{K}_\lambda(\rho, \nabla, g) = \tau[\mathfrak{D}_\lambda(\rho)^* K_{\text{rec}} \mathfrak{D}_\lambda(\rho)]$. Inertia is therefore the resistance of the paired regional disturbance to identity-preserving algebraic support derivations, weighted by K_{rec} .

13.6 Junction recovery loss and record susceptibility

Let $\mathcal{R}_\lambda : \mathcal{A}_\lambda \rightarrow \mathcal{W}$ be an admissible recovery channel. The junction recovery loss is

$$\mathcal{L}_\lambda(\rho) := S(\rho \| \mathcal{R}_\lambda E_\lambda(\rho)).$$

For a support derivation δ with $\rho_s = \alpha_s(\rho)$, the record susceptibility Hessian is

$$\chi_\lambda(\delta, \delta) := \left. \frac{d^2}{ds^2} \mathcal{L}_\lambda(\rho_s) \right|_{s=0}.$$

This is the operational meaning of M_{rec} at the observer cut. The record-bandwidth law of Section 11.2 reads

$$c_*^2(\lambda) = \inf_{\delta \neq 0} \frac{\mathcal{I}_\lambda(\delta)}{\chi_\lambda(\delta, \delta)} = \lambda_{\min} \left(M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2} \right) = \lambda_{\min}(G_{\text{acc}}).$$

13.7 Temperature as a scaling-limit observer parameter

A thermal interpretation appears only when a stable record cut admits a modular scale with respect to which the restricted state has the KMS form [14]

$$\omega_\lambda(ab) = \omega_\lambda(b \sigma_{i\Theta^{-1}}^{\omega_\lambda}(a)).$$

Θ is not a primitive substrate temperature; it is the conversion factor between modular record-time and exterior-accessible entropy production, calibrated by the paired geometry $(K_{\text{rec}}, M_{\text{rec}})$. Unruh- [16] and Hawking-type temperatures are candidates for such observer-junction modular parameters.

13.8 Powers-type product states as the model example

Let $\mathcal{W}_n = M_2(\mathbb{C})^{\otimes n}$, $\tau_n = 2^{-n} \text{Tr}$, and $\omega_n = \bigotimes_{k=1}^n (p_k|0\rangle\langle 0| + (1-p_k)|1\rangle\langle 1|)$. This is the Powers product-state construction [15] of inequivalent factor representations of the UHF algebra. The finite modular spectrum is generated by $r_k = \log(p_k/(1-p_k))$. The same finite trace substrate supports three qualitatively different observer limits:

$$\begin{aligned} \text{tracial observer} &\Rightarrow \text{Type II-like / no modular arrow,} \\ \text{lattice modular ratios} &\Rightarrow \text{Type III}_\lambda\text{-like periodic modularity,} \\ \text{dense modular ratios} &\Rightarrow \text{Type III}_1\text{-like modularity.} \end{aligned}$$

13.9 Junction principle

Principle 13.2 (Type II₁ substrate, Type III₁ observer junction). The substrate disturbance geometry is computed in a tracial Type II₁-like presentation through the paired $(K_{\text{rec}}, M_{\text{rec}})$ bookkeeping. An observer is represented by a modular embedding

$$(\mathcal{W}, \tau, K_{\text{rec}}, M_{\text{rec}}) \xrightarrow{E_\lambda} (\mathcal{A}_\lambda, \phi_\lambda) \xrightarrow{\pi_{\phi_\lambda}} \mathcal{M}_\lambda^{\text{obs}} = \pi_{\phi_\lambda}(\mathcal{A}_\lambda)''.$$

The junction is where time appears: the effective time of the observer is $\sigma_t^{\phi_\lambda}$, not a primitive flow on \mathcal{W} . If the modular spectrum becomes dense in the scaling limit, the observer representation may be Type III₁-like.

13.10 Scaling-limit observer-recoverability speed: proposal status

For a finite approximant, the formal quantity is

$$c_{\mathcal{O},n} := \sup_\lambda \frac{\Delta_{\lambda,n}^{\text{rec}}}{\Delta\tau_{\lambda,n}},$$

where $\Delta_{\lambda,n}^{\text{rec}}$ is the largest algebraic record distance and $\Delta\tau_{\lambda,n}$ is the modular record increment. A scaling-limit light speed would be $c_* := \lim_{n \rightarrow \infty} c_{\mathcal{O},n}$, equal in the paired-operator language to the limiting lowest generalized eigenvalue of G_{acc} .

13.11 Block-cut diagnostic on the Powers family

For the product-state family, the new one-site modular increment is

$$\Delta\tau_{m,n} := \sqrt{\text{Var}_{\omega_n}(K_{m+1,n} - K_{m,n})} = \left| \log \frac{p_{m+1}}{1-p_{m+1}} \right| \sqrt{p_{m+1}(1-p_{m+1})}.$$

With $\Delta_{m,n}^{\text{shell}} = 1$, the candidate finite block-cut speed is $c_{\mathcal{O},n}^{\text{shell}} := \sup_{0 \leq m < n} 1/\Delta\tau_{m,n}$. Representative results for $n \leq 8$:

family	modular-density behavior	$c_{\mathcal{O},n}^{\text{shell}}$
$p_k = 1/2$	0 for all n	undefined
$p_k = 0.7$	lattice saturation	2.575457
$p_k = \sigma(0.35\sqrt{q_k})$	$\rightarrow 1$ by $n \approx 7$	4.164987

The modular-density column is a robust finite signal for Type III₁-like observer modularity; the product block-cut speed is a candidate construction, not yet a light-speed derivation.

13.12 Correlated observer-state record-bandwidth diagnostic

To obtain a nontrivial record-distance diagnostic, replace the product state by

$$\rho_n(h, J, \beta) = Z_n^{-1} e^{-\beta H_n(h, J)}, \quad H_n(h, J) = h \sum_{i=1}^n Z_i + J \sum_{i=1}^{n-1} Z_i Z_{i+1}.$$

Define $C_{m,m+1}^{(n)} = |\langle Z_m Z_{m+1} \rangle_{\rho_n} - \langle Z_m \rangle_{\rho_n} \langle Z_{m+1} \rangle_{\rho_n}|$, $\Delta_{m,n}^{\text{rec}} := -\log C_{m,m+1}^{(n)}$, $\Delta\tau_{m,n} := \sqrt{\text{Var}_{\rho_n}(I_{m+1})}$. The finite correlated observer-cone speed is

$$c_{\mathcal{O},n}^{\text{corr}}(h, J, \beta) := \sup_{1 \leq m < n} \frac{\Delta_{m,n}^{\text{rec}}}{\Delta\tau_{m,n}}.$$

This uses both sides required by the framework: algebraic record-distance from correlations (the K_{rec} side) and modular record-time cost (the M_{rec} side).

Gauge saturation. The gauge/light sector is identified with the limiting speed only when it saturates the same recoverability cone: $\lim_{n \rightarrow \infty} c_{\mathcal{O},n}^{\text{gauge}} = \lim_{n \rightarrow \infty} c_{\mathcal{O},n}^{\text{corr}} = c_* = \lambda_{\min}(G_{\text{acc}})$.

13.13 Finite diagnostic checklist

For a sequence of accessible blocks $(\mathcal{A}_n, \omega_n)$, compute:

1. the modular-ratio group Γ_n ;
2. the spectral spread of $H_n = -\log \rho_n$;
3. the stability of candidate record cuts under $\sigma_t^{\omega_n}$;
4. whether the paired trace-defect disturbance $(K_{\text{rec}}, M_{\text{rec}})$ remains stable under τ_n while modular access data becomes nontrivial;
5. a recoverability speed using an algebraic record distance and a modular increment.

The expected signature of the observer junction is the coexistence of stable paired substrate disturbance geometry under τ_n with a densifying observer modular spectrum under ω_n .

14 Horizon-cut thermodynamics and black-hole radiation

This section is a framework-native interpretation, not a derivation of Hawking's formula [34, 33]. Its purpose is to locate black-hole radiation in the same two-layer language used above for thermodynamics and observer recoverability through the paired $(K_{\text{rec}}, M_{\text{rec}})$ structure.

14.1 High-disturbance recoverability traps

At the substrate level a black hole is not first a region of spacetime. It is a high-disturbance recoverability trap: a region for which the exterior access map loses stable invertibility on interior record data. The region $\mathcal{B} \subset \mathcal{W}$ is black-hole-like for an exterior observer \mathcal{O} with cut $E_{\text{out},\lambda} : \mathcal{W} \rightarrow \mathcal{A}_{\text{out},\lambda}$ when three conditions hold:

1. large localized regional paired disturbance, $\mathcal{K}_{\mathcal{B}} \gg \mathcal{K}_{\text{ambient}}$;
2. high erasure threshold, $\Delta\mathcal{K}^{\text{erase}}(\mathcal{B}) \gg \Theta_{\text{rec}}$;
3. exterior nonrecoverability: the exterior recovery error $\varepsilon_{\mathcal{B}}(\mu) := \inf_{\mathcal{R}} \|\mathcal{R}(E_{\text{out},\mu}(\rho_{\mathcal{B}})) - \rho_{\mathcal{B}}\|_1 \geq \varepsilon_0 > 0$ for all admissible exterior cuts within the observer's record budget.

Equivalently, the interior-to-exterior recoverability rate vanishes:

$$c_*^{\mathcal{B} \rightarrow \text{out}}(\mathcal{O}) = 0, \quad \text{i.e.,} \quad \lambda_{\min} \left((M_{\text{rec},\mathcal{B}}^{\text{ext}})^{-1/2} K_{\text{rec},\mathcal{B}}^{\text{ext}} (M_{\text{rec},\mathcal{B}}^{\text{ext}})^{-1/2} \right) = 0$$

on interior modes. Thus: black hole = high-disturbance region whose internal records are substrate-real but exterior-inaccessible.

14.2 Horizons as recoverability boundaries

A horizon is not primitive geometry. It is the boundary of stable exterior recoverability. A boundary is horizon-like for \mathcal{O} when interior records fail the observer-cone and disturbance-barrier tests:

$$d_{\omega}(R_{\text{in}}, R_{\text{out}}) > c_* \Delta\tau \quad \text{or} \quad \Delta\mathcal{K}^{\text{erase}}(R_{\text{in}} \rightarrow R_{\text{out}}) > \Theta_{\text{rec}},$$

while the exterior cut still detects a boundary disturbance flux through the paired $(K_{\text{rec}}, M_{\text{rec}})$ geometry. Equivalently, $\partial\mathcal{B}$ = the cut where internal records cease to be stably recoverable outside.

14.3 Radiation as exterior disturbance flux

Radiation is the exterior observer's record of hidden paired disturbance relaxation. The substrate state remains ω on \mathcal{W} . The exterior observer sees only $\omega_{\text{out},\lambda} = \omega \circ E_{\text{out},\lambda}$. Define the outgoing disturbance heat by

$$\delta Q_{\partial\mathcal{B}}^{\text{out}} := -\delta\Phi_{\text{out}}(\mathcal{K}_{\mathcal{B}}),$$

with the convention that $\delta Q_{\partial\mathcal{B}}^{\text{out}} > 0$ means paired disturbance leaves the trapped region. Black-hole radiation is therefore outgoing disturbance flux across a horizon cut, recorded after exterior coarse-graining. At the primitive level, evaporation is a change in the exterior algebra's recovery and entropy bookkeeping; the particle language is an observer-effective decomposition of boundary correlations across an exterior cut.

14.4 Cut temperature

Temperature is not assigned to \mathcal{B} as a primitive substrate scalar. It is a cut quantity. Let $S_{\partial\mathcal{B}}$ denote the entropy of exterior-accessible boundary records. The horizon-cut thermodynamic balance is

$$\delta Q_{\partial\mathcal{B}}^{\text{out}} = \Theta_{\partial\mathcal{B}} \delta S_{\partial\mathcal{B}}.$$

Thus $\Theta_{\partial\mathcal{B}}^{-1} = \partial S_{\partial\mathcal{B}} / \partial Q_{\partial\mathcal{B}}^{\text{out}}$. A natural framework-native analogue of surface gravity is

$$\kappa_{\text{rec}} := \|\nabla_{\text{rec}} \Phi_{\text{rec}}(\mathcal{K}_{\mathcal{B}})\|_{\partial\mathcal{B}},$$

with the conjectural relation $\Theta_{\partial\mathcal{B}} \propto \kappa_{\text{rec}}$: the sharper the recoverability barrier, the hotter the boundary appears to an exterior observer.

14.5 Thermality and modular appearance

The exterior state can look thermal because the exterior observer has coarse-grained over inaccessible interior correlations. The exterior marginal is naturally described by a modular Hamiltonian $\rho_{\text{out}} \sim e^{-K_{\text{out}}}$. Thermal radiation is the modular exterior description of hidden boundary-disturbance correlations. Pair creation is an observer-effective factorization of a non-separable boundary correlation, not a fundamental substrate process.

14.6 Information loss and Page-like recovery

There is no substrate-level information destruction. The total state remains ω on \mathcal{W} . What changes is the exterior restriction $\omega_{\text{out}} = \omega \circ E_{\text{out}}$. Information loss is therefore relative to the exterior algebra:

information loss = failure of exterior records to recover substrate correlations.

Evaporation is the redistribution of paired regional disturbance from inaccessible support to accessible outgoing records:

$$\frac{d}{d\tau} \Phi_{\text{out}}(\mathcal{K}_{\mathcal{B}}) < 0, \quad \delta Q_{\partial\mathcal{B}}^{\text{out}} > 0.$$

A Page-like transition [35, 36] becomes a record-recoverability transition. Early radiation appears thermal because the relevant correlations are not yet recoverable through \mathcal{A}_{out} . Later recovery occurs if the accumulated exterior record algebra becomes sufficient to reconstruct those correlations: $\mathcal{M}_{\text{early rad}} \subseteq \mathcal{M}_{\text{late out}}$ up to controlled error. The recent island/quantum-extremal-surface program [37, 38, 39] supplies a holographic dictionary in which such a recovery transition is realized; in the present language the dictionary entry is record-algebra inclusion of $\mathcal{M}_{\text{early rad}}$ into a sufficient exterior cut. The structural reading of horizon entropy as entanglement entropy across an exterior cut has a long history [78, 79], and the generalized second law as recoverability monotonicity along the cut has been proved in the AQFT setting in [80].

A minimal Page-curve functional: let the full substrate state on the joined black-hole/radiation sector be globally recoverable in the substrate representation, while the exterior observer has only $\rho_{\text{rad},\lambda} = E_{\text{rad},\lambda}(\rho_{\mathcal{B}\cup\text{rad}})$. Define

$$S_{\text{ext}}(\lambda) := -\text{Tr} \rho_{\text{rad},\lambda} \log \rho_{\text{rad},\lambda}.$$

The framework-native Page behavior is

$$\frac{dS_{\text{ext}}}{d\lambda} > 0 \quad \text{before recovery completion}, \quad \frac{dS_{\text{ext}}}{d\lambda} \leq 0 \quad \text{after recovery completion},$$

while the full substrate description remains globally non-destroying. This is a Page-curve placement theorem-schema, not a numerical Page curve.

14.7 Summary of the framework-native reading

Inside the present framework the dictionary is:

- black hole = high-disturbance recoverability-trapping region,
- horizon = observer cut where internal records cease to be exterior-recoverable,
- radiation = exterior-accessible disturbance flux across that cut,
- temperature = modular conversion between paired disturbance flux and record entropy,
- information loss = loss relative to \mathcal{A}_{out} , not destruction in \mathcal{W} .

This is a placement result: black-hole radiation belongs to the same substrate–observer thermodynamic layer as $\delta Q_R = \Theta_R \delta S_R$, with the horizon interpreted as a paired-recoverability boundary and the radiation as recorded disturbance flux.

15 Internal charged sector: identification, minimality, and observer-junction diagnostics

The framework's substrate layer fixes a tenseless algebraic totality with paired regional disturbance $(K_{\text{rec}}, M_{\text{rec}})$; the observer layer fixes record-ordered cuts. Up to this point the internal automorphism sector of the substrate has been treated as additional declared structure. This section asks how far the framework can constrain that sector by substrate-style principles, what charged-sector structure follows, and what numerical diagnostics emerge at the observer junction through the paired bookkeeping.

The result is not a derivation of the Standard Model. It is a chain of conditional consequences. If physical charge is identified with the minimal observer-stable internal holonomy obstruction visible to admissible cuts, and if the access-mixing structure is built from the same observer cut and the same triple-overlap obstruction by the smallest trace-norm preserving recombinations, then the minimal internal algebra is fixed, the minimal chiral content is fixed, the hypercharge values are fixed by recoverability-obstruction cancellation, the trace ratio $\text{Tr}(Y^2)/\text{Tr}(T_3^2) = 5/3$ and the substrate-scale weak angle $\sin^2 \theta_W = 3/8$ follow, and a numerical fine-structure diagnostic and a charged-lepton ladder become available through the paired $(K_{\text{rec}}, M_{\text{rec}})$ structure.

The guiding methodological rule throughout is that no measured value of $\alpha, m_e, m_\mu, m_\tau$, or m_p is used to choose a coefficient. The electron mass enters only as a normalization unit when reporting dimensionless lepton ratios.

The fine-structure diagnostic output of Section 15.9, $A = \alpha^{-1}$, is the access capacity input used by the mass-shape sector of Section 5. The chain of dependencies is therefore: the substrate normal form of Section 4 supplies the Schur-complement machinery; the present section produces A as a Level-1 invariant via the Schur-trace identity; the mass-shape sector of Section 5 consumes A together with the four-axis observer-access lattice of Section 2.3 to predict the tau and baryonic mean access ratios; and the metrological completion of Section 7 carries those predictions into laboratory-comparable mass ratios. The present section is therefore upstream of the mass-shape sector in the framework's predictive chain.

The section is organized as a logical chain. Sections 15.1 through 15.6 establish the substrate-side primitives of the charged sector: charge as minimal observer-stable holonomy, the minimal three-branch closure, the minimal chiral content, hypercharge fixing, and the Maxwell-type identification. Sections 6.8 through 15.8 establish the three mechanism theorems that govern the access dressing: the Schur-complement theorem (sign and form of the dressing), the access-recovery composition theorem (primitive scalar access operator), and the projective access-cover theorem (termination of the polarization-order expansion at minimal closure). Section 15.9 assembles these results into the single Schur-trace identity for the fine-structure diagnostic. Section 15.11 records the corresponding strong-sector adjoint holonomy diagnostic built from the fixed $b = 3$ observer geometry and the non-projective record-boundary trace. Sections 15.12 through 15.14 carry out consistency checks: observer-access dimension ablation and the substrate-observer duality, the duality reversal that derives $T_{\text{bare}} = 21/2$ from observer-side primitives, and the substrate content realization that saturates the duality value. Section 15.15 collects independent consequences of the projective-readout axiom. The remaining sections cover the charged-lepton ladder, charge-structure consequences, what is not recovered, and the framework's status.

15.1 Charge as minimal observer-stable internal holonomy

A candidate charged sector is taken to be a nontrivial, observer-stable, internal holonomy label of a persistent regional defect. Four requirements are imposed:

- *internality*: not a displacement in substrate time or observer-effective spacetime;
- *record stability*: survives admissible observer cuts through the paired barrier;

- *relationality*: compares internal frames between access algebras;
- *non-degeneracy*: not a trivial identity or backtracking inverse.

Connection-block representation lemma. A gauge-sector readout in this framework is not a trace over matter support labels. It is a readout of the stiffness of an internal connection or holonomy variation of the paired regional disturbance. Matter support states may carry the corresponding internal label, but the gauge coupling itself is controlled by the response of the disturbance functional to changes of the internal transport data.

Let the admissible internal transition maps between observer-access algebras be regarded as local holonomy data. An infinitesimal variation of such holonomy data is an element of the Lie algebra of the corresponding internal automorphism group. Under a change of internal frame, Lie-algebra-valued connection variations transform by the adjoint action. Therefore, when a non-Abelian gauge-sector block is Schur-reduced, the hidden internal block is the adjoint connection block, not the fundamental matter-support block.

For the strong-sector readout, this means that the primary $SU(3)$ block is the adjoint color-holonomy block. The fundamental triplet describes color-carrying matter support, but it is not the connection block whose stiffness defines the strong-sector gauge readout. Thus the internal dimension and stiffness weight used in the strong-sector diagnostic are those of the $SU(3)$ adjoint sector.

Let $\mathcal{A}_i, \mathcal{A}_j, \mathcal{A}_k \subset \mathcal{W}$ be local access algebras and $U_{ij} : \mathcal{A}_j \rightarrow \mathcal{A}_i$ admissible internal transition maps. A one-frame loop $U_{ii} = 1$ carries no charged distinction. A two-frame loop $U_{ij}U_{ji} = 1$ tests reversibility. The first nondegenerate compatibility obstruction is a triple comparison:

$$\Omega_{ijk} = U_{ij}U_{jk}U_{ki}.$$

Under the primitive scalar-isotropic charged-cell assumption (the observer-readable part of the primitive charged obstruction lies in a central Abelian phase sector; the three edge transports are equivalent under cell symmetry; the lowest-disturbance charged readout is the scalar component of the obstruction), each edge carries the same central phase z :

$$U_{ij} = U_{jk} = U_{ki} = z, \quad \Omega_{ijk} = z^3.$$

Record-stable closure requires $z^3 = 1$, giving three solutions: $z_0 = 1$, $z_1 = e^{2\pi i/3}$, $z_2 = e^{4\pi i/3}$. The primitive charged sector has exactly three stable branches $\mathcal{D}_0, \mathcal{D}_1, \mathcal{D}_2$.

Higher N -frame loops decompose into triple-overlap pieces; thus $N \geq 4$ closures may represent composite or excited charged structures but do not define new primitive elementary charged sectors.

Proposition 15.1 (Primitive isotropic triple-overlap charged sector). *Let a charged sector be defined as a nontrivial, observer-stable, internal holonomy label of a persistent regional defect. Assume the substrate admits local access algebras with invertible pairwise internal transition maps U_{ij} . Then the first nondegenerate charged obstruction is $\Omega_{ijk} = U_{ij}U_{jk}U_{ki}$. Under the primitive scalar-isotropic charged-cell assumption, record closure gives $z^3 = 1$ and the primitive charged sector has exactly three stable branches.*

15.2 From discrete branches to continuous mixing

The triple-overlap closure delivers three discrete charged branches. The passage to the familiar non-Abelian factors requires two further moves grounded in substrate/access primitives.

Two-state access mixing as cut/complement. The substrate-derived two-state structure is the cut/complement decomposition at a fixed access label. The observer cut $E_\lambda : \mathcal{W} \rightarrow \mathcal{A}_\lambda$ defines, given a local charged disturbance D , two primitive components $D_\parallel = E_\lambda(D)$, $D_\perp = D - E_\lambda(D)$, so the minimal access-discrimination object is the pair

$$\Psi_D = \begin{pmatrix} D_\parallel \\ D_\perp \end{pmatrix}.$$

The trace-quadratic disturbance norm (weighted by the paired $(K_{\text{rec}}, M_{\text{rec}})$ at the same cut) is

$$\|\Psi_D\|_\tau^2 = \tau(D_\parallel^* D_\parallel) + \tau(D_\perp^* D_\perp).$$

Admissible internal recombinations preserving this norm form a $U(2)$ -type group on the two-component access space. Removing the central Abelian phase (already counted by the holonomy sector) leaves

$$SU(2).$$

Continuous mixing of the three primitive branches. The triple-overlap closure produces a three-dimensional branch space $\mathcal{H}_3 = \text{span}\{\mathcal{D}_0, \mathcal{D}_1, \mathcal{D}_2\}$. Admissible internal recombinations preserving the trace-quadratic branch norm act by $U(3)$. Removing the separately counted central $U(1)$ phase leaves

$$SU(3).$$

The corrected chain is: primitive triple-overlap obstruction \Rightarrow 3 branches $\Rightarrow U(3) \Rightarrow SU(3)$ after removing $U(1)$. The discrete closure $z^3 = 1$ supplies the primitive three-dimensional branch space; the continuous $SU(3)$ -like factor is the minimal trace-norm preserving non-Abelian mixing.

15.3 Minimal charged-sector selection principle

Principle 15.2 (Minimal charged-sector selection). The physical charged sector is the smallest nontrivial observer-stable internal automorphism sector that simultaneously satisfies:

1. a central Abelian phase readout;
2. minimal non-Abelian two-state access mixing;
3. minimal non-Abelian three-frame holonomy-index structure;
4. chiral doublet/singlet access splitting;
5. cancellation of global recoverability obstructions.

A pure $U(1)$ fails (2), (3), (4). $SU(2) \times U(1)$ fails (3). $SU(3) \times U(1)$ fails (2), (4). $SU(2) \times SU(2) \times U(1)$ fails (3). $SU(5)$ or $SU(4) \times SU(2) \times U(1)$ contain extra branches; they fail minimality.

The minimal product is therefore the Standard Model gauge group [49, 50, 51]

$$SU(3) \times SU(2) \times U(1).$$

Factor	Framework origin
Abelian charge	Central phase of primitive internal holonomy obstruction
Two-state access mixing	Cut/complement pair (D_\parallel, D_\perp) at one cut
Three-branch holonomy mixing	Primitive triple-overlap obstruction; trace-norm preserving mixing of three

Two constraints — chiral doublet/singlet access splitting and cancellation of global recoverability obstructions — are imported as structural requirements in the present treatment.

Central Abelian footprint rule. The central Abelian factor is not an optional exterior addition to the charged sector. It is the scalar phase reference common to the charged internal holonomy readout. Consequently, a non-Abelian charged boundary that is read through the full internal algebra carries a residual Abelian footprint whenever the observer cut does not isolate the non-Abelian sector from the central phase sector. Once the Abelian baseline has been fixed by the electromagnetic Schur-trace diagnostic, this residual footprint enters subsequent charged-sector diagnostics as the inverse of that baseline readout. In this sense the Abelian footprint is not a second electromagnetic calculation appended to a non-Abelian trace; it is the central-phase remnant of reading a charged non-Abelian boundary through the full minimal internal algebra.

15.4 Minimal chiral content and hypercharge fixing

Once the minimal internal algebra is $SU(3) \times SU(2) \times U(1)$, the minimal chiral content is selected by representation theory:

$$Q_L : (3, 2), \quad u_R : (3, 1), \quad d_R : (3, 1), \quad L_L : (1, 2), \quad e_R : (1, 1),$$

up to neutral singlets. Let Y_Q, Y_u, Y_d, Y_L, Y_e be unknown. The framework analogue of anomaly cancellation [54, 55, 56] reads:

$$\begin{aligned} SU(3)^2 U(1) : \quad & 2Y_Q - Y_u - Y_d = 0, \\ SU(2)^2 U(1) : \quad & 3Y_Q + Y_L = 0, \\ U(1)^3 : \quad & 6Y_Q^3 - 3Y_u^3 - 3Y_d^3 + 2Y_L^3 - Y_e^3 = 0. \end{aligned}$$

With $Q_{\text{em}} = T_3 + Y$ and normalization $Y_e = -1$, the minimal rational solution is

$$Y_Q = \frac{1}{6}, \quad Y_u = \frac{2}{3}, \quad Y_d = -\frac{1}{3}, \quad Y_L = -\frac{1}{2}, \quad Y_e = -1.$$

The one-generation Abelian trace is $T_{\text{gen}} = \text{Tr}(Y^2) = 6(\frac{1}{6})^2 + 3(\frac{2}{3})^2 + 3(-\frac{1}{3})^2 + 2(-\frac{1}{2})^2 + (-1)^2 = 10/3$. The one-generation weak trace is $\text{Tr}(T_3^2) = 3 \cdot \frac{1}{2} + 1 \cdot \frac{1}{2} = 2$.

15.5 Hypercharge normalization and conditional weak angle

The ratio is $\text{Tr}(Y^2)/\text{Tr}(T_3^2) = (10/3)/2 = 5/3$. In partition-strain language $\text{Tr}(Y^2)$ is the Abelian internal holonomy stiffness trace through K_{rec} and $\text{Tr}(T_3^2)$ is the weak doublet stiffness trace. Trace-normalized stiffness equality defines the canonically normalized Abelian generator

$$Y_{\text{can}} = \sqrt{\frac{3}{5}} Y, \quad \text{Tr}(Y_{\text{can}}^2) = \text{Tr}(T_3^2),$$

equivalently $g_1^2 = (5/3)g_Y^2$. This is the usual 5/3 hypercharge normalization, obtained as a trace-ratio consequence of the minimal one-generation charged content rather than as an $SU(5)$ [52, 53] embedding postulate.

If the canonically normalized primitive Abelian and weak sectors share the same leading paired stiffness coefficient through K_{rec} , $g_1 = g_2$ at the substrate normalization scale, then

$$\sin^2 \theta_W = \frac{g_Y^2}{g_2^2 + g_Y^2} = \frac{3/5}{1 + 3/5} = \frac{3}{8}.$$

This is the unification-scale value, not the observed low-energy weak angle. The 3/8 value is a structural consistency check on the trace logic, not a low-energy prediction.

15.6 Conditional identification with the Maxwell readout and the lepton ladder

Proposition 15.3 (Conditional charged-sector identification). *Assume the observer-effective charged sector is the minimal internal sector satisfying: internality, observer-record stability, relationality, central Abelian phase readout, entry into the regional paired disturbance functional through the same internal connection whose stationary variation defines the Abelian observer field readout, and use of its stable branches as the charged identity sectors for the record-preserving strain ladder. Then the primitive charged obstruction is Ω_{ijk} . Isotropic Abelian closure gives $z^3 = 1$, hence three stable charged branches. The same internal connection produces the Maxwell-type observer readout through stationarity of \mathcal{K}_R with the matched pair $(K_{\text{rec}}^{\text{em}}, M_{\text{rec}}^{\text{em}}) = g_{\text{eff}}^{-2}(K_{\text{rec}}, M_{\text{rec}})$, and the same three branches supply the charged strain ladder used in the lepton diagnostic.*

Proof sketch. The triple-overlap argument supplies the three branches. For the Maxwell readout, perturb the internal Abelian connection by A :

$$\mathcal{K}_R(\nabla + A) = \mathcal{K}_R(\nabla) + \frac{1}{4g_{\text{eff}}^2} \langle F_A, F_A \rangle_{\tau, R} + \langle A, J \rangle + O(A^3),$$

where $F_A = dA$. Stationarity $\delta_A \mathcal{K}_R = 0$ gives $\delta F_A = g_{\text{eff}}^2 J$. Combined with $dF = 0$:

$$dF = 0, \quad \delta F = g_{\text{eff}}^2 J.$$

The matched pair $(K_{\text{rec}}^{\text{em}}, M_{\text{rec}}^{\text{em}})$ ensures the electromagnetic record cone equals the universal c_* . The lepton ladder follows from using the three stable branches as charged identity sectors and the bi-Laplacian stiffness $E_k = k^4$, $S_n = \sum_{k=1}^n k^4$. \square

15.7 Schur-complement and access–recovery application to the charged sector

The Schur-complement reduction (Theorem 6.11 of Section 6.8) and the access–recovery composition theorem (Theorem 6.12 of Section 6.9) of the observer layer apply directly to the fine-structure calculation. At the observer junction, let x denote the scalar charged-readout amplitude and y the recoverable access-sector amplitudes. The framework’s paired $(K_{\text{rec}}, M_{\text{rec}})$ structure organizes the joint stiffness as the positive-definite coupled quadratic form $Q(x, y) = K_{\text{bare}} x^2 + 2xBy + \langle y, Cy \rangle$ of Section 6.8, with $K_{\text{bare}} > 0$ the undressed charged stiffness, $C > 0$ the access-sector stiffness (a restriction of K_{rec} to the recoverable access channel), and B the off-diagonal block. Theorem 6.11 fixes the effective scalar stiffness as $K_{\text{eff}} = K_{\text{bare}} - BC^{-1}B^\dagger$, with the access-dressing operator $BC^{-1}B^\dagger \succeq 0$ positive semidefinite.

Consequences for the fine-structure formula. Theorem 6.11 fixes three structural features of the diagnostic calculation independently of the specific operator that produces the access dressing.

First, the sign of every access-dressing correction is negative: charged record modes coupled to recoverable access modes always lower the observer-effective inverse stiffness, never raise it. The minus signs in $-3/(8\pi^4)$ and $-21/(32\pi^{12})$ are forced, not chosen.

Second, the access dressing is a single operator $BC^{-1}B^\dagger$, not a sum of independently-postulated corrections. The decomposition of this operator into a single-channel piece and a projective-cover piece is a spectral decomposition of one access-dressing object, not a postulate that several independent corrections exist.

Third, the mechanism is exact at the level of the coupled quadratic form. There is no perturbative truncation in the Schur step itself; truncation enters only when $BC^{-1}B^\dagger$ is expanded over the access-mode spectrum. The fine-structure diagnostic is therefore a calculation of the spectrum of one operator, not a perturbation series at arbitrary order.

Inter-sector Schur overlap in the charged sector. Section 6.8 also establishes the minimal inter-sector overlap correction, $\Delta K_{ij}^{\text{overlap}} = -1/(T_i T_j)$ of Equation (1), applicable whenever two internal holonomy sectors are simultaneously unresolved by one observer cut. The framework's neutral electroweak visibility readout of Section 15.10 uses this rule for the simultaneously unresolved $(SU(3), SU(2))$, $(SU(3), U(1))$, and $(SU(2), U(1))$ overlap pairs at the electroweak cut.

First single-channel access-dressing contribution. Theorem 6.12 of Section 6.9 identifies the primitive scalar access-dressing kernel as L_{acc}^{-2} , the squared inverse access Laplacian, with first inverse eigenvalue π^{-4} on the normalized Dirichlet access interval. Applied to the Schur-complement structure with $C = L_{\text{acc}}^2$ on the recoverable access channel and the off-diagonal coupling carrying the charged projector P_{ch} , the first single-channel contribution to $BC^{-1}B^\dagger$ is the charged-projector trace of L_{acc}^{-2} on the first access mode,

$$\Pi_{\text{access}} = \text{Tr}_{\text{acc}}(P_{\text{ch}} L_{\text{acc}}^{-2}) = \frac{3}{8\pi^4}.$$

The factor $3/8$ is the charged trace fraction derived from the minimal chiral content and $\text{Tr}(Y^2)/\text{Tr}(T_3^2) = 5/3$, and the factor π^{-4} is the first inverse eigenvalue of L_{acc}^{-2} on the normalized Dirichlet interval. This is the third term of the fine-structure formula, recovered as the leading single-channel scalar access dressing through the observer-layer Schur reduction.

The Schur-complement theorem and the access-recovery composition theorem together determine the first nonzero correction up to its sign and its leading π^{-4} order, and identify its operator origin. The remaining content of the diagnostic value at the single-channel order is the charged-projector trace, which is fixed by the substrate-side minimal-content analysis.

15.8 Projective access-cover operator and polarization-order expansion

The Schur-complement theorem fixes the sign and form of the access dressing, and the access-recovery composition theorem fixes the primitive scalar kernel L_{acc}^{-2} . The remaining task is to define the projective-cover part of the same dressing operator and compute its trace. This subsection does that explicitly.

Spatial projective cover. The projective closure acts on the three spatial access directions Σ_{obs}^3 , not on the full observer-access space O_{acc} . The record-ordering coordinate τ_{obs} sequences records, but it does not generate independent spatial projective sectors. Let

$$b := \dim \Sigma_{\text{obs}}^3, \quad d := \dim O_{\text{acc}}.$$

For a b -branch spatial cover, the projective record space is the direct sum over nonempty branch subsets,

$$\mathcal{H}_{\text{proj}}^{(b)} := \bigoplus_{\emptyset \neq S \subseteq \{1, \dots, b\}} \mathcal{H}_S.$$

The empty subset is removed because a projective readout has no zero-amplitude representative. The number of retained cover sectors is therefore $2^b - 1$.

Tensor access kernel. On each recoverable access branch, the primitive closed access-recovery kernel is $\Lambda_{\text{acc}} := L_{\text{acc}}^{-2}$. At the leading Dirichlet mode in an access dimension d , its inverse eigenvalue is π^{-d} . For a b -branch projective cover the closed kernel is the tensor product

$$\Lambda^{(b)} := \Lambda_{\text{acc},1} \otimes \cdots \otimes \Lambda_{\text{acc},b},$$

so the leading inverse eigenvalue is multiplicative:

$$\lambda_1^{(b)} = \pi^{-bd}.$$

At the minimal observer cut, $b = 3$ and $d = 4$, so the projective-cover kernel is suppressed by π^{-12} .

Projective normalization and charged projector. The projective-readout axiom quotients overall access amplitude. The remaining normalization is the average over the full observer-access volume, giving the factor $1/d$. Equivalently, define the normalized projective-cover propagator

$$G_{\text{proj}}^{(b)} := \frac{1}{d} \Lambda^{(b)}$$

on $\mathcal{H}_{\text{proj}}^{(b)}$. The charged part of the recoverable access channel is selected by the same charged projector used in the single-channel term,

$$P_{\text{ch}} = w_Y P_1, \quad w_Y = \frac{3}{8},$$

where w_Y is fixed by the minimal charged-content trace ratio $\text{Tr}(Y^2)/\text{Tr}(T_3^2) = 5/3$ at the shared substrate/junction normalization.

The projective-cover projector is

$$P_{\text{proj}}^{(b)} := P_{\text{ch}} \otimes \bigoplus_{\emptyset \neq S \subseteq \{1, \dots, b\}} I_S,$$

where I_S is the identity on the leading mode of the branch-subset sector \mathcal{H}_S . This makes the multiplicity explicit: the trace over the cover counts exactly the $2^b - 1$ nonempty projective sectors.

Theorem 15.4 (Projective access-cover trace). *For a b -branch observer projective cover in access dimension d , with primitive access-recovery kernel $\Lambda_{\text{acc}} = L_{\text{acc}}^{-2}$ and charged projector $P_{\text{ch}} = (3/8)P_1$, the leading projective-cover contribution to the Schur dressing is*

$$\Pi_{\text{proj}}^{(b)} = \text{Tr}_{\mathcal{H}_{\text{proj}}^{(b)}} \left(P_{\text{proj}}^{(b)} G_{\text{proj}}^{(b)} \right) = \frac{3}{8} \frac{2^b - 1}{d} \frac{1}{\pi^{bd}}.$$

In particular, the minimal memory-bearing projective cut has $(b, d) = (3, 4)$ and therefore

$$\Pi_{\text{proj}}^{(3)} = \frac{3}{8} \cdot \frac{7}{4} \cdot \frac{1}{\pi^{12}} = \frac{21}{32\pi^{12}}.$$

Proof. The leading mode of each closed access-recovery branch contributes π^{-d} . The tensor product over b independent spatial projective branches contributes π^{-bd} . The direct sum over nonempty projective branch subsets contributes the multiplicity $2^b - 1$. The charged projector contributes the trace weight $w_Y = 3/8$. The projective normalization divides by the observer-access dimension d . Multiplying these four independent factors gives

$$\text{Tr} \left(P_{\text{proj}}^{(b)} G_{\text{proj}}^{(b)} \right) = \frac{3}{8} \frac{2^b - 1}{d} \pi^{-bd}.$$

Setting $b = 3$ and $d = 4$ gives $21/(32\pi^{12})$. □

This result upgrades the projective-cover term from a structural decomposition to an operator trace. The numerator $21 = 3(2^3 - 1)$ is not chosen after the fact: the factor 3 is the charged trace weight numerator and $2^3 - 1$ is the number of nonempty projective cover sectors. The denominator $32 = 8 \cdot 4$ is likewise not an imposed rational: 8 is the denominator of the charged trace weight and 4 is the observer-access normalization. The power π^{12} is the tensor product of three closed access-recovery kernels over a four-dimensional observer-access interface.

Schur-complement realization. The single-channel and projective-cover terms are two spectral blocks of one positive Schur dressing operator. Decompose the recoverable access space as

$$\mathcal{Y}_{\text{acc}} = \mathcal{Y}_1 \oplus \mathcal{Y}_{\text{proj}},$$

with

$$C^{-1} = L_{\text{acc}}^{-2} \oplus G_{\text{proj}}^{(3)}, \quad B^\dagger B = P_{\text{ch}} \oplus P_{\text{proj}}^{(3)}.$$

Then the access dressing appearing in the Schur complement is

$$BC^{-1}B^\dagger = \text{Tr}_{\mathcal{Y}_1}(P_{\text{ch}}L_{\text{acc}}^{-2}) + \text{Tr}_{\mathcal{Y}_{\text{proj}}}(P_{\text{proj}}^{(3)}G_{\text{proj}}^{(3)}),$$

namely

$$BC^{-1}B^\dagger = \frac{3}{8\pi^4} + \frac{21}{32\pi^{12}}.$$

Substitution into the exact Schur complement gives the observer-effective charged stiffness

$$K_{\text{eff}} = K_{\text{bare}} - \frac{3}{8\pi^4} - \frac{21}{32\pi^{12}}.$$

The two negative corrections therefore arise from one positive access-dressing operator, not from independently appended numerical terms.

Termination of the polarization-order expansion. The polarization expansion is finite under the minimal projective-cover rule. The $b = 1$ block is the ordinary single-channel access polarization,

$$\Pi^{(1)} = \frac{3}{8\pi^d},$$

which at $d = 4$ gives $3/(8\pi^4)$. The $b = 2$ block carries only a removable relative phase; after the projective quotient it has no non-removable holonomy and contributes no primitive charged correction. The $b = 3$ block is the first non-removable projective closure and gives $21/(32\pi^{12})$. Covers with $b \geq 4$ decompose into triple-overlap pieces and do not define new primitive charged sectors. Thus the fine-structure access dressing terminates at the minimal triple cover:

$$\sum_b \Pi^{(b)} = \frac{3}{8\pi^4} + \frac{21}{32\pi^{12}}.$$

15.9 Fine-structure diagnostic as a single Schur-trace identity

The preceding subsections established the charged-sector machinery used by the Abelian readout: the charged-holonomy selection principle, the connection-block representation lemma, the central Abelian footprint rule, the Schur-complement theorem (Section 6.8), the access-recovery composition theorem (Section 6.9), and the projective access-cover theorem (Section 15.8). This subsection assembles those ingredients into the fine-structure diagnostic in a form sharper than a four-term sum. The four numerical pieces are presented as the four exposed components of a single Schur-complement trace on one substrate-cut block operator. The unit of derivational claim is one operator, not four addends.

The framework-native meaning of α inside the two-layer organization should be recalled briefly. Within the substrate-observer language, the fine-structure constant is the readout ratio between the minimal Abelian charged-holonomy unit visible to an admissible observer cut and the isotropic record-cone access capacity supplied by the paired baseline susceptibility, and α^{-1} counts the number of isotropic record-access units required to resolve one central phase unit at the observer junction. In this role, the Abelian fine-structure diagnostic also fixes the baseline central phase footprint used by later charged-sector readouts: non-Abelian holonomy boundaries

do not replace the $U(1)$ reference; they are read through the same full charged-sector algebra and inherit its scalar Abelian baseline. The diagnostic question is then whether this central Abelian readout ratio can be computed from one operator on the substrate–cut Hilbert space. The intended chain is

$$\boxed{\delta S_{\text{sub}} = 0 \implies (K_{\text{rec}}, M_{\text{rec}}) \implies L_{\text{acc}} \implies \mathcal{H}_{\text{cut}} \implies \text{Tr Schur}_C(\mathcal{A}_{\text{sub/cut}}) = \alpha_{\text{eff}}^{-1}} \quad (5)$$

The first arrow is a proposal: a substrate variational principle is exhibited below in a form general enough to admit the required normal form, but the specific choice of two scalar functionals inside that form is left open. The second arrow is conditional on the boundary-condition choice for the access operator, which is adopted and flagged rather than derived. The third arrow is constructive: the cut Hilbert space and the block operator on it are written out explicitly. The fourth arrow is the proposition of this subsection. Its content is that, once the block operator is fixed, the four numerical pieces of the diagnostic are forced.

Substrate variational principle: proposal, not theorem. The timeless substrate is not specified by a metric, a background time, or a fixed spacetime manifold. It is specified by a positive record susceptibility M_{rec} , a positive record stiffness K_{rec} , and a locality–recoverability constraint. It is proposed that these objects are selected by stationarity of a substrate functional of the form

$$S_{\text{sub}}[K, M; \Lambda] = \tau \left[KM^{-1} + \Phi(M) + \Psi(M^{-1/2}KM^{-1/2}) + \Lambda C_{\text{loc}}(K, M) \right], \quad (6)$$

on the cone $\{K > 0, M > 0\}$, with Φ and Ψ scalar trace functionals to be specified, Λ a self-adjoint Lagrange multiplier, and C_{loc} encoding finite-access locality together with refinement covariance of the substrate–cut graph. Admissibility requires positivity of the stationary pair, locality of the normalized operator $M_{\text{rec}}^{-1/2}K_{\text{rec}}M_{\text{rec}}^{-1/2}$ on the access graph, and invariance of its spectrum under refinement of the graph representative.

The role of (6) in this paper is structural. It is not claimed that an explicit pair (Φ, Ψ) has been identified for which the stationary record pair satisfies the metric-compatible normal form

$$K_{\text{rec}} = M_{\text{rec}}^{1/2} L_{\text{acc}} M_{\text{rec}}^{1/2}, \quad (7)$$

with L_{acc} a positive self-adjoint access operator. What is claimed is the weaker statement that (7) is a natural normal form to be selected by such a principle: once M_{rec} is used to define the record inner product, the only refinement-covariant positive stiffness with the same generalized spectrum as L_{acc} is conjugate to L_{acc} by $M_{\text{rec}}^{1/2}$. Specifying Φ and Ψ so that (7) emerges as the unique stationary point on the admissible cone is the central substrate-side open problem of the framework, and is identified as such in the closing status statement of this chapter.

Given (7) as a working hypothesis, the generalized record-cone diagnostic is $c_*^2 = \lambda_{\min}(M_{\text{rec}}^{-1}K_{\text{rec}}) = \lambda_{\min}(L_{\text{acc}})$, so the access cone is determined by the spectrum of the substrate–cut operator L_{acc} rather than by a background spacetime speed.

Access spectrum from the stable cut. The observer cut is modeled as a finite-access boundary between the timeless substrate and a recordable sector. After normalization of the access measure the primitive access coordinate is $u \in [0, 1]$. The framework needs two endpoint conditions on L_{acc} , one on each side of this interval. The observer-side endpoint at $u = 1$ is the natural one: record closure means that an access channel returns a determinate recorded value, which fixes the boundary mode to zero. The substrate-side endpoint at $u = 0$ is the asymmetric one. The choice made here is that both endpoints absorb,

$$L_{\text{acc}} = -\frac{d^2}{du^2}, \quad f(0) = f(1) = 0, \quad (8)$$

with the substrate-side absorbing condition interpreted as the statement that an unrecorded variation is not yet accessible to the observer cut and therefore contributes no boundary amplitude. This is the same Dirichlet problem used throughout the framework, and it is reproduced here without claiming that the substrate-side absorbing condition has been derived from substrate primitives. A reflecting (Neumann) substrate-side endpoint would shift every π -power below; this dependence is flagged in the chapter's closing status statement.

Under (8) the spectrum and eigenfunctions are $L_{\text{acc}}e_n = (n\pi)^2e_n$ with $e_n(u) = \sin(n\pi u)$, normalized so that the access measure $d\mu_{\text{acc}}(u) = 2 du$ on $[0, 1]$ gives $\int_0^1 e_n(u)^2 d\mu_{\text{acc}} = 1$. This convention is adopted at the outset, so that the first-mode access amplitude $a_1 = \int_0^1 e_1(u) d\mu_{\text{acc}} = 2 \int_0^1 \sin(\pi u) du = 2/\pi$ emerges without an after-the-fact rescaling. The first-mode contribution to the bare junction stiffness is therefore

$$4\pi a_1^2 = 4\pi \frac{4}{\pi^2} = \frac{16}{\pi}. \quad (9)$$

The appearance of π is a spectral residue of the unique normalized absorbing access interval, not an inserted convention.

Closed access-loop recoverability: substrate reading of imported constraints. The charged sector is not derived from the substrate alone in this framework, and the status is stated explicitly. The cut/complement argument of Section 15.2 supplies the two-state access mixing structure $SU(2)$, but it does not supply the chiral split between left- and right-access sectors; chirality is asserted here as a primitive observer-cut constraint. The cubic-trace conditions

$$\sum_i Y_i = 0, \quad \sum_i Y_i^3 = 0, \quad \sum_{\text{weak}} Y_i = 0, \quad \sum_{\text{color}} Y_i = 0, \quad (10)$$

are imported from the Standard Model anomaly-cancellation pattern. The substrate reading of (10) is that a charged record transported around a closed access loop must return as a single-valued record, so each closed loop carries a recoverability obstruction that must vanish; the equations are then the obstruction-vanishing conditions for the minimal loops. This reading is interpretive. Deriving both the form and the multiplicities of (10) from substrate primitives remains open and is listed in the chapter's status statement.

Under the imported constraints, the minimal chiral solution with nontrivial color multiplicity, weak pairing, and one access-stabilizing scalar doublet has hypercharges $(Y_Q, Y_u, Y_d, Y_L, Y_e) = (1/6, 2/3, -1/3, -1/2, -1)$ and $|Y_H| = 1/2$, with per-generation charge-square trace $T_{\text{gen}} = 10/3$. The three-generation fermionic charge-square trace and the scalar contribution combine to

$$T_{\text{bare}} = 3T_{\text{gen}} + T_H = 10 + \frac{1}{2} = \frac{21}{2}. \quad (11)$$

The trace ratio $\text{Tr}(Y^2)/\text{Tr}(T_3^2) = 5/3$ then gives the charged access weight

$$w_Y = \frac{g_Y^2}{g_2^2 + g_Y^2} = \frac{3}{8}, \quad (12)$$

where the second equality uses $g_1 = g_2$ at the shared junction trace-normalization, itself a normal-form assertion that is flagged in the chapter's status statement rather than derived here.

The same number $T_{\text{bare}} = 21/2$ also arises from an observer-side duality independent of any specific substrate content. The substrate-observer duality

$$d \cdot T_{\text{bare}} = \binom{d}{2}(2^b - 1) \implies T_{\text{bare}} = \frac{d-1}{2}(2^b - 1) \quad (13)$$

gives $T_{\text{bare}} = (3/2) \cdot 7 = 21/2$ at the minimal memory-bearing projective cut $(d, b) = (4, 3)$. The two routes (substrate-content counting on (11) and observer-side duality, treated in detail

in Section 6.10) are consistent. Consistency is recorded as an internal compatibility condition between substrate realization and observer projective access, not as a proof of either route; the substrate-content route inherits whatever fragility chirality and the cubic-trace constraints carry.

The corresponding bare junction inverse coupling is

$$\alpha_{\text{junction}}^{-1} = 4\pi T_{\text{bare}} + 4\pi a_1^2 = 42\pi + \frac{16}{\pi} = 137.0398496297\dots \quad (14)$$

Primitive access kernel: definitional minimality. The access dressing is obtained by Schur reduction of the coupled observer-junction quadratic form (Section 6.8), with effective scalar stiffness $K_{\text{eff}} = K_{\text{bare}} - BC^{-1}B^\dagger$. The Schur-complement theorem fixes the sign of every access correction: access dressing lowers the inverse effective coupling, never raises it.

The specific kernel that realizes C^{-1} on the Dirichlet access interval was taken in Section 6.9 to be $G_{\text{acc}} = L_{\text{acc}}^{-2}$. The kernel is fixed here by access–recovery composition together with a definitional minimality condition. Call an access-dressing kernel G *primitive* if (i) G is positive self-adjoint, (ii) the trace of G against the charged projector on the leading mode is finite, (iii) G is compatible with the absorbing endpoint conditions of (8), and (iv) G contains no extraneous higher-order smoothing beyond what is needed for finiteness of the leading trace. Under this definition, the first primitive kernel that meets (i)–(iv) for the access operator (8) is L_{acc}^{-2} . The kernel L_{acc}^{-1} is endpoint-sensitive in a sense that does not define a finite charged trace against the leading mode without auxiliary regularization; the informal statement is that $L_{\text{acc}}^{-1}\delta_u$ produces a logarithmically singular response at the boundary, but a fully rigorous statement of the exclusion of L_{acc}^{-1} as a primitive kernel on the Dirichlet interval is not given here. Kernels L_{acc}^{-p} with integer $p \geq 3$ are admissible under (i)–(iii) but violate (iv): they add extra smoothing that is not forced by the structure of the substrate–cut junction. The primitiveness condition is therefore selecting L_{acc}^{-2} as the minimal kernel meeting the first three conditions; the condition (iv) is doing the selection work and is stated as a definition rather than as a theorem.

On the normalized Dirichlet interval the leading inverse eigenvalue of L_{acc}^{-2} is π^{-4} , and the charged single-channel Schur dressing is

$$\Pi_{\text{access}} = w_Y \pi^{-4} = \frac{3}{8\pi^4}. \quad (15)$$

The higher-mode tail is bounded by $\sum_{n \geq 2} (n\pi)^{-4} = \pi^{-4}(\zeta(4) - 1) \approx 8.3 \times 10^{-4}$, so the first-mode truncation has relative error $\zeta(4) - 1 \approx 0.082$ at the level of the single-channel trace. This truncation is therefore not a 10^{-9} statement on its own; the 10^{-9} -level agreement of the full four-part diagnostic with experiment is a property of the four contributions taken together, and the closing comparison records what this combination does and does not test.

Projective-cover trace at the minimal observer cut. The minimal memory-bearing projective cut has spatial branching multiplicity $b = 3$ and access dimension $d = 4$. The projective record space at this cut is the direct sum over nonempty branch subsets, $\mathcal{H}_{\text{proj}}^{(3)} = \bigoplus_{\emptyset \neq S \subseteq \{1,2,3\}} \mathcal{H}_S$, with $2^3 - 1 = 7$ retained sectors. On the triple cover the closed access–recovery kernel is $\Lambda^{(3)} = \Lambda_{\text{acc}} \otimes \Lambda_{\text{acc}} \otimes \Lambda_{\text{acc}}$ with leading inverse eigenvalue $\pi^{-bd} = \pi^{-12}$. The projective normalization carries the inverse observer-access dimension $1/d = 1/4$ together with the inverse projective-sheet count $1/2^b = 1/8$, giving the propagator normalization $G_{\text{proj}}^{(3)} = (1/32)\Lambda^{(3)}$, and the projective-cover dressing is

$$\Pi_{\text{proj}}^{(3)} = w_Y \cdot \frac{2^b - 1}{2^b \cdot d} \cdot \pi^{-bd} = \frac{3}{8} \cdot \frac{7}{32} \cdot \frac{1}{\pi^{12}} = \frac{21}{32\pi^{12}}. \quad (16)$$

The $b = 2$ block carries only a removable relative phase and does not contribute a primitive non-removable obstruction; covers with $b \geq 4$ decompose into triple-overlap pieces and do not generate new primitive sectors. The polarization expansion in b therefore terminates at $b = 3$, and the full access dressing is the two-term sum $\sum_b \Pi^{(b)} = 3/(8\pi^4) + 21/(32\pi^{12})$.

The single substrate–cut operator. The operator whose Schur complement reproduces the four-part diagnostic is written explicitly. Let $\mathcal{H}_{\text{cut}} = \mathcal{H}_{\text{ch}} \oplus \mathcal{H}_{\text{acc}} \oplus \mathcal{H}_{\text{proj}}$, with \mathcal{H}_{ch} one-dimensional (the recorded charged amplitude), \mathcal{H}_{acc} the recoverable access channel restricted to its leading Dirichlet mode, and $\mathcal{H}_{\text{proj}}$ the projective triple cover at $(b, d) = (3, 4)$. On \mathcal{H}_{cut} define the block operator

$$\mathcal{A}_{\text{sub/cut}} = \begin{pmatrix} 4\pi T_{\text{bare}} + 4\pi a_1^2 & B_Y & B_{\text{proj}} \\ B_Y^\dagger & C_Y & 0 \\ B_{\text{proj}}^\dagger & 0 & C_{\text{proj}} \end{pmatrix}, \quad (17)$$

with diagonal access-sector blocks $C_Y = L_{\text{acc}}^2|_{\mathcal{H}_{\text{acc}}}$ and $C_{\text{proj}} = (\Lambda^{(3)})^{-1}|_{\mathcal{H}_{\text{proj}}}$, and off-diagonal couplings B_Y, B_{proj} chosen so that

$$B_Y C_Y^{-1} B_Y^\dagger = w_Y \pi^{-4} = \frac{3}{8\pi^4}, \quad B_{\text{proj}} C_{\text{proj}}^{-1} B_{\text{proj}}^\dagger = \frac{21}{32\pi^{12}}. \quad (18)$$

The choice of B_Y and B_{proj} at this stage is dictated by the single-channel and projective-cover trace calculations of (15) and (16), not by independent postulate; what is asserted is that the same coupling structure appears in both blocks, and that the substrate–cut operator therefore has one off-diagonal B with two spectral channels rather than two unrelated couplings.

The Schur complement of $\mathcal{A}_{\text{sub/cut}}$ against the unrecorded access sectors $C_Y \oplus C_{\text{proj}}$ is

$$\begin{aligned} \text{Schur}_{C_Y \oplus C_{\text{proj}}}(\mathcal{A}_{\text{sub/cut}}) &= 4\pi T_{\text{bare}} + 4\pi a_1^2 - B_Y C_Y^{-1} B_Y^\dagger - B_{\text{proj}} C_{\text{proj}}^{-1} B_{\text{proj}}^\dagger \\ &= 42\pi + \frac{16}{\pi} - \frac{3}{8\pi^4} - \frac{21}{32\pi^{12}}. \end{aligned} \quad (19)$$

The recorded sector is one-dimensional, so the Schur complement is a scalar and its trace is itself. The single identity is therefore

$$\boxed{\alpha_{\text{eff}}^{-1} = \text{Tr Schur}_{C_Y \oplus C_{\text{proj}}}(\mathcal{A}_{\text{sub/cut}}) = 42\pi + \frac{16}{\pi} - \frac{3}{8\pi^4} - \frac{21}{32\pi^{12}}.} \quad (20)$$

The four visible terms in α_{eff}^{-1} are therefore not independent choices: they are the four exposed components of one Schur trace. The bare-trace term 42π and the first-access term $16/\pi$ are the diagonal contributions of \mathcal{H}_{ch} ; the dressings $-3/(8\pi^4)$ and $-21/(32\pi^{12})$ are the two spectral channels of $-BC^{-1}B^\dagger$ on \mathcal{H}_{acc} and $\mathcal{H}_{\text{proj}}$. The Schur-complement theorem fixes both signs to be negative.

Numerical comparison and the CODATA 2022 anchor. Evaluating (20) numerically gives

$$\alpha_{\text{eff}}^{-1} = 137.035\,999\,176\,3\dots \quad (21)$$

The current CODATA 2022 adjusted value of the inverse fine-structure constant is

$$\alpha_{\text{CODATA 2022}}^{-1} = 137.035\,999\,177(21), \quad (22)$$

with quoted one-sigma standard uncertainty $\sigma_{\text{CODATA 2022}} = 2.1 \times 10^{-8}$ [58]. The residual is

$$\Delta_{\text{CODATA 2022}} = \alpha_{\text{eff}}^{-1} - \alpha_{\text{CODATA 2022}}^{-1} \approx -7 \times 10^{-10}, \quad (23)$$

which is approximately 0.03σ and is therefore comfortably inside the recommended uncertainty.

This comparison is anchored on the CODATA 2022 adjustment. The CODATA 2018 value $\alpha_{\text{CODATA 2018}}^{-1} = 137.035\,999\,084(21)$ [57] is superseded by the 2022 adjustment, and the comparison anchor used here is recorded explicitly so the reader can reproduce the calculation.

The CODATA 2022 adjustment is itself a weighted reconciliation of two direct atom-recoil measurements that are in roughly 5σ tension with one another. The Parker–Yu–Zhong–Estey–Müller measurement using a caesium recoil in 2018 gives $\alpha^{-1} = 137.035\,999\,046(27)$ [60]; the

Morel–Yao–Cladé–Guellati–Khélifa measurement using a rubidium recoil in 2020 gives $\alpha^{-1} = 137.035\,999\,206(11)$ [59]. An independent precision input to the CODATA chain is the electron magnetic-moment measurement of Fan, Myers, Sukra, and Gabrielse [61], which, combined with the QED theory calculation of the electron anomalous magnetic moment [82], supplies a third determination of α in approximate consistency with the rubidium-recoil result. Against Parker the present formula sits at $\sim 4.8\sigma$; against Morel it sits at $\sim 2.7\sigma$; against the CODATA 2022 midpoint it sits at $\sim 0.03\sigma$. The diagnostic is therefore consistent with the recommended adjustment but is not simultaneously consistent with both direct measurements, in the same sense that any prediction landing near the CODATA midpoint inherits the Rb/Cs tension. This is recorded as a property of the comparison, not as a feature of the formula: future adjustments that resolve the Rb/Cs tension will sharpen the test of (20) in one direction or the other.

What this subsection closes and what it does not. The contribution of this subsection is the single Schur-trace identity (20). What it closes is the unit of derivational claim. The four-part value of the fine-structure diagnostic is the trace of a single Schur complement on the substrate–cut block operator (17), with the Schur-complement theorem fixing the signs of the two access dressings and the spectrum of one operator fixing their values. It also fixes the Abelian baseline used by later charged-sector readouts: the scalar $U(1)$ footprint is not recalculated in each non-Abelian sector but inherited from this central phase Schur trace.

What this subsection does not close is the substrate uniqueness of that operator’s blocks. The calculation is machinery-fixed once the charged-sector selection, the access-boundary condition, the primitive access kernel, the chiral/cubic trace constraints, and the shared junction normalization are accepted; what remains open is the substrate-level uniqueness of those inputs. Five items remain open and are listed explicitly so the status of the diagnostic is unambiguous. First, the substrate variational principle (6) is exhibited as a form rather than a specification; the scalar functionals Φ and Ψ that would produce (7) as the unique stationary point on the admissible cone have not been identified. Second, the Dirichlet boundary condition on the access interval is adopted rather than derived; the substrate-side absorbing condition at $u = 0$ is the asymmetric choice. Third, chirality is asserted as a primitive observer-cut constraint, and the cubic-trace conditions (10) are imported from the Standard Model anomaly-cancellation pattern with a substrate reading rather than a substrate derivation. Fourth, the primitiveness condition selecting L_{acc}^{-2} as the leading access kernel is definitional; the exclusion of L_{acc}^{-1} on the Dirichlet interval is stated informally and would benefit from a precise endpoint-regularity argument. Fifth, the equality $g_1 = g_2$ at the shared junction trace-normalization is asserted as a normal-form condition under canonical trace normalization of primitive internal generators, rather than as a theorem from substrate content. These five items are the substantive remaining program for the substrate–cut foundation of the diagnostic; they are collected in the chapter’s status table together with the other items not recovered.

The status of the section, in compact form, is therefore the following:

Numerical match to CODATA 2022: within quoted uncertainty. Structural compression: four terms unified into one Schur trace. Abelian baseline footprint fixed for later charged-sector readouts. Substrate uniqueness of the block operator: not yet established.	(24)
---	------

The framework’s claim is therefore that, within the stated charged-sector and access-normalization inputs, the Abelian fine-structure readout is fixed by the Schur complement of a single block operator on the substrate–cut Hilbert space. The remaining question is not whether the four terms are independently adjustable inside this diagnostic; they are not. The remaining question is whether the substrate machinery uniquely forces the block operator and boundary data from which the diagnostic is built.

15.10 Neutral electroweak visibility and the SLD/LEP- b channel split

The neutral electroweak diagnostic is a visibility readout of an observer-access boundary. It is not an inverse-stiffness readout. Inverse couplings are obtained from effective stiffness after hidden-sector elimination; a weak mixing angle is instead a scalar share of the neutral boundary visible through the central Abelian phase. The object being read is therefore the neutral $SU(2)$ – $U(1)$ access boundary, and the result depends on the observer support through which that boundary is resolved.

The Abelian trace entering this readout is fixed upstream by the fine-structure Schur-trace diagnostic of Section 15.9. In the notation of that subsection,

$$T_1 := \alpha_{\text{eff}}^{-1} = 137.035999176\dots \quad (25)$$

This is the central Abelian phase footprint carried forward into charged-boundary readouts. It is not replaced by a separate electroweak input in this subsection. The same primitive Abelian footprint is read through different neutral observer-access supports.

Neutral-boundary visibility lemma. For a neutral electroweak observer cut with $SU(2)$ trace weight T_2 and central Abelian trace weight T_1 , the effective weak-mixing readout is the scalar visibility of the central Abelian phase inside the jointly accessible neutral boundary. Trace normalization assigns reciprocal visibility weights to the direct sector channels. The minimal unresolved $SU(2)$ – $U(1)$ boundary contact contributes the product of the two reciprocal weights. Thus the direct leptonic or polarized neutral-current readout is

$$\sin^2 \theta_{\text{eff},\ell} = \frac{1}{T_2} + \frac{1}{T_1} + \frac{1}{T_2 T_1}. \quad (26)$$

The first term is the direct $SU(2)$ neutral visibility, the second term is the direct central Abelian footprint, and the third term is the minimal shared neutral-boundary visibility.

The $SU(2)$ trace weight is also fixed by the framework bookkeeping. The neutral connection block is an adjoint holonomy block. For $SU(2)$ the adjoint internal dimension is 3 and the adjoint Casimir weight is 2. With the observer projective multiplicity $b = 3$, the trace weight is

$$T_2 = 3 \cdot \frac{3}{2} = \frac{9}{2}. \quad (27)$$

This value is also recovered by the closure-normalized observation-support theorem of Appendix 6.11 (Theorem 6.16) with branch count $d_{\text{supp}} = 2$ for the weak-isospin doublet observation, giving $T_2 = b^2 d_{\text{supp}}/d = 9 \cdot 2/4 = 9/2$ on the substrate $(b, d) = (3, 4)$. The framework reads T_2 through either route consistently. Combining (25) and (27) in the neutral-boundary visibility lemma gives

$$\sin^2 \theta_{\text{eff},\ell} = \frac{2}{9} + \frac{1}{137.035999176\dots} + \frac{1}{(9/2)(137.035999176\dots)} = 0.2311412087\dots \quad (28)$$

This is the leptonic or polarized neutral-boundary readout. More precisely, it is the framework's value for the direct polarized leptonic neutral-current readout, the observer-access cut in which both the initial-state and the final-state supports are purely leptonic and no color-carrying support participates in the access path at any vertex. The polarized e^+e^- measurement (SLD A_{LR}) realizes this access cut. Other channel classes correspond to distinct observer-access cuts: the b -quark forward-backward readout (below) reads the same neutral boundary through a color-carrying final-state support; the hadron-collider Drell-Yan readout (Section 15.10) reads the same boundary through a color-carrying initial-state support. The framework's prediction for each channel class is determined by which supports participate in the access path; equation (28) is specifically the polarized-leptonic case.

Sign convention. The sign rule follows from the type of readout. Schur elimination lowers an inverse-stiffness readout, so unresolved hidden overlap enters inverse-coupling diagnostics with a negative sign. Visibility readouts count accessible boundary channels, so shared boundary contact enters with a positive sign. Thus Section 15.11 subtracts unresolved electroweak overlap in an inverse-stiffness diagnostic, while the neutral-boundary visibility readout in (26) adds the shared $SU(2)$ – $U(1)$ visibility:

$$\begin{aligned} \text{inverse stiffness readout:} & \quad \text{hidden overlap subtracts,} \\ \text{visibility/share readout:} & \quad \text{accessible overlap adds.} \end{aligned} \tag{29}$$

Connection blocks and observation-channel supports. The representation rule is fixed by what the diagnostic measures. If the readout measures gauge stiffness or internal holonomy response, the hidden block is Lie-algebra-valued and transforms in the adjoint representation. If the readout measures an observer channel passing through a matter support, the relevant block is the representation carried by that support:

$$\begin{aligned} \text{connection-stiffness readout:} & \quad \text{use adjoint blocks,} \\ \text{observation-channel readout:} & \quad \text{use the support representation.} \end{aligned} \tag{30}$$

This is why the strong coupling diagnostic uses the $SU(3)$ adjoint holonomy block, while the b -quark final-state channel uses the $SU(3)$ fundamental color-support block.

The SLD left–right asymmetry and the LEP b -quark forward–backward asymmetry need not correspond to the same observer-access cut. The SLD left–right asymmetry is a leptonic neutral-boundary readout. The LEP b -quark forward–backward asymmetry reads the same neutral boundary through a color-carrying final-state support. In the framework these are two observer-access readouts of the neutral electroweak boundary through different supports. A leptonic or polarized readout sees the neutral boundary directly. A b -quark forward–backward readout sees it through a fundamental color-carrying support, producing an additional final-state boundary contribution.

For the fundamental representation of $SU(3)$, the internal dimension is 3 and the quadratic Casimir is $4/3$. With the same observer projective multiplicity $b = 3$, the fundamental color-support trace weight is

$$T_{b,\text{col}} = 3 \cdot \frac{3}{4/3} = \frac{27}{4}. \tag{31}$$

This is the trace weight for *any* fundamental-representation $SU(3)$ color-carrying support read through the $b = 3$ observer projective cut; the framework does not distinguish heavy-quark from light-quark contributions at this structural level. The match to the LEP A_{FB}^b measurement is to the specific experimentally tagged final state of that measurement, not a framework-side flavor selection. The value is also recovered by the closure-normalized observation-support theorem of Appendix 6.11 (Theorem 6.16) with branch count $d_{\text{supp}} = 3$, giving $T_{b,\text{col}} = b^2 d_{\text{supp}}/d = 27/4$ on $(b, d) = (3, 4)$. The framework’s $4/3$ closure unit coincides numerically with the $SU(3)$ fundamental quadratic Casimir; the substrate origin is primary, with the gauge-theoretic re-expression as a secondary identification valid for gauge-charged supports but not generally.

The derivation of δ_b uses no fermion mass input. The channel split arises from color-support trace structure, not from a mass hierarchy: $T_{b,\text{col}}$ is fixed by the fundamental color-support branch count on the $(b, d) = (3, 4)$ substrate, T_1 is fixed by the fine-structure Schur identity (Section 15.9), and the product $T_{b,\text{col}} T_1$ is mass-independent. The framework does not encode the top-Yukawa coupling, the $Zb\bar{b}$ vertex correction, or any fermion mass ratio in this derivation.

Because the b -quark final state is electrically charged, the color-support channel is read through the central Abelian phase footprint. The minimal final-state boundary contribution is the reciprocal overlap of the fundamental color-support trace and the Abelian footprint:

$$\delta_b = \frac{1}{T_{b,\text{col}} T_1} = \frac{1}{(27/4)(137.035999176\dots)} = 0.0010810893\dots \tag{32}$$

The b -channel effective readout is therefore

$$\sin^2 \theta_{\text{eff},b} = \sin^2 \theta_{\text{eff},\ell} + \delta_b = 0.2322222980\dots \quad (33)$$

The channel split is consequently carried by the access support: direct neutral visibility for the leptonic channel, and neutral visibility through a fundamental color-support boundary for the b -quark channel.

Numerical channel comparison. Using the historical channel values quoted in precision electroweak summaries, the two framework readouts sit inside the quoted one-sigma intervals for the two separate channels [62]:

Channel	Framework	Experiment	1σ band	Deviation
Lepton channel (SLD A_{LR})	0.23114	0.23098 ± 0.00026	[0.23072, 0.23124]	$+0.62\sigma$
b -quark channel (LEP A_{FB}^b)	0.23222	0.23221 ± 0.00029	[0.23192, 0.23250]	$+0.04\sigma$
Splitting	0.00108	0.00123 ± 0.00039	[0.00084, 0.00162]	-0.38σ

Conditional on the neutral-boundary visibility lemma, the sign convention, and the connection/support representation rule, the SLD/LEP- b split is a native observer-access feature of the framework: the same neutral electroweak boundary is resolved through two different supports, and those supports carry different visibility contributions.

Hadron-collider Drell-Yan readouts as a distinct observer-access cut. Recent precision measurements of the leptonic effective weak-mixing angle at hadron colliders, primarily Drell-Yan dilepton production at the Tevatron and the LHC, provide additional inputs to the precision-electroweak landscape. Representative values include CMS at 8 TeV (0.23101 ± 0.00053), LHCb at 13 TeV (0.23147 ± 0.00050), and CMS at 13 TeV with the largest current Drell-Yan dataset (0.23157 ± 0.00031); a Bodek-Seo-Yang reanalysis with profiled PDFs has reported 0.23156 ± 0.00024 , currently the most precise single determination. The framework's leptonic-channel prediction $\sin^2 \theta_{\text{eff},\ell} = 0.23114$ is within 1σ of the CMS 8 TeV and LHCb 13 TeV measurements, but sits at approximately $+1.4\sigma$ and $+1.7\sigma$ from the two CMS 13 TeV results.

This is not a tension with the leptonic visibility lemma as stated. Hadron-collider Drell-Yan production has a color-carrying initial state ($q\bar{q}$ from the proton parton distributions) and a leptonic final state. The observer-access cut therefore differs from the polarized e^+e^- case (where no color support enters the path) and from the LEP b -quark case (where the color support is in the final state). The relevant access geometry has color in the initial state and leptons in the final state. Under the connection/support representation rule (30), this is a third channel class with its own framework prediction.

The Drell-Yan retrieval support R_3 . The first ingredient of the Drell-Yan readout is a definite operator-algebraic construction of the relevant retrieval support, traceable to existing framework primitives. Starting from the minimal observer cut $\Sigma_{\text{obs}} = S_{\text{obs}} \oplus \tau_{\text{obs}}$ with $\dim S_{\text{obs}} = 3$ and $\dim \Sigma_{\text{obs}} = 4$ (Section 6.6), and from the paired record bookkeeping ($K_{\text{rec}}, M_{\text{rec}}$) controlling the generalized record operator $G_{\text{acc}} = M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2}$, the spatial-access component of the recoverable structure carries three primitive projective branches induced by the projective cover of S_{obs} . Denote the corresponding minimal stable spectral support projections of $G_{\text{acc}}|_{S_{\text{obs}}}$ by e_1, e_2, e_3 , with $e_i e_j = 0$ for $i \neq j$ and $e_1 + e_2 + e_3 = \mathbf{1}_{R_3}$. The finite trace algebra they generate is $R_3 := \text{Alg}\{e_1, e_2, e_3\}$. This is the retrieval-support algebra introduced in Section 4.2. It is not identical to S_{obs} (the geometric/projective access layer) and not identical to the seven-element projective cover of Section 15.8 (used in fine-structure dressing).

Conditional theorem. Conditional on three operator-algebra hypotheses about the framework's existing structures, the Drell-Yan channel weight is fixed:

Theorem 15.5 (Drell-Yan projection theorem, conditional form). *Let $R_3 = \text{Alg}\{e_1, e_2, e_3\}$ be the primitive retrieval-support algebra defined above. Assume:*

- (H1) Equal-trace condition. *The three primitive support projections satisfy $\tau(e_1) = \tau(e_2) = \tau(e_3)$.*
- (H2) Symmetric hadronic averaging. *The unresolved hadronic-initial-state averaging map H is the trace-preserving conditional expectation of R_3 onto its central scalar subalgebra, acting symmetrically on the three primitive support projections.*
- (H3) Rank-one leptonic projection. *The clean leptonic neutral-current final-state conditional expectation E_ℓ , composed with H , projects the centralized retrieval support onto a single primitive trace unit.*

Then the Drell-Yan channel weight is

$$w_{\text{DY}} = \frac{\tau(P_{\text{DY}}R_3)}{\tau(R_3)} = \frac{1}{3}, \quad P_{\text{DY}} := E_\ell \circ H.$$

The arithmetic of the conclusion is immediate once the three hypotheses are granted. The substance of the theorem is the three hypotheses, each of which is an independently checkable operator-algebra claim about the framework's existing structures.

Status of (H1), (H2), and (H3): all three closed conditional on operator-algebraic conditions. All three hypotheses of Theorem 15.5 admit conditional derivations from existing framework structure. Hypothesis (H1) is closed by Theorem 4.3 of Section 4.2: if the spatial Lorentz-recovery subgroup $G_{\text{sp}} \subset G_{\text{Lor}}(\Sigma_{\text{obs}})$ exists, acts irreducibly on S_{obs} , and preserves the record-support trace, then $\tau(e_1) = \tau(e_2) = \tau(e_3)$; this routes (H1) to the existing Lorentz-recovery program of Section 11.3 with no new substrate primitive. Hypothesis (H2) is closed by Lemma 15.7 below: an operator-algebraic statement of what it means for the prepared channel state to be unresolved with respect to the primitive retrieval algebra, together with the trace-equivalence supplied by (H1), forces the hadronic-initial-state map on R_3 to be the symmetric conditional expectation. Hypothesis (H3) is closed by Lemma 15.8 below: a forward trace-preservation condition on the retrieval inclusion of R_3 into the leptonic final-state algebra, together with a Murray–von Neumann identification of clean leptonic records as the image of one primitive retrieval projection, forces the leptonic visible support to carry exactly one primitive retrieval trace unit.

Hypothesis (H2): from unresolvedness to symmetric averaging. The initial-state observable algebra of the channel, A_{init} , contains R_3 as a distinguished subalgebra: $R_3 \hookrightarrow A_{\text{init}}$. There is no claim of a global tensor factorization $A_{\text{init}} = A_{\text{had}} \otimes R_3$; the relation needed is only the inclusion of R_3 together with the existence of an S_3 action on R_3 that commutes with the experimentally resolved hadronic observables of the channel (flavor, PDF, kinematics, QCD-sector data). In the minimal model used here, R_3 is commutative, $R_3 \cong \mathbb{C}^3$.

Let S_3 act on R_3 by permutation of the primitive projections, $\sigma(e_i) = e_{\sigma(i)}$. By Theorem 4.3, $\tau(e_1) = \tau(e_2) = \tau(e_3)$, so this action is trace-preserving with respect to the substrate trace τ . Without (H1), the S_3 action would still exist as an algebraic permutation of R_3 , but it would not in general preserve τ ; trace-preservation of the action is the load-bearing point at which (H1) is used.

Definition 15.6 (R_3 -unresolved channel state). A channel state ω on A_{init} is *unresolved with respect to R_3* if its restriction to R_3 is invariant under the S_3 branch-permutation action:

$$\omega(\sigma(x)) = \omega(x) \quad \text{for every } \sigma \in S_3 \text{ and every } x \in R_3.$$

This is the operator-algebraic statement that the channel state carries no observable label distinguishing the three primitive retrieval branches. It does not require flavor symmetry of the proton, equal parton distribution functions, symmetric kinematics, or symmetric detector acceptance: those structures live in operators outside R_3 and may be arbitrarily asymmetric. It requires only that the channel state's restriction to R_3 be permutation-invariant.

Define the symmetric expectation on R_3 as the finite-group average,

$$E_{\text{sym}} : R_3 \rightarrow R_3^{S_3}, \quad E_{\text{sym}}(x) = \frac{1}{6} \sum_{\sigma \in S_3} \sigma(x).$$

Under $R_3 \cong \mathbb{C}^3$, the S_3 -fixed subalgebra is one-dimensional, generated by $\mathbf{1}_{R_3} = e_1 + e_2 + e_3$. The trace-preservation of the S_3 action (supplied by Theorem 4.3) makes E_{sym} a trace-preserving conditional expectation onto $R_3^{S_3}$; it is the unique such expectation. On the primitive generators,

$$E_{\text{sym}}(e_i) = \frac{1}{3} (e_1 + e_2 + e_3) \quad (i = 1, 2, 3).$$

Lemma 15.7 (Unresolved hadronic conditional expectation; (H2)). *Assume Theorem 4.3, so the primitive retrieval projections are trace-equivalent. Let $R_3 \cong \mathbb{C}^3$ be the primitive commutative retrieval subalgebra of A_{init} , and let ω be an R_3 -unresolved channel state in the sense of Definition 15.6. Then the restriction of ω to R_3 factors through the symmetric conditional expectation:*

$$\omega(x) = \omega(E_{\text{sym}}(x)) \quad \text{for every } x \in R_3.$$

Proof. By Theorem 4.3, $\tau(e_1) = \tau(e_2) = \tau(e_3)$, so the S_3 action on R_3 is trace-preserving and E_{sym} is the trace-preserving conditional expectation onto $R_3^{S_3}$. For an R_3 -unresolved state ω , Definition 15.6 gives $\omega(\sigma(x)) = \omega(x)$ for every $\sigma \in S_3$ and every $x \in R_3$. Therefore

$$\omega(E_{\text{sym}}(x)) = \frac{1}{6} \sum_{\sigma \in S_3} \omega(\sigma(x)) = \frac{1}{6} \sum_{\sigma \in S_3} \omega(x) = \omega(x). \quad \square$$

Lemma 15.7 converts (H2) of Theorem 15.5 into the conditional statement: *if the prepared channel state is R_3 -unresolved in the sense of Definition 15.6, then its action on R_3 is the symmetric averaging map E_{sym} . The hypothesis to be justified for any specific channel is therefore Definition 15.6, a checkable operator-algebraic condition on the channel state rather than a verbal claim about hadronic symmetry. The Drell-Yan initial state, in particular, is expected to be R_3 -unresolved because the experimentally resolved hadronic observables of the channel do not carry a primitive spatial-retrieval branch index by construction: the branch index is internal to the framework's observer-access bookkeeping, not a measured quantum number.*

Hypothesis (H3): from a trace-preserving retrieval inclusion to one primitive retrieval unit. After (H2) has reduced the hadronic initial-state input on R_3 to the scalar unresolved support $\mathbf{1}_{R_3} \in R_3^{S_3}$, the leptonic final-state image cannot live inside R_3 itself: under $R_3 \cong \mathbb{C}^3$ with the branch-permutation S_3 action, the only S_3 -invariant projections of R_3 are 0 and $\mathbf{1}_{R_3}$, and neither is a rank-one branch-blind visible support. The clean leptonic visible support therefore lives in a distinct algebra: the leptonic final-state algebra A_ℓ , related to R_3 by an inclusion

$$\iota_\ell : R_3 \rightarrow A_\ell$$

that brings primitive retrieval projections into the leptonic visible setting. The condition required on ι_ℓ for the trace to inherit correctly is forward trace-preservation: $\tau_\ell(\iota_\ell(x)) = \tau(x)$ for every

$x \in R_3$, where τ_ℓ denotes the trace on A_ℓ and τ the substrate trace inherited by R_3 . This is the forward direction of a Markov inclusion; only this forward direction is used in what follows, and no reverse conditional expectation or Jones-index normalization is invoked.

Sub-trace projections may exist in the ambient algebra A_ℓ (the substrate Type II₁ structure permits projections of arbitrary trace value in $[0, 1]$), but those projections are not Murray–von Neumann equivalent to the image of any primitive retrieval projection under ι_ℓ and therefore do not represent clean leptonic retrieval-supported records.

Lemma 15.8 (Leptonic scalar trace-inheritance; (H3)). *Assume Theorem 4.3, so the primitive retrieval projections satisfy $\tau(e_1) = \tau(e_2) = \tau(e_3)$. Let A_ℓ be the clean leptonic neutral-current final-state algebra with trace τ_ℓ , and let*

$$\iota_\ell : R_3 \rightarrow A_\ell$$

be a trace-preserving retrieval inclusion:

$$\tau_\ell(\iota_\ell(x)) = \tau(x) \quad \text{for every } x \in R_3.$$

Suppose a clean leptonic neutral-current record is represented by a visible scalar projection $q_\ell \in A_\ell$ that is Murray–von Neumann equivalent to the image of one primitive retrieval projection,

$$q_\ell \sim \iota_\ell(e_i),$$

in the sense that there exists a partial isometry $v \in A_\ell$ with $v^*v = \iota_\ell(e_i)$ and $vv^* = q_\ell$. Then

$$\tau_\ell(q_\ell) = \tau(e_i).$$

Proof. Murray–von Neumann equivalent projections in a von Neumann algebra carry equal trace under any tracial state [4, 5]: $\tau_\ell(vv^*) = \tau_\ell(v^*v)$ by the trace property, hence $\tau_\ell(q_\ell) = \tau_\ell(\iota_\ell(e_i))$. Forward trace-preservation of ι_ℓ gives $\tau_\ell(\iota_\ell(e_i)) = \tau(e_i)$. Composing the two equalities, $\tau_\ell(q_\ell) = \tau(e_i)$. \square

Lemma 15.8 converts (H3) of Theorem 15.5 into two operator-algebraic conditions on the leptonic final-state side: (i) the retrieval inclusion $\iota_\ell : R_3 \rightarrow A_\ell$ is trace-preserving in the forward direction; (ii) a clean leptonic neutral-current record is the Murray–von Neumann equivalence class of one primitive retrieval image under ι_ℓ . The equivalence $q_\ell \sim \iota_\ell(e_i)$ does not identify q_ℓ with a labeled branch of R_3 ; doing so would place the visible support back inside R_3 and break branch-blindness. The identification is at the level of trace, not at the level of projection identity: q_ℓ lives in A_ℓ and carries the same trace as the image of one primitive retrieval unit. The branch symmetry of the leptonic readout is preserved at the trace level rather than at the projection-identity level.

What the conditional theorem now establishes. Theorem 15.5 is the form of the Drell-Yan derivation, with all three of its hypotheses now established conditional on operator-algebraic conditions on the framework’s existing structures: (H1) via Theorem 4.3, conditional on a trace-preserving transitive spatial group action on the primitive retrieval projections; (H2) via Lemma 15.7, conditional on S_3 -invariance of the prepared channel state on R_3 (Definition 15.6); (H3) via Lemma 15.8, conditional on a forward-trace-preserving retrieval inclusion $\iota_\ell : R_3 \rightarrow A_\ell$ together with the Murray–von Neumann identification of clean leptonic records as primitive retrieval images. The Drell-Yan weight follows arithmetically: Theorem 4.3 gives $\tau(\mathbf{1}_{R_3}) = 3\tau(e_i)$; Lemma 15.7 reduces the hadronic initial-state input on R_3 to the scalar support $\mathbf{1}_{R_3}$; Lemma 15.8 gives the leptonic visible image q_ℓ with $\tau_\ell(q_\ell) = \tau(e_i)$. Hence

$$w_{\text{DY}} = \frac{\tau_\ell(q_\ell)}{\tau(\mathbf{1}_{R_3})} = \frac{\tau(e_i)}{3\tau(e_i)} = \frac{1}{3},$$

without reference to any measured collider value. The conditional Drell-Yan prediction $w_{\text{DY}} = 1/3$ is therefore a framework prediction, conditional on each of the three operator-algebraic conditions of Theorem 4.3 and Lemmas 15.7, 15.8 being satisfied by the framework’s existing structures and by the prepared channel state of the relevant experiment. Each condition is a checkable algebraic statement rather than a verbal claim about hadronic symmetry, color, or final-state structure. A successful identification of A_ℓ and verification of the trace-preservation and Murray–von Neumann equivalence conditions inside the framework’s existing observer-cut machinery would lift the conditional status of (H3) entirely; an inability to verify any of these conditions, or a requirement of ex post fitting choices, would falsify the framework’s claim to capture the Drell-Yan channel structure of these measurements.

Resolved Drell-Yan kernel and the isotropic limit. The conditional Drell-Yan theorem (Theorem 15.5) gives the channel weight $w_{\text{DY}} = 1/3$ as a single value for the entire hadron-collider Drell-Yan class. Precision measurements at hadron colliders, however, extract $\sin^2 \theta_{\text{eff},\ell}^{\text{DY}}$ values that differ between experiments and analysis channels (CDF vs. D0, ATLAS vs. CMS vs. LHCb, electron vs. muon, central vs. forward acceptance) by amounts comparable to the experimental precision profiles. The conditional theorem in its current form does not distinguish among these combinations. The present block recasts $w_{\text{DY}} = 1/3$ as the *isotropic limit* of a more general visibility formula, built on the resolved retrieval algebra \tilde{R}_3 constructed below as an admissible refinement of R_3 in the sense of Definition 3.1. The recasting does not modify the conditional theorem: $w_{\text{DY}} = 1/3$ remains the framework’s prediction under the operator-algebraic conditions of Theorem 4.3 and Lemmas 15.7, 15.8, and is now recovered exactly through Corollary 3.5.

Resolved retrieval algebra as an admissible refinement. The resolved retrieval algebra is constructed as an admissible refinement of R_3 in the sense of Section 3.3:

$$\tilde{R}_3 = R_3 \otimes \mathcal{F}_{\text{DY}}, \quad \mathcal{F}_{\text{DY}} = \mathcal{F}_{\text{flavor}} \otimes \mathcal{F}_\xi,$$

where $\mathcal{F}_{\text{flavor}}$ is the finite-dimensional commutative W^* -algebra spanned by orthogonal flavor projections $\{P_u, P_d\}$ with $P_u + P_d = \mathbf{1}_{\text{flavor}}$ and democratic trace $\tau_{\text{flavor}}(P_u) = \tau_{\text{flavor}}(P_d) = 1/2$, and \mathcal{F}_ξ is an abstract abelian W^* -algebra of kinematic support coordinates (parton momentum fractions, dilepton rapidity, invariant mass, transverse momentum, detector acceptance projections) with a normalized faithful trace τ_ξ . The product trace is $\tau_{\text{DY}} = \tau_{\text{flavor}} \otimes \tau_\xi$.

Verification of the four assumptions of Definition 3.1 is direct: (A1) holds by construction since the lifted primitive projections are $e_i \otimes \mathbf{1}_{\text{DY}}$; (A2) is the product-trace definition above; (A3) holds because the S_3 permutation action on R_3 extends to \tilde{R}_3 by acting only on the first tensor factor (flavor and kinematic labels are inert under branch permutation); (A4) is realized by the canonical lift of Definition 3.3,

$$\tilde{K}_{\text{rec}} = K_{\text{rec}}^{R_3} \otimes \mathbf{1}_{\text{DY}}, \quad \tilde{M}_{\text{rec}} = M_{\text{rec}}^{R_3} \otimes \mathbf{1}_{\text{DY}}.$$

Theorem 3.4 therefore applies: trace factorization, equivariant lift of E_{sym} to the first tensor factor, and access-geometry factorization $\tilde{G}_{\text{acc}} = G_{\text{acc}}^{R_3} \otimes \mathbf{1}_{\text{DY}}$. The resolved branch projections are

$$e_{i,q,\xi} = e_i \otimes P_q \otimes f_\xi.$$

The forward trace-preservation and equivariance properties that the resolved DY block relies on are immediate consequences of Theorem 3.4, with no new substrate input required.

Experiment-and-channel projection and visibility weight. Each precision experiment publishes a specific kinematic acceptance, channel selection, and extraction methodology. Encode these as an element $P_{\text{exp},\ell} \in \tilde{R}_3$:

$$P_{\text{exp},\ell} = \sum_q \int d\xi \rho_q(\xi; \text{exp}) \Pi_{\text{exp},\ell}(\xi) \mathbf{S}_{\text{exp},q}(\xi) \sum_i e_{i,q,\xi}, \quad (34)$$

where $\rho_q(\xi; \text{exp})$ is the parton-distribution weighting of flavor q at kinematic point ξ for the experiment's beam configuration, $\Pi_{\text{exp},\ell}(\xi)$ is the experiment's published acceptance, and $\mathbf{S}_{\text{exp},q}(\xi)$ is the experiment's signed weak-angle extraction score on the resolved kinematic factor \mathcal{F}_ξ :

$$\mathbf{S}_{\text{exp},q}(\xi) = \frac{\partial \log \sigma_{\text{exp},q}(\xi)}{\partial \sin^2 \theta_{\text{eff}}}. \quad (35)$$

The score is the bin-resolved log-derivative that controls the experimental extraction of $\sin^2 \theta_{\text{eff},\ell}^{\text{DY}}$ from the angular asymmetry coefficient $A_4 = (8/3) A_{FB}$ in the Collins–Soper-frame template fit. It is required because the measured observable in every modern Drell–Yan effective-angle extraction (CDF, D0, ATLAS, CMS, LHCb) is a signed angular asymmetry rather than a rate: a rate-weighted projection ($\mathbf{S}_{\text{exp},q} \equiv 1$) reproduces the framework's earlier toy form but cannot in general distinguish experiments with similar initial-state luminosities and different angular-template responses, the limiting case being CDF ee versus D0 $\mu\mu$ at the Tevatron.

The element $P_{\text{exp},\ell}$ is fully external phenomenological data: it introduces no framework-internal free parameter, and its construction from PDF sets, from the experiment's published acceptance, and from the experiment's published template (or its angular-coefficient $A_4(M, y)$ analysis) is the same input that a Standard-Model template fit uses. The score $\mathbf{S}_{\text{exp},q}(\xi)$ carries a sign and makes $P_{\text{exp},\ell}$ a generically signed (rather than positive) element of \tilde{R}_3 ; every framework consequence below uses only trace pairings, which remain well defined as real numbers. By construction $P_{\text{exp},\ell}$ remains S_3 -invariant on the coarse retrieval index, so the isotropic limit $w_{\text{exp},\ell} \rightarrow 1/3$ of Corollary 3.5 is preserved unchanged: when $\mathbf{S}_{\text{exp},q}$ is flavor-symmetric on the up-class and down-class projectors, the trace pairing of $P_{\text{exp},\ell}$ with K_Y vanishes.

The experiment-and-channel-specific Drell-Yan visibility weight is the trace ratio

$$w_{\text{exp},\ell} := \frac{\tau\left(P_{\text{exp},\ell} \tilde{K}_{\text{DY}}^{\text{int}}\right)}{\tau(P_{\text{exp},\ell})},$$

where $\tilde{K}_{\text{DY}}^{\text{int}} \in \tilde{R}_3$ is the *resolved retrieval kernel* constructed below. The corresponding framework-predicted effective angle is

$$\sin^2 \theta_{\text{eff},\ell}^{\text{DY}}(\text{exp}, \ell) = V_\ell^0 + w_{\text{exp},\ell} \delta_{\text{init}},$$

with $V_\ell^0 = 1/T_2 + 1/T_1 + 1/(T_2 T_1)$ from the neutral-boundary visibility lemma (26) and $\delta_{\text{init}} = 1/(T_{q,\text{col}} T_1)$ the initial-state fundamental color-support increment, structurally the same as the δ_b of Equation (32) with $T_{q,\text{col}} = 27/4$ from Equation (31) for any fundamental-representation quark.

Explicit construction of the resolved kernel. As clarified in Section 3.3 (the two-distinct-objects discussion), the substrate-level paired bookkeeping ($\tilde{K}_{\text{rec}}, \tilde{M}_{\text{rec}}$) is flat on the \mathcal{F}_{DY} factor by Definition 3.3, while the measurement kernel $\tilde{K}_{\text{DY}}^{\text{int}}$ is a derived object that may carry non-trivial flavor structure. The framework-internal resolved retrieval kernel is the sum of two structurally distinct pieces:

$$\tilde{K}_{\text{DY}}^{\text{int}} = K_{\text{iso}} + K_Y, \quad (36)$$

where

$$K_{\text{iso}} = \frac{1}{3} \mathbf{1}_{\tilde{R}_3}, \quad K_Y = \lambda_Y \left[Y_u^2 P_u + Y_d^2 P_d - \langle Y^2 \rangle_{\text{iso}} \mathbf{1}_{\tilde{R}_3} \right]. \quad (37)$$

The first piece K_{iso} is the scalar isotropic baseline supplying the 1/3 trace at the unrefined level. The second piece K_Y is the trace-centered Abelian-footprint correction in the flavor factor; here λ_Y is a framework-internal coupling whose explicit numerical value is part of the deferred Drell–Yan computation, and the isotropic subtraction constant $\langle Y^2 \rangle_{\text{iso}}$ is the democratic flavor average

$$\langle Y^2 \rangle_{\text{iso}} = \tau_{\text{flavor}}(Y_u^2 P_u + Y_d^2 P_d) = \frac{1}{2}(Y_u^2 + Y_d^2) = \frac{1}{2}\left(\frac{4}{9} + \frac{1}{9}\right) = \frac{5}{18}, \quad (38)$$

which is the unique value of the subtraction constant that makes the correction K_Y trace-centered in \mathcal{F}_{DY} ($\tau_{\text{DY}}(K_Y) = 0$). The choice is not a normalization convention: it is the algebraic condition required for Corollary 3.5 to apply, with the democratic flavor weighting of $\mathcal{F}_{\text{flavor}}$ fixing the numerical value 5/18.

Theorem 15.9 (Drell-Yan isotropic limit, exact form). *For the resolved retrieval kernel $\tilde{K}_{\text{DY}}^{\text{int}}$ of Equations (36)–(38) and for any experiment-and-channel projection $P_{\text{exp},\ell}$ constructed as above, the unrefined-limit projection of $\tilde{K}_{\text{DY}}^{\text{int}}$ under $\mathcal{E}_{\text{DY}} := \text{id}_{R_3} \otimes \tau_{\text{DY}}$ is*

$$\mathcal{E}_{\text{DY}}(\tilde{K}_{\text{DY}}^{\text{int}}) = \frac{1}{3} \mathbf{1}_{R_3},$$

and the visibility weight in the unrefined limit is

$$w_{\text{exp},\ell}|_{\text{iso}} = \frac{1}{3}.$$

Proof. The kernel $\tilde{K}_{\text{DY}}^{\text{int}} = K_{\text{iso}} + K_Y$ has K_{iso} scalar on \tilde{R}_3 and, by the choice of $\langle Y^2 \rangle_{\text{iso}} = 5/18$ in Equation (38), $\tau_{\text{DY}}(K_Y) = 0$. Corollary 3.5 therefore gives $\mathcal{E}_{\text{DY}}(\tilde{K}_{\text{DY}}^{\text{int}}) = K_{\text{iso}}|_{R_3} = \frac{1}{3} \mathbf{1}_{R_3}$. For the visibility weight, the conditional expectation \mathcal{E}_{DY} commutes with multiplication by S_3 -symmetric elements such as $P_{\text{exp},\ell}$ (which is constructed to be S_3 -invariant on the coarse retrieval index), so $\tau(P_{\text{exp},\ell} \tilde{K}_{\text{DY}}^{\text{int}}) = \tau(P_{\text{exp},\ell} \mathcal{E}_{\text{DY}}(\tilde{K}_{\text{DY}}^{\text{int}})) = \frac{1}{3} \tau(P_{\text{exp},\ell})$ in the unrefined limit, hence $w_{\text{exp},\ell} = 1/3$. \square

Theorem 15.5’s $w_{\text{DY}} = 1/3$ is therefore not a contingent claim but an exact consequence of the refinement-compatibility architecture: it is what Corollary 3.5 produces when applied to the explicit kernel of Equations (36)–(38). The recovery is exact, not asymptotic, because the factorization theorem makes the trace product structurally exact rather than approximately separable.

Anchor-projection reading of the visibility weight. The visibility-weight formula $\sin^2 \theta_{\text{eff},\ell}^{\text{DY}}(\text{exp}, \ell) = V_\ell^0 + w_{\text{exp},\ell} \delta_{\text{init}}$ admits a structural reading as a two-anchor projection. The direct leptonic neutral-boundary readout

$$\sin^2 \theta_L := V_\ell^0 = \frac{1}{T_2} + \frac{1}{T_1} + \frac{1}{T_2 T_1} = 0.2311412087\dots$$

is the low anchor: the value the framework returns when no color-carrying final support participates in the observer-access cut, identified at SLD through A_{LR} in Section 15.10. The final-support trace overlap $\delta_{\text{init}} = 1/(T_{q,\text{col}} T_1)$ defines a high anchor

$$\sin^2 \theta_H := V_\ell^0 + \delta_{\text{init}} = 0.2322222980\dots,$$

structurally identical to the LEP $A_{\text{FB}}^{0,b}$ value of Equation (32). The visibility weight $w_{\text{exp},\ell}$ is the experiment-and-channel projection weight onto the high anchor, with substrate-derived bounds given by the algebraic refinement below. The isotropic value $w = 1/3$ of Theorem 15.9 corresponds to the trace-democratic projection of an unresolved channel state. The score-projected pairing constructed below is the framework’s substrate-internal mechanism for $w_{\text{exp},\ell}$ to deviate from 1/3 when the channel state resolves the up-class versus down-class projectors asymmetrically; the structural object being projected is the anchor pair $(\sin^2 \theta_L, \sin^2 \theta_H)$, and the SLD/LEP- b split of Section 15.10 is the same anchor pair read at the endpoint values $w = 0$ and $w = 1$.

Primary substrate architecture: anchor amplitude, residual amplitude, Pythagorean budget. The framework's substrate-derived DY readout is fixed by the closure-normalized observation-support theorem of Appendix 6.11 together with the operator-algebraic class structure introduced here. Let P_{\min} denote the minimal observation support projector for the charged-lepton final-state record at the neutral-current angular-asymmetry vertex (Definition 6.15), with $d_{\text{supp}} = 3$ branch coordinates (charge orientation, projective direction, kinematic-closure magnitude) and trace weight $T_{\text{ch}} = 27/4$ from Theorem 6.16. Let $P_{\text{exp},\ell}$ denote the experiment-and-channel projection operator. Decompose

$$P_{\text{exp},\ell} = P_{\text{exp},\ell}^{\parallel} + P_{\text{exp},\ell}^{\perp}, \quad P_{\text{exp},\ell}^{\parallel} := P_{\text{exp},\ell} \wedge P_{\min}, \quad P_{\text{exp},\ell}^{\perp} := P_{\text{exp},\ell} - P_{\text{exp},\ell}^{\parallel}, \quad (39)$$

with $P_{\text{exp},\ell}^{\perp} P_{\min} = 0$ (the residual is operator-orthogonal to the minimal support). Define the substrate-internal trace fractions and the corresponding anchor-coordinate amplitude and residual amplitude:

$$v_{\text{exp},\ell} := \frac{\tau(P_{\text{exp},\ell}^{\parallel})}{\tau(P_{\min})}, \quad u_{\text{exp},\ell} := \frac{\tau(P_{\text{exp},\ell}^{\perp})}{\tau(P_{\min})}, \quad w_{\text{exp},\ell} := \sqrt{v_{\text{exp},\ell}}, \quad r_{\text{exp},\ell} := \sqrt{u_{\text{exp},\ell}}. \quad (40)$$

The amplitudes w and r are the primary substrate-internal quantities entering the DY readout; the trace fractions v and u are derived auxiliary quantities. Let $\sigma_{\text{exp},\ell}, \sigma'_{\text{exp},\ell} \in \{+1, -1\}$ denote the visibility/Schur signs of Equation (29) applied to the anchor coordinate and the orthogonal coordinate respectively: each sign is $+1$ for the visibility-positive readout class and -1 for the Schur-eliminated transport readout class, fixed by readout structure rather than fitted. The two signs are structurally independent.

The substrate-derived DY central-value shift is the sum of two orthogonal contributions to the same observable

$$\sin^2 \theta_{\text{eff,DY}}(\text{exp}, \ell) = V_{\ell}^0 + \delta_{\text{anchor}}(\text{exp}, \ell) + \delta_{\text{projection}}(\text{exp}, \ell), \quad (41)$$

$$\delta_{\text{anchor}}(\text{exp}, \ell) := \sigma_{\text{exp},\ell} w_{\text{exp},\ell} \delta_{\text{init}}, \quad \delta_{\text{projection}}(\text{exp}, \ell) := \sigma'_{\text{exp},\ell} r_{\text{exp},\ell} \delta_{\text{init}}, \quad (42)$$

with $\delta_{\text{init}} = 1/(T_{\text{ch}} T_1)$.

Pythagorean substrate-budget bound (primary). The amplitudes are constrained by the substrate-internal Pythagorean bound

$$w_{\text{exp},\ell}^2 + r_{\text{exp},\ell}^2 \leq 1, \quad (43)$$

equivalently $\delta_{\text{anchor}}^2 + \delta_{\text{projection}}^2 \leq \delta_{\text{init}}^2$. The bound is the trace-content statement $\tau(P_{\text{exp},\ell}) \leq \tau(P_{\min})$: the experiment-and-channel projection cannot fill more substrate-trace than the minimal observation support carries. The anchor amplitude and the residual amplitude share a single substrate-trace budget; they cannot both saturate. This is the operator-algebraic statement that prevents the residual from acting as an independent second anchor-scale contribution, and it is the structurally primary form of the bound.

Simple per-component cap (derived corollary). Equation (43) implies the per-component cap

$$|w_{\text{exp},\ell}| \leq 1, \quad |r_{\text{exp},\ell}| \leq 1, \quad \iff \quad |\delta_{\text{anchor}}| \leq \delta_{\text{init}}, \quad |\delta_{\text{projection}}| \leq \delta_{\text{init}}.$$

The simple per-component cap alone would not prevent the residual from becoming a second independent anchor-scale contribution and would re-introduce the overfitting risk that the v19 architecture removes. The Pythagorean form (43) is the substrate-level statement; the per-component cap is its derived consequence.

Observer-access classification. Measurements satisfying the operator-algebraic condition

$$P_{\text{exp},\ell} \leq P_{\text{min}} \quad (\text{operator ordering on } \tilde{R}_3), \quad (44)$$

equivalently $P_{\text{exp},\ell}^\perp = 0$, hence $r_{\text{exp},\ell} = 0$ and $\delta_{\text{projection}} = 0$, form the *Tevatron class*: the projection onto charge \times direction \times kinematic-closure-magnitude is a subprojection of P_{min} , with detector-level coordinates (isolation, shower shape, track quality, fiducial geometry) entering only as event-selection thresholds applied before $P_{\text{exp},\ell}$ is defined. The class is delineated structurally by the feature that the support projection and the parton weighting enter the extraction as external phenomenological inputs, not as joint fit objects against the measured A_{FB} template.

Measurements with $r_{\text{exp},\ell} > 0$ form the *LHC joint-fit class*: parton-distribution constraints derived from the measured A_{FB} distribution itself project onto support outside P_{min} , so $P_{\text{exp},\ell}^\perp \neq 0$. The class is delineated structurally by the feature that the proton parton distributions and the weak-angle parameter are extracted jointly against a shared A_{FB} template. The framework’s primary statement on the joint-fit class is the Pythagorean bound (43) on amplitudes, not a row-level central-value prediction; full row-level prediction would require reconstructing the joint-fit machinery (PDF inference, quark-direction dilution, channel acceptance, rapidity weighting, template covariance structure), which the framework’s substrate-derived primary architecture does not undertake.

Secondary residual apparatus: structural diagnostic of K_Y (demoted from primary in v19)

The remainder of this subsection retains the structural identification of K_Y within the resolved retrieval algebra together with the score-projected experiment-and-channel pairing $\mathcal{A}_{\text{exp},\ell}$. In the v19 architecture this material is a *bounded residual diagnostic*, not a primary mechanism for predicting row-level central values: the primary DY architecture (41) is fixed by the closure-normalized observation-support theorem and the Pythagorean budget bound (43), both of which are insensitive to K_Y and λ_Y . The constructions below are retained because real Drell–Yan extractions depend on parton distributions, channel acceptance, rapidity weighting, mass-window choices, and template covariance structure, and the framework should be able to discuss those objects structurally without pretending experiments are projection-free.

Two things are explicitly *removed as primary* in v19. First, row-by-row fits of the coupling λ_Y inside K_Y against individual experiment central values are *structurally inadmissible*: a row-level λ_Y fit would absorb measurement variation into a framework parameter and remove the falsifiability gain of the primary anchor-projection architecture. Second, prediction of row-level joint-fit class central values through any K_Y -driven score-projection construction is removed: the joint-fit class is bounded by (43), not predicted row by row. The score-projected pairing of Equation (34) below remains framework-internal structural content; closure of λ_Y as a single substrate-internal coupling (not as per-experiment fits) is deferred as a residual-diagnostic continuation work, and the Pythagorean bound caps every residual-diagnostic contribution by construction.

Source of S_3 -symmetry breaking: per-flavor central Abelian footprint. Experiment-and-channel-specific deviations from $w_{\text{exp},\ell} = 1/3$ enter through the K_Y term when an experiment-specific projection $P_{\text{exp},\ell}$ samples up-type and down-type partons asymmetrically. The fundamental color-support trace weight $T_{q,\text{col}} = 27/4$ used in the SLD/LEP- b channel split is identical for every fundamental-representation quark under $SU(3)$; Equation (31) carries no flavor index at the framework level. The central Abelian phase footprint, however, scales as $|Y_q|^2$ in the hypercharge normalization fixed by the minimal chiral content of Section 15.4. With the assignments

$$(Y_Q, Y_u, Y_d, Y_L, Y_e) = \left(\frac{1}{6}, \frac{2}{3}, -\frac{1}{3}, -\frac{1}{2}, -1\right)$$

of right-handed singlet hypercharges from Section 15.4, the squared footprints of up-type and down-type quark singlets satisfy

$$|Y_u|^2 = |Y_c|^2 = |Y_t|^2 = \frac{4}{9}, \quad |Y_d|^2 = |Y_s|^2 = |Y_b|^2 = \frac{1}{9},$$

giving the framework-native ratio

$$\boxed{|Y_u|^2 : |Y_d|^2 = 4 : 1.}$$

This ratio enters K_Y in Equation (37) as the framework-internal source of S_3 -symmetry breaking when an experiment-specific projection $P_{\text{exp},\ell}$ samples flavors asymmetrically. The ratio is fixed by the minimal-chiral-content recoverability-obstruction equations of Section 15.4 and not by any external fit; it is the same input that fixes the $T_{\text{bare}} = 21/2$ bare trace and the $5/3$ hypercharge normalization, used here in its squared form because the resolved retrieval kernel reads the Abelian phase footprint at the quark-Z vertex rather than the unsquared hypercharge.

Experiments and channels that select different up-type vs. down-type quark mixtures — through the PDF weighting $\rho_q(\xi; \text{exp})$ and the kinematic projection $\Pi_{\text{exp},\ell}(\xi)$ — therefore sample different effective values of the Abelian footprint and receive different framework-predicted $w_{\text{exp},\ell}$ values, away from the isotropic value $1/3$. The structural prediction is that experiment dependence enters entirely through this flavor-and-kinematic support weighting and not through any modification of the primitive observer branch structure: the layer-protection content of Theorem 3.4 (the untwisted-extension assumption (A1) of Definition 3.1) guarantees that no measurement can require a coupling between kinematics and the projective branch index. The associated falsifier is Remark 3.6, specialized to the Drell–Yan context.

Structural identification of K_Y within the framework algebra. The coupling λ_Y in Equation (37) is the trace-centered Abelian-footprint coupling of the resolved retrieval kernel. Its magnitude is not pinned in the present block: the framework-internal derivation that closes λ_Y requires the score-projected experiment-and-channel pairing of Equation (34), whose explicit form depends on external template inputs and is deferred to the Drell–Yan companion. What can be settled here is the *structural class* of K_Y within the framework algebra: which primitive composition K_Y is an instance of. Three structurally distinct identifications are formally available, and we work them out explicitly to identify which class is admissible.

The eigenvalues of K_Y on the flavor support are calculable from the substrate-fixed hypercharge assignments. From the trace-centered form of Equation (37),

$$\text{spec}(K_Y) = \{ \lambda_Y(Y_u^2 - \langle Y^2 \rangle_{\text{iso}}), \lambda_Y(Y_d^2 - \langle Y^2 \rangle_{\text{iso}}) \} = \{ \lambda_Y/6, -\lambda_Y/6 \},$$

where the prefactor $1/6$ is substrate-derived: $Y_u^2 - \langle Y^2 \rangle_{\text{iso}} = 4/9 - 5/18 = 3/18 = 1/6$, and analogously $Y_d^2 - \langle Y^2 \rangle_{\text{iso}} = -1/6$. The leading nonzero eigenvalue magnitude is therefore $|\lambda_Y|/6$.

(P1) Stiffness identification:

K_Y enters the substrate stiffness K_{rec} directly, so its eigenvalues modify G_{acc} and the access spectrum. Matching the leading eigenvalue to the primitive access gap gives $|\lambda_Y|/6 = \pi^{-4}$, hence $\lambda_Y = 6\pi^{-4}$. This identification is *ruled out* by the canonical-lift commitment of Definition 3.3: the substrate stiffness lifts strictly as $\tilde{K}_{\text{rec}} = K_{\text{rec}}^{R_3} \otimes \mathbf{1}_{\text{DY}}$, with no contribution from K_Y .

(P2) Linear-composition identification:

K_Y plays the role of $\Lambda_{\text{acc}}^{-1}$ (single primitive access step without recovery), and K_Y^2 is the analog of the squared primitive composition $\Lambda_{\text{acc}}^{-2}$. Matching the leading eigenvalue of K_Y^2 to π^{-4} gives $(\lambda_Y/6)^2 = \pi^{-4}$, hence $\lambda_Y = 6\pi^{-2}$. This identification is also *ruled out*: as established in Section 6.9, a single primitive access operator without recovery is not closed under record recovery and is not admissible as a primitive access dressing. There is no framework-admitted $\Lambda_{\text{acc}}^{-1}$ object for K_Y to be the analog of.

(P3) Primitive-composition identification:

K_Y is itself the primitive flavor-side access-recovery composition, structurally analogous to $\Lambda_{\text{acc}}^{-2}$. The squared-hypercharge form Y_q^2 in Equation (37) encodes the two-vertex composition: the Abelian phase enters at one quark– Z vertex (access) and exits at the conjugate vertex (recovery), giving the squared structure $Y_q \cdot Y_q = Y_q^2$. This identification places K_Y in the primitive access-recovery class with no canonical-lift conflict and no $\Lambda_{\text{acc}}^{-1}$ -analog requirement. It does *not*, in the present block, fix the coupling λ_Y to a specific value: the prefactor 6 in $\text{spec}(K_Y) = \{\pm\lambda_Y/6\}$ is the algebraic inverse of $|Y_q^2 - \langle Y^2 \rangle_{\text{iso}}|$ fixed by substrate-derived hypercharge assignments, but the overall scale λ_Y is the framework-internal coupling of K_Y within the primitive access-recovery class, to be closed by a derivation that involves the score-projected experiment-and-channel pairing.

Only identification (P3) is structurally admissible: it respects the canonical-lift commitment, requires no $\Lambda_{\text{acc}}^{-1}$ -analog object, and matches the substrate-derived squared-hypercharge structure of K_Y to the squared structure of $\Lambda_{\text{acc}}^{-2}$ in the fine-structure access dressing. The numerical value of λ_Y within the admissible (P3) class is a single substrate-internal coupling whose closure requires the score-projected pairing of Equation (34) and is deferred to the Drell–Yan companion.

Score-projected experiment-and-channel asymmetry. The trace pairing of $P_{\text{exp},\ell}$ with K_Y defines the experiment-and-channel-specific deviation from the isotropic visibility weight as the score-projected up–down asymmetry of the centered hypercharge kernel:

$$\Delta w_{\text{exp},\ell} = \frac{\tau(P_{\text{exp},\ell} K_Y)}{\tau(P_{\text{exp},\ell})} = \frac{\lambda_Y}{6} \mathcal{A}_{\text{exp},\ell}, \quad (45)$$

with

$$\mathcal{A}_{\text{exp},\ell} := \frac{\int d\xi \Pi_{\text{exp},\ell}(\xi) [\rho_u(\xi; \text{exp}) \mathbf{S}_{\text{exp},u}(\xi) - \rho_d(\xi; \text{exp}) \mathbf{S}_{\text{exp},d}(\xi)]}{\int d\xi \Pi_{\text{exp},\ell}(\xi) \sum_q \rho_q(\xi; \text{exp}) \mathbf{S}_{\text{exp},q}(\xi)}, \quad (46)$$

where the up-class projector covers $\{u, c, t\}$ and the down-class projector covers $\{d, s, b\}$ in the framework’s class assignment. The score-projected asymmetry decomposes naturally as

$$\mathcal{A}_{\text{exp},\ell} = \mathcal{A}_{\text{exp},\ell}^{\text{rate}} + \mathcal{A}_{\text{exp},\ell}^{\text{score}}, \quad (47)$$

where $\mathcal{A}_{\text{exp},\ell}^{\text{rate}} = (\langle \rho_u \rangle - \langle \rho_d \rangle) / \langle \rho_{\text{total}} \rangle$ is the rate-weighted up–down asymmetry (the form obtained by setting $\mathbf{S}_{\text{exp},q} \equiv 1$, which corresponds to a $\sin^2 \theta_{\text{eff}}$ extraction from event counts) and $\mathcal{A}_{\text{exp},\ell}^{\text{score}}$ is the signed angular-template correction, sourced by the deviation of the actual extraction score $\mathbf{S}_{\text{exp},q}$ from a flavor-symmetric flat weighting.

Substrate-internal mechanism for the visibility weight. The two pieces of Equation (47) behave very differently across experiments. The rate-weighted piece $\mathcal{A}_{\text{exp},\ell}^{\text{rate}}$ depends only on PDF luminosities and acceptance support; for two experiments at the same beam configuration and similar acceptance it returns nearly the same value. The score piece $\mathcal{A}_{\text{exp},\ell}^{\text{score}}$, by contrast, picks up the experiment’s signed angular response: it differs between mass-window choices that sample different ratios of γ – Z interference structure, between angular acceptance projections, and between template-fit and angular-coefficient-fit methodologies. The score-projected pairing of Equation (34) is therefore the framework’s substrate-internal mechanism for the visibility weight $w_{\text{exp},\ell}$ to deviate from the isotropic value $1/3$ within the anchor-projection architecture. It does not enlarge the structural content of the anchor pair ($\sin^2 \theta_L, \sin^2 \theta_H$); it specifies how an experiment-and-channel-resolved measurement projects onto that pair.

Tevatron-class corollary (primary). Within the Tevatron class, $r_{\text{exp},\ell} = 0$ and the primary DY shift (41) reduces to

$$\sin^2 \theta_{\text{eff,DY}}(\text{exp}, \ell) = V_\ell^0 + \sigma_{\text{exp},\ell} w_{\text{exp},\ell} \delta_{\text{init}}, \quad |w_{\text{exp},\ell}| \leq 1.$$

The framework predicts that every Tevatron-class measurement central value lies in the anchor-bounded band

$$\sin^2 \theta_{\text{eff,DY}}(\text{exp}, \ell) \in [V_\ell^0 - \delta_{\text{init}}, V_\ell^0 + \delta_{\text{init}}] = [0.230\,06, 0.232\,22]. \quad (48)$$

This is the framework’s primary, substrate-derived bound on Tevatron-class central values, free of any score-projection or row-level coupling input. The lower edge coincides numerically with the leptonic value minus the substrate increment; the upper edge coincides with the leptonic-plus-final-color anchor identified at LEP $A_{\text{FB}}^{0,b}$ in Section 15.10. Within the band, the framework’s structural prediction is a candidate channel sign/order with sign fixed by $\sigma_{\text{exp},\ell}$ and magnitude bounded by δ_{init} ; row-level central values are not predicted by the primary architecture.

The CDF and D0 channel-separated Run II measurements sit consistent with this band. CDF reports $\sin^2 \theta_{\text{eff}}(ee) = 0.23248 \pm 0.00053$ and $\sin^2 \theta_{\text{eff}}(\mu\mu) = 0.23150 \pm 0.00046$; D0 reports $\sin^2 \theta_{\text{eff}}(ee) = 0.23106 \pm 0.00057$ and $\sin^2 \theta_{\text{eff}}(\mu\mu) = 0.23016 \pm 0.00067$. The D0 $\mu\mu$ value at $V_\ell^0 - 0.91 \delta_{\text{init}}$ corresponds to $\sigma = -1$ (Schur-eliminated transport class) with $w \approx 0.91$. The CDF ee central value lies at $V_\ell^0 + 1.24 \delta_{\text{init}}$, 0.5σ above the upper edge of the primary band; this does not falsify the band at the present statistical level (the bound holds within experimental uncertainty), but a refined Tevatron-class re-analysis with central value outside the band at $> 3\sigma$ would falsify the primary architecture (criterion F1 below).

LHC joint-fit class (bound only). Within the LHC joint-fit class, $r_{\text{exp},\ell} > 0$ is allowed and the Pythagorean bound (43) constrains each amplitude individually to $|w|, |r| \leq 1$, with joint constraint $w^2 + r^2 \leq 1$. The framework’s primary statement on joint-fit class measurements is the bound

$$|\sin^2 \theta_{\text{eff,DY}}(\text{exp}, \ell) - V_\ell^0| \leq |\delta_{\text{anchor}}| + |\delta_{\text{projection}}| \leq 2 \delta_{\text{init}} \approx 0.002\,16, \quad (49)$$

with the tighter Pythagorean form $\delta_{\text{anchor}}^2 + \delta_{\text{projection}}^2 \leq \delta_{\text{init}}^2$ holding at the substrate level.

The framework does *not* predict row-level central values for individual joint-fit class measurements such as CMS at $\sqrt{s} = 8$ TeV ($\sin^2 \theta_{\text{eff}}^\ell(\text{combined}) = 0.23101 \pm 0.00053$), LHCb at 13 TeV (0.23147 ± 0.00050), or CMS at 13 TeV (0.23157 ± 0.00031). Reconstructing row-level central values requires modeling the joint-fit machinery (PDF inference from the measured A_{FB} distribution, quark-direction dilution, channel acceptance, rapidity weighting, template covariance structure), and the framework’s substrate-derived primary architecture does not undertake that reconstruction. Predicting joint-fit class central values through the secondary K_Y /score-projection apparatus without reconstructing the joint-fit machinery would re-introduce the score-projection knob under another name and is structurally inadmissible at the v19 level.

Forward-looking falsifiers. The primary DY architecture supports the following falsifiers, ordered by structural priority:

(F1) Tevatron-class anchor-bound test (primary).

Any future neutral-current Drell–Yan extraction in $p\bar{p}$ or comparable valence-dominated kinematics with externally-fixed parton distributions, in any leptonic channel, with central value outside the band $[V_\ell^0 - \delta_{\text{init}}, V_\ell^0 + \delta_{\text{init}}] = [0.230\,06, 0.232\,22]$ at $> 3\sigma$ falsifies the primary anchor-projection bound (48) and, with it, the closure-normalized observation-support theorem (Theorem 6.16) interpretation of the charged-lepton final-state record.

(F2) Externally-fixed-PDF LHC anchor-bound test (primary).

An LHC neutral-current Drell–Yan extraction performed against externally-fixed parton distributions only — without joint A_{FB}/PDF constraint — structurally defines a Tevatron-class measurement at LHC kinematics. Any such extraction with central value outside $[V_\ell^0 - \delta_{\text{init}}, V_\ell^0 + \delta_{\text{init}}]$ at $> 3\sigma$ falsifies the framework’s substrate-derived bound jointly with the Tevatron-class/joint-fit class distinction. Either CMS or ATLAS can in principle perform this re-analysis at the cost of larger PDF uncertainties.

(F3) Tevatron-class channel sign/order test (secondary).

Within the Tevatron-class anchor-bound band, the secondary residual apparatus predicts $\sin^2 \theta_{\text{eff}}(ee) > \sin^2 \theta_{\text{eff}}(\mu\mu)$ for the visibility-positive readout class. A future Tevatron-class measurement showing $\sin^2 \theta_{\text{eff}}(ee) < \sin^2 \theta_{\text{eff}}(\mu\mu)$ at $> 3\sigma$, with both values inside the anchor-bounded band of (F1), falsifies the channel sign/order content of the secondary residual apparatus (specifically the (P3) structural identification of K_Y) but not the primary anchor-projection bound.

(F4) Pythagorean substrate-budget test (secondary).

Any LHC joint-fit class measurement with reconstructed amplitudes violating $w_{\text{exp},\ell}^2 + r_{\text{exp},\ell}^2 > 1$ (equivalently, a central value outside the loose corollary band $[V_\ell^0 - 2\delta_{\text{init}}, V_\ell^0 + 2\delta_{\text{init}}]$ at $> 3\sigma$ for the linear-sum corollary, or a tighter Pythagorean violation when the residual amplitude is structurally accessible) falsifies the Pythagorean substrate-budget bound (43) for the joint-fit class.

The priority ordering reflects the v19 architectural commitment. (F1) and (F2) test the framework’s primary substrate-derived bound and are structurally clean. (F3) tests the secondary residual diagnostic; a future measurement that passes (F1) but fails (F3) would identify the (P3) structural class of K_Y as a deficient residual identification — a continuation-work target on the residual-diagnostic side — without falsifying the primary architecture. The layer-protection falsifier of the factorization theorem (Remark 3.6) remains in force across all classes.

Scope of the present block. The v19 architecture introduced above is a *hierarchical re-placement*, not a full additive layering. The primary content is fixed by three structural objects: the anchor-projection theorem of the present block, the closure-normalized observation-support trace weight $T_{\text{ch}} = 27/4$ from Theorem 6.16, and the operator-algebraic observer-access class distinction (44) between Tevatron-class measurements (where the framework predicts the substrate-bounded band $[V_\ell^0 \pm \delta_{\text{init}}]$ and a candidate channel sign/order at the secondary residual level) and LHC joint-fit class measurements (where the framework provides a Pythagorean substrate-budget bound only and explicitly does not predict row-level central values). The Pythagorean bound (43) is the substrate-internal statement; the simple per-component cap is its derived corollary.

The secondary residual apparatus retains the structural identification of K_Y as the primitive flavor-side access-recovery composition in the (P3) class and the score-projected experiment-and-channel pairing $\mathcal{A}_{\text{exp},\ell}$ as a structural diagnostic. It does not predict row-level central values, and row-level fits of λ_Y against individual measurement values are structurally inadmissible because they would absorb measurement variation into a framework parameter and remove the falsifiability gain of the primary architecture. Closure of λ_Y as a single substrate-internal coupling (one number, not per-experiment) within the (P3) class is identified as a residual-diagnostic continuation work, separate from the row-level fitting that the v19 architecture forbids.

The hadron-collider Drell–Yan extension at the v19 level fixes (a) the framework-internal substrate scale δ_{init} via T_{ch} , (b) the operator-algebraic Tevatron-class / joint-fit class distinction, (c) the Pythagorean substrate-budget bound on amplitudes, and (d) the candidate channel sign/order within the Tevatron class at the secondary residual level. It does not pin row-level

central values for any individual measurement; the primary architecture is bounded and class-distinguishing, not row-level predictive. The conditional theorem of the resolved Drell–Yan section continues to claim $w_{\text{DY}} = 1/3$ in the isotropic limit; the present subsection adds the v19 primary architecture above the isotropic limit, with row-level visibility-weight prediction deliberately removed. Falsification proceeds primarily through (F1)–(F2) above, with (F3)–(F4) testing the secondary residual apparatus and the Pythagorean budget independently.

15.11 Strong-sector adjoint holonomy and electroweak-resolution overlap

The same substrate–cut bookkeeping gives a non-Abelian strong-sector diagnostic when the observer geometry is kept fixed and the internal holonomy block is changed. The observer-access data remain

$$b = 3, \quad d = 4.$$

The number b is the spatial projective multiplicity of the observer cut; it is not a gauge-sector parameter. Moving from the electromagnetic readout to the strong readout therefore does not mean varying the number of projective branches. It means replacing the Abelian charged block by the internal color-holonomy block while keeping the same memory-bearing observer geometry.

The internal block is fixed by the connection-block representation lemma of Section 15.1. A gauge coupling readout is a stiffness readout of internal connection or holonomy variation, not a trace over matter support labels. Infinitesimal non-Abelian connection variations are Lie-algebra-valued and transform by the adjoint action. Hence the primary strong-sector block is the adjoint $SU(3)$ color-holonomy block rather than the fundamental color-support triplet. The fundamental triplet describes color-carrying matter support; it is not the hidden connection block whose stiffness defines the strong gauge readout.

For $SU(3)$, the adjoint internal dimension is

$$d_{\text{int}} = \dim(\text{adj}_{SU(3)}) = 8,$$

and the quadratic stiffness normalization is the adjoint Casimir

$$C_2(\text{adj}_{SU(3)}) = 3.$$

Writing the paired regional disturbance block in observer/internal form,

$$\mathcal{A}_{\text{sub/cut}}^{(3)} = \begin{pmatrix} P & Q \\ R & S \end{pmatrix},$$

the observer block P is governed by the fixed three-branch projective layer, the internal block S is the unrecorded adjoint color-holonomy sector, and Q, R are the disturbance couplings between spatial projective branches and internal color-holonomy modes. Under uniform spatial-branch coupling and Casimir-normalized internal stiffness, the primary strong-sector Schur trace weight is

$$\Delta\chi_{SU(3),\text{adj}} = b \frac{d_{\text{int}}}{C_2(\text{adj})} = 3 \cdot \frac{8}{3} = 8. \quad (50)$$

This value is also stated in Appendix 6.11 as a corollary of the adjoint connection block theorem (Theorem 6.17); $T_s = 8$ does not admit a closure-normalized observation-support derivation, since solving $T_s = b^2 d_{\text{supp}}/d = 8$ for d_{supp} on $(b, d) = (3, 4)$ gives the non-integer $32/9$. The adjoint connection block theorem is the only route to T_s within the framework. This is the adjoint color-holonomy trace seen through the fixed three-branch observer geometry. It is not obtained by changing the observer-access dimension and not by promoting the record-order coordinate to a fourth spatial branch.

The observer cut also contains the record-order coordinate. That coordinate is not a fourth projective branch, because it does not generate independent spatial localization sectors. Its

contribution is instead boundary-like. The elementary inclusion boundary carries one scalar trace unit before memory orientation. In the pre-oriented substrate, the two formal sides of that boundary are trace-equivalent by isotropy and Type II₁ trace uniqueness. A memory-bearing observer selects one oriented side as the recorded boundary. Therefore the observer-visible record-order contribution is one half of the elementary scalar boundary unit,

$$\tau_{\partial\text{rec}} = \frac{1}{2}.$$

The first static strong-sector inverse-coupling diagnostic therefore has two structural contributions,

$$\alpha_{s,\text{stat}}^{-1} \Big|_{SU(3),\text{adj}} = \Delta\chi_{SU(3),\text{adj}} + \tau_{\partial\text{rec}} = 8 + \frac{1}{2} = \frac{17}{2}. \quad (51)$$

The first term is the adjoint color-holonomy trace. The second term is the one-sided non-projective record-boundary trace. This differs from the fine-structure access corrections. In the electromagnetic diagnostic, the subleading terms are Schur access dressings and therefore enter with the negative sign fixed by Schur reduction. The strong-sector one-half term is not such a dressing and carries no access-spectrum factor of π ; it is the scalar boundary contribution of the record-order coordinate after the pre-oriented unit boundary has been split by trace uniqueness and one oriented side has been selected by memory.

At electroweak access resolution, the primary $SU(3)$ adjoint boundary need not remain perfectly isolated from neighboring internal holonomy sectors. The observer's conditional expectation onto the accessible algebra can overlap multiple internal sector boundaries at the same resolution. The first adjacent non-Abelian overlap is the $SU(3)$ – $SU(2)$ overlap.

The reciprocal-overlap rule (1) applies to two simultaneously unresolved internal holonomy sectors. For the strong–weak overlap, the strong primary trace weight is

$$T_s = 8.$$

The weak adjoint block has internal dimension 3 and adjoint Casimir 2. With the same fixed projective multiplicity $b = 3$, its primary trace weight is

$$T_w = 3 \cdot \frac{3}{2} = \frac{9}{2}.$$

Therefore the minimal $SU(3)$ – $SU(2)$ overlap correction is

$$\Delta K_{sw}^{\text{overlap}} = -\frac{1}{T_s T_w} = -\frac{1}{8 \cdot (9/2)} = -\frac{1}{36}. \quad (52)$$

This term is a Schur dressing: it lowers the observer-effective inverse stiffness because the strong readout is not completely isolated from the neighboring weak holonomy boundary at the same access resolution.

The remaining internal sector is the Abelian phase background. Since the color-charged boundary is read through the full minimal charged-sector algebra, it retains the central Abelian phase footprint described in Section 15.3. In the static electroweak-resolution diagnostic, this residual scalar contribution is represented by the inverse Abelian baseline fixed by the fine-structure Schur trace of Section 15.9,

$$+\frac{1}{137}.$$

Combining the primary adjoint trace, the one-sided record-boundary trace, the minimal strong–weak overlap, and the residual Abelian footprint gives the *leading* Schur expansion at the electroweak observer cut,

$$\alpha_s^{-1}(\mu_{\text{EW}}) = 8 + \frac{1}{2} - \frac{1}{36} + \frac{1}{137}. \quad (53)$$

Numerically,

$$\alpha_s^{-1}(\mu_{\text{EW}}) \approx 8.4795. \quad (54)$$

This is the leading-order Schur expansion at the access depth where $SU(3)$, $SU(2)$, and $U(1)$ are simultaneously resolvable as separate boundaries (Section 6.7); it is not a static commitment for all energies, and it is not an accidental match at one human-selected scale. The PDG 2024 world average for the strong coupling at the Z pole, $\alpha_s(M_Z) = 0.1180 \pm 0.0009$ [62], corresponds to $\alpha_s^{-1}(M_Z) = 8.4746 \pm 0.0646$. The leading-expansion prediction $\alpha_s^{-1} = 8.4795$ deviates by $+0.08\sigma$, $\alpha_s(M_Z) = 0.117931 \dots$ vs. measured 0.1180 ± 0.0009 .

Sub-leading terms and the convergence of the Schur expansion. The reciprocal-overlap rule (1) applies to every pair of simultaneously unresolved internal holonomy sectors. At the electroweak observer cut all three pairs — $(SU(3), SU(2))$, $(SU(3), U(1))$, $(SU(2), U(1))$ — are present, and higher-order overlaps among more than two sectors also enter at progressively subleading orders. The natural sub-leading terms suggested by the framework’s machinery are evaluated below.

Sub-leading term	Value	Cumulative $\alpha_s(M_Z)$
Leading 4-term (53)	—	0.117931
+ $SU(3)$ – $U(1)$ Schur: $-1/(T_s T_1)$	-9.12×10^{-4}	0.117944
+ Triple overlap: $-1/(T_s T_2 T_1)$	-2.03×10^{-4}	0.117934

The leading-order formula sits at -0.08σ from the PDG world-average. Including the $SU(3)$ – $U(1)$ pairwise Schur overlap moves the prediction to -0.06σ . Including the triple overlap moves it to -0.07σ . All three sit comfortably inside the experimental one-sigma band. The $SU(2)$ – $U(1)$ pairwise overlap is not added separately because it is already absorbed into the visibility-type Abelian footprint $+1/T_1$ inherited from the fine-structure Schur trace, where it enters with the visibility sign convention; double-counting would violate (29). The prediction is therefore robust against the inclusion question. The framework predicts that the natural sub-leading corrections decay as inverse powers of the large trace weight $T_1 \sim 137$ and are not resolvable until experimental precision improves by roughly an order of magnitude.

Why this matches at the electroweak scale. The four-term formula (53) is precisely the leading Schur expansion at the access depth where $SU(3)$, $SU(2)$, and $U(1)$ are simultaneously resolvable as separate boundaries (Section 6.7). The experimental anchor of that depth is the electroweak scale $\mu \sim M_Z$, with the $SU(3)$ – $SU(2)$ overlap term entering because $SU(2)$ is resolvable and the Abelian footprint entering because $U(1)$ is resolvable. At deep-QCD access depths ($\mu \ll M_{W,Z}$), $SU(2)$ is no longer resolvable as a propagating mode; the framework therefore drops the $-1/(T_s T_2)$ overlap term and recomputes the Schur trace with only $SU(3)$ and $U(1)$ contributions. At deep-UV access depths additional internal sectors may become resolvable and add new overlap terms. The prediction (53) is the leading-order value at the specific access depth corresponding to the electroweak measurement; explicit computation of the access-depth profile is part of the continuation work declared in Section 21.

Rigidity of the strong-coupling prediction. The PDG world-average uncertainty $\sigma_{\alpha_s} = 0.0009$ is large enough that some sub-corrections sit below current experimental resolution. The structural choices that are sharply tested by this prediction are the top-level ones.

Trace-weight ablation on T_s	$\alpha_s(M_Z)$	Deviation
Reference: $T_s = 8$ (adjoint $SU(3)$)	0.117931	-0.08σ
$T_s = 27/4$ (fundamental $SU(3)$)	0.1384	$+22.7\sigma$
$T_s = 3$ (internal dimension only)	0.2913	$+192\sigma$

The connection/support representation rule’s selection of adjoint over fundamental for the strong inverse-coupling readout is discriminated against representation alternatives at the $\sim 23\sigma$ level. The record-order boundary contribution $\tau_{\partial\text{rec}} = 1/2$ is also sharply tested:

Boundary-trace ablation on $\tau_{\partial\text{rec}}$	$\alpha_s(M_Z)$	Deviation
Reference: $\tau_{\partial\text{rec}} = 1/2$	0.117 931	-0.08σ
$\tau_{\partial\text{rec}} = 0$ (no record boundary)	0.125 3	$+8.1\sigma$
$\tau_{\partial\text{rec}} = 1$ (no half-splitting)	0.111 4	-7.4σ

The half-unit value forced by Type II₁ trace uniqueness on a memory-oriented inclusion boundary is therefore sharply tested at the $\sim 8\sigma$ level against the trivial alternatives. By contrast, the two finer corrections (the $-1/36$ overlap and the $+1/T_1$ Abelian footprint) sit just below current experimental precision and would require a $\sim 3\times$ tightening of the world average to be discriminated individually. The framework’s commitment is that the leading-order match at -0.08σ is the visible part of a structurally larger agreement that future precision will continue to test.

Thus the electroweak-resolution strong-sector readout is organized as four leading structural contributions, with a calculable hierarchy of sub-leading terms and a calculable access-depth profile. The interpretation is not temporal fluctuation of virtual particles in a vacuum. The resolution dependence comes from how the observer’s memory-bearing access cut partitions, overlaps, and Schur-reduces internal holonomy sectors at a given access scale.

15.12 Observer-access dimension ablation and substrate–observer duality

The minimal stable observer-access dimension $d = \dim O_{\text{acc}} = 4$ enters the fine-structure formula through the π -powers of the polarization corrections. Holding substrate quantities fixed and varying only the access dimension d produces a sharp ablation of the $3+1$ observer-cut structure.

Dimension-dependent four-part formula. With T_{bare} , a_1 , $3/8$, and P_{ch} held fixed at their substrate-derived values, the polarization corrections at single-channel and minimal projective-cover orders become

$$\Pi^{(1)}(d) = \frac{3}{8\pi^d}, \quad \Pi^{(3)}(d) = \frac{3}{8} \cdot \frac{7}{d} \cdot \frac{1}{\pi^{3d}} = \frac{21}{8d\pi^{3d}}.$$

The dimension-dependent four-part diagnostic value is

$$\alpha_{\text{eff}}^{-1}(d) = 42\pi + \frac{16}{\pi} - \frac{3}{8\pi^d} - \frac{21}{8d\pi^{3d}}.$$

Numerical ablation across $d \in \{2, 3, 4, 5\}$.

d	$\Pi^{(1)}$	$\Pi^{(3)}$	$\alpha_{\text{eff}}^{-1}(d)$	deviation from CODATA
2	3.80×10^{-2}	1.37×10^{-3}	137.0006	3.5×10^{-2}
3	1.21×10^{-2}	2.93×10^{-5}	137.0278	8.2×10^{-3}
4	3.85×10^{-3}	7.10×10^{-7}	137.0359992	8×10^{-10}
5	1.22×10^{-3}	1.83×10^{-8}	137.0387	2.7×10^{-3}

The formula agrees with the CODATA value $\alpha^{-1} = 137.035\,999\,177(21)$ only at $d = 4$, at the 10^{-9} level. The deviations at adjacent access dimensions are large: $d = 3$ misses by 0.008, $d = 5$ misses by 0.003, both many thousand standard deviations from the CODATA value. No nearby access dimension reproduces the measured constant.

Substrate–observer duality at $d = 4$. The dimensional sensitivity is reinforced by an internal duality between the substrate’s trace arithmetic and the observer’s access combinatorics. The bare-trace integer 42 admits the substrate reading $d \cdot T_{\text{bare}}$ and the observer reading $\binom{d}{2}(2^b - 1)$, and these are equal precisely when

$$d \cdot T_{\text{bare}} = \binom{d}{2}(2^b - 1).$$

At $T_{\text{bare}} = 21/2$ and $b = 3$ this reduces to $21d/2 = 7d(d - 1)/2$, that is $21 = 7(d - 1)$, with unique algebraic solution $d = 4$. The first record-projection integer 16 admits a parallel duality. The substrate reading is $d \cdot (a_1\pi)^2$ with $a_1 = 2/\pi$, giving $4d$. The observer reading is the Clifford-grade total $\sum_{k=0}^d \binom{d}{k} = 2^d$. The equality

$$d \cdot (a_1\pi)^2 = 2^d$$

reduces at $a_1 = 2/\pi$ to $4d = 2^d$, with unique solution $d = 4$ among integers $d \geq 2$.

Internal-consistency reading of the 3 + 1 axiom. The two duality equations together establish: at $d = 4$, the substrate trace counting and the observer access combinatorics produce the same integers 42 and 16 that enter the bare junction stiffness, and at no other access dimension does this internal consistency hold. The 3 + 1 observer-access structure is therefore consistent on two reinforcing grounds. First, it is the minimal memory-bearing projective cut: three projective localization directions plus one record-ordering direction. Second, the substrate–observer duality between trace arithmetic and access combinatorics closes only at $d = 4$. The four-part diagnostic value matching CODATA only at $d = 4$ is then an empirical consistency check, not the sole reason for choosing the dimension. The two duality equations are internal-consistency conditions among the framework’s chosen primitives ($T_{\text{bare}} = 21/2$, $a_1 = 2/\pi$, $b = 3$) rather than external selection rules; their closure at $d = 4$ is non-trivial because the equations $21d/2 = 7d(d - 1)/2$ and $4d = 2^d$ each have unique integer solutions, but it would be an overstatement to claim that $d = 4$ is derived from outside the framework. Combined with the substrate-side derivations of $T_{\text{bare}} = 21/2$, $a_1 = 2/\pi$, $3/8$ from $\text{Tr}(Y^2)/\text{Tr}(T_3^2) = 5/3$, P_{ch} from minimal chiral content, the three-branch closure from minimal non-removable holonomy, and the projective-cover trace theorem above, the four-part fine-structure formula is determined by substrate primitives plus the minimal observer cut.

Spatial-access reading of the duality solution. The first duality $21 = 7(d - 1)$ rearranges to $d - 1 = 3 = b$, that is

$$\dim \Sigma_{\text{obs}}^3 = d - 1 = b,$$

which states that the number of spatial access directions of the observer interface equals the minimal projective-closure multiplicity. The duality therefore does more than fix $d = 4$: it identifies the spatial-access dimension of Σ_{obs}^3 with the branching multiplicity of the minimal projective cover. The +1 in the observer-access decomposition $O_{\text{acc}} = \Sigma_{\text{obs}}^3 \oplus \tau_{\text{obs}}$ is the record-ordering coordinate that supplies sequencing but does not participate in the spatial projective closure, and the duality closure $d - 1 = b$ is the algebraic statement of exactly that asymmetry. The empirical 3+1 observer-access decomposition, the minimal non-removable projective closure at $b = 3$, and the substrate–observer trace duality are three different routes to the same access geometry, and they agree at $d = 4, b = 3$.

15.13 Observer-side bare-trace duality applied to the charged sector

The observer-side duality of the bare trace is established in Section 6.10 of the observer-layer chapter: at the minimal stable observer-access dimension $d = 4$ and minimal projective-closure

multiplicity $b = 3$, Theorem 6.13 fixes $T_{\text{bare}} = (d - 1)(2^b - 1)/2 = 21/2$ from the duality $d \cdot T_{\text{bare}} = \binom{d}{2}(2^b - 1)$, with the bare junction stiffness $4\pi T_{\text{bare}} = 42\pi$ as an observer-side prediction independent of any specific substrate content realization. The framework therefore does not claim that the fine-structure diagnostic depends on the Standard Model chiral content as an input. The bare-trace value is fixed by the observer-side duality at the framework's selected access dimension and minimal closure multiplicity; the substrate's chiral content is then constrained to realize this trace value through admissible representations. That the minimal realization coincides with three-generation Standard Model charged content is established in the next subsection as an internal consistency between observer-layer geometry and admissible substrate content, not as an imported assumption.

15.14 Bare-trace content realization: minimal chiral content saturates the duality value

The previous subsection derives $T_{\text{bare}} = 21/2$ from the observer-side duality with $d = 4$ and $b = 3$, leaving the question of which substrate content realizes this trace value. The realization is sharply constrained: only the minimal three-generation chiral charged content with one access-stabilizing doublet at $|Y_H| = 1/2$ saturates the duality target. The following ablation exhibits the rigidity of the realization against alternative content choices.

Generation count. Let N_g denote the number of generations. With one-generation trace $T_{\text{gen}} = 10/3$ and Higgs trace $T_H = 1/2$,

$$T_{\text{bare}}(N_g) = N_g \cdot \frac{10}{3} + \frac{1}{2}, \quad 4\pi T_{\text{bare}}(N_g) = \frac{40\pi}{3} N_g + 2\pi.$$

For $N_g = 1$: $T_{\text{bare}} = 23/6$, giving $4\pi T_{\text{bare}} \approx 48.17$. For $N_g = 2$: $T_{\text{bare}} = 43/6$, giving ≈ 90.06 . For $N_g = 3$: $T_{\text{bare}} = 21/2$, giving $42\pi \approx 131.95$. Only the three-generation case approaches the observed $\alpha^{-1} \approx 137$, and single-generation or two-generation frameworks would miss the fine-structure diagnostic by several tens of inverse-coupling units at the bare-trace level alone.

Higgs doublet count. Let N_H denote the number of minimal access-stabilizing doublets, each with $|Y_H| = 1/2$ fixed by the projective-readout axiom. With three generations,

$$T_{\text{bare}}(N_H) = 10 + \frac{N_H}{2}.$$

For $N_H = 0$: $T_{\text{bare}} = 10$, $4\pi T_{\text{bare}} = 40\pi \approx 125.66$. For $N_H = 1$: $T_{\text{bare}} = 21/2$, $4\pi T_{\text{bare}} = 42\pi \approx 131.95$. For $N_H = 2$: $T_{\text{bare}} = 11$, $4\pi T_{\text{bare}} = 44\pi \approx 138.23$. Only $N_H = 1$ is consistent with the bare-trace step of the diagnostic; a two-Higgs-doublet variant would shift the bare value by approximately 2π , far beyond what the polarization corrections can absorb.

Hypercharge normalization. The trace ratio $\text{Tr}(Y^2)/\text{Tr}(T_3^2) = 5/3$ uses the unnormalized hypercharge convention $(Y_Q, Y_u, Y_d, Y_L, Y_e) = (1/6, 2/3, -1/3, -1/2, -1)$ with $Y_e = -1$ as the unit. Canonical 5/3-normalized hypercharge $Y' = \sqrt{3/5} Y$ would give $\text{Tr}(Y'^2) = (3/5) \text{Tr}(Y^2)$ per generation. Under that rescaling the bare trace would become $(3/5) \cdot 42\pi \approx 79.17$ and the fine-structure diagnostic would fail by a factor of 5/3. The unnormalized substrate convention is therefore the unique one consistent with the diagnostic and is fixed by the recoverability-obstruction equations, not by canonical-coupling preferences.

Conclusion of the trace-content ablation. The bare value $T_{\text{bare}} = 21/2$ survives only under three-generation minimal chiral content with exactly one minimal access-stabilizing doublet at $|Y_H| = 1/2$, in the unnormalized hypercharge convention selected by the recoverability-obstruction equations. None of these are tunable parameters of the calculation. Variations in

any of them shift 42π by amounts much larger than the polarization corrections can absorb, so the diagnostic agreement at the bare-trace level is a separate evidential signal for the minimal chiral content.

15.15 Independent consequences of the projective-readout axiom

The projective-readout axiom does not only enter the fine-structure diagnostic. It carries independent consequences for the charged-sector structure that can be checked against the rest of the framework.

Higgs trace contribution fixed. The minimal access-stabilizing doublet $H = (H^+, H^0) : (1, 2)$ has a hypercharge label $|Y_H|$ that the projective-readout structure fixes uniquely. For an $SU(2)$ doublet, the projective quotient $U(2)/U(1)_{\text{central}} = SU(2)/\mathbb{Z}_2$ requires that the central $U(1)$ phase be single-valued under the doublet representation. This is consistent only when the hypercharge label takes half-integer values, with the minimal nontrivial access-stabilizing assignment $|Y_H| = 1/2$. The resulting Higgs trace contribution is

$$T_H = 2|Y_H|^2 = 2\left(\frac{1}{2}\right)^2 = \frac{1}{2},$$

which is the value used in the bare-trace calculation $T_{\text{bare}} = 10 + 1/2 = 21/2$. The projective readout fixes the $1/2$ in the bare trace rather than leaving it as a convention.

Charged-sector gauge structure as projective symmetry. The projective-readout axiom implies that the observer-effective gauge group of the charged sector is the projective symmetry group of the minimal charged record space. For the minimal three-branch closure with internal algebra $M_3(\mathbb{C})$, this is

$$PU(3) = U(3)/U(1)_{\text{central}} = SU(3)/\mathbb{Z}_3,$$

with the central $U(1)$ factored out by the projective quotient and the remaining \mathbb{Z}_3 the center of $SU(3)$. The projective axiom therefore selects the same gauge structure that the minimal-content argument independently selects through the eightfold internal-automorphism closure, an internal consistency check between the two routes into the charged sector.

Recoverability-obstruction equations as projective single-valuedness. The hypercharge values $(Y_Q, Y_u, Y_d, Y_L, Y_e) = (1/6, 2/3, -1/3, -1/2, -1)$ are fixed in the framework by recoverability-obstruction equations that mirror anomaly cancellation. In the projective-readout language, these equations are the single-valuedness conditions for the projective phase under closed chiral access transformations: the projective record readout must return to itself after a complete loop in the charged-access geometry, which forces precisely the cubic and quadratic trace cancellations. The hypercharge solution is therefore the unique projective single-valued assignment for the minimal chiral content, an independent confirmation of the trace-cancellation result.

Mass ratios as projective invariants. The charged-lepton mass ratios computed in the framework are projective invariants by construction: the singular values of the effective Yukawa bridge $Y_{\text{eff}} = M_{0,R}^{-1/2} B_{LR} M_{0,L}^{-1/2}$ are invariant under the projective rescaling $Y_{\text{eff}} \mapsto e^{i\theta} \lambda Y_{\text{eff}}$. The mass ratios m_n/m_e therefore depend only on the support-deformation labels and on the sectorwise-isotropic baseline M_0 , and not on any overall normalization of the bridge. The projective axiom guarantees that lepton ratios are framework-native predictions in the projective sense, not normalization-dependent diagnostics.

15.16 Charged-lepton ladder

The three stable triple-overlap branches $\mathcal{D}_0, \mathcal{D}_1, \mathcal{D}_2$ are now used as charged identity sectors of the record-preserving access ladder

$$\mathcal{D}_0 \xrightarrow{A_1} \mathcal{D}_1 \xrightarrow{A_2} \mathcal{D}_2.$$

The observer cannot access \mathcal{D}_2 while erasing the record of the path through \mathcal{D}_1 , so the activated set is the prefix set $A_n = \{1, 2, \dots, n\}$, and the strain readout is cumulative: $S_n = \sum_{k=1}^n E_k$. The single-mode energy uses the bi-Laplacian stiffness $L\psi_k = k^2\psi_k, L^2\psi_k = k^4\psi_k$, so $E_k = k^4$.

$$S_n = \sum_{k=1}^n k^4, \quad S_0 = 0, \quad S_1 = 1, \quad S_2 = 17.$$

The ladder mobility-suppression coefficient is a counted ratio of charged holonomy branches to observer access boundaries:

$$C_{\text{ladder}} = \frac{N_{\text{hol}}}{N_{\text{bdry}}} = \frac{3}{2}.$$

Using the corrected observer-effective fine-structure diagnostic,

$$B_{\text{eff}} = C_{\text{ladder}} \alpha_{\text{eff}}^{-1} = \frac{3}{2} \alpha_{\text{eff}}^{-1} = 205.5539998 \dots,$$

and the charged-lepton ladder readout is

$$\frac{m_n}{m_e} = 1 + \frac{3}{2} \alpha_{\text{eff}}^{-1} \sum_{k=1}^n k^4, \quad n = 0, 1, 2.$$

The three values are $m_0/m_e = 1$, $m_1/m_e \approx 206.5540$, $m_2/m_e \approx 3495.42$. Observed values from the Particle Data Group [62]: $m_\mu/m_e \approx 206.7683$, $m_\tau/m_e \approx 3477.15$. Relative errors: -0.10% and $+0.53\%$. These are catastrophic at experimental precision; the ladder formula reproduces the structural form but the underlying mass mechanism has not been correctly identified.

15.17 Charge structure consequences

Without further parameters, the minimal chiral content gives, via $Q_{\text{em}} = T_3 + Y$:

$$Q_u = +\frac{2}{3}, \quad Q_d = -\frac{1}{3}, \quad Q_e = -1, \quad Q_\nu = 0.$$

Valence accounting: $Q_p = +1$, $Q_n = 0$, $Q_H = Q_p + Q_e = 0$. These are charge-algebra consequences of the minimal chiral content.

15.18 What is not recovered

Proton-electron mass ratio. The proton is a confined composite, not a charged ladder mode. Its mass requires a baryonic confinement/binding readout not constructed in the framework.

Low-energy weak mixing angle. $\sin^2 \theta_W = 3/8$ is the substrate-scale unification value, not the measured low-energy value.

Generation count. Whether the same \mathbb{Z}_3 supplies both color and generations is not resolved.

Chiral access splitting. The cut/complement argument supplies $SU(2)$ but not chirality; chirality is asserted as a primitive constraint.

Recoverability-obstruction equation form. The form of the cubic-trace constraints is imported as the framework analogue of anomaly cancellation.

Fine-structure and lepton-ladder values at experimental precision. The single Schur-trace identity of Section 15.9 brings the fine-structure diagnostic to within $\sim 0.03\sigma$ of the CODATA 2022 recommended value $\alpha^{-1} = 137.035\,999\,177(21)$, but does so under five named structural inputs: the substrate variational principle exhibited only as a form, the Dirichlet boundary condition on the access interval, the assertion of chirality as a primitive observer-cut constraint, the import of the cubic-trace recoverability conditions from anomaly cancellation, and the definitional minimality of L_{acc}^{-2} as the primitive access kernel. The framework therefore presents a compressed diagnostic, not a closed substrate derivation; deriving any of the five inputs would tighten the claim. The charged-lepton ladder remains at the 10^{-3} relative-error level. Note also that the CODATA 2022 adjustment reconciles Rb (Morel 2020) and Cs (Parker 2018) atom-recoil measurements that are themselves in $\sim 5\sigma$ tension, and the diagnostic sits near the recommended midpoint rather than near either direct measurement.

15.19 Status and remaining problems

Item	Status
Primitive charged holonomy	Conditionally triple-overlap governed.
Three charged branches	From triple-overlap closure plus primitive scalar-isotropic charged-cell assumption.
Two-state access mixing	$SU(2)$ from cut/complement modulo central phase.
Three-branch holonomy mixing	$SU(3)$ as continuous trace-norm preserving mixing modulo central phase.
Minimal charged algebra	$SU(3) \times SU(2) \times U(1)$ from selection principle.
Minimal chiral content	Determined up to neutral singlets.
Hypercharge values	From recoverability-obstruction cancellation.
Hypercharge normalization	From $\text{Tr}(Y^2)/\text{Tr}(T_3^2)$.
5/3	
Weak angle 3/8	Substrate-scale value only, conditional on $g_1 = g_2$.
Identification with Maxwell sector	Conditional theorem with matched pair $(K_{\text{rec}}^{\text{em}}, M_{\text{rec}}^{\text{em}})$.
Charged-lepton ladder	Structural diagnostic, not at experimental precision.
Fine-structure diagnostic	Bare $\alpha_{\text{junction}}^{-1} = 42\pi + 16/\pi$; with Schur access dressing on a single substrate-cut block operator $\mathcal{A}_{\text{sub/cut}}$, $\alpha_{\text{eff}}^{-1} = 42\pi + 16/\pi - 3/(8\pi^4) - 21/(32\pi^{12}) = 137.0359991763\dots$, agreeing with CODATA 2022 $137.035\,999\,177(21)$ at $\sim 0.03\sigma$. Single Schur-trace diagnostic; substrate uniqueness of the block operator open.
Electric charge pattern	Charge-algebra consequence.
Proton-electron mass ratio	Not recovered; baryonic sector required.
Three generations	\mathbb{Z}_3 double duty not resolved.
Chiral access splitting	Imported as primitive constraint.
Recoverability-obstruction form	Imported as framework analogue of anomaly cancellation.

The framework has not proved the physical charged sector from the substrate alone. It has localized the remaining assumptions inside a single explicit principle: the physical charged sector is the minimal observer-stable internal automorphism sector satisfying central Abelian phase

readout, cut/complement two-state access mixing, three-branch triple-overlap holonomy mixing, chiral doublet/singlet splitting, and global recoverability-obstruction cancellation. Under this principle the minimal charged algebra, the hypercharge pattern, the 5/3 trace normalization, the conditional unification-scale weak angle, the fine-structure diagnostic, its access-polarization correction, and the charged-lepton ladder all follow as one conditional chain through the paired $(K_{\text{rec}}, M_{\text{rec}})$ bookkeeping rather than as independent sector choices.

The remaining substrate-level work consists of four explicit problems. First, derive the primitive scalar-isotropic charged-cell assumption from the substrate. Second, derive the chiral access splitting from substrate/access primitives. Third, derive the specific form and multiplicities of the recoverability-obstruction equations. Fourth, distinguish the \mathbb{Z}_3 producing color from the \mathbb{Z}_3 producing generations. The access-polarization operator has been localized to an explicit Schur-complement trace, but it still deserves ablation against nearby symbolic alternatives. Together these problems define the remaining technical program for the charged-sector layer.

16 Layer allocation and status

The framework has enough moving parts that the status of each claim must be kept separate. The following table is part of the claim discipline of the paper.

Phenomenon	Layer	Status
Conditional-expectation defect	Substrate	Proved as a Type II ₁ local normal form under positivity, differentiability, covariance, and isotropy hypotheses.
Paired regional disturbance $(K_{\text{rec}}, M_{\text{rec}})$	Substrate	Proposed compressed paired generator with structural decomposition $K_{\text{rec}} = K_0 + \kappa_{\bullet} K_{\bullet}$, $M_{\text{rec}} = M_0 + \eta_{\bullet} M_{\bullet}$. Entropy, inertia, gravity-like response, and connection stationarity are readouts of this paired defect geometry through $G_{\text{acc}} = M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2}$.
Shared-Connection principle	Substrate/readout bridge	Derived as a shared-generator result from one paired regional disturbance, timeless path-independence, minimal bookkeeping, and trace uniqueness.
Composition independence	Substrate	Internal sources change defect shape but not the canonical leading paired disturbance geometry.
Recovered thermodynamics	Junction	Recovered only after regional disturbance flux through K_{rec} is paired with observer-access entropy through M_{rec} and modular scale.
Black-hole radiation	Junction/future direction	Horizon-cut paired disturbance thermodynamics: $\delta Q_{\partial\mathcal{B}} = \Theta_{\partial\mathcal{B}} \delta S_{\partial\mathcal{B}}$. Placement and formalization, not a derivation of Hawking's formula.
Light-speed limitation	Junction	Defined as $c_*^2(R) = \lambda_{\min}(M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2}) = \lambda_{\min}(G_{\text{acc}})$, the lowest generalized eigenvalue of the paired operator. Universal value not assigned.
Lorentzian kinematics and invariant speed	Junction	Conditional on a genuine observer-access cone. Correlated finite-cone diagnostic supplies the observer interval $ds_{\text{rec}}^2 = c_*^2 d\tau^2 - d\omega^2$.

Phenomenon	Layer	Status
Newtonian/free-fall motion	mo- Junction	Recovered as least-disturbance record continuation plus weak-limit shared-generator cancellation.
Non-Abelian response	Substrate/connection sector	Selection criterion stated by noncommuting coarsening transitions.
Maxwell, weak-field, and continuum gravity schema	Junction	Recovered as stationary holonomy disturbance, stationary accessible connection response, and conditional Einstein-translation when continuum hypotheses hold. Electromagnetic sector uses matched pair $K_{\text{rec}}^{\text{em}} = g_{\text{eff}}^{-2} K_{\text{rec}}$, $M_{\text{rec}}^{\text{em}} = g_{\text{eff}}^{-2} M_{\text{rec}}$ so that $c_{\text{em}} = c_*$.
Finite graph or script validations	External validation	Useful implementation checks; not premises.
Finite QEC recovery-disturbance diagnostics	Finite trace diagnostic	Stabilizer erasure checks test the Hessian law near exact Knill–Laflamme surfaces; amplitude-damping code tests disturbance tracking away from exact recovery.
Type II ₁ → Type III ₁ observer junction	Substrate/observer bridge	Bridge proposal with modular-density as the finite Type III-like signal and a correlated-state cone-speed diagnostic.

The shortest spine of the framework is

$$\mathfrak{D}_R \Rightarrow (K_{\text{rec}}, M_{\text{rec}}) \Rightarrow \left\{ \begin{array}{l} \text{entropy readout,} \\ \text{inertial support Hessian via } K_{\text{rec}}, \\ \text{gravity-like connection variation,} \\ \text{gauge/field stationarity,} \\ \text{record cone via } G_{\text{acc}} = M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2} \end{array} \right\} \Rightarrow \delta Q_R = \Theta_R \delta S_R,$$

where the first two arrows are substrate-side claims, the readouts are substrate–observer bridge claims, and the final arrow is a junction claim requiring the observer layer. The horizon-cut reading applies the same final arrow to black-hole-like regions by replacing R with $\partial\mathcal{B}$: $\delta Q_{\partial\mathcal{B}} = \Theta_{\partial\mathcal{B}} \delta S_{\partial\mathcal{B}}$.

17 Worked finite trace diagnostics: Petz recovery as paired-regional-disturbance diagnostics

This section gives a finite test of whether the paired $(K_{\text{rec}}, M_{\text{rec}})$ framework computes anything useful in a finite trace setting, rather than merely renaming familiar continuum equations. The constructions remain finite matrix approximants but use the same trace geometry as the Type II₁ normal form and connect conditional-expectation defect to a standard information-theoretic recovery task.

The finite task is erasure recovery for stabilizer codes: the perfect five-qubit code $[[5, 1, 3]]$ [42, 43, 44] with $d_C = 2$, the four-qubit detection code $[[4, 2, 2]]$ with $d_C = 4$, and the six-qubit detection code $[[6, 4, 2]]$ with $d_C = 16$. The progression matters because a claimed Petz–disturbance coefficient should scale with code dimension.

Let $P_C = VV^\dagger$ be the code projector, with V the isometry from logical space to physical Hilbert space. For erased set E , let \mathcal{P}_E^0 be the non-identity Pauli operators on E . The Knill–Laflamme erasure condition [42] is $P_C O_E P_C = \alpha(O_E) P_C$ for $O_E \in \mathcal{P}_E^0$. This is naturally a

trace-defect condition. Define

$$B_a(V) := V^\dagger O_a V - \frac{\text{Tr}(V^\dagger O_a V)}{d_C} I_C, \quad \sigma_E(V) := \sum_{O_a \in \mathcal{P}_E^0} \|B_a(V)\|_F^2.$$

Thus $\sigma_E = 0$ iff the erased subsystem is exactly correctable. In the present language, σ_E is the finite trace defect measuring whether an observer cut that loses E still leaves the encoded record recoverable from the retained algebra.

17.1 Exact threshold as a sanity check

The exact erasure threshold is a sanity check, not the principal result. At $\theta = 0$, $\sigma_E = 0$ is the Knill–Laflamme condition written as a trace defect. The finite computations recover the expected thresholds.

For $[[5, 1, 3]]$ with stabilizers $XZZXI, IXZZX, XIXZZ, ZXIXZ$:

Retained	Erased	Petz fidelity / Defect disturbance
5,4,3	0,1,2	1.0 / ≈ 0 (correctable)
2,1,0	3,4,5	0.25 / 6, 24, 96

For $[[4, 2, 2]]$: correctable for one erasure with infidelity ≈ 0 ; for two erasures fidelity 0.25, defect 12. For $[[6, 4, 2]]$ with $d_C = 16$: correctable for one erasure; two-erasure fidelity 0.25, defect 48. These tables verify the code implementation and normalization.

17.2 Perturbative Petz–paired-disturbance normal form

Let $V_\theta = e^{-i\theta H} V$. At an exactly correctable erasure, $B_a(V) = 0$ for all $O_a \in \mathcal{P}_E^0$. The first variation of Petz entanglement fidelity [46, 47, 48] vanishes, and the leading term is quadratic in the defect variables B_a :

$$1 - F_{\text{Petz}}(V_\theta) = \frac{1}{4d_C} \sum_a \|B_a(V_\theta)\|_F^2 + O(\|B(\theta)\|^3) = \frac{1}{4d_C} \sigma_E(V_\theta) + O(\|D_E\|^3).$$

The predicted ratio is $\sigma_E/(1 - F_{\text{Petz}}) = 4d_C + O(\theta)$. For $d_C = 2$, ratio is 8. For $d_C = 4$, ratio 16. For $d_C = 16$, ratio 64.

The numerical diagnostic confirms:

Code	d_C	Predicted ratio	Observed at $\theta = 10^{-3}$
$[[5, 1, 3]]$	2	8	7.99984 (1 erased), 8.00053 (2 erased)
$[[4, 2, 2]]$	4	16	15.99991
$[[6, 4, 2]]$	16	64	64.00012

The origin of the coefficient is explicit: trace subtraction removes the identity component; Pauli orthogonality turns the second-order defect into a Frobenius sum; the Petz expansion normalizes by the logical trace average through the paired geometry, yielding the factor $1/(4d_C)$. This is the local Hessian statement of the paired regional disturbance at the exact-recovery surface.

17.3 Sector splits: useful diagnostic, limited analogy

The five-qubit erased-pair disturbance admits a Pauli-type split, $\sigma_E = \sigma_E^X + \sigma_E^Y + \sigma_E^Z + \sigma_E^{\text{mixed}}$. The shared-coefficient statement is $1 - F_{\text{Petz}} = \frac{1}{8}(\sigma_E^X + \sigma_E^Y + \sigma_E^Z + \sigma_E^{\text{mixed}}) + O(\|D_E\|^3)$. A closer analogue of physical composition universality would split by deformation origin — logical, stabilizer, gauge-fixing, and generic physical — and check the same coefficient.

17.4 What the example establishes

The example shows that the conditional-expectation defect language and the paired $(K_{\text{rec}}, M_{\text{rec}})$ bookkeeping compute something concrete: near an exact recovery surface, the Petz recovery loss and the trace-defect disturbance have the same Hessian up to a computable normalization. The progression through three codes provides a pressure test: $[[4, 2, 2]]$ gives ratio 16, forcing the coefficient $1/(4d_C)$; $[[6, 4, 2]]$ with $d_C = 16$ confirms the corrected scaling at 64.

17.5 Approximate recovery away from the exact-erasure surface

Modern recovery theory [75, 76, 77] has tightened the Petz framework [46, 47, 48] considerably since the original Petz sufficiency results, producing universal recovery maps and approximate-sufficiency bounds that translate directly into the present paper's $(K_{\text{rec}}, M_{\text{rec}})$ language: stiffness controls the obstruction to exact recovery, susceptibility controls the bound on the recovery error. The four-qubit Leung–Nielsen–Chuang–Yamamoto (LNCY) amplitude-damping code [45],

$$|0_L\rangle = (|0000\rangle + |1111\rangle)/\sqrt{2}, \quad |1_L\rangle = (|0011\rangle + |1100\rangle)/\sqrt{2},$$

serves as an approximate-channel pressure test. For one qubit, the amplitude-damping channel has $A_0 = |0\rangle\langle 0| + \sqrt{1-\gamma}|1\rangle\langle 1|$, $A_1 = \sqrt{\gamma}|0\rangle\langle 1|$. Let $K_s(\gamma) = A_s V$ for $s \in \{0, 1\}^4$. The channel–Knill–Laflamme defect is

$$B_{st}(\gamma) := K_s(\gamma)^\dagger K_t(\gamma) - \frac{\text{Tr}(K_s(\gamma)^\dagger K_t(\gamma))}{d_C} I_C, \quad \sigma_{\text{AD}}^{\text{raw}}(\gamma) := \sum_{s,t} \|B_{st}(\gamma)\|_F^2.$$

Direct expansion gives $\sigma_{\text{AD}}^{\text{raw}}(\gamma) = 2\gamma^2 - 4\gamma^3 + 13\gamma^4 - 20\gamma^5 + 18\gamma^6 - 8\gamma^7 + 2\gamma^8$. The Petz entanglement infidelity has leading expansion $1 - F_{\text{Petz}}(\gamma) = \frac{7}{4}\gamma^2 + O(\gamma^3)$. Hence

$$\frac{\sigma_{\text{AD}}^{\text{raw}}(\gamma)}{1 - F_{\text{Petz}}(\gamma)} \longrightarrow \frac{8}{7}.$$

The coefficient $8/7$ is a structural normalization difference: both quantities are built from the same defect operators B_{st} , but the Petz infidelity weights them through the recovery metric $\mathcal{N}(\sigma)^{-1/2}$ (the operational M_{rec} -side weighting), while $\sigma_{\text{AD}}^{\text{raw}}$ uses the bare Frobenius metric.

Higher-order drift. The raw disturbance and Petz loss do not track each other exactly as functions of γ :

γ	$1 - F_{\text{Petz}}$	$\sigma^{\text{raw}}/(1 - F)$
10^{-4}	1.75×10^{-8}	1.1426
10^{-3}	1.75×10^{-6}	1.1404
10^{-2}	1.75×10^{-4}	1.1190
10^{-1}	1.76×10^{-2}	0.9713

Recovered-channel disturbance identity. The Petz-corrected disturbance is the traceless Frobenius variance of the recovered channel $\mathcal{R}_{\text{Petz}} \circ \mathcal{N}$. For any trace-preserving channel \mathcal{C} on a d -dimensional code space with Kraus operators $\{A_i\}$, define $B_i := A_i - \text{Tr}(A_i)I/d$ and $\sigma_{\mathcal{C}} := \sum_i \|B_i\|_F^2$. Then $\sigma_{\mathcal{C}} = d(1 - F_e(\mathcal{C}))$. Therefore

$$\sigma_{\text{AD}}^{\text{Petz}} = d_C(1 - F_{\text{Petz}}).$$

This identity is exact because the recovered Petz channel is trace-preserving on the code.

Layer-safe interpretation. The lesson is not that an observer freely chooses a disturbance functional. The substrate theorem gives a paired trace-defect normal form $(K_{\text{rec}}, M_{\text{rec}})$. A channel or access map pulls that defect geometry back to a particular defect space, inducing a specific channel metric. The erasure-Pauli metric is orthogonal, unital, and Pauli complete, giving the coefficient $1/(4d_C)$. The amplitude-damping metric is non-unital and Kraus-weighted, giving $8/7$. The framework supports a family of structurally analogous channel-disturbance functionals, one per access/noise model; their leading coefficients are computable from the channel geometry, not arbitrary new substrate couplings.

Status of the finite QEC diagnostic. Pauli-erasure disturbance provides a tangent law near exact Knill–Laflamme surfaces, $1 - F_{\text{Petz}} = \sigma_E/(4d_C) + O(\|D_E\|^3)$. Raw amplitude-damping disturbance provides a channel-specific tangent diagnostic with finite- γ drift. Recovered-channel disturbance is exactly the standard entanglement infidelity written as a traceless Kraus-variance disturbance, and its sector decomposition identifies the branches responsible for residual infidelity. This closes the finite QEC diagnostic as a useful worked example of the paired $(K_{\text{rec}}, M_{\text{rec}})$ framework rather than an additional foundational axiom.

18 Paired record-pair audit: universal stiffness–susceptibility consistency across sectors

The preceding section diagnosed paired regional disturbance through finite Petz recovery in a single QEC channel. This section runs a complementary finite audit on the universal record pair $(K_{\text{rec}}, M_{\text{rec}})$ itself, against the multi-sector consistency the framework’s compression principle demands. The same finite $N = 8$ approximant is used for record-cone recovery, Maxwell-type stationarity, the fine-structure access-polarization correction, and the charged-lepton ladder. The purpose is not to add an independently tuned sector. It is to verify that one paired substrate object can pass several tests simultaneously, and to localize where the freedom in choosing $(K_{\text{rec}}, M_{\text{rec}})$ does and does not survive consistency.

18.1 Meaning of the susceptibility operator

The generalized record cone is controlled by the positive pair $(K_{\text{rec}}, M_{\text{rec}})$, not by a stiffness operator alone. The cone speed is the bottom of the generalized spectrum,

$$c_*^2 = \lambda_{\min} \left(M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2} \right) = \inf_{v \neq 0} \frac{\langle v, K_{\text{rec}} v \rangle}{\langle v, M_{\text{rec}} v \rangle}.$$

K_{rec} is the universal record stiffness: the substrate’s resistance to admissible access or record deformation. M_{rec} is the universal record susceptibility: the record-loading, record-inertia, or access response carried by the same deformation channel. High stiffness raises the record speed; high susceptibility lowers it. A horizon-like region is represented by susceptibility growth: if $M_{\text{rec}} \mapsto s M_{\text{rec}}$ with K_{rec} held fixed, then

$$c_{*,\text{ext}} \sim s^{-1/2} c_*.$$

M_{rec} is therefore not an electromagnetic-only object. It is part of the universal record geometry seen by the observer layer, and every sector — Maxwell-type, classical support motion, charged-sector access polarization, horizon exterior — must inherit the same pair, up to sector-specific scalar coupling factors.

18.2 The identity-susceptibility toy and its limitation

A frequently used finite toy sets

$$K_{\text{rec}} = L_{\text{acc}}, \quad M_{\text{rec}} = I,$$

with the matched electromagnetic prescription of the Maxwell readout,

$$K_{\text{rec}}^{\text{em}} = g_{\text{eff}}^{-2} L_{\text{acc}}, \quad M_{\text{rec}}^{\text{em}} = g_{\text{eff}}^{-2} I.$$

The electromagnetic record cone equals the base cone because the g_{eff}^{-2} factor cancels:

$$(M_{\text{rec}}^{\text{em}})^{-1/2} K_{\text{rec}}^{\text{em}} (M_{\text{rec}}^{\text{em}})^{-1/2} = L_{\text{acc}}.$$

In the finite $N = 8$ audit this gives $c_* = c_{\text{em}} = 3.1256671980$, with the Maxwell-type stationarity equation

$$\mathcal{K}_R^{\text{em}}[A; J] = \frac{1}{2g_{\text{eff}}^2} A^T L_{\text{acc}} A - J^T A$$

yielding $L_{\text{acc}} A = g_{\text{eff}}^2 J$ with residual of order 10^{-15} .

This is internally consistent but it makes $M_{\text{rec}} = I$ a normalization choice rather than a physical substrate response. Such a choice is too narrow if the same paired structure is meant to organize classical support motion, Maxwell-type fields, horizon-like recovery collapse, charged-sector access polarization, and quantum non-separation from one substrate.

18.3 First nontrivial universal susceptibility attempt

A natural nontrivial susceptibility built from substrate ingredients already present in the toy is

$$M_{\text{rec}} = I + \eta \hat{L}_{\text{acc}}^{d-1} + \mu \hat{L}_{\text{acc}}^d + \lambda \hat{P}_{\text{ns}},$$

where the hats denote trace-mean normalization and P_{ns} is a low-rank nonseparable or sector-memory projector. The finite $N = 8$ audit uses the simple first-channel version

$$M_{\text{rec}} = I + 0.25 \hat{L}_{\text{acc}}^{d-1} + 0.25 \hat{L}_{\text{acc}}^d + 0.25 \hat{P}_{c1}.$$

This matrix is symmetric positive definite and genuinely nontrivial:

$$\lambda_{\min}(M_{\text{rec}}) = 1.2750000000, \quad \lambda_{\max}(M_{\text{rec}}) = 4.2587057502, \quad \|M_{\text{rec}} - I\|_F = 3.4268896978.$$

Keeping the matched electromagnetic prescription $K_{\text{rec}}^{\text{em}} = g_{\text{eff}}^{-2} K_{\text{rec}}$, $M_{\text{rec}}^{\text{em}} = g_{\text{eff}}^{-2} M_{\text{rec}}$ still cancels the coupling out of the record cone. With $K_{\text{rec}} = L_{\text{acc}}$ left unchanged, the finite audit gives

$$c_* = c_{\text{em}} = 1.5146208051,$$

with cone matching to numerical precision and Maxwell stationarity still passing.

However, the charged-sector fine-structure correction breaks. The continuum correction of Section 15 used

$$\Pi_{\text{access}} = \frac{3}{8\pi^4},$$

and at finite $N = 8$ this becomes

$$\Pi_{\text{access}}^{(N=8),\text{old}} = \frac{3}{8} \lambda_1(L_{\text{acc}})^{-2} = 0.003928803674,$$

giving a corrected value $\alpha^{-1} = 137.0359208260$. The naive M_{rec} -aware correction would use

$$\Pi_{\text{access}}^{M_{\text{rec}}} = \frac{3}{8} \lambda_1 \left(M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2} \right)^{-2}.$$

With $K_{\text{rec}} = L_{\text{acc}}$ unchanged, the generalized first eigenvalue rises to

$$\lambda_1\left(M_{\text{rec}}^{-1/2} L_{\text{acc}} M_{\text{rec}}^{-1/2}\right) = 2.2946129466,$$

so the correction blows up to $\Pi_{\text{access}}^{M_{\text{rec}}} = 0.071255041188$ and

$$\alpha_{M_{\text{rec}}}^{-1} = 136.9685945885.$$

This is worse than the bare junction value $42\pi + 16/\pi = 137.0398496297$. The lesson is that universalizing M_{rec} alone is incomplete: it changes the record-loading side of the generalized eigenproblem while leaving the stiffness side as a bare access Laplacian. The charged-sector audit then breaks because the access-polarization correction was calibrated against $M_{\text{rec}} = I$.

18.4 Paired universal stiffness

The consistent correction is to universalize the pair, not only the susceptibility. The base object is

$$(K_{\text{rec}}, M_{\text{rec}}),$$

with both entries representing the same substrate record geometry. The finite toy implements this by keeping the nontrivial M_{rec} above and defining a metric-compatible stiffness

$$K_{\text{rec}} = M_{\text{rec}}^{1/2} L_{\text{acc}} M_{\text{rec}}^{1/2}.$$

The generalized record operator is then

$$G_{\text{acc}} = M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2} = L_{\text{acc}}.$$

This does not set $K_{\text{rec}} = L_{\text{acc}}$. It says the universal stiffness is the lift of the access stiffness into the universal record-susceptibility metric. The stiffness matrix is nontrivial in Euclidean coordinates, but its generalized spectrum relative to M_{rec} recovers the access spectrum.

The matched electromagnetic sector remains

$$K_{\text{rec}}^{\text{em}} = g_{\text{eff}}^{-2} K_{\text{rec}}, \quad M_{\text{rec}}^{\text{em}} = g_{\text{eff}}^{-2} M_{\text{rec}},$$

and the Maxwell-type stationarity equation becomes

$$\mathcal{K}_R^{\text{em}}[A; J] = \frac{1}{2g_{\text{eff}}^2} A^T K_{\text{rec}} A - J^T A, \quad K_{\text{rec}} A = g_{\text{eff}}^2 J.$$

In the $N = 8$ paired audit the stationarity residual is of order numerical roundoff, and the record cone returns to

$$c_* = c_{\text{em}} = 3.1256671980,$$

with

$$\max_i |\lambda_i(G_{\text{acc}}) - \lambda_i(L_{\text{acc}})| = 1.48 \times 10^{-12}.$$

This is a finite- N realization of the substrate-side weaker compatibility condition $\text{spec}(M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2}) \rightarrow \text{spec}(G_{\text{acc}})$ stated in Section 3: M_{rec} is nontrivial, K_{rec} is its metric-compatible lift, and the paired generalized spectrum is the access spectrum.

18.5 Fine-structure and lepton-ladder audit after pairing

Under the paired definition, the access-polarization correction no longer blows up because $\lambda_1(G_{\text{acc}}) = \lambda_1(L_{\text{acc}})$. The finite $N = 8$ correction is

$$\Pi_{\text{access}}^{(N=8),\text{pair}} = \frac{3}{8}\lambda_1(G_{\text{acc}})^{-2} = 0.003928803674,$$

and

$$\alpha_{\text{pair},(N=8)}^{-1} = 137.0359208260.$$

This agrees with the finite- N corrected value obtained under $M_{\text{rec}} = I$, now reproduced from a universal pair rather than from a universal susceptibility against a bare stiffness. The continuum limit recovers

$$\alpha_{\text{eff}}^{-1} = 42\pi + \frac{16}{\pi} - \frac{3}{8\pi^4} = 137.035999886\dots$$

of Section 15.

Using the same charged-lepton ladder diagnostic of Section 15.16,

$$\frac{m_n}{m_e} = 1 + \frac{3}{2}\alpha^{-1} \sum_{k=1}^n k^4,$$

the paired finite audit gives

$$\frac{m_\mu}{m_e} = 206.5538812391, \quad \frac{m_\tau}{m_e} = 3495.4159810640.$$

These are structural charged-sector diagnostics, not precision mass predictions, and the precision miss already noted in Section 15 remains.

The finite generalized second-sector energy is

$$E_2 = \left(\frac{\lambda_2(G_{\text{acc}})}{\lambda_1(G_{\text{acc}})} \right)^2,$$

which in the $N = 8$ Dirichlet audit is the discrete value rather than the continuum idealization. In the continuum/index-normalized ladder the idealized value is $E_2 = 2^4 = 16$.

18.6 Audit table

Item	Result	Meaning
Nontrivial M_{rec}	PASS	M_{rec} is SPD and not the identity.
Paired K_{rec}	PASS	$K_{\text{rec}} = M_{\text{rec}}^{1/2} L_{\text{acc}} M_{\text{rec}}^{1/2}$ is SPD and not bare L_{acc} in Euclidean coordinates.
Generalized spectrum	PASS	$M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2}$ recovers the L_{acc} spectrum to 1.48×10^{-12} .
Matched EM cone	PASS	$c_{\text{em}} = c_* = 3.1256671980$; g_{eff} cancels from the record cone.
Maxwell stationarity	PASS	$g_{\text{eff}}^{-2} K_{\text{rec}} A = J$ holds to numerical roundoff.
Bare α^{-1}	STRUCTURAL	$42\pi + 16/\pi = 137.0398496297$.
Continuum correction	LEGACY	$42\pi + 16/\pi - 3/(8\pi^4) = 137.0359998864$.
Naive universal- M_{rec} with bare K_{rec}	FAIL	Gives 136.9685945885 because first-mode loading is counted without matching stiffness.
Paired M_{rec} -aware finite correction	PASS (toy)	Gives 137.0359208260, restoring the finite continuum correction under the universal pair.
Charged-lepton ladder	STRUCTURAL ONLY	$m_\mu/m_e = 206.5538812391$ and $m_\tau/m_e = 3495.4159810640$ in the paired finite audit.

18.7 Status of the paired audit

The paired construction is the correct finite-toy consistency move for the universal record pair. The framework should not make M_{rec} universal while leaving K_{rec} bare. The physical object is the generalized pair $(K_{\text{rec}}, M_{\text{rec}})$, and every sector inherits the same pair up to sector-specific scalar coupling factors such as g_{eff}^{-2} in the electromagnetic amplitude.

The construction should not be oversold. The metric-compatible lift

$$K_{\text{rec}} = M_{\text{rec}}^{1/2} L_{\text{acc}} M_{\text{rec}}^{1/2}$$

is one realization of the substrate-side weaker compatibility condition $\text{spec}(M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2}) \rightarrow \text{spec}(G_{\text{acc}})$ stated in Section 3. It repairs the finite-toy inconsistency and preserves the charged-sector audit, but a paper-level derivation must explain why the universal stiffness must take this form, or must derive an equivalent paired structure from a substrate variational principle. The technical problem is therefore sharper: derive the universal record pair $(K_{\text{rec}}, M_{\text{rec}})$ from the substrate so that classical motion, Maxwell stationarity, finite record cones, horizon-like susceptibility growth, quantum/access non-separation, and the charged-sector access-polarization and lepton-ladder diagnostics all use the same generalized stiffness/susceptibility geometry.

The role of this audit in the fine-structure diagnostic of Section 15 is consistency-side, not derivation-side. It rules out one nearby alternative — universalizing only M_{rec} while leaving K_{rec} bare — as failing the finite consistency test, and confirms that the paired metric-compatible lift restores the continuum correction value. It does not resolve the ansatz $\alpha_{\text{junction}}^{-1} = 4\pi T_{\text{bare}}$, the unnormalized-versus-canonical Y -trace question in $T_{\text{bare}} = 21/2$, the substrate motivation for the Dirichlet boundary condition, or the lepton-ladder precision miss. Those remain in the remaining-problems list of Section 15.

19 Relation to neighboring frameworks

The framework sits near several established approaches. Its proposed contribution is a different bookkeeping discipline: substrate-side structure (carried by the paired $(K_{\text{rec}}, M_{\text{rec}})$ operator) and observer-side reconstruction are kept explicitly separate, and the finite validations track which assumptions are substrate axioms, which are observer ansätze, and which results are constructed consequences.

Page–Wootters. Page–Wootters models [20, 21, 22] recover dynamics from correlations inside a globally stationary quantum state by using an internal clock subsystem. The present framework is close in spirit but does not identify the observer merely with a clock. It gives the observer a record algebra $\mathcal{M}_\lambda \subseteq \mathcal{A}_\lambda$ and derives the access order from record recoverability through the paired generator $G_{\text{acc}} = M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2}$. This record-recoverability order is the main distinguishing bookkeeping move.

Relational quantum mechanics. Relational QM [23] emphasizes that states are relative to observers or systems. The present framework accepts observer-relative state assignment but adds a fixed location map: stable separability, the paired disturbance bookkeeping, stiffness, and connection data are substrate-side; time, collapse, motion, field evolution, and Lorentz kinematics are observer-side.

Connes–Rovelli thermal time. The thermal-time hypothesis [24] extracts a time flow from the modular automorphism group of a state on a von Neumann algebra. The present framework differs in aim: it treats experienced time as record-recoverability order, not as a canonical modular flow. The bandwidth radius and access speed are induced by the lowest generalized eigenvalue of the paired G_{acc} .

Jacobson thermodynamics of spacetime. Jacobson [25, 26] derives the Einstein equation as an equation of state from the Clausius relation applied to local causal horizons. The present paper agrees with the structural moral that gravity-like dynamics may be thermodynamic rather than microscopic. It differs in primitive data: Jacobson’s local horizons are spacetime objects; the present framework’s cuts are algebraic conditional expectations selected by modular stability through $(K_{\text{rec}}, M_{\text{rec}})$. Thermodynamics is recovered only from substrate–observer coupling: $\delta Q_R = \Theta_R \delta S_R$, where δQ_R is paired-disturbance-cost variation across a cut and δS_R is accessible record-entropy variation through the same cut.

Decoherent histories, causal sets, and tensor networks. The decoherent-histories program [27, 28, 29] also locates the operational arrow in record-bearing structure, but stays inside ordinary quantum mechanics on a spacetime background; here the substrate is tenseless and the record order is intrinsic to the observer junction. Causal set theory [30] replaces spacetime with a discrete partial order; the present framework keeps the algebraic side continuous and lets order arise from record inclusion rather than from a primitive causal order. Tensor-network and emergent-geometry constructions [31, 32] build spacetime from entanglement structure; in the present language those constructions can be read as observer-side reconstruction of access geometry from a substrate-side correlation pattern, with the paired $(K_{\text{rec}}, M_{\text{rec}})$ supplying the cone speed.

Crossed-product observer programs. Recent work on observer algebras and the crossed product in semiclassical gravity [17, 18, 19] shows how Type III algebras become Type II once an observer with a clock is included. The bookkeeping here is structurally parallel — a tracial substrate with a non-tracial observer cut — but starts from the substrate side and treats the Type II/III question as a junction question rather than a gravity-input question. The general algebraic-QFT background [68, 69, 70], including the geometric content of modular flow established by Bisognano and Wichmann [66, 67] and the Connes classification of injective factors [65], supplies the operator-algebraic background against which the present framework’s substrate–observer split is formulated.

Operational reconstructions of quantum theory. A separate but related class of reconstructions derives quantum theory itself from operational axioms [83, 84]. The present framework is not a reconstruction of quantum theory: it accepts quantum theory’s observer-side mathematics (states, channels, conditional expectations) and asks what substrate carries it. The two programs are orthogonal but compatible: an operational-reconstruction observer layer could in principle sit on top of the paired-bookkeeping substrate proposed here.

Induced and entropic gravity. The framework’s structural reading of gravity as boundary-access stiffness has a long lineage in the induced-gravity program of Sakharov [71, 72] and in the more recent entropic-gravity proposals [73, 74]. The substrate-internal route taken here differs in primitive data: the paired regional disturbance $(K_{\text{rec}}, M_{\text{rec}})$ is the substrate primitive, rather than a vacuum-fluctuation density or a holographic-screen entropy, and the heat-kernel reduction is performed on the generalized record operator G_{acc} in the spectral-action sense [63, 64, 81] rather than on a matter sector embedded in a background metric.

Reconstruction-class positioning. The framework is not a competitor to the Standard Model or general relativity at the level of Feynman rules, loop diagrams, or local field equations. It is a reconstruction-class proposal in the same architectural genus as the algebraic-QFT crossed-product program [17, 18, 19], causal-set foundations [30], and Connes–Rovelli modular thermal time [24]. It asks whether established phenomenological structures, including some precision observables, can be read as effective consequences of one substrate–observer architecture. The

framework does not replace QFT computation. It claims that several QFT-computed quantities, at their experimental anchor scales, agree with the leading Schur expansion at the corresponding observer cut, and offers the architecture in which that agreement is not coincidental. The QFT calculation of $\sin^2 \theta_{\text{eff},\ell}$ from loop diagrams and the framework’s visibility-lemma readout are two languages for the same observable at one access depth; the framework adds the claim that the latter language is the substrate-native one and that access-depth variation is the substrate-native counterpart of renormalization-group flow.

The value of the framework is diagnostic: it marks which effective laws are substrate-only, which are observer-only, and which are recovered only at their junction. Newtonian mechanics is recovered as least-disturbance record continuation; Maxwell-type equations as stationary holonomy disturbance with the matched pair $K_{\text{rec}}^{\text{em}} = g_{\text{eff}}^{-2} K_{\text{rec}}$, $M_{\text{rec}}^{\text{em}} = g_{\text{eff}}^{-2} M_{\text{rec}}$; weak-field gravity as stationary accessible connection response to paired regional disturbance; and Lorentz kinematics as a conditional invariant-access-cone recovery.

20 Modular flow, sector visibility, and the gravitational coupling

The preceding sections developed the substrate–observer architecture, the paired bookkeeping $(K_{\text{rec}}, M_{\text{rec}})$, the charged-sector diagnostics, and the algebraic necessity program for the gauge sectors. This section collects three further results that the framework reaches by the same operator-algebraic machinery: an access-flow ordering principle that identifies the substrate-internal coordinate on which any framework-derived running coupling must live (Theorem 20.1), a sector-visibility theorem that fixes the functional form of the framework’s analog of the renormalization-group profile (Theorem 20.2), and a structurally complete operator-algebraic identification of Newton’s constant in terms of one substrate-level dimensional scale (Theorem 20.3). Each is established through a closed lemma cascade rather than a postulate.

These results were carried as “open questions” in earlier versions of this manuscript; they are now collected here, between the comparison to neighboring frameworks and the genuinely open continuation work of Section 21. The open continuation work specific to each — numerical determination of the modular generator’s spectrum, of the sector-overlap coefficients, and of the boundary-curvature stiffness κ_{∂} — is listed in Section 21.

20.1 Modular access-flow principle

The framework reads at specific observer cuts, producing values at the access depths of those cuts (Section 21.5). The Standard Model’s $\alpha^{-1}(\mu)$ and $\alpha_s^{-1}(\mu)$ curves, by contrast, are dynamical objects derived from the renormalization-group equations: a continuous running between cuts. The structural question is therefore: what substrate-internal mechanism does the framework provide for the ordering between cuts, and is it consistent with the framework’s commitment to substrate-internal-only readings (no external scale, no external clock, no postulated correspondence)? This subsection identifies the unique substrate-internal mechanism, derives the form the access-depth ordering coordinate must take, and identifies that coordinate with the modular parameter of the framework’s Type III₁ observer cut. It does not derive the QED or QCD beta function; it identifies the structural primitive any framework derivation must use.

Four substrate-axiom conditions on access-flow. An access-flow ordering between observer cuts at distinct invariant resolutions must satisfy four substrate-axiom requirements:

1. *Static substrate.* The Type II₁ substrate carries the tracial state, and the modular automorphism group of a tracial state is trivial. No flow arises from the substrate itself; any non-trivial ordering must arise from the observer cut.

2. *No external scale.* The framework reads only substrate-internal primitives. The ordering of two cuts therefore depends only on dimensionless ratios of invariant resolutions, not on either resolution separately.
3. *Path independence of refinement.* Refining the observer cut from invariant resolution μ_0 to μ_1 and then from μ_1 to μ_2 must be equivalent to refining directly from μ_0 to μ_2 . This is the framework's canonical-lift commitment on refinements (Section 3.3) applied to the access hierarchy.
4. *Algebraic covariance.* The ordering of cuts is implemented by the intrinsic automorphism group of the observer algebra, not by any external structure. The framework's observer cut is a Type III₁ factor (Section 13); the intrinsic one-parameter automorphism group of a Type III₁ factor with a faithful normal state is its modular automorphism group σ_t by Tomita–Takesaki, unique up to inner cocycles for any choice of state.

Theorem 20.1 (Modular access-flow principle). *Under the four conditions above, the access-depth ordering coordinate λ between cuts at invariant resolutions μ, μ_0 satisfies*

$$\lambda = \log(\mu^2/\mu_0^2), \quad (55)$$

and is the modular parameter of the framework's Type III₁ observer cut.

Sketch. Conditions 2 and 3 require the ordering coordinate to be a continuous function λ of the resolution ratio $R = \mu^2/\mu_0^2$ satisfying the Cauchy functional equation $\lambda(R_{01}R_{12}) = \lambda(R_{01}) + \lambda(R_{12})$. The unique continuous solution on $\mathbb{R}_{>0}$ is $\lambda(R) = c \log R$ for constant c ; the canonical normalization $c = 1$ (one access-depth unit equals one e-fold of invariant resolution) gives Equation (55). Condition 1 says the ordering does not arise from the tracial substrate. Condition 4 says it is implemented by the intrinsic automorphism group of the observer algebra; on a Type III₁ factor this group is the modular automorphism group σ_t generated by any faithful normal state, with the choice of state determining the flow up to inner cocycle. The two together identify the access-depth coordinate λ with the modular parameter t of σ_t , with the canonical normalization fixing the proportionality between the two. \square

The principle has two structural consequences.

(C1) Modular parameter and log-scale variable are the same object. The framework does not posit a separate modular time and renormalization-group scale and then map between them. The modular parameter of the Type III₁ observer cut, which generates the framework's intrinsic flow at the observer junction, is the same object as $\log(\mu^2/\mu_0^2)$ in renormalization-group flow. This is a structural identification forced by the four substrate axioms, not a postulated correspondence.

(C2) Reference scale μ_0 is sector-dependent. The four axioms fix the form $\log(\mu^2/\mu_0^2)$ but not the choice of μ_0 . The natural sector-specific reference is the lowest invariant resolution present in that sector's framework reading: $\mu_0 = m_e$ for the QED sector (the electron threshold), $\mu_0 = \Lambda_{\text{QCD}}$ or a comparable infrared scale for the QCD sector. Identifying these reference scales substrate-internally is part of the continuation-work program below.

First falsifiable target. The empirical QED running between Thomson and the Z -pole is

$$\Delta\alpha^{-1} := \alpha^{-1}(M_Z) - \alpha^{-1}(0) \approx 127.95 - 137.036 \approx -9.09,$$

integrated over the access-depth interval

$$\Delta\lambda = \log(M_Z^2/m_e^2) \approx 24.2$$

at $\mu_0 = m_e$. Any substrate-derived expression for $d\alpha^{-1}/d\lambda$ produced by the framework's continuation-work program must, when integrated over this interval with the appropriate threshold-counting structure, reproduce $\Delta\alpha^{-1} \approx -9.09$ within experimental uncertainty. This is the first concrete falsifiable target of the modular-flow program. A second target is the QCD running between Λ_{QCD} and M_Z , integrated over the corresponding access-depth interval. Both targets are structurally well-defined; neither is derived by the present paper.

Continuation-work deliverables. Theorem 20.1 identifies the substrate-internal coordinate on which the framework's running-coupling derivation must live; it does not constitute the derivation. The remaining items, in priority order, are:

1. Construction of the sector turning-on function $f_{\text{sector}}(\lambda)$. The formal definition is supplied by Theorem 20.2 of Section 20.2 as the normalized modular spectral visibility $f_s(\lambda) = \tau(P_s \chi_{(-\infty, \lambda]}(K_{\text{mod}}))/\tau(P_s)$, where P_s is the substrate-derived sector support projection and K_{mod} is the modular generator of the observer cut. The function is framework-internal (no fitted sigmoid or phenomenological interpolation), and its derivative is the modular spectral density $\rho_s(\lambda)$ of the sector. What remains as continuation work is the explicit computation of K_{mod} for the framework's observer junction, which would yield the numerical $\rho_s(\lambda)$ profile.
2. Derivation of $d\alpha^{-1}/d\lambda$ and $d\alpha_s^{-1}/d\lambda$ from substrate primitives, with leading-order targets the Standard Model one-loop coefficients at each threshold-crossing access depth.
3. Higher-loop coefficients from higher-order substrate Schur expansions.
4. Substrate-internal derivation of the sector-specific reference scales μ_0 .
5. Verification that the integrated running between the framework's existing readings (Thomson and electroweak access cuts) reproduces $\Delta\alpha^{-1} \approx -9.09$ and the analogous QCD target within experimental uncertainty.

Status. The present subsection does not close the gap identified at the end of Section 21.5. It identifies the substrate-internal primitive — the modular parameter of the Type III₁ observer cut — that closure would use, derives the substrate-axiom form $\lambda = \log(\mu^2/\mu_0^2)$ of the access-depth coordinate, identifies the first falsifiable target, and lists the continuation-work deliverables. The framework's claim is that the structural route to running couplings is unique under the four substrate axioms: any framework-internal derivation of running couplings must use the modular parameter of the Type III₁ observer cut as its scale variable, and any deviation from this form would violate substrate-axiom requirements the framework already has in force. Whether the continuation-work program completes successfully is open; whether it could close to anything other than modular flow on the Type III₁ observer cut is structurally forbidden.

20.2 Modular sector-visibility functions

Theorem 20.1 of the previous subsection identifies the access-depth coordinate λ with the modular parameter of the Type III₁ observer cut and forces the form $\lambda = \log(\mu^2/\mu_0^2)$. Item (1) of the continuation-work list there asked for the sector turning-on function $f_{\text{sector}}(\lambda)$ encoding how each framework sector contributes to a reading at access depth λ . The present subsection supplies the formal definition: $f_{\text{sector}}(\lambda)$ is the normalized modular spectral visibility of the sector under the Type III₁ observer cut, defined as a trace ratio of operator-algebraic projections. The construction is framework-internal — not a fitted sigmoid, not a phenomenological interpolation, not a time evolution of the substrate — and it gives the framework's analog of the renormalization-group derivative as a sum of modular spectral densities weighted by substrate-derived Schur coefficients.

Substrate decomposition into sectors. The substrate record algebra decomposes into sector-supported pieces

$$R = R_0 + \sum_s R_s + \sum_{s < t} R_{st} + \dots, \quad (56)$$

where R_0 is the sector-neutral base record, R_s is the pure contribution from sector s (s ranging over the Abelian sector $U(1)_Y$, the weak sector $SU(2)$, the strong sector $SU(3)$, the charged final-support sector, the color final-support sector, and analogous primitive supports), R_{st} is the contact contribution between sectors s and t , and higher-order terms encode higher-order Schur/modular overlaps. Each sector has a support projection P_s in the resolved retrieval algebra, with finite substrate trace $\tau(P_s) > 0$, and analogously for overlap projections P_{st} .

Modular access cut. Let $K_{\text{mod}} = -\log \Delta_\omega$ be the modular generator of the Type III₁ observer cut, where Δ_ω is the modular operator associated with the framework's substrate-derived faithful normal state ω . The access-depth cut at modular parameter λ is represented by the spectral projector

$$Q(\lambda) := \chi_{(-\infty, \lambda]}(K_{\text{mod}}), \quad (57)$$

which selects all modular modes of depth at most λ in the modular spectral decomposition of the observer cut. The family $\{Q(\lambda)\}_{\lambda \in \mathbb{R}}$ is monotone non-decreasing in λ and converges to the identity on the cut as $\lambda \rightarrow \infty$. On the Type III₁ observer cut, K_{mod} has continuous spectrum \mathbb{R} , and $Q(\lambda)$ is a non-trivial spectral projection for every λ .

Theorem 20.2 (Modular sector visibility). *Let P_s be the support projection for substrate sector s in the resolved retrieval algebra, and K_{mod} the modular generator of the Type III₁ observer cut. The access-depth visibility of sector s at modular depth λ is*

$$f_s(\lambda) := \frac{\tau(P_s \chi_{(-\infty, \lambda]}(K_{\text{mod}}))}{\tau(P_s)}, \quad (58)$$

where τ is the substrate trace lifted to the Type II _{∞} core of the Type III₁ observer cut via the modular crossed product (Section 13). Then:

1. $0 \leq f_s(\lambda) \leq 1$ for all $\lambda \in \mathbb{R}$;
2. f_s is monotone non-decreasing in λ when the access family $\{Q(\lambda)\}$ is nested;
3. the derivative is the normalized modular spectral density of the sector,

$$f'_s(\lambda) = \rho_s(\lambda) := \frac{\tau(P_s \delta(\lambda - K_{\text{mod}}))}{\tau(P_s)}. \quad (59)$$

The corresponding statements hold for overlap visibility functions $f_{st}(\lambda)$ with P_s replaced by P_{st} .

Sketch. (1) Boundedness. From $0 \leq \chi_{(-\infty, \lambda]}(K_{\text{mod}}) \leq \mathbf{1}$ in the operator ordering, the positive square root gives $P_s \chi_{(-\infty, \lambda]}(K_{\text{mod}}) P_s \leq P_s$, and faithfulness of τ yields $0 \leq \tau(P_s Q(\lambda)) \leq \tau(P_s)$.

(2) Monotonicity. For $\lambda_1 \leq \lambda_2$ in nested cuts, the spectral projection family satisfies $\chi_{(-\infty, \lambda_1]}(K_{\text{mod}}) \leq \chi_{(-\infty, \lambda_2]}(K_{\text{mod}})$ as projections, and positivity of τ preserves the inequality.

(3) Derivative. Using the spectral decomposition $K_{\text{mod}} = \int \lambda' dE(\lambda')$, the cumulative spectral measure $\sigma_s(\lambda) = \tau(P_s \chi_{(-\infty, \lambda]}(K_{\text{mod}}))$ has derivative $d\sigma_s/d\lambda$ that, by definition of the spectral measure, equals the formal pairing of P_s with $\delta(\lambda - K_{\text{mod}})$ under τ . Normalizing by $\tau(P_s)$ gives Equation (59). \square

Cut-dependent inverse-coupling readout. A framework-internal cut-dependent readout decomposes as

$$G(\lambda) = G_{\text{base}} + \sum_s f_s(\lambda) \Delta_s + \sum_{s < t} f_{st}(\lambda) \Delta_{st} + \dots, \quad (60)$$

where G_{base} is the cut-independent base reading and Δ_s, Δ_{st} are the framework's substrate-derived sector and overlap coefficients (the $T_1, T_2, T_s, T_{b,\text{col}}, w_Y$, and inter-sector overlaps assembled in Appendix 6.11). Differentiating in λ and using Theorem 20.2 gives the framework's analog of the renormalization-group derivative

$$\frac{dG}{d\lambda} = \sum_s \rho_s(\lambda) \Delta_s + \sum_{s < t} \rho_{st}(\lambda) \Delta_{st} + \dots. \quad (61)$$

With the identification $\lambda = \log(\mu^2/\mu_0^2)$ from Theorem 20.1, Equation (61) is the framework's analog of the renormalization-group derivative $dG/d\log\mu^2$. The right-hand side is a sum of modular spectral densities weighted by substrate-derived Schur coefficients — the framework's structural form of a beta function. Both factors on the right are framework-internal: the spectral densities $\rho_s(\lambda)$ are determined by the modular generator K_{mod} of the observer cut, and the coefficients Δ_s, Δ_{st} are the framework's existing trace-weight content.

Boundary conditions at the framework's existing access cuts. The Thomson and electroweak readings of Section 15.9 and Section 15.11 are boundary values of the sector visibility functions at two specific access depths λ_0 (Thomson) and λ_Z (electroweak):

Thomson cut (Abelian boundary). At λ_0 the Abelian sector is fully resolved and the non-Abelian and inter-sector contributions are absent from the QED reading,

$$f_Y(\lambda_0) = 1, \quad f_2(\lambda_0) = f_3(\lambda_0) = f_{Y2}(\lambda_0) = f_{Y3}(\lambda_0) = f_{23}(\lambda_0) = 0.$$

The resulting reading is $\alpha^{-1}(0) = T_1 = 137.035\,999\,176\dots$

Electroweak cut (multi-sector boundary). At λ_Z the relevant electroweak, strong, and inter-sector contributions are resolved with their cut-specific visibilities,

$$f_Y(\lambda_Z) = 1, \quad f_2(\lambda_Z), f_3(\lambda_Z), f_{Y2}(\lambda_Z), f_{Y3}(\lambda_Z), \dots \in [0, 1],$$

with each $f_s(\lambda_Z)$ encoding partial or full resolution of sector s at the electroweak cut. The resulting reading is $\alpha_s^{-1}(M_Z) = 8.4795$.

The continuous curves of $f_s(\lambda)$ between λ_0 and λ_Z are determined by the modular spectral profile of K_{mod} on the framework's observer cut. Explicit computation of K_{mod} has not been carried out in the present paper.

Status. Theorem 20.2 closes item (1) of the continuation-work list of Section 20.1: the formal definition of $f_{\text{sector}}(\lambda)$ as a normalized modular spectral visibility is in place. The remaining items are unchanged in form but now have concrete operator-algebraic targets:

- explicit computation of K_{mod} for the framework's observer junction, which gives the numerical spectral profile $\rho_s(\lambda)$ for each sector;
- identification of the sector and overlap coefficients Δ_s, Δ_{st} in Equation (60) from the framework's existing Schur-expansion content, with the Standard Model one-loop coefficients b_0 as the leading-order targets;
- verification that integration of Equation (61) between λ_0 and λ_Z reproduces $\Delta\alpha^{-1} \approx -9.09$ within experimental uncertainty.

The structural content the present subsection adds is that $f_{\text{sector}}(\lambda)$ is not a free function to be fitted: it is determined by the substrate-derived sector projection P_s , the framework's modular generator K_{mod} , and the substrate trace τ , with the form fixed by Equation (58). The framework's beta-function analog (61) is therefore also not a fitted function: it is a sum of modular spectral densities weighted by substrate Schur coefficients, with both factors fixed framework-internally once K_{mod} is computed.

20.3 Gravitational access and the unified modular-flow architecture

The modular flow architecture established in the previous two subsections gives a substrate-internal access-depth coordinate λ (Theorem 20.1) and sector visibility functions $f_{\text{sector}}(\lambda)$ (Theorem 20.2) governing how Standard Model couplings depend on the observer cut. The same architecture admits a natural extension to gravitational access without introducing additional substrate primitives. This subsection identifies the structural unification: the same modular flow coordinate λ orders both gauge-coupling resolution and gravitational area resolution, with G_N entering not through a new sector visibility function but through dimensional analysis of the dimensionless area-resolution ratio at the observer cut. The framework does not derive G_N or M_{Pl} from substrate primitives; it claims that gravitational access is governed by the same single underlying modular-flow effect that governs gauge access.

The 3 + 1 split: spatial support and modular ordering. The framework's observer readout is 3 + 1-dimensional but does not interpret the four directions as four spacetime coordinates in the usual sense. The three spatial dimensions are the framework's spatial support primitives, providing the boundary-area resolution at the observer cut. The additional ordering parameter is the modular access depth λ of Theorem 20.1, not a fourth substrate dimension; λ labels the refinement of observer cuts through the static substrate, not an evolution within the substrate. Time, like sector running, is emergent from modular flow at the observer junction. This is consistent with the framework's existing position that the substrate is timeless and that observer-effective time emerges from the modular automorphism of the Type III₁ observer cut of Section 13.

Area-like resolution variable. In a substrate with three spatial support dimensions, a boundary cut is codimension-one in space, with boundary-area resolution measured by an invariant length-squared scale. The framework's reference observer cuts have invariant resolutions μ^2 (dimension $[\text{length}]^{-2} = [\text{energy}]^2$ in natural units) ranging from the Thomson resolution at small μ to the electroweak resolution at $\mu = M_Z$. The modular access depth $\lambda = \log(\mu^2/\mu_0^2)$ is the framework-internal logarithmic ordering of these resolutions. The spatial support explains why the natural resolution variable is area-like (μ^2 rather than μ); the modular ordering explains why the refinement parameter is logarithmic (λ in μ^2).

Gravitational coupling as dimensionless area-resolution ratio. The gravitational coupling G_N has dimension $[\text{length}]^2 = [\text{energy}]^{-2}$ in natural units; equivalently $M_{\text{Pl}}^2 = 1/G_N$ has dimension $[\text{energy}]^2$. The natural dimensionless combination with the framework's boundary-area resolution μ^2 is

$$\alpha_G(\mu) := G_N \mu^2 = \frac{\mu^2}{M_{\text{Pl}}^2}, \quad (62)$$

the ratio of the boundary-area resolution at the cut to the Planck area. This dimensionless gravitational coupling is determined by the same μ^2 that characterizes the observer cut for gauge readouts, with G_N supplying the external dimensional reference M_{Pl} . No new substrate primitive is required: α_G is built from the boundary-area resolution that the cut already carries and the dimensional input G_N that the framework treats as external.

Unified modular-flow architecture. The same access-depth coordinate λ that orders the gauge-coupling sector visibility functions $f_s(\lambda)$ also orders the gravitational coupling $\alpha_G(\mu)$. With the gravitational reference scale $\mu_0 = M_{\text{Pl}}$,

$$\lambda = \log(\mu^2/M_{\text{Pl}}^2), \quad \alpha_G(\lambda) = e^\lambda, \quad \alpha_G^{-1}(\lambda) = e^{-\lambda}. \quad (63)$$

Gravitational coupling has exponential dependence on the modular access depth, in contrast with the logarithmic dependence that gauge couplings exhibit through their sector visibility functions $f_s(\lambda)$. The exponential form is dimensional, fixed by G_N being a dimensionful constant; the logarithmic form of λ in μ^2 is the substrate-axiom form derived in Theorem 20.1. Both readings live on the same modular flow coordinate. They differ in whether the coupling is dimensionless and runs perturbatively (gauge sectors) or carries a dimensional reference and scales as the boundary-area ratio itself (gravitational sector).

Holographic and area-resolution interpretation. The dimensionless combination $G_N\mu^2 = \mu^2/M_{\text{Pl}}^2$ that fixes the gravitational coupling at the framework’s observer cut coincides with the dimensionless area-resolution ratio that enters holographic descriptions of black-hole entropy and area-law entropy bounds. The Bekenstein–Hawking entropy of a horizon of area A is $S_{\text{BH}} = A/(4G_N)$, equivalently $A \cdot \mu^2/(4\alpha_G)$ at boundary-area resolution μ^2 , so that one unit of area at resolution μ^2 contributes $1/(4\alpha_G)$ to the entropy. The framework’s substrate-internal interpretation is that gravitational coupling is fundamentally an area-resolution effect of the observer cut: at $\mu = M_{\text{Pl}}$, the boundary area is resolved at Planck-scale cells and $\alpha_G = 1$; below the Planck scale ($\mu < M_{\text{Pl}}$), $\alpha_G \ll 1$ and the boundary cut has not resolved enough area for gravitational coupling to dominate the framework’s readouts. This is consistent with the framework’s Einstein–Hilbert continuation-work program of Section 12.2, which seeks to derive Einstein’s field equations from the spectral continuum limit of the generalized record operator G_{acc} .

What the unification claims and does not claim. The structural unification claim of this subsection is that gravitational access and gauge access in the framework are governed by a single underlying effect: modular flow on the Type III₁ observer cut, with the single access-depth coordinate $\lambda = \log(\mu^2/\mu_0^2)$. Gauge couplings depend logarithmically in this coordinate through the sector visibility functions of Theorem 20.2; gravitational coupling depends exponentially on the same coordinate through the dimensionless ratio $\alpha_G = G_N\mu^2$ with reference $\mu_0 = M_{\text{Pl}}$. The framework does not numerically derive G_N or M_{Pl} in the present paper; both quantities are structurally tied to the framework’s paired bookkeeping by Theorem 20.3 of the next subsection (Section 20.4), which identifies $G_N = 1/(16\pi C_1)$ with C_1 a substrate-internal spectral coefficient of $(K_{\text{rec}}, M_{\text{rec}})$ at boundary modes, but the leading coefficient C_1 is not computed in closed form here. The framework does not derive Einstein’s field equations from $\alpha_G(\mu)$; that derivation is the framework’s existing continuation-work program for the spectral continuum limit of the generalized record operator. What the framework claims is that the structural ordering of gravitational access cuts is the same modular flow that orders gauge access cuts: one substrate-internal flow effect, not two parallel mechanisms.

Status. The unification stated here is a structural identification, not a derivation. It rests on three facts already established in the framework: (i) the same Type III₁ observer cut underlies both gauge readouts (Section 15) and gravitational readouts (Section 12.2); (ii) the access-depth coordinate is uniquely $\lambda = \log(\mu^2/\mu_0^2)$ under the substrate axioms of Theorem 20.1; (iii) dimensional analysis of G_N against the boundary-area resolution μ^2 produces the dimensionless coupling $\alpha_G = G_N\mu^2$ as the unique combination respecting the framework’s no-external-scale axiom up to the external dimensional input M_{Pl} . The structural payoff is that the framework’s modular-flow architecture admits gravitational access as another readout on the same underlying

coordinate, supporting the framework’s broader claim that quantum, gauge, and gravitational structures are observer-effective readouts of one substrate–observer architecture rather than parallel mechanisms requiring independent foundational treatments.

20.4 Boundary-stiffness Planck coefficient theorem

The previous subsection unified gauge and gravitational access on the same modular flow coordinate λ but explicitly carved out G_N and M_{Pl} as external dimensional inputs not derived by the framework. This subsection refines that carve-out: G_N is structurally determined by a specific spectral coefficient C_1 of the framework’s paired bookkeeping $(K_{\text{rec}}, M_{\text{rec}})$ at boundary-support modes, conditional on the framework’s existing Einstein–Hilbert continuation-work hypotheses (locality and diffeomorphism-covariance of the long-wavelength limit). The framework therefore has a candidate derivation of G_N from substrate primitives that is not yet a numerical derivation — the leading coefficient C_1 has not been computed in closed form — but is structurally complete in the sense that G_N is no longer a free external parameter once the existing reduction-target hypotheses are verified.

What deriving G_N requires. In natural units $G_N = 1/M_{\text{Pl}}^2$, equivalently $M_{\text{Pl}}^2 = 1/G_N$. Deriving G_N therefore means deriving an absolute access-area scale $A_* := G_N = 1/M_{\text{Pl}}^2$. The framework needs a structural object playing the role of a minimal nontrivial boundary-refinement unit, whose stiffness fixes this absolute area scale. No such object can appear among the gauge-sector trace weights $T_1, T_2, T_{b,\text{col}}, T_s, w_Y$ of Appendix 6.11, because those are dimensionless ratios derived from the substrate’s $(b, d) = (3, 4)$ combinatorics and carry no absolute scale. The framework’s natural candidate for the boundary-refinement stiffness is the paired bookkeeping $(K_{\text{rec}}, M_{\text{rec}})$ projected to long-wavelength boundary-deformation modes.

Gravity as universal boundary-access stiffness. The framework’s structural reading of gravitational response is that gravity is the cost of resolving an observer-access boundary refinement. Where gauge couplings are sector-visibility readouts ordered by the modular flow coordinate λ through the visibility functions $f_s(\lambda)$ of Theorem 20.2, gravity is the universal stiffness of boundary refinement itself: the response of the paired bookkeeping $(K_{\text{rec}}, M_{\text{rec}})$ to a long-wavelength deformation of the observer cut. Newton’s constant measures the inverse of this stiffness, $G_N \sim 1/K_{\text{grav}}$; the Planck mass squared is the corresponding stiffness scale, $M_{\text{Pl}}^2 \sim K_{\text{grav}}$. The candidate theorem below makes this stiffness identification operator-algebraically precise.

Theorem 20.3 (Boundary-stiffness Planck coefficient). *Let P_∂ project the paired record bookkeeping $(K_{\text{rec}}, M_{\text{rec}})$ onto long-wavelength observer-boundary deformation modes, and define the boundary-projected paired operators*

$$K_\partial := P_\partial K_{\text{rec}} P_\partial, \quad M_\partial := P_\partial M_{\text{rec}} P_\partial. \quad (64)$$

For a boundary metric perturbation mode h_q at four-momentum q , define the boundary stiffness response

$$\Pi_{\text{grav}}(q) := \frac{\langle h_q, K_\partial h_q \rangle}{\langle h_q, M_\partial h_q \rangle}. \quad (65)$$

Suppose the long-wavelength continuum limit of $\Pi_{\text{grav}}(q)$ is local and diffeomorphism-covariant. Then the leading nontrivial curvature term in the corresponding observer-effective action is Einstein–Hilbert,

$$S_{\text{eff}}[g] = \frac{M_*^2}{2} \int d^4x \sqrt{-g} R + \dots, \quad (66)$$

with reduced Planck stiffness

$$M_*^2 = 2C_1, \quad C_1 := \lim_{q \rightarrow 0} \frac{\Pi_{\text{grav}}(q)}{q^2}, \quad (67)$$

and Newton's constant

$$G_N = \frac{1}{8\pi M_*^2} = \frac{1}{16\pi C_1}. \quad (68)$$

Sketch via four-step lemma chain. The theorem decomposes into a four-step lemma chain that makes the derivation structure explicit, separating definitional steps (Lemmas 20.4, 20.5, 20.7) from the single conditional step (Lemma 20.6).

Lemma 20.4 (Boundary restriction). *The gravitational access sector is the restriction of the paired record bookkeeping $(K_{\text{rec}}, M_{\text{rec}})$ to observer-boundary deformation modes. Explicitly, $K_{\partial} = P_{\partial} K_{\text{rec}} P_{\partial}$ and $M_{\partial} = P_{\partial} M_{\text{rec}} P_{\partial}$ where P_{∂} projects onto long-wavelength observer-boundary deformation modes (transverse geometric perturbations $h_{ab}(x)$ with gauge and coordinate artifacts projected out). The framework's gravitational object is not a new coupling but the boundary-mode specialization of the existing paired bookkeeping.*

This lemma is definitional within the framework's operator-algebraic structure: it identifies the gravitational sector as a specific projection of $(K_{\text{rec}}, M_{\text{rec}})$ rather than as an independent operator.

Lemma 20.5 (Gravitational response). *The long-wavelength gravitational stiffness is the q^2 coefficient in the boundary-deformation response $\Pi_{\text{grav}}(q) = \langle h_q, K_{\partial} h_q \rangle / \langle h_q, M_{\partial} h_q \rangle$. The local expansion at small q*

$$\Pi_{\text{grav}}(q) = C_0 + C_1 q^2 + C_2 q^4 + \dots$$

separates the constant term C_0 (cosmological/volume-type stiffness) from the curvature stiffness C_1 and higher derivative-sensitive coefficients C_2, \dots . The leading curvature coefficient is

$$C_1 = \lim_{q \rightarrow 0} \frac{\Pi_{\text{grav}}(q) - \Pi_{\text{grav}}(0)}{q^2}, \quad (69)$$

with the explicit subtraction of $\Pi_{\text{grav}}(0)$ separating cosmological-constant stiffness from curvature stiffness.

This lemma is also definitional: it identifies which coefficient in the local expansion of $\Pi_{\text{grav}}(q)$ corresponds to the gravitational coupling and gives the subtracted extraction formula for C_1 that cleanly separates the two leading scalar stiffness contributions.

Lemma 20.6 (Spectral/Laplace Einstein–Hilbert reduction). *Suppose the boundary-deformation response of $(K_{\text{rec}}, M_{\text{rec}})$ admits a continuum limit satisfying the following four substrate-level hypotheses:*

1. *long-wavelength locality of the continuum limit;*
2. *diffeomorphism covariance of the continuum effective action;*
3. *spectral/Laplace generation of continuum geometry from the framework's generalized record operator $G_{\text{acc}} = M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2}$;*
4. *no preferred background direction in the boundary-deformation sector.*

Then the continuum effective action takes the form

$$S_{\text{eff}}[g] = C_{\text{vol}} \int d^4x \sqrt{-g} + C_1 \int d^4x \sqrt{-g} R + (\text{higher-curvature}), \quad (70)$$

where the higher-curvature terms include R^2 , $R_{ab}R^{ab}$, $R_{abcd}R^{abcd}$, and higher-derivative corrections, and C_1 is the boundary-stiffness coefficient of Lemma 20.5. The leading nontrivial two-derivative scalar functional is uniquely the Einstein–Hilbert term: it is the unique local diffeomorphism-covariant two-derivative scalar in four dimensions, up to boundary terms and normalization.

Lemma 20.6 is the single conditional step of the chain. The four substrate-level hypotheses are precisely the framework’s Einstein–Hilbert continuation-work program of Section 12.2 (spectral continuum limit of G_{acc} , the heat-kernel / Laplace structure addressed in Section 12.1); the present theorem introduces no new conditionality beyond those existing framework commitments. The lemma’s content is that, once those hypotheses hold, the Einstein–Hilbert form is not inserted by hand but emerges as the unique allowed leading scalar by representation-theoretic uniqueness of two-derivative diffeomorphism-covariant scalars on four-dimensional manifolds.

Lemma 20.7 (Planck stiffness identification). *Matching the continuum effective action (70) to the Einstein–Hilbert action in reduced Planck convention $S_{\text{EH}} = (M_*^2/2) \int d^4x \sqrt{-g} R$ gives the Planck stiffness identification $M_*^2 = 2C_1$ and Newton’s constant $G_N = 1/(8\pi M_*^2) = 1/(16\pi C_1)$, with C_1 extracted from $\Pi_{\text{grav}}(q)$ by Equation (69).*

This lemma is coefficient-matching: it identifies the framework’s C_1 with the reduced Planck stiffness through the standard Einstein–Hilbert kinetic-term coefficient relation $M_*^2/2$ and the standard reduced Planck mass identification $G_N = 1/(8\pi M_*^2)$.

Combining the four lemmas: Lemma 20.4 defines K_∂, M_∂ ; Lemma 20.5 extracts C_1 from $\Pi_{\text{grav}}(q)$ with the subtraction (69); Lemma 20.6 establishes the Einstein–Hilbert form of the continuum effective action under the four substrate-level hypotheses; Lemma 20.7 identifies $G_N = 1/(16\pi C_1)$. This completes the proof. \square

Explicit construction of P_∂ . The boundary projection of Lemma 20.4 admits a framework-internal construction as a composite of four admissibility filters, each removing a distinct class of non-gravitational modes from the raw boundary record-deformation space. Let \mathcal{H}_{raw} denote the space of linearized observer-boundary record deformations. The gravitational boundary sector is the intersection of four admissibility subspaces:

$$\mathcal{H}_\partial = \mathcal{H}_{\text{metric}} \cap \mathcal{H}_{\text{comp}} \cap \mathcal{H}_{\text{diff}} \cap \mathcal{H}_{\text{IR}}, \quad (71)$$

with the four constituent subspaces characterized as follows.

Metric-deformation subspace $\mathcal{H}_{\text{metric}}$. The space of symmetric boundary tensor perturbations $h_{ab} = \delta\gamma_{ab}$ that act as variations of the induced boundary metric γ_{ab} on the observer-cut surface ∂O . This filter removes non-geometric record disturbances: internal-sector phase changes, gauge-sector deformations, matter-record fluctuations, and pure detector-support modes that do not modify boundary geometry.

Record-compatibility subspace $\mathcal{H}_{\text{comp}}$. Deformations $h_{ab} \in \mathcal{H}_{\text{metric}}$ such that $\delta_h K_{\text{rec}}$ and $\delta_h M_{\text{rec}}$ remain paired under the framework’s record-admissibility (paired bookkeeping) condition: the metric deformation must preserve the conjugate stiffness–susceptibility relation that defines $G_{\text{acc}} = M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2}$. This filter removes deformations that would break stable recordability of the paired bookkeeping under boundary metric variation.

Diffeomorphism-quotient subspace $\mathcal{H}_{\text{diff}}$. The quotient of $\mathcal{H}_{\text{comp}}$ by infinitesimal boundary diffeomorphisms acting as $h_{ab} \rightarrow h_{ab} + \nabla_a \xi_b + \nabla_b \xi_a$, equivalently the orthogonal complement of pure diffeomorphism modes with respect to the M_{rec} -induced inner product. This filter removes coordinate-redefinition modes that are not physical boundary deformations.

Long-wavelength spectral subspace \mathcal{H}_{IR} . The low-spectral band of the boundary Laplace operator, $\mathcal{H}_{\text{IR}} = \text{ran } \chi_{[0, \Lambda^2]}(\Delta_\partial)$, where Δ_∂ is the boundary Laplacian induced on metric deformations and Λ^2 is a spectral cutoff to be removed as $q^2 \rightarrow 0$ in the long-wavelength limit. This filter selects derivative-sensitive curvature modes (those that contribute to the $C_1 q^2$ coefficient of Lemma 20.5).

The boundary projector P_∂ is the orthogonal projection of \mathcal{H}_{raw} onto \mathcal{H}_∂ with respect to the M_{rec} -induced inner product

$$\langle h, k \rangle_M := \langle h, M_{\text{rec}} k \rangle, \quad (72)$$

i.e. P_∂ is the M_{rec} -orthogonal projection onto the intersection (71). When the four constituent filter-projections $P_{\text{metric}}, P_{\text{comp}}, P_{\text{diff}}, P_{\text{IR}}$ commute, P_∂ admits the ordered product form

$$P_\partial = P_{\text{IR}} P_{\text{diff}} P_{\text{comp}} P_{\text{metric}} \quad (73)$$

(equivalent forms by symmetry of orthogonal projection onto an intersection of subspaces). In the general case where the filters do not commute exactly, the safe formal definition is the M_{rec} -orthogonal projection onto the intersection \mathcal{H}_∂ , which is unambiguous regardless of filter ordering.

Framework-native inner product. Use of the M_{rec} -induced inner product (72) makes the orthogonality structure of P_∂ compatible with the susceptibility content of M_{rec} already used in $\Pi_{\text{grav}}(q)$ of Lemma 20.5, and imports no external Hilbert-space metric: M_{rec} is the framework's record susceptibility and supplies the natural notion of orthogonality for physical record deformations. The boundary Laplace eigenmode h_q used in $\Pi_{\text{grav}}(q)$ is normalized in this inner product,

$$\Delta_\partial h_q = q^2 h_q, \quad h_q \in \mathcal{H}_\partial, \quad \langle h_q, h_q \rangle_M = 1. \quad (74)$$

The boundary stiffness response $\Pi_{\text{grav}}(q)$ of Equation (65) is then a ratio of M_{rec} -inner products in which the normalization of h_q cancels.

Structural status of the four-step chain. The chain decomposes the theorem into one conditional step (Lemma 20.6) and three definitional or coefficient-matching steps (Lemmas 20.4, 20.5, 20.7). The conditional step's four substrate-level hypotheses — long-wavelength locality, diffeomorphism covariance, spectral/Laplace continuum generation, no preferred background direction — are precisely the framework's existing Einstein–Hilbert reduction continuation-work commitments; the four-step chain introduces no new structural conditionality beyond those already in force. The boundary projector P_∂ used in Lemma 20.4 is explicitly constructed above (Equations (71)–(73)) as the M_{rec} -orthogonal projection onto the intersection of four admissibility subspaces using only framework-internal objects, so no abstract definition is left implicit in the chain. What remains open is the *numerical evaluation* of C_1 from the framework's paired operators projected through the explicit P_∂ , not the structural derivation of G_N from C_1 : the structural derivation is complete via the four lemmas with P_∂ now constructively defined, and the conditionality has been absorbed into the framework's existing reduction-program hypotheses rather than introduced fresh.

The theorem does not constitute a numerical derivation of G_N for two operationally distinct reasons. First, the leading coefficient C_1 has not been computed in closed form for the framework's paired operators; this computation requires evaluating the four admissibility filters $P_{\text{metric}}, P_{\text{comp}}, P_{\text{diff}}, P_{\text{IR}}$ on specific substrate operators and an explicit boundary-mode basis $\{h_q\}$, neither of which the present paper provides. Second, while the four substrate-level hypotheses of Lemma 20.6 are derived consequences of the closed cascade of Section 20.6 (combining the boundary-deformation scaffold of Section 20.5 with the framework's existing Shared-Connection principle), their operator-algebraic verification on the framework's specific paired operators — e.g. the $\mathcal{H}_{\text{comp}}$ closure, the analyticity of $\Pi_{\text{grav}}(q)$ at $q^2 = 0$, the absence of any preferred background tensor surviving P_∂ — has not been carried out in the present paper. These verifications remain part of the framework's Einstein–Hilbert continuation-work program of Section 12.2, which seeks to derive Einstein's field equations from the spectral continuum limit of G_{acc} ; the framework has not verified these hypotheses within the present paper. Third, the theorem identifies the dimensional combination C_1 but does not by itself fix an absolute scale: C_1 inherits the framework's primitive scale structure, which would need separate substrate-level identification.

What the theorem changes is the structural status of G_N . Before Theorem 20.3, G_N was an external dimensional input with no framework derivation. After it, G_N is conditional on the existing Einstein–Hilbert reduction succeeding: if that reduction succeeds, G_N is determined by the specific framework-internal coefficient C_1 through Equation (68), no longer a free input. The framework’s scope statement is correspondingly tightened: “ G_N is not derived numerically by the present paper” is preserved, but G_N is no longer an independent external parameter — it is structurally tied to the framework’s existing reduction-target program through the boundary-stiffness coefficient.

Connection to the Einstein–Hilbert continuation work. The theorem is structurally aligned with the framework’s existing Einstein–Hilbert reduction program. Where the existing program asks whether the continuum spectral limit of G_{acc} yields a local diffeomorphism-covariant action with an Einstein–Hilbert leading term, the present theorem identifies the specific spectral coefficient C_1 that determines the magnitude of G_N once that reduction succeeds. The two pieces of work are complementary: the existing continuation-work program addresses the *form* of the long-wavelength action (does it land on Einstein–Hilbert?); the present theorem identifies the *magnitude* of the coefficient G_N once the form is fixed. Both are required for a full substrate-internal derivation of gravity; the framework now has both pieces stated structurally, with the form-question and the magnitude-question both reduced to concrete operator-algebraic targets on the same paired bookkeeping $(K_{\text{rec}}, M_{\text{rec}})$.

Continuation-work targets for the Planck-coefficient theorem. Closing the derivation of G_N from substrate primitives requires:

1. Numerical evaluation of the composite projector P_∂ for the framework’s specific paired operators $(K_{\text{rec}}, M_{\text{rec}})$. The structural construction of P_∂ as the M_{rec} -orthogonal projection onto $\mathcal{H}_\partial = \mathcal{H}_{\text{metric}} \cap \mathcal{H}_{\text{comp}} \cap \mathcal{H}_{\text{diff}} \cap \mathcal{H}_{\text{IR}}$ is supplied by Equations (71)–(73) above; what remains is the operator-algebraic evaluation of the four constituent admissibility filters on specific substrate operators.
2. Derivation of the primitive boundary curvature stiffness κ_∂ from a microscopic substrate variational principle. Under the minimal substrate variational principle of Lemma 20.8 of Section 20.5, three of the four leading expansion coefficients are fixed by minimality and normalization ($M_0 = 1$, $M_2 = 0$, $K_0 = \Lambda_\partial$ subtracted in curvature extraction), and only the single substrate-level scale κ_∂ remains as the framework’s primitive gravitational stiffness; $G_N = 1/(16\pi\kappa_\partial)$ then follows from Equation (96).
3. Operator-algebraic verification on the framework’s specific paired operators $(K_{\text{rec}}, M_{\text{rec}})$ of the structural hypotheses identified by the cascade of Section 20.6: the $\mathcal{H}_{\text{comp}}$ closure of Lemma 20.9, the self-adjointness and analytic spectral expansion of Lemma 20.10, the analyticity of $\Pi_{\text{grav}}(q)$ at $q^2 = 0$ of Lemma 20.11, the absence of any preferred background tensor surviving the boundary projection of Lemma 20.12, and the matter-sector universality closed by the framework’s Shared-Connection principle. The cascade reduces these to closed subsidiary statements with no new structural conditionality beyond existing framework commitments; what remains is operator-algebraic evaluation.
4. Numerical evaluation of $G_N = 1/(16\pi C_1)$ for the framework’s specific substrate, with the empirical target $G_N \approx 6.674 \times 10^{-11} \text{ m}^3 \text{ kg}^{-1} \text{ s}^{-2}$ as the verification anchor.

The first three items are operator-algebraic; the fourth is the framework’s first numerical falsifiable target for gravity. A computed value of G_N from the framework’s primitive bookkeeping that differs from the empirical value by more than the precision of the computation would falsify the framework’s claim that gravity is a readout of the same paired record geometry as the gauge sector.

20.5 Boundary-deformation reduction and the Einstein–Hilbert coefficient

The previous subsection established Theorem 20.3 via the four-lemma chain and supplied the explicit composite construction of P_∂ as the M_{rec} -orthogonal projection onto $\mathcal{H}_\partial = \mathcal{H}_{\text{metric}} \cap \mathcal{H}_{\text{comp}} \cap \mathcal{H}_{\text{diff}} \cap \mathcal{H}_{\text{IR}}$. The present subsection supplies the explicit mathematical scaffold needed for the operator-algebraic evaluation of the four-lemma chain: explicit forms of the boundary deformation variables, the record-compatibility constraint, the boundary Laplace operator, the mode normalization, and the local expansion of K_∂, M_∂ . The end result is a closed-form expression for the curvature stiffness coefficient C_1 in terms of four leading expansion coefficients of K_∂, M_∂ , completing the operational specification of the Einstein–Hilbert reduction program.

Boundary deformation variables. Let $\partial\mathcal{O}$ denote the codimension-one observer-access boundary, with induced boundary metric γ_{ab} . A boundary deformation is a small perturbation of the induced metric,

$$\gamma_{ab} \mapsto \gamma_{ab} + h_{ab}, \quad h_{ab} = h_{ba}, \quad (75)$$

so the deformation variable is the variation $h_{ab} = \delta\gamma_{ab}$. The raw deformation space is the space of symmetric rank-two tensor fields on the observer boundary,

$$\mathcal{H}_{\text{raw}} = \Gamma(\text{Sym}^2 T^* \partial\mathcal{O}), \quad (76)$$

and the metric-deformation filter P_{metric} of the previous subsection selects this geometric subspace from the larger raw record-deformation space (which also contains internal-sector, gauge-sector, and matter-record deformations).

Record-compatibility constraint. A boundary deformation is admissible only if it preserves the paired record structure of $(K_{\text{rec}}, M_{\text{rec}})$. Writing the first-order deformations

$$K_{\text{rec}}(h) = K_{\text{rec}} + \delta_h K_{\text{rec}} + O(h^2), \quad M_{\text{rec}}(h) = M_{\text{rec}} + \delta_h M_{\text{rec}} + O(h^2),$$

the compatibility condition is that the generalized eigenvalue structure of the pair remains M_{rec} -self-adjoint and positive. In normalized form, define

$$A_{\text{rec}} := M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2}. \quad (77)$$

A deformation h is record-compatible iff the first variation $\delta_h A_{\text{rec}}$ remains self-adjoint on the record Hilbert space and the deformed operators preserve positivity:

$$M_{\text{rec}} + \delta_h M_{\text{rec}} > 0, \quad K_{\text{rec}} + \delta_h K_{\text{rec}} \geq 0 \quad (78)$$

to first order in h . Equivalently, the generalized eigenproblem $K_{\text{rec}} u = \lambda M_{\text{rec}} u$ must remain a valid paired stiffness/susceptibility problem under deformation. The record-compatibility condition defines the compatible subspace $\mathcal{H}_{\text{comp}}$ from the metric-deformation subspace.

Diffeomorphism quotient. Infinitesimal boundary diffeomorphisms act on metric perturbations as

$$h_{ab} \rightarrow h_{ab} + \nabla_a \xi_b + \nabla_b \xi_a, \quad (79)$$

generating coordinate-redundancy modes that are not physical boundary deformations. The diffeomorphism-quotient subspace is the equivalence-class quotient

$$\mathcal{H}_{\text{diff}} = \mathcal{H}_{\text{comp}} / \{\nabla_a \xi_b + \nabla_b \xi_a\}, \quad (80)$$

equivalently the M_{rec} -orthogonal complement of the pure-diffeomorphism modes within $\mathcal{H}_{\text{comp}}$. The explicit gauge choice within the quotient (e.g. transverse-traceless conditions $\nabla^a h_{ab} = 0$ and $h^a_a = 0$ for isolating propagating tensor modes) is a representative-selection issue that does not affect gauge-invariant content; the framework uses the equivalence-class formulation of Equation (80) as the structural prescription.

Boundary Laplace operator. The boundary Laplace operator on symmetric tensor fields is the Lichnerowicz Laplacian

$$(\Delta_L h)_{ab} = -\nabla^2 h_{ab} - 2R_{acbd}h^{cd} + R_a{}^c h_{cb} + R_b{}^c h_{ac}, \quad (81)$$

where the curvature quantities R_{abcd}, R_{ab} are intrinsic to the observer-boundary geometry (or to the induced continuum geometry in the boundary sector). In the nearly flat long-wavelength limit, $(\Delta_L h)_{ab} \simeq -\partial^2 h_{ab}$, and spectral eigenmodes satisfy

$$\Delta_L h_{ab}^{(q)} = q^2 h_{ab}^{(q)}. \quad (82)$$

The long-wavelength projector entering P_∂ is the spectral cutoff

$$P_{\text{IR}} = \chi_{[0, \Lambda^2]}(\Delta_L), \quad (83)$$

with the Einstein–Hilbert reduction obtained in the $q^2 \rightarrow 0$ continuum limit on spectral modes within this band.

Mode normalization. The boundary modes are normalized using the susceptibility inner product induced by M_∂ :

$$\langle h, k \rangle_{M_\partial} = \int_{\partial\mathcal{O}} d\Sigma \sqrt{\gamma} h_{ab} (M_\partial)^{ab, cd} k_{cd}, \quad (84)$$

where $d\Sigma \sqrt{\gamma}$ is the boundary volume element. A normalized physical boundary mode h_q satisfies the three conditions

$$P_\partial h_q = h_q, \quad \Delta_L h_q = q^2 h_q, \quad \langle h_q, h_q \rangle_{M_\partial} = 1. \quad (85)$$

This explicit normalization fixes the mode amplitude before extracting the stiffness response of Lemma 20.5.

Local expansion of K_∂ and M_∂ . Under the framework’s spectral/Laplace continuation-work hypotheses (Lemma 20.6, conditions 1–3), the boundary-projected paired operators admit local expansions in powers of the Lichnerowicz Laplacian on the long-wavelength sector:

$$K_\partial = K_0 + K_2 \Delta_L + K_4 \Delta_L^2 + \dots, \quad M_\partial = M_0 + M_2 \Delta_L + M_4 \Delta_L^2 + \dots, \quad (86)$$

with K_0, K_2, K_4, \dots and M_0, M_2, M_4, \dots scalar coefficients fixed by the substrate-internal structure of the paired bookkeeping at the observer boundary. For a normalized eigenmode h_q with $\Delta_L h_q = q^2 h_q$,

$$\langle h_q, K_\partial h_q \rangle = K_0 + K_2 q^2 + K_4 q^4 + \dots, \quad \langle h_q, M_\partial h_q \rangle = M_0 + M_2 q^2 + M_4 q^4 + \dots,$$

and the stiffness response is

$$\Pi_{\text{grav}}(q) = \frac{\langle h_q, K_\partial h_q \rangle}{\langle h_q, M_\partial h_q \rangle} = \Pi_0 + \Pi_2 q^2 + O(q^4), \quad (87)$$

with leading coefficients

$$\Pi_0 = \frac{K_0}{M_0}, \quad \Pi_2 = \frac{K_2 M_0 - K_0 M_2}{M_0^2}. \quad (88)$$

The constant term Π_0 is the cosmological/volume-type stiffness; the q^2 term Π_2 is the curvature stiffness. The curvature stiffness coefficient of Lemma 20.5 is therefore

$$C_1 = \Pi_2 = \frac{K_2 M_0 - K_0 M_2}{M_0^2}, \quad (89)$$

equivalently the subtraction form $C_1 = \lim_{q \rightarrow 0} [\Pi_{\text{grav}}(q) - \Pi_{\text{grav}}(0)]/q^2$ of Equation (69). The closed-form expression (89) reduces the computation of C_1 to the determination of four leading coefficients K_0, K_2, M_0, M_2 of the local expansions of K_∂, M_∂ on the boundary sector.

Einstein–Hilbert reduction and Newton’s constant. Under the four hypotheses of Lemma 20.6 — long-wavelength locality, diffeomorphism covariance, spectral/Laplace continuum generation, and no preferred background direction — the continuum effective action takes the form

$$S_{\text{eff}}[g] = C_{\text{vol}} \int d^4x \sqrt{-g} + C_1 \int d^4x \sqrt{-g} R + O(R^2, \nabla R, \dots), \quad (90)$$

with the volume coefficient C_{vol} related to Π_0 (cosmological-constant-like piece) and the curvature coefficient C_1 given by Equation (89). Matching to the Einstein–Hilbert action in reduced Planck convention $S_{\text{EH}} = (M_*^2/2) \int d^4x \sqrt{-g} R$ gives the reduced Planck stiffness

$$M_*^2 = 2 C_1 = \frac{2(K_2 M_0 - K_0 M_2)}{M_0^2}, \quad (91)$$

and Newton’s constant

$$G_N = \frac{1}{8\pi M_*^2} = \frac{1}{16\pi C_1} = \frac{M_0^2}{16\pi (K_2 M_0 - K_0 M_2)}. \quad (92)$$

Equation (92) is the framework’s closed-form expression for Newton’s constant in terms of four leading coefficients of the boundary-projected paired operators K_∂, M_∂ . The gravitational coupling G_N is not a new gauge-like trace value: it is the inverse of the universal long-wavelength boundary-deformation stiffness of the observer-access record pair.

Minimal substrate variational principle. Equation (92) expresses Newton’s constant in terms of four leading expansion coefficients K_0, K_2, M_0, M_2 of K_∂, M_∂ . A natural refinement that further specifies these coefficients is to impose a minimality principle on the boundary record pair: at long wavelengths, use the lowest-order local, diffeomorphism-covariant, record-compatible quadratic form on boundary metric deformations, with canonical susceptibility normalization. The four conditions of the minimal principle are:

1. locality of the continuum effective action on boundary deformations;
2. diffeomorphism covariance of the boundary functional;
3. record-compatibility, i.e. preservation of the paired $(K_{\text{rec}}, M_{\text{rec}})$ admissibility under boundary deformation (Equation (78));
4. canonical leading-order normalization of the susceptibility, $\langle h_q, M_\partial h_q \rangle = 1$ for boundary eigenmodes.

Conditions 1–3 are exactly the substrate-level hypotheses of Lemma 20.6; condition 4 is the normalization convention of Equation (85). No new substrate-level hypothesis is introduced beyond those already in force.

Lemma 20.8 (Minimal boundary spectral coefficients). *Under the four conditions of the minimal substrate variational principle above, the leading expansion coefficients of K_∂, M_∂ in Equation (86) take the form*

$$K_0 = \Lambda_\partial, \quad K_2 = \kappa_\partial, \quad M_0 = 1, \quad M_2 = 0, \quad (93)$$

where Λ_∂ is the boundary-volume (cosmological-type) stiffness and κ_∂ is the primitive boundary curvature stiffness. Higher-order coefficients K_4, K_6, \dots and M_4, M_6, \dots correspond to higher-curvature and higher-derivative invariants, suppressed at leading order in the long-wavelength limit.

Sketch. $M_0 = 1$: The susceptibility M_∂ defines the normalization of boundary eigenmodes through the inner product (84). The canonical normalization $\langle h_q, M_\partial h_q \rangle = 1$ at leading order (Equation (85)) directly fixes $M_0 = 1$. This is a normalization convention, not a physical assumption.

$M_2 = 0$: Minimality of the susceptibility content places derivative-dependence first in the stiffness operator K_∂ rather than in the susceptibility metric M_∂ . This is the framework analog of the standard field-theoretic convention separating the kinetic stiffness operator (which contains ∂^2 -type terms) from the inner-product metric (which is taken ultralocal at leading order). Under this convention, $M_\partial = M_0 + O(\Delta_L^2)$ with $M_2 = 0$.

$K_0 = \Lambda_\partial$: The zero-derivative scalar functional of the boundary metric is the boundary volume $\int_{\partial\mathcal{O}} d\Sigma \sqrt{\gamma}$ (the unique local diffeomorphism-covariant scalar with no derivatives of the metric). Therefore K_0 is the boundary-volume stiffness Λ_∂ , the boundary analog of a cosmological-constant-type term.

$K_2 = \kappa_\partial$: The leading two-derivative diffeomorphism-covariant scalar functional of the boundary metric is the Ricci scalar of the boundary geometry (the same uniqueness argument used in Lemma 20.6: the Ricci scalar is the unique local diff-covariant two-derivative scalar in four dimensions, up to boundary terms and normalization). Therefore K_2 is the primitive boundary curvature stiffness κ_∂ . \square

Specialized closed-form for G_N . Substituting the minimal coefficients (93) into the general closed-form Equation (89) gives

$$C_1 = \frac{K_2 M_0 - K_0 M_2}{M_0^2} = \frac{\kappa_\partial \cdot 1 - \Lambda_\partial \cdot 0}{1^2} = \kappa_\partial, \quad (94)$$

so the curvature stiffness coefficient equals the primitive boundary curvature stiffness under the minimal principle. The full stiffness response is

$$\Pi_{\text{grav}}(q) = \Lambda_\partial + \kappa_\partial q^2 + O(q^4), \quad (95)$$

and the reduced Planck stiffness and Newton's constant specialize to

$$M_*^2 = 2 \kappa_\partial, \quad G_N = \frac{1}{16\pi \kappa_\partial}. \quad (96)$$

The boundary-volume stiffness Λ_∂ is a separate framework primitive (the boundary analog of a cosmological-constant-type term); it enters $\Pi_{\text{grav}}(0)$ but is subtracted out in the curvature extraction and does not contribute to G_N . The full structural derivation of Newton's constant under the minimal substrate variational principle therefore reduces to determining a single substrate-level dimensional scale: the primitive boundary curvature stiffness κ_∂ .

What is derived versus what remains. Lemma 20.8 derives the *form* of the leading expansion coefficients in Equation (93) from locality, diffeomorphism covariance, record-compatibility, and canonical susceptibility normalization — conditions the framework already commits to. It does not derive the *numerical value* of κ_∂ . Numerical determination of κ_∂ requires a microscopic substrate variational principle supplying an absolute boundary-stiffness scale, which is where M_* structurally lives:

$$\kappa_\partial = \frac{M_*^2}{2} = \frac{1}{16\pi G_N}.$$

After Lemma 20.8, the framework's outstanding work for gravity is sharply localized: derive *one* substrate-level dimensional constant κ_∂ from microscopic substrate physics, with $M_0 = 1$ fixed by normalization, $M_2 = 0$ by minimality of susceptibility, $K_0 = \Lambda_\partial$ subtracted in the curvature extraction, and $K_2 = \kappa_\partial$ the single number remaining as the framework's primitive gravitational scale.

Operational status. Equation (92) closes the structural derivation of G_N at the closed-form level. Under the minimal substrate variational principle of Lemma 20.8, this specializes to the single-coefficient form $G_N = 1/(16\pi\kappa_\partial)$ of Equation (96): three of the four expansion coefficients are fixed by minimality and normalization ($M_0 = 1$ by canonical normalization, $M_2 = 0$ by ultralocal susceptibility minimality, $K_0 = \Lambda_\partial$ subtracted in the curvature extraction), and only the primitive boundary curvature stiffness κ_∂ remains as a substrate-level dimensional scale to be derived. The framework's outstanding work for gravity is therefore sharply localized to one substrate-level dimensional constant, not four expansion coefficients. The four substrate-level hypotheses of Lemma 20.6 are derived consequences of the closed cascade of Section 20.6, no longer independent continuation-work commitments; their operator-algebraic verification on the framework's specific paired operators remains continuation work, but no new structural conditionality has been introduced. The empirical check $G_N \approx 6.674 \times 10^{-11} \text{ m}^3 \text{ kg}^{-1} \text{ s}^{-2}$ remains a numerical target requiring derivation of κ_∂ from microscopic substrate physics to close first.

20.6 Substructure of the Einstein–Hilbert reduction hypotheses

Lemma 20.6 of Section 20.4 states the Einstein–Hilbert reduction under four substrate-level hypotheses: long-wavelength locality, diffeomorphism covariance, spectral/Laplace continuum generation, and no preferred background direction. These four hypotheses were carried as framework continuation-work commitments through the closed-form C_1 derivation of Section 20.5 and the minimal-coefficient specialization $G_N = 1/(16\pi\kappa_\partial)$ of Lemma 20.8. The present subsection breaks the four-hypothesis bundle into a closed cascade of six subsidiary lemmas, each of whose residual condition is resolved by the next, bottoming out in the framework's existing single-observer-boundary construction (with the Shared-Connection principle providing independent confirmation). After the cascade, the four hypotheses of Lemma 20.6 are no longer independent continuation-work items: they are derived consequences of the boundary-deformation scaffold of Section 20.5 together with the framework's existing single-observer-boundary construction (which also underwrites the Shared-Connection principle).

Lemma 20.9 (Boundary deformation sector exists). *Assume M_{rec} is positive on admissible boundary record deformations. Let \mathcal{H}_{raw} be the M_{rec} -Hilbert completion of symmetric observer-boundary metric perturbations $h_{ab} = \delta\gamma_{ab}$, and let $\mathcal{H}_\partial = \mathcal{H}_{\text{metric}} \cap \mathcal{H}_{\text{comp}} \cap \mathcal{H}_{\text{diff}} \cap \mathcal{H}_{\text{IR}}$ as in Equation (71). If each of these four subspaces is closed in the M_{rec} -topology, then \mathcal{H}_∂ is closed and the M_{rec} -orthogonal projection $P_\partial : \mathcal{H}_{\text{raw}} \rightarrow \mathcal{H}_\partial$ exists and is unique.*

Sketch. The M_{rec} -inner product $\langle h, k \rangle_M = \langle h, M_{\text{rec}} k \rangle$ is positive on the admissible deformations by hypothesis, and \mathcal{H}_{raw} is its Hilbert completion. The metric-deformation, diffeomorphism-quotient, and IR-spectral subspaces $\mathcal{H}_{\text{metric}}, \mathcal{H}_{\text{diff}}, \mathcal{H}_{\text{IR}}$ are closed by standard arguments: symmetric tensor sections form a closed subspace of $\Gamma(T^*\partial\mathcal{O} \otimes T^*\partial\mathcal{O})$; the diffeomorphism-quotient subspace is the orthogonal complement of the closed range of the symmetrized gradient operator; the IR-spectral band is the range of the bounded spectral projection $\chi_{[0, \Lambda^2]}(\Delta_L)$ of the self-adjoint Lichnerowicz operator. The intersection of finitely many closed subspaces is closed, and the orthogonal projection onto a closed subspace of a Hilbert space exists and is unique. \square

Residual condition. The closedness of $\mathcal{H}_{\text{comp}}$ for the framework's specific paired operators $(K_{\text{rec}}, M_{\text{rec}})$ is the one nontrivial closure condition. This is resolved by Lemma 20.10: if the boundary-projected operators are self-adjoint with analytic spectral expansion, their domain of self-adjointness within \mathcal{H}_∂ is closed, forcing $\mathcal{H}_{\text{comp}}$ closure on the boundary sector.

Lemma 20.10 (Spectral/Laplace expansion exists). *Assume the boundary-projected operators K_∂ and M_∂ are self-adjoint on \mathcal{H}_∂ with respect to the M_{rec} -inner product, and in the long-wavelength regime are functions of the Lichnerowicz operator Δ_L whose spectral functions are*

analytic at the bottom of the spectrum. Then

$$K_{\partial} = F_K(\Delta_L), \quad M_{\partial} = F_M(\Delta_L), \quad (97)$$

with analytic Taylor expansions

$$F_K(x) = K_0 + K_2 x + O(x^2), \quad F_M(x) = M_0 + M_2 x + O(x^2), \quad (98)$$

yielding the local expansion of Equation (86) of Section 20.5.

Sketch. By spectral calculus on the self-adjoint Δ_L , any operator that commutes with Δ_L on the long-wavelength sector admits the form $F(\Delta_L)$ via the spectral theorem; the self-adjointness hypothesis on $K_{\partial}, M_{\partial}$ together with their generation by Δ_L in the IR sector gives the spectral functional form (97). Analyticity of F_K, F_M at the spectral bottom yields the Taylor expansions (98). \square

Residual condition. That F_K, F_M are in fact analytic at $q^2 = 0$ for the framework's specific paired operators after boundary projection. This is resolved by Lemma 20.11: if $\Pi_{\text{grav}}(q)$ is analytic at $q^2 = 0$ with no nonanalytic terms, then $\langle h_q, K_{\partial} h_q \rangle$ and $\langle h_q, M_{\partial} h_q \rangle$ on Δ_L -eigenmodes are themselves analytic, forcing F_K, F_M analyticity.

Lemma 20.11 (Long-wavelength limit is local). *If the boundary-projected record kernels of K_{∂} and M_{∂} are quasi-local on the observer boundary — meaning their position-space kernel decays sufficiently fast at large boundary separation, equivalently their small- q spectral expansion contains no nonanalytic terms $|q|$, $q^2 \log q^2$, or $1/q^2$ — then the gravitational stiffness response is analytic at $q^2 = 0$:*

$$\Pi_{\text{grav}}(q) = \Pi_0 + C_1 q^2 + O(q^4), \quad (99)$$

and the corresponding continuum effective functional is expandable in local curvature invariants.

Sketch. Nonanalytic terms $|q|$, $q^2 \log q^2$, $1/q^2$ in $\Pi_{\text{grav}}(q)$ correspond by Fourier transform to position-space kernels with $1/r^2$, $1/r$, or growing tails — nonlocal long-range integrated contributions. Their absence forces $\Pi_{\text{grav}}(q)$ to be analytic at $q^2 = 0$, with the small- q expansion controlled by the local Taylor coefficients of the spectral functions F_K, F_M . By the polynomial-in- Δ_L form on the long-wavelength sector, the corresponding continuum action is expandable in local diffeomorphism-covariant invariants of the boundary geometry. \square

Residual condition. The absence of nonlocal IR singularities in the boundary-projected record pair, equivalently the absence of an unaccounted massless non-geometric sector surviving inside P_{∂} . This is resolved by Lemma 20.12: a surviving massless non-geometric sector would correspond to a preferred background tensor or field structure transforming non-covariantly under diffeomorphisms, and Lemma 20.12's diffeomorphism-covariance hypothesis forbids such surviving structures.

Lemma 20.12 (Diffeomorphism covariance of the continuum limit). *Assume P_{∂} quotients pure infinitesimal boundary diffeomorphisms (Equation (80)), and assume $K_{\partial}, M_{\partial}$ are constructed from intrinsic boundary-geometric data only, with no preferred coordinate frame or background tensor surviving the projection. Then the long-wavelength local continuum functional is diffeomorphism-covariant, and consequently the leading nontrivial two-derivative scalar term is proportional to*

$$\int d^4x \sqrt{-g} R, \quad (100)$$

with higher-order corrections beginning at curvature-squared or higher-derivative order.

Sketch. Quotienting by pure diffeomorphisms ensures that \mathcal{H}_∂ contains only physical (non-gauge) deformation classes. Constructing K_∂, M_∂ from intrinsic boundary geometry ensures no preferred coordinate or background tensor enters the response. The combination yields a continuum functional invariant under boundary diffeomorphisms. In four dimensions, the unique local diffeomorphism-covariant two-derivative scalar functional of the metric is, up to a topological boundary term and an overall coefficient, $\int d^4x \sqrt{-g} R$ (standard representation-theoretic uniqueness, used also in the proof of Lemma 20.6). \square

Residual condition. That no preferred background tensor survives the boundary projection P_∂ in the framework's specific construction. This is resolved by Lemma 20.13: a surviving preferred tensor would provide a sector-specific coupling channel through which different matter sectors could acquire independent gravitational charges, and Lemma 20.13's universality hypothesis forbids sector-specific couplings, forcing the absence of any preferred tensor that could mediate them.

Lemma 20.13 (Universal stress-readout coupling). *Let h_{ab} be an admissible observer-boundary metric deformation in \mathcal{H}_∂ . If all matter-sector record supports couple to h_{ab} only through the variation of the common observer-boundary metric, then their first-order response has the universal form*

$$\delta S_{\text{matter}} = \frac{1}{2} \int d^4x \sqrt{-g} T^{ab} h_{ab}, \quad (101)$$

and the same boundary curvature stiffness $M_*^2 = 2\kappa_\partial$ controls the geometric response to every matter stress readout. The resulting continuum equation has the universal form

$$M_*^2 G_{ab} = T_{ab}, \quad G_{ab} = 8\pi G_N T_{ab}, \quad G_N = \frac{1}{8\pi M_*^2} = \frac{1}{16\pi\kappa_\partial}, \quad (102)$$

with G_N universal across matter sectors.

Sketch. Under the universality hypothesis, every matter sector couples to boundary geometry through the common metric variation h_{ab} , giving the matter-side stress response $\delta S_{\text{matter}}/\delta g^{ab} = -\frac{1}{2}\sqrt{-g}T^{ab}$ in Equation (101). The same boundary curvature stiffness $M_*^2 = 2\kappa_\partial$ from Lemma 20.8 controls the metric response, yielding Einstein's equation with universal Newton's constant $G_N = 1/(8\pi M_*^2)$. \square

Residual condition. That matter sectors couple only through the common boundary metric deformation h_{ab} , with no sector-specific independent gravitational couplings. This is closed internally by Lemma 20.14 below: the framework's single-observer-boundary construction forces \mathcal{H}_∂ to be matter-sector-independent and P_∂ to be unique, so all matter sectors couple to the same projected $h_{ab} = P_\partial \delta\gamma_{ab}$ by structural necessity rather than by hypothesis. The same conclusion is independently confirmed by the framework's Shared-Connection principle (Section 9 establishing the alignment of inertial and gravity-like readouts as two variational readouts of the same paired regional disturbance $(K_{\text{rec}}, M_{\text{rec}})$) together with the composition-insensitivity theorem (Section 9.2 ensuring internal energy fractions across binding, kinetic, electromagnetic, and rest-like contributions do not introduce independent canonical couplings). Under the framework's existing equivalence-principle content, all matter supports couple through the same paired book-keeping $(K_{\text{rec}}, M_{\text{rec}})$, with the only available gravitational coupling channel being the boundary metric deformation. The universality hypothesis of Lemma 20.13 is therefore not an additional commitment; it is a consequence both of the single-boundary construction (Lemma 20.14) and of the framework's existing Shared-Connection principle applied to the boundary-deformation context of Section 20.5.

Lemma 20.14 (Uniqueness of the gravitational boundary projector). *Suppose the framework has a single observer-access boundary $\partial\mathcal{O}$ with one induced boundary metric γ_{ab} , and matter sectors may have different internal supports but are all read through this same observer-access boundary. Then the gravitational boundary sector*

$$\mathcal{H}_\partial = \mathcal{H}_{\text{metric}} \cap \mathcal{H}_{\text{comp}} \cap \mathcal{H}_{\text{diff}} \cap \mathcal{H}_{\text{IR}}$$

is matter-sector-independent, and the M_{rec} -orthogonal projection $P_\partial : \mathcal{H}_{\text{raw}} \rightarrow \mathcal{H}_\partial$ is unique and common to all matter sectors. Consequently, all matter supports couple to the same projected metric deformation

$$h_{ab} = P_\partial \delta\gamma_{ab},$$

and matter sectors differ only in their stress readouts T_s^{ab} , not in the gravitational coupling channel.

Sketch. Each of the four filters defining \mathcal{H}_∂ is constructed from observer-boundary geometric data, not from any matter-sector internal label:

- $\mathcal{H}_{\text{metric}}$ selects variations of the single induced boundary metric γ_{ab} on the single observer boundary $\partial\mathcal{O}$;
- $\mathcal{H}_{\text{comp}}$ imposes the record-compatibility condition on the total paired bookkeeping $(K_{\text{rec}}, M_{\text{rec}})$, not separately on any matter-sector restriction (Equation (78));
- $\mathcal{H}_{\text{diff}}$ quotients by infinitesimal boundary diffeomorphisms of the single observer boundary (Equation (80));
- \mathcal{H}_{IR} is the low-spectral band of the Lichnerowicz operator on the single boundary geometry (Equation (83)).

None of these conditions depends on a matter-sector label. The intersection \mathcal{H}_∂ is therefore one shared geometric subspace, common to all matter supports. The M_{rec} -orthogonal projection onto a closed subspace of a Hilbert space is unique (Lemma 20.9); since \mathcal{H}_∂ is defined once, P_∂ is defined once. A sector-specific projector $P_{\partial,s}$ would require either a sector-specific boundary metric $\gamma_{ab}^{(s)}$ or a sector-specific compatibility condition $\mathcal{H}_{\text{comp},s}$, neither of which the framework introduces.

For the matter-coupling consequence: each matter-sector stress response to the common deformation $h_{ab} = P_\partial \delta\gamma_{ab}$ has the standard form $\delta S_s = \frac{1}{2} \int d^4x \sqrt{-g} T_s^{ab} h_{ab}$. The stress tensor T_s^{ab} is sector-dependent; the deformation field h_{ab} is not. Summing over sectors gives a single total source $T^{ab} = \sum_s T_s^{ab}$, which is exactly the universality hypothesis of Lemma 20.13. \square

Single-boundary framework commitment. The premise of Lemma 20.14 — that the framework has one observer-access boundary $\partial\mathcal{O}$ with one induced boundary metric γ_{ab} , through which all matter sectors are read — is not a new commitment. It is built into the framework’s substrate-observer construction from the outset: the timeless partition-strain architecture has a single substrate-observer junction with a single access cut, and matter sectors are represented by different internal supports read through the same junction, not by separate junctions with separate boundaries. Lemma 20.14 makes this commitment structurally explicit and shows that it forces the matter-coupling universality of Lemma 20.13. Failure modes that would invalidate the cascade — sector-specific boundary metrics $\gamma_{ab}^{(s)}$ or sector-specific gravitational projectors $P_{\partial,s}$ — are excluded by the single-boundary construction.

Cascade closure and consequences for Lemma 20.6. The six-lemma chain forms a closed cascade:

- Lemma 20.9’s residual ($\mathcal{H}_{\text{comp}}$ closure) is resolved by Lemma 20.10’s self-adjointness/analyticity;
- Lemma 20.10’s residual (F_K, F_M analyticity at $q^2 = 0$) is resolved by Lemma 20.11’s small- q analyticity of Π_{grav} ;
- Lemma 20.11’s residual (no nonlocal IR singularities) is resolved by Lemma 20.12’s diffeomorphism-covariance;
- Lemma 20.12’s residual (no preferred background tensor surviving) is resolved by Lemma 20.13’s universal coupling;
- Lemma 20.13’s residual (no sector-specific gravitational charges) is resolved internally by Lemma 20.14’s uniqueness-of-projector argument from the framework’s single-observer-boundary construction;
- Lemma 20.14’s premise (single observer boundary with single induced metric, all matter sectors read through it) is the framework’s existing substrate-observer construction, independently confirmed by the Shared-Connection principle’s identification of inertial and gravity-like readouts as two variational readouts of the same paired regional disturbance.

The four substrate-level hypotheses of Lemma 20.6 — long-wavelength locality (Lemma 20.11), diffeomorphism covariance (Lemma 20.12), spectral/Laplace continuum generation (Lemma 20.10), and no preferred background direction (Lemma 20.12’s no-preferred-tensor clause, closed by Lemma 20.13) — are derived consequences of the cascade, conditional only on the framework’s existing single-observer-boundary construction (the cascade’s structural foundation, also underwriting the Shared-Connection principle) plus the explicit boundary-deformation scaffold of Section 20.5. The framework’s gravity program is correspondingly tightened: what remains open is the derivation of the single dimensional scale κ_{∂} from microscopic substrate physics (Lemma 20.8), with the structural cascade from $(K_{\text{rec}}, M_{\text{rec}})$ to $G_N = 1/(16\pi\kappa_{\partial})$ now fully closed at the operator-algebraic level.

21 Open questions

21.1 Type II₁ to Type III₁ observer junction

The central structural problem is whether local entropy-maximizing observer states on increasing access cuts, constrained by stable records and denied the global trace, can converge to Type III₁-like modular algebras. The first target is a rigorous version of the modular-density diagnostic: construct a sequence $(\mathcal{A}_n, \omega_n)$ over finite trace approximants such that the paired $(K_{\text{rec}}, M_{\text{rec}})$ disturbance normal form remains stable while the modular ratio groups of ω_n become dense in \mathbb{R} .

The framework leaves several questions open.

1. *Continuum and refinement.* Implement algebraic shell growth and subalgebra joins on genuine von Neumann subalgebra systems without embedded coordinate assumptions.
2. *Justifying the primitive principles.* Derive the paired master regional disturbance functional $(K_{\text{rec}}, M_{\text{rec}})$, its admissible readout maps, and the structural decomposition $K_{\text{rec}} = K_0 + \sum_{\bullet} \kappa_{\bullet} K_{\bullet}$, $M_{\text{rec}} = M_0 + \sum_{\bullet} \eta_{\bullet} M_{\bullet}$ from a more primitive operator-algebraic condition.
3. *Operator-algebraic implementation.* A genuine implementation on a von Neumann algebra is the natural next mathematical step.

4. *Observer reconstruction at scale.* A larger-scale demonstration with realistic record-bearing dynamics is required.
5. *Observer-effective quantum reconstruction.* A larger operator-algebraic reconstruction would need non-commuting accessible algebras, POVMs, entanglement across cuts, decohered record subalgebras, and a principled account of Born weight selection.
6. *Bandwidth-derived Lorentz observer kinematics.* A stronger account would prove convergence of the generalized spectrum $K_{\text{rec}}v = c^2 M_{\text{rec}}v$ from the Type II₁ trace geometry and derive Lorentz transformations between different observer cuts.
7. *Relativistic extension.* A genuine relativistic extension requires tensor stress-energy with the right signature, a metric or richer connection sector, a variational principle, and an account of diffeomorphism redundancy.
8. *Refinement-program hardening.* Replace embedded-coordinate benchmarks with a substrate-intrinsic refinement category; construct transfer operators from algebraic conditional expectations; extend metric action beyond a quadratic positive-definite perturbative model; couple metric and gauge fields; run Level-1 non-Abelian validations.
9. *Comparison with alternative framings.* A systematic comparison with relational QM, Page–Wootters, thermal time, decoherent histories, causal sets, and tensor-network geometries is open work.
10. *Modular thermodynamics.* Make $\delta Q_R = \Theta_R \delta S_R$ intrinsic using modular relative entropy rather than finite approximants.
11. *Type III₁ lift.* Lift the shared-generator composition-independence result from Type II₁ trace geometry to a Type III₁ modular setting.

21.2 Einstein-reduction proof obligations

The continuum Einstein-translation theorem schema reduces the gravity problem to four precise proof obligations. First, the generalized record operator G_{acc} must lie in a Laplace-type spectral universality class with nonzero scalar-curvature coefficient. Second, admissible metric variations must be lifted from algebraic automorphisms. Third, the observer access-cut and recovery-loss functional must supply the boundary term. Fourth, the coefficient of the scalar-curvature term must be normalized against $c_*^2 = \lambda_{\min}(G_{\text{acc}})$ so that $\kappa_{\text{eff}} = 8\pi G_{\text{obs}}/c_*^4$.

21.3 Horizon-cut radiation

Compute $\Theta_{\partial\mathcal{B}}$ from modular data; derive or falsify $\Theta_{\partial\mathcal{B}} \propto \kappa_{\text{rec}}$; define a finite model exhibiting a Page-like recoverability transition.

21.4 Cosmological-constant dictionary entry

A cosmological-constant-like term would be placed as a global paired disturbance baseline or vacuum offset, not as localized matter disturbance. The present work does not derive such a term; it only identifies its placement.

21.5 Multi-cut readings at distinct access depths

The framework does not produce a continuous running function $\alpha(\mu)$ or $\alpha_s(\mu)$ from a single observer cut. It reads at specific access cuts, producing values at the access depths of those cuts. The fine-structure result $T_1^{-1} = \alpha_{\text{eff}}^{-1} = 137.035\,999\,176\dots$ is the framework’s reading at the

Thomson access cut (Section 15.9), where the central Abelian phase footprint is resolved without internal-sector interference. The strong-coupling result $\alpha_s^{-1}(M_Z) = 8.4795$ is the framework’s reading at the electroweak access cut (Section 15.11), where the $SU(3)$ adjoint block is resolved against the neighboring electroweak adjoint and Abelian sectors. The two cuts are different observer-access objects, returning different framework-internal trace expansions; neither is the “default” scale, and neither requires the other to be evaluated separately. The framework’s analog of perturbative renormalization-group running is access-depth dependence: as the observer cut moves through the framework’s access hierarchy, the relevant Schur-expansion sectors change, and the resulting trace expansion evaluates to a different value at each depth.

The mapping between the framework’s effective access depth (derived from the modular automorphism structure of the observer cut, cf. Section 13) and the physical momentum scale μ at which standard couplings are measured is not constructed in the present paper. The framework reads at two specific access cuts (Thomson and electroweak); the continuous access-depth profile that would map the framework’s reading onto the measured $\alpha_s(\mu)$ curve between low- μ and high- μ scales is identified as continuation work and is required for direct comparison with standard renormalization-group running profiles.

21.6 Continuation work flagged by the relocated structural results

The modular access-flow principle (Theorem 20.1), the modular sector-visibility theorem (Theorem 20.2), and the boundary-stiffness Planck-coefficient theorem (Theorem 20.3) are now collected in Section 20. Each carries its own open numerical-determination program: explicit computation of the modular generator K_{mod} on the framework’s observer junction (and thus of the spectral densities $\rho_s(\lambda)$ entering the framework’s beta-function analog), identification of the sector and overlap coefficients Δ_s, Δ_{st} from the framework’s existing Schur-expansion content, verification that integration of the framework’s beta-function analog between λ_0 (Thomson) and λ_Z (electroweak) reproduces the measured $\Delta\alpha^{-1} \approx -9.09$ within experimental uncertainty, and derivation of the single dimensional scale κ_∂ from microscopic substrate physics so that $G_N = 1/(16\pi\kappa_\partial)$ becomes a numerical prediction rather than a structural identification.

21.7 Projective-access proof obligations

The mass-shape sector of Section 5 carries five named derivation targets. Each is a substrate-axiom-level open problem whose resolution would convert one of the section’s selection rules into a derived theorem and tighten the predictive standing of the two-mode projective-access core lemma (Theorem 5.8).

- (O1) *Projective-access ladder derivation.* Identify the projective access space \mathcal{P} , the invariant measure $d\mu_{\mathcal{P}}$, and the averaged Schur kernel κ_{proj} such that $\int_{\mathcal{P}} \kappa_{\text{proj}} d\mu_{\mathcal{P}} = 1/\pi$ (Principle 5.4). This converts Lemma 5.3 from a normalization lemma to a derived theorem and locks the projective normalization on which the rest of Section 5 depends. Of the five obligations this is the highest-leverage target: pinning O1 closes the section’s primary structural gap.
- (O2) *Closure condition.* Derive Principle 5.5 (second-order projective closure) from a substrate-side minimality principle on the visible–hidden decomposition. The mass-shape sector should be identified as the unique closed visible–hidden block under a specified completeness criterion, rather than chosen by convention.
- (O3) *Support-type assignment.* Derive Principle 5.6: show that the tau visible support is ordered (4!) and the baryonic mean visible support is subset $(2^4 - 1)$ from the framework’s existing sector classification (Section 15), rather than imposed by the multiplicity rule.

- (O4) *Schur sign rule.* Derive Principle 5.7 from a substrate-level positivity criterion. The framework should explain why charged-leptonic dressings are additive while baryonic mean defects are subtractive, in terms of the Schur-complement structure on the visible block of Definition 5.1.
- (O5) *Branch-pair averaging rule.* Derive Principle 5.9 (the muon coefficient 1/36) from the $b = 3$ branch structure of Section 15 and the four-axis observer-access lattice of Section 2.3, or determine that no such derivation exists and treat the muon line as a separate sector requiring independent structural input.

Pinning O1 alone elevates Lemma 5.3 to a derived ladder. Pinning O2 through O4 converts the core lemma into a theorem with all premises derived from substrate axioms. Pinning O5 promotes the muon prediction from the provisional extension of Section 5.7 to the core. The five obligations are independent and may be pursued separately.

21.8 Relation to perturbative electroweak corrections

The Standard Model derives the channel split between $\sin^2 \theta_{\text{eff},\ell}$ and $\sin^2 \theta_{\text{eff},b}$ through specific loop-level radiative corrections: the $Zb\bar{b}$ vertex correction (with internal W -boson and top-quark exchange), ρ -parameter shifts from the W/Z mass-splitting, and form factor structures that depend on the momentum transfer q^2 . The numerical agreement of the framework's $\delta_b = 1/(T_{b,\text{col}} T_1)$ with the measured channel split arises from a different starting point: an algebraic trace overlap between the fundamental color-support trace weight and the central Abelian phase footprint, with no loop integral, no fermion-mass input, and no diagrammatic vertex-correction structure.

The framework is therefore offering an *alternative algebraic derivation* of the same numerical result, not a re-expression of the Standard Model loop calculation in different language. Whether a structural dictionary exists between the algebraic primitives of the substrate–observer architecture (paired bookkeeping, observer-access cuts, Schur-trace identities, observation supports and adjoint connection blocks of Appendix 6.11) and the diagrammatic primitives of perturbative quantum field theory (propagators, vertices, loop integrals over internal momenta) is genuine open work and is identified here as a separate research target. The present paper does not claim such a bridge; pretending to one would be overreach.

The mass-hierarchy question requires a parallel acknowledgment. The Standard Model's loop derivation depends critically on the top-quark mass: the $Zb\bar{b}$ vertex correction scales with m_t^2/M_W^2 . The framework's algebraic derivation of δ_b uses no mass input: $T_{b,\text{col}}$ is fixed by the fundamental color-support branch count $d_{\text{supp}} = 3$ on the $(b, d) = (3, 4)$ substrate (Appendix 6.11), T_1 is fixed by the fine-structure Schur identity (Section 15.9), and the product $T_{b,\text{col}} T_1$ is mass-independent. That the two routes converge to the same value is itself a non-trivial structural feature of the framework's substrate-derived bookkeeping.

The framework does not derive the lepton and quark mass ratios as part of its scope (Section 10). The minimal-chiral-content construction of Section 15.4 produces the algebraic ratio $|Y_u|^2 : |Y_d|^2 = 4 : 1$ in the squared-hypercharge form needed for the resolved Drell–Yan kernel of Section 15.10; it does not produce the Yukawa hierarchy. A framework derivation of the Yukawa hierarchy would require a separate substrate-internal mechanism that the present paper neither constructs nor relies on.

21.9 Falsification criteria

Beyond the technical open questions above, the framework commits to specific empirical falsification targets. None can be evaded by adjusting a free parameter, because there are no free parameters to adjust. The architecture either continues to match tightened precision measurements or it does not.

1. *Leptonic effective weak-mixing angle (polarized e^+e^- class).* If future precision tightens the SLD/Tevatron-class polarized leptonic effective angle to a one-sigma band that does not contain $\sin^2 \theta_{\text{eff},\ell} = 0.231\,141\,2087$, the polarized-leptonic visibility readout of Section 15.11 (and its source visibility lemma) is falsified.
2. *b -quark channel.* If future precision tightens the LEP/SLD b -quark effective angle to a one-sigma band that does not contain $\sin^2 \theta_{\text{eff},b} = 0.232\,222\,30$, the representation-rule choice of fundamental color-support for the b -quark is falsified.
3. *Fine-structure precision.* If CODATA α^{-1} at the 10^{-12} relative level deviates from $137.035\,999\,176\dots$ by more than experimental uncertainty, the single Schur-trace identity is falsified.
4. *Strong coupling at M_Z .* If improved $\alpha_s(M_Z)$ precision at $\sigma_{\alpha_s} \lesssim 10^{-4}$ deviates from the leading four-term Schur expansion of Section 15.11, either the leading expansion or the framework's prediction of sub-leading corrections is falsified.
5. *Drell-Yan projection theorem, verification of operator-algebraic conditions.* As discussed in Section 15.10, the framework's prediction for the hadron-collider Drell-Yan channel readout at the isotropic level depends on three operator-algebraic conditions on the framework's existing structures: a trace-preserving transitive spatial group action on the primitive retrieval projections (closing the hypothesis of Theorem 4.3); S_3 -invariance of the prepared channel state on R_3 in the sense of Definition 15.6 (closing the hypothesis of Lemma 15.7); and a forward-trace-preserving retrieval inclusion $\iota_\ell : R_3 \rightarrow A_\ell$ together with Murray-von Neumann equivalence of clean leptonic records to primitive retrieval images (closing the hypotheses of Lemma 15.8). When all three conditions are verified inside the framework's existing observer-cut machinery, $w_{\text{DY}} = 1/3$ becomes a free-standing framework prediction in the isotropic limit, and the corresponding isotropic central value of $\sin^2 \theta_{\text{eff,DY}}$ must agree with a flavor-symmetric score-projected average of the measured Drell-Yan values. A verification of all three conditions that nevertheless produces an isotropic w_{DY} disagreeing with the data, or a verification that requires ex post fitting choices to match the data, falsifies the framework's claim that channel readouts are observer-access-cut features of one architecture. Above the isotropic limit, Section 15.10 casts the experiment-and-channel-specific deviation as a signed projection of the centered hypercharge kernel onto each experiment's weak-angle extraction score $\mathbf{S}_{\text{exp},q}(\xi) = \partial \log \sigma_{\text{exp},q} / \partial \sin^2 \theta_{\text{eff}}$, with magnitude set by the substrate-internal coupling λ_Y in the (P3) structural class of Equation (37) and sign set by the substrate-fixed weight $Y_u^2 - Y_d^2 > 0$. A reconstructed score-projected pattern within the Tevatron class (CDF and D0 channel-separated values; the CMS measurement at $\sqrt{s} = 8$ TeV and several subsequent LHC measurements employ joint A_{FB} -PDF extraction and lie in a structurally distinct observer-access class for which the framework currently does not derive a prediction, as discussed in Section 15.10) that does not track the observed CDF and D0 channel ordering in sign falsifies the (P3) structural identification of K_Y , since (P1) and (P2) are independently ruled out and no other class of K_Y is admissible. The three forward-looking falsifiers (F1)–(F3) stated in Section 15.10 bound the framework's Tevatron-class prediction independently: future Tevatron-class channel-separated measurements, an LHC re-analysis against externally-fixed PDFs, and the magnitude of the channel separation bounded by δ_{init} .
6. *Access-depth running profile.* As discussed in Section 6.7, the framework's analog of perturbative running is the access-depth dependence of which sectors enter the Schur expansion. When the framework's access-depth-dependent Schur trace is computed and compared to the measured $\alpha_s(\mu)$ profile between $\mu \sim 2$ GeV and $\mu \sim 1$ TeV, agreement at the level of the experimental precision profile is required. A confirmed disagreement that cannot be absorbed into the framework's access-depth dependence falsifies the access-depth analog of running.

7. *Type II₁ normal-form scalar-reduction step.* Theorem 4.1's proof contains one step where the scalar reduction via isotropy plus unitary covariance is asserted in compressed form rather than expanded via explicit Schur-lemma analysis of the adjoint unitary representation. A rigorous mathematical treatment that finds the step cannot be closed under stated hypotheses would falsify the foundational normal-form claim and, with it, the shared-generator proposition of Section 9 and the cone-recovery lemma of Section 4.2.
8. *Tau mass ratio.* Section 5 predicts $r_\tau = m_\tau/m_e = 4!A + 4\Pi_1 + \Pi_2 = 3477.22828$ under the two-mode projective-access core lemma (Theorem 5.8), where $A = \alpha^{-1}$ is the framework's fine-structure prediction of Section 15.9. If improved measurements of m_τ/m_e place its central value outside the interval $[3477.225, 3477.231]$, the combination of selection rules in the core lemma is falsified. A failure does not invalidate the substrate normal form of Section 4 or the metrological completion of Section 7; it falsifies the specific combination of multiplicity, closure, and sign rules in Theorem 5.8.
9. *Baryon mean mass ratio.* Section 5 also predicts $r_B = (m_p + m_n)/(2m_e) = (2^4 - 1)A - 5\Pi_1 = 1837.44042$. The current measured value 1837.41817 leaves a relative residual of $\sim 10^{-5}$; improved precision tightening the residual below the propagated α^{-1} uncertainty would force a refinement of the multiplicity assignment for the baryonic mean sector or a recognition of additional projective modes beyond the second-order closure of Principle 5.5.
10. *Layer protection of substrate-level support primitives.* As discussed in Section 3.3 and made explicit by Remark 3.6, the refinement-factorization architecture forbids direct coupling between observer-resolvable refinement labels and the substrate-level support index. Any measurement or analysis requiring a readout whose support projection on a refined retrieval algebra cannot be written as $P_{\mathcal{R}} \otimes P_{\mathcal{F}}$ within experimental or analytic uncertainty falsifies the untwisted-extension assumption (A1) of Definition 3.1, and with it the canonical-lift commitment (Definition 3.3) on which the refinement compatibility of paired bookkeeping rests. This is a framework-wide structural prediction; it applies wherever the framework refines a substrate algebra by attaching branch-neutral labels.

These targets are named here so that future tightening of measurement precision, mathematical scrutiny, and continuation work each have specific deliverables against which framework agreement or disagreement can be checked.

22 Conclusion

The thesis is reconstructive and organizational. Two layers separate cleanly: a tenseless algebraic substrate carries stable relational structure through the paired $(K_{\text{rec}}, M_{\text{rec}})$ bookkeeping, and a memory-bearing observer carries record-ordered access. The compressed substrate object is the paired regional disturbance: the universal record stiffness K_{rec} and the universal record susceptibility M_{rec} , with the observable record cone governed by the generalized record operator $G_{\text{acc}} = M_{\text{rec}}^{-1/2} K_{\text{rec}} M_{\text{rec}}^{-1/2}$. Locality, non-separation, persistence, support-stiffness, gravity-like response, and gauge connection are not treated as separate primitive disturbance sectors; they are readouts or stationarity conditions of this paired disturbance under different probes. Experienced time, motion, measurement update, and field evolution remain observer-side effects of stable record access.

The operational replacement for primitive time is the inclusion index of stable records. A modular time scale appears when the local entropy-maximizing observer state on the recoverable algebra is faithful and nonracial. This keeps the substrate tenseless while giving the observer layer enough structure to support thermal flow, finite access cones, and relativistic kinematics as reconstruction targets.

The central gain of the compressed formulation is that entropy, inertia, and gravity-like response no longer need to be engineered as parallel functionals. Entropy is scalar unrecoverability through M_{rec} . Inertia is the algebraic support-deformation Hessian through K_{rec} . Gravity-like response is the connection variation of the same paired disturbance. The light-speed limit is $c_*^2 = \lambda_{\min}(G_{\text{acc}})$. The electromagnetic sector matches the same paired structure, $(K_{\text{rec}}^{\text{em}}, M_{\text{rec}}^{\text{em}}) = g_{\text{eff}}^{-2}(K_{\text{rec}}, M_{\text{rec}})$, ensuring that the coupling controls amplitude without altering the universal record cone.

The minimal stable observer-access dimension is $3 + 1$: three projective localization directions plus one record-ordering coordinate. The structural function of this decomposition is not numerical convenience. The three spatial directions Σ_{obs}^3 are the smallest cover that closes orientation, angular separation, occlusion, and parallax without an exterior reference frame; the record-ordering coordinate τ_{obs} supplies sequencing along the inclusion index of stable records. This $3 + 1$ observer-access structure enters consistently into the Lorentz-recovery target, the access-dimension ablation of the fine-structure diagnostic, and the substrate–observer duality that fixes $T_{\text{bare}} = 21/2$ at $(d, b) = (4, 3)$.

As a concrete observer-junction deliverable inside this organization, the fine-structure access-polarization diagnostic compresses to a single Schur-complement trace on one substrate-cut block operator. The four numerical components of the diagnostic are not four independent choices: they are the diagonal contributions of the charged-readout sector and the two spectral channels of one positive access-dressing operator. The Schur-complement theorem fixes the signs of the dressings; the access–recovery composition theorem fixes the kernel as L_{acc}^{-2} ; the projective access-cover theorem fixes the multiplicity $2^b - 1$ at the minimal projective closure. The resulting value agrees with the CODATA 2022 recommended fine-structure constant within experimental uncertainty under five explicitly named structural inputs that remain part of the open program.

The Type II₁ normal-form theorem remains the substrate anchor. Under explicit covariance and isotropy assumptions, local conditional-expectation defects have one canonical quadratic trace geometry, and the paired $(K_{\text{rec}}, M_{\text{rec}})$ structure inherits the same canonical trace.

The observer junction remains essential. A purely tracial substrate has no experienced time, no nontrivial thermal flow, and no primitive light speed. Nontracial observer states and scaling limits supply modular time, finite recoverability cones, temperature-like parameters, and Type III-like behavior. The black-hole reading uses the same separation: horizon entropy, exterior information loss, and radiation-like flux are observer-access readouts of large boundary paired disturbance, while the full substrate description is not destroyed.

The strongest continuation is to replace the remaining constructive choices with primitive principles. The paired master disturbance functional, admissible regional category, stiffness and susceptibility operators with their structural decompositions, source readouts, and refinement protocols are still specified rather than uniquely derived. The refinement compatibility theorem of Section 3.3 closes one structural gap: it shows that under four explicit assumptions, admissible refinements of a substrate retrieval algebra factor as a tensor product, the canonical lift is forced by access-geometry preservation, and the unrefined readout is recovered exactly under coarse-graining over the refinement labels. The finite QEC examples show how conditional-expectation defect diagnoses Petz recoverability and channel recovery loss in concrete finite trace settings consistent with the paired framework. The paired master disturbance and observer-access category define the present reconstruction core; the remaining open work is to sharpen their uniqueness and empirical dictionary without changing the substrate–junction spine.

A Validation material

Validation protocols include a six-region Level-0 construction, Newtonian and Lorentz observer finite-model models, sourced $U(1)$ solves through the matched pair $(K_{\text{rec}}^{\text{em}}, M_{\text{rec}}^{\text{em}})$, weak-field scalar solves, coordinate-graph refinement tests, topological stress tests, and finite non-Abelian pilot

scripts. The sourced $U(1)$ solve is summarized in the main Maxwell-type section only as a stationarity benchmark for the holonomy readout. The remaining calculations are useful as consistency checks, but they are not premises of the foundation argument.

The main claim is therefore not that finite scripts reproduce selected formulas. The main claim is that conditional-expectation defect supplies the local geometry of a paired master regional disturbance ($K_{\text{rec}}, M_{\text{rec}}$); entropy, recoverability loss, inertial support stiffness, gravity-like connection response, and gauge stationarity are readouts of that one paired disturbance through observer-access cuts.

Acknowledgements

The author thanks the LLM assistants who served as drafting and critique collaborators throughout the development of this work. Each was used substantively for technical sharpening, restructuring, prose drafting, and consistency checking across multiple iterations. These systems do not bear authorship, but their contribution to the development of this manuscript was material and is acknowledged here. The author also thanks Nicolas Joly for the idea on black-hole evaporation.

Copyright

Copyright © 2026 Daniel de Souza Casali. All rights reserved.

No part of this manuscript — including the text, figures, tables, equations, and any accompanying source files — may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the author, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law.

For permission requests, contact the author.

Suggested citation. Daniel de Souza Casali, *Timeless Partition-Strain on a Structured Algebraic Universe: A Unisound Substrate–Observer Reconstruction*, May 2026.

Disclaimer. This work is provided in good faith based on the author’s current understanding of the framework it describes. The picture is one possible reading and may be incorrect. The author makes no factual claims about ultimate physical reality and disclaims liability for damages resulting from the use of this material to the extent permitted by applicable law.

References

- [1] B. S. DeWitt, “Quantum theory of gravity. I. The canonical theory”, *Phys. Rev.* **160** (1967) 1113.
- [2] K. V. Kuchař, “Time and interpretations of quantum gravity”, in *Proc. 4th Canadian Conf. on General Relativity and Relativistic Astrophysics*, eds. G. Kunstatter et al. (World Scientific, 1992); reprinted *Int. J. Mod. Phys. D* **20** (2011) 3.
- [3] C. J. Isham, “Canonical quantum gravity and the problem of time”, in *Integrable Systems, Quantum Groups, and Quantum Field Theories*, eds. L. A. Ibort and M. A. Rodríguez (Kluwer, 1993), pp. 157–287, arXiv:gr-qc/9210011.
- [4] F. J. Murray and J. von Neumann, “On rings of operators”, *Ann. Math.* **37** (1936) 116.
- [5] M. Takesaki, *Theory of Operator Algebras I*, Springer (2002).

- [6] M. Takesaki, *Theory of Operator Algebras II*, Springer (2003).
- [7] A. Connes, *Noncommutative Geometry*, Academic Press (1994).
- [8] A. Connes, “Une classification des facteurs de type III”, *Ann. Sci. École Norm. Sup. (4)* **6** (1973) 133.
- [9] M. Tomita, “Standard forms of von Neumann algebras”, *The Fifth Functional Analysis Symposium of the Math. Soc. of Japan, Sendai* (1967).
- [10] M. Takesaki, *Tomita’s Theory of Modular Hilbert Algebras and its Applications*, Lecture Notes in Math. 128, Springer (1970).
- [11] M. Takesaki, “Duality for crossed products and the structure of von Neumann algebras of type III”, *Acta Math.* **131** (1973) 249.
- [12] H. Umegaki, “Conditional expectation in an operator algebra”, *Tohoku Math. J.* **6** (1954) 177.
- [13] M. Takesaki, “Conditional expectations in von Neumann algebras”, *J. Funct. Anal.* **9** (1972) 306.
- [14] R. Haag, N. M. Hugenholtz and M. Winnink, “On the equilibrium states in quantum statistical mechanics”, *Commun. Math. Phys.* **5** (1967) 215.
- [15] R. T. Powers, “Representations of uniformly hyperfinite algebras and their associated von Neumann rings”, *Ann. Math.* **86** (1967) 138.
- [16] W. G. Unruh, “Notes on black-hole evaporation”, *Phys. Rev. D* **14** (1976) 870.
- [17] E. Witten, “Gravity and the crossed product”, *JHEP* **10** (2022) 008, arXiv:2112.12828.
- [18] V. Chandrasekaran, G. Penington and E. Witten, “Large N algebras and generalized entropy”, *JHEP* **04** (2023) 009, arXiv:2209.10454.
- [19] V. Chandrasekaran, R. Longo, G. Penington and E. Witten, “An algebra of observables for de Sitter space”, *JHEP* **02** (2023) 082, arXiv:2206.10780.
- [20] D. N. Page and W. K. Wootters, “Evolution without evolution: dynamics described by stationary observables”, *Phys. Rev. D* **27** (1983) 2885.
- [21] W. K. Wootters, “ ‘Time’ replaced by quantum correlations”, *Int. J. Theor. Phys.* **23** (1984) 701.
- [22] V. Giovannetti, S. Lloyd and L. Maccone, “Quantum time”, *Phys. Rev. D* **92** (2015) 045033, arXiv:1504.04215.
- [23] C. Rovelli, “Relational quantum mechanics”, *Int. J. Theor. Phys.* **35** (1996) 1637, arXiv:quant-ph/9609002.
- [24] A. Connes and C. Rovelli, “Von Neumann algebra automorphisms and time-thermodynamics relation in generally covariant quantum theories”, *Class. Quantum Grav.* **11** (1994) 2899, arXiv:gr-qc/9406019.
- [25] T. Jacobson, “Thermodynamics of spacetime: the Einstein equation of state”, *Phys. Rev. Lett.* **75** (1995) 1260, arXiv:gr-qc/9504004.
- [26] T. Jacobson, “Entanglement equilibrium and the Einstein equation”, *Phys. Rev. Lett.* **116** (2016) 201101, arXiv:1505.04753.

- [27] R. B. Griffiths, “Consistent histories and the interpretation of quantum mechanics”, *J. Stat. Phys.* **36** (1984) 219.
- [28] M. Gell-Mann and J. B. Hartle, “Quantum mechanics in the light of quantum cosmology”, in *Complexity, Entropy and the Physics of Information*, ed. W. H. Zurek (Addison-Wesley, 1990) p. 425.
- [29] R. Omnès, “Consistent interpretations of quantum mechanics”, *Rev. Mod. Phys.* **64** (1992) 339.
- [30] L. Bombelli, J. Lee, D. Meyer and R. D. Sorkin, “Space-time as a causal set”, *Phys. Rev. Lett.* **59** (1987) 521.
- [31] B. Swingle, “Entanglement renormalization and holography”, *Phys. Rev. D* **86** (2012) 065007, arXiv:0905.1317.
- [32] M. Van Raamsdonk, “Building up spacetime with quantum entanglement”, *Gen. Relativ. Gravit.* **42** (2010) 2323, arXiv:1005.3035.
- [33] J. D. Bekenstein, “Black holes and entropy”, *Phys. Rev. D* **7** (1973) 2333.
- [34] S. W. Hawking, “Particle creation by black holes”, *Commun. Math. Phys.* **43** (1975) 199.
- [35] D. N. Page, “Average entropy of a subsystem”, *Phys. Rev. Lett.* **71** (1993) 1291, arXiv:gr-qc/9305007.
- [36] D. N. Page, “Time dependence of Hawking radiation entropy”, *JCAP* **09** (2013) 028, arXiv:1301.4995.
- [37] A. Almheiri, N. Engelhardt, D. Marolf and H. Maxfield, “The entropy of bulk quantum fields and the entanglement wedge of an evaporating black hole”, *JHEP* **12** (2019) 063, arXiv:1905.08762.
- [38] G. Penington, “Entanglement wedge reconstruction and the information paradox”, *JHEP* **09** (2020) 002, arXiv:1905.08255.
- [39] A. Almheiri, T. Hartman, J. Maldacena, E. Shaghoulian and A. Tajdini, “The entropy of Hawking radiation”, *Rev. Mod. Phys.* **93** (2021) 035002, arXiv:2006.06872.
- [40] J. W. York, “Role of conformal three-geometry in the dynamics of gravitation”, *Phys. Rev. Lett.* **28** (1972) 1082.
- [41] G. W. Gibbons and S. W. Hawking, “Action integrals and partition functions in quantum gravity”, *Phys. Rev. D* **15** (1977) 2752.
- [42] E. Knill and R. Laflamme, “Theory of quantum error-correcting codes”, *Phys. Rev. A* **55** (1997) 900, arXiv:quant-ph/9604034.
- [43] R. Laflamme, C. Miquel, J. P. Paz and W. H. Zurek, “Perfect quantum error correcting code”, *Phys. Rev. Lett.* **77** (1996) 198, arXiv:quant-ph/9602019.
- [44] C. H. Bennett, D. P. DiVincenzo, J. A. Smolin and W. K. Wootters, “Mixed-state entanglement and quantum error correction”, *Phys. Rev. A* **54** (1996) 3824, arXiv:quant-ph/9604024.
- [45] D. W. Leung, M. A. Nielsen, I. L. Chuang and Y. Yamamoto, “Approximate quantum error correction can lead to better codes”, *Phys. Rev. A* **56** (1997) 2567, arXiv:quant-ph/9704002.
- [46] D. Petz, “Sufficient subalgebras and the relative entropy of states of a von Neumann algebra”, *Commun. Math. Phys.* **105** (1986) 123.

- [47] D. Petz, “Sufficiency of channels over von Neumann algebras”, *Quart. J. Math. Oxford* **39** (1988) 97.
- [48] H. Barnum and E. Knill, “Reversing quantum dynamics with near-optimal quantum and classical fidelity”, *J. Math. Phys.* **43** (2002) 2097, arXiv:quant-ph/0004088.
- [49] S. L. Glashow, “Partial-symmetries of weak interactions”, *Nucl. Phys.* **22** (1961) 579.
- [50] S. Weinberg, “A model of leptons”, *Phys. Rev. Lett.* **19** (1967) 1264.
- [51] A. Salam, “Weak and electromagnetic interactions”, in *Elementary Particle Theory*, ed. N. Svartholm (Almqvist & Wiksell, 1968) p. 367.
- [52] H. Georgi and S. L. Glashow, “Unity of all elementary-particle forces”, *Phys. Rev. Lett.* **32** (1974) 438.
- [53] H. Georgi, H. R. Quinn and S. Weinberg, “Hierarchy of interactions in unified gauge theories”, *Phys. Rev. Lett.* **33** (1974) 451.
- [54] S. L. Adler, “Axial-vector vertex in spinor electrodynamics”, *Phys. Rev.* **177** (1969) 2426.
- [55] J. S. Bell and R. Jackiw, “A PCAC puzzle: $\pi^0 \rightarrow \gamma\gamma$ in the σ -model”, *Nuovo Cim. A* **60** (1969) 47.
- [56] H. Georgi and S. L. Glashow, “Gauge theories without anomalies”, *Phys. Rev. D* **6** (1972) 429.
- [57] E. Tiesinga, P. J. Mohr, D. B. Newell and B. N. Taylor, “CODATA recommended values of the fundamental physical constants: 2018”, *Rev. Mod. Phys.* **93** (2021) 025010.
- [58] P. J. Mohr, D. B. Newell, B. N. Taylor and E. Tiesinga, “CODATA recommended values of the fundamental physical constants: 2022”, arXiv:2409.03787 (2024); see also the NIST CODATA reference at physics.nist.gov/constants, giving $\alpha^{-1} = 137.035\,999\,177(21)$.
- [59] L. Morel, Z. Yao, P. Cladé and S. Guellati-Khélifa, “Determination of the fine-structure constant with an accuracy of 81 parts per trillion”, *Nature* **588** (2020) 61.
- [60] R. H. Parker, C. Yu, W. Zhong, B. Estey and H. Müller, “Measurement of the fine-structure constant as a test of the Standard Model”, *Science* **360** (2018) 191.
- [61] X. Fan, T. G. Myers, B. A. D. Sukra and G. Gabrielse, “Measurement of the electron magnetic moment”, *Phys. Rev. Lett.* **130** (2023) 071801.
- [62] S. Navas et al. (Particle Data Group), “Review of Particle Physics”, *Phys. Rev. D* **110** (2024) 030001.
- [63] A. H. Chamseddine and A. Connes, “The spectral action principle”, *Commun. Math. Phys.* **186** (1997) 731, arXiv:hep-th/9606001.
- [64] A. H. Chamseddine and A. Connes, “Resilience of the spectral standard model”, *JHEP* **09** (2012) 104, arXiv:1208.1030.
- [65] A. Connes, “Classification of injective factors. Cases II_1 , II_∞ , III_λ , $\lambda \neq 1$ ”, *Ann. Math.* **104** (1976) 73.
- [66] J. J. Bisognano and E. H. Wichmann, “On the duality condition for a Hermitian scalar field”, *J. Math. Phys.* **16** (1975) 985.

- [67] J. J. Bisognano and E. H. Wichmann, “On the duality condition for quantum fields”, *J. Math. Phys.* **17** (1976) 303.
- [68] R. Haag, *Local Quantum Physics: Fields, Particles, Algebras*, 2nd ed., Springer (1996).
- [69] O. Bratteli and D. W. Robinson, *Operator Algebras and Quantum Statistical Mechanics, Vol. I*, 2nd ed., Springer (1987).
- [70] O. Bratteli and D. W. Robinson, *Operator Algebras and Quantum Statistical Mechanics, Vol. II*, 2nd ed., Springer (1997).
- [71] A. D. Sakharov, “Vacuum quantum fluctuations in curved space and the theory of gravitation”, *Sov. Phys. Dokl.* **12** (1968) 1040.
- [72] M. Visser, “Sakharov’s induced gravity: a modern perspective”, *Mod. Phys. Lett. A* **17** (2002) 977, arXiv:gr-qc/0204062.
- [73] E. Verlinde, “On the origin of gravity and the laws of Newton”, *JHEP* **04** (2011) 029, arXiv:1001.0785.
- [74] T. Padmanabhan, “Thermodynamical aspects of gravity: new insights”, *Rep. Prog. Phys.* **73** (2010) 046901, arXiv:0911.5004.
- [75] O. Fawzi and R. Renner, “Quantum conditional mutual information and approximate Markov chains”, *Commun. Math. Phys.* **340** (2015) 575, arXiv:1410.0664.
- [76] M. Junge, R. Renner, D. Sutter, M. M. Wilde and A. Winter, “Universal recovery maps and approximate sufficiency of quantum relative entropy”, *Ann. Henri Poincaré* **19** (2018) 2955, arXiv:1509.07127.
- [77] D. Sutter, M. Berta and M. Tomamichel, “Multivariate trace inequalities”, *Commun. Math. Phys.* **352** (2017) 37, arXiv:1604.03023.
- [78] L. Bombelli, R. K. Koul, J. Lee and R. D. Sorkin, “Quantum source of entropy for black holes”, *Phys. Rev. D* **34** (1986) 373.
- [79] M. Srednicki, “Entropy and area”, *Phys. Rev. Lett.* **71** (1993) 666, arXiv:hep-th/9303048.
- [80] A. C. Wall, “A proof of the generalized second law for rapidly changing fields and arbitrary horizon slices”, *Phys. Rev. D* **85** (2012) 044011, arXiv:1105.3445.
- [81] D. V. Vassilevich, “Heat kernel expansion: user’s manual”, *Phys. Rep.* **388** (2003) 279, arXiv:hep-th/0306138.
- [82] T. Aoyama, T. Kinoshita and M. Nio, “Theory of the anomalous magnetic moment of the electron”, *Atoms* **7** (2019) 28.
- [83] L. Hardy, “Quantum theory from five reasonable axioms”, arXiv:quant-ph/0101012 (2001).
- [84] L. Masanes and M. P. Müller, “A derivation of quantum theory from physical requirements”, *New J. Phys.* **13** (2011) 063001, arXiv:1004.1483.