

The GRIP on GDOS

GRIP: the Geometric Engine of Symmetry Breaking, from the Big Bang to Mind

Katsushi Arisaka

*Department of Physics and Astronomy, Electrical and Computer Engineering,
University of California, Los Angeles, Los Angeles, CA 90095, USA*

Contact: arisaka@physics.ucla.edu

Abstract

Nothing in the universe stayed symmetric. It began hot and nearly featureless, and everything in it now — the particles, the elements, the molecules, the living cell, the mind reading this sentence — is what was left standing after one symmetry broke, and then another, and then another.

This paper proposes that a single mechanism drives all of it, and writes that mechanism down. We call it GRIP, the Gauge-Runtime Invariance Pipeline: a four-move cascade in which a system first holds every gauge equivalence open and commits to none (DGI); a geometric free-energy landscape (GFF) then fixes which way is downhill; the descent selects a branch at a saddle and the symmetry breaks (GSSB); and the system is captured in a basin — a Geometric Attractor (GA), or, when the basin is the deepest one available, a *Ground* Geometric Attractor (GGA), which is why matter is stable rather than merely long-lived.

The pipeline runs in exactly two regimes, and the distinction is the paper's second claim. GRIP-In is descent on a landscape that does not move, and it halts. GRIP-Out is the transient event that reorganizes the landscape itself and hands the next descent its starting point. GRIP-Out opens the journey; GRIP-In finishes it. The weak interaction is Nature's implementation of GRIP-In on the strata of matter; the electroweak transition was a GRIP-Out.

We give the geometry that makes the pipeline exact — four Geometric Invariance Pillars of the open-system boundary law, of which GRIP is the fourth and the other three are its machinery — and with it the family count $N_g \leq 3$ follows as a theorem rather than as an observation. Every prediction is a closed-form function of five integers with no free parameter left to tune, and each pillar is handed an experiment that can end it.

Executive Summary

Three moves, and this paper is the third. The GG-Theory program rests on three ideas, and only three. The first is a six-dimensional geometry in which the four dimensions we inhabit are a boundary and the two extra ones are not further directions in space but a scale and a phase (Arisaka, 2026, A2-GG6) — a geometry that proves anomaly-free (Arisaka, 2026, A3-BI6) and leaves no continuous parameter free to adjust (Arisaka, 2026, A4-Euclid), on a constitutional charter of four axioms and seven postulates (Arisaka, 2026, A1-Constitution). The second is that no system in nature is closed: the bulk is projecting onto the boundary at every instant, and what we call the laws of physics is what that projection looks like from inside (Arisaka, 2026, A5-GDOS). The third is the subject of this paper. A geometry says what *can* exist and an open-system law says how it is *seen*; neither says why the universe, having begun almost featureless, did not stay that way. Something has to break the symmetry, choose among the pieces, and stop. That something is what we name and write down here.

The pipeline, in one line. We call it GRIP, the Gauge-Runtime Invariance Pipeline, and it has four moves:

$$\text{DGI} \longrightarrow (\text{GFF}) \longrightarrow \text{GSSB} \longrightarrow \text{GA} / \text{GGA}.$$

Read it left to right. DGI, Dynamical Gauge Invariance, is the symmetric top of the descent: every gauge equivalence is still open and nothing has been committed. GFF, the Geometric Free-Energy Functional, is the landscape — written in parentheses because it is not a move but the terrain the moves happen on, the thing that says which way is downhill and how steep. GSSB, Geometric Spontaneous Symmetry Breaking, is the moment of choice: the descent reaches a saddle, one branch is taken, the others are not, and a symmetry that was real a moment ago is gone. GA, the Geometric Attractor, is where the system is captured. That is the whole engine. Everything else in this paper exists to make those four moves exact enough to be wrong.

Two regimes: one stroke opens, the other closes. The engine runs in exactly two ways, and the difference is a single question — *does the landscape itself change?* GRIP-In is descent on a landscape that stays put: the system rolls into a basin, tunnels between basins, may rest a long while at a way station, and halts for good when it reaches a basin with no exit. GRIP-Out is the opposite: a transient, out-of-equilibrium event that reorganizes the landscape and hands the next descent a new starting point. In one sentence: GRIP-Out opens the journey, GRIP-In finishes it. They are not competing accounts of one thing but alternating strokes of one engine, and the alternation is what lets a universe go on producing new structure instead of settling once and stopping. The weak interaction is Nature’s implementation of GRIP-In on the strata of matter; the electroweak transition was a GRIP-Out. What is invariant under a descent and what is destroyed by it are questions of symmetry and of holonomy, and both are treated at cohort level (Arisaka, 2026, W1-Symmetry; Arisaka, 2026, W2-Holonomy); that a descent runs one way at all is the same arrow the mirrors of matter carry (Arisaka, 2026, P5-CPT).

Why matter is stable. The terminus deserves a sharper name than “attractor,” because not all basins are alike. A GA is any stable basin; a GGA, a *Ground* Geometric Attractor, is the deepest one on its stratum, with nothing below it to fall to. That distinction is not bookkeeping — it is the reason matter persists. The electron is stable because it is the ground attractor of its stratum (Arisaka, 2026, P3-Electron). So is the proton, whose refusal to decay is then a fact about geometry rather than an unexplained bound (Arisaka, 2026, P6-Proton). So, if the framework is right, is the dark-matter candidate this program predicts (Arisaka, 2026, C1-Kaxion). It is the same reason a water molecule is not a temporary arrangement of atoms that will eventually find something better to be: it is already at the bottom. A universe whose descents merely slowed down would have no permanent objects in it, and nothing to build with.

From the first microsecond to the mind. The claim this paper is written to support is that the four moves are not a mechanism for one epoch but the mechanism, running over and over at every scale. The grand-unified transition in the first instant is one turn of the crank, and the resonance it leaves near the lock scale is a standing test of it (Arisaka, 2026, P1-KK). So is the electroweak transition, which gives the particles their masses and fixes the Standard Model’s parameters (Arisaka, 2026, A9-SM); so is the settling of quarks into protons and neutrons, and of nucleons into helium, carbon, oxygen, iron; so is the binding of atoms into water, carbon dioxide, ammonia; of small molecules into amino acids and peptides; of those into the molecular machinery of the living cell — ATP, RNA, DNA; and, the program contends, of cells into organisms, of organisms into animals, and of animals into the kind that can ask why any of it happened. The same descent read at cosmological scale gives the boundary cosmology (Arisaka, 2026, A10-Cosmos), the two ends of the expansion corridor (Arisaka, 2026, C2-Hubble), and the horizon’s entropy (Arisaka, 2026, C3-BH); read all the way up, it is the capstone’s subject (Arisaka, 2026, A11-Origin). It is worth being exact about what that sentence is. The pipeline, its two regimes and its halting theorem are proved here. Their deployment through the full ladder of cascade layers is carried by the capstone, and the identification of the higher layers — chemical, biochemical, biological, cognitive — with the same four moves is the cohort’s central wager rather than a theorem of this paper. What this paper contributes is the thing that wager needs: a pipeline stated sharply enough that the identification could be shown to fail.

The geometry underneath, briefly. Making the pipeline exact takes machinery, and the machinery is four Geometric Invariance Pillars of which GRIP is the fourth. The first gathers gravity and the three gauge forces into a single composite connection on the six-dimensional bulk, so that there is one object to descend rather than four. The second is the conservation ledger that object must satisfy, and its arithmetic is unexpectedly rigid: closure forces a unique set of five integers, $(1, 3, 5; 7, 35)$, and everything downstream is a closed-form function of them. The third is the free-energy landscape itself, together with the proof that descent on it runs one way — which is where the arrow of time enters the runtime, and where the rules of quantum mechanics stop being a list to assume and become consequences (Arisaka, 2026, A8-QM). The fourth is the pipeline. That the boundary’s single imaginary unit is itself derived rather than posited (Arisaka, 2026, P4-Real) is what lets the descent be written on real geometry throughout. The four pillars are not four


constructions but four faces of one geometric object, and their mutual forcing is what makes the family count come out at $N_g \leq 3$ instead of being read off the data.

What is established here, and what would refute it. The boundary of a claim matters more than the claim. Derived in this paper: the pipeline and its four moves, the two regimes and their necessity, the halting theorem at the ground attractor, the conservation ledger, the one-way descent, and the family bound $N_g \leq 3$ — the last conditional on the capacity theorem of the three-generations companion (Arisaka, 2026, P2-N3), whose integrality residue is discharged on a benchmark background with two refinements named and left open. Taken as given and not re-derived here: the constitutional charter, the six-dimensional geometry, the anomaly-freedom proof and the open-system boundary law, each the subject of a companion preprint and each cited at the step that uses it. Against that, three experiments are named with the values they must return: a modified equation of stellar structure in compact stars, equipartition together with the fluctuation–dissipation relation in any boundary observable, and a generation-ordered ladder of decay widths falling by a fixed exponential factor from one family to the next. Any one of them can end the program, and the geometry leaves nowhere to retreat: the dimension is fixed, the integers are fixed, and there is no adjustable landscape to hide in.

A closing thought. Physics has been very good at describing what is stable and much less comfortable with how anything came to be stable in the first place. The symmetric state is always the easy one to write down; the world we have is the one that stopped being symmetric, repeatedly, in a particular order, and produced observers late in the sequence. If the four moves proposed here are right, then that sequence is not a history of accidents but the same operation applied again and again to its own output — and the cosmos, life, and mind are not three separate puzzles but three depths of one descent. That is a large claim, and it is offered in the only form that makes a large claim respectable: as something that can be measured, and found wanting.


Keywords: GG-Theory; GDOS; Gi-4 Pillars; TGG (Total Geometric-Gauge object); CCI (Covariant Conservation Identity); GFF (Geometric Free-Energy Functional); GRIP (Gauge-Runtime Invariance Pipeline); GEP (Geometry-Gauge Energy Principle); CEP (Curvature-Entropy Principle); DGI (Dynamical Gauge Invariance); GSSB (Geometric Spontaneous Symmetry Breaking); GA (Geometric Attractor); GGA (Ground Geometric Attractor); GRIP-In (fixed-landscape regime); GRIP-Out (landscape-reorganizing regime); OGGE; MSMF; BI-6; OGN-5; Spin(1,5); SL(2,H); composite connection; boundary stress tensor; master conservation law; modified Tolman–Oppenheimer–Volkoff equation; anomaly polynomial; Hopkins–Singer differential cohomology; Morse theory on stratified spaces; boundary projector; bulk-to-boundary descent; family-projector moduli; Wilson–Hosotani (Hosotani, 2005); instanton-action budget; magic-dimension ladder; foundational open-system framework

Four Geometric Pillars, and the One Law They Hold Up



What this paper is. GDOS is the boundary law; the four Gi-4 pillars are what forces it to be that law and not another. Each pillar inherits its constitutional content from GDOS above and its bulk-side content from the 6D geometry below, and the horizontal arrows are implications, not analogies: $\text{TGG} \Rightarrow \text{CCI} \Rightarrow \text{GFF} \Rightarrow \text{GRIP}$. The four named compatibility relations tie each pillar's algebraic layer to its operational one. Read this page and the shape of the argument is fixed; the twelve Parts are the work of earning it. In full: Figure 2, p. 33.

How Structure Settles: One Cascade, and the Two Regimes It Runs In



The pillar the paper is named for. The central column is the runtime: from a fully symmetric top where no gauge equivalence has been spent, the system descends the fixed landscape, branches at the symmetry-breaking saddles, and is captured into a stable basin. The two-regime story is the point. At a way station the basin still has exit channels and the descent continues — that is GRIP-In, and the weak interaction is Nature’s implementer of it. At the deepest basin the exit set is empty, GRIP-In halts, and only GRIP-Out can act: a transient out-of-equilibrium event that reorganizes the landscape itself and supplies the initial datum of the next descent. GRIP-Out opens the landscape; GRIP-In finishes the journey. In full: Figure 8, p. 146.

Every Prediction the Four Pillars Make That Could Kill the Framework

Where the four pillars are exposed. One row per pillar sub-theorem, each naming the prediction, the observable that would falsify it, and where the measurement stands today. The fourth column is the honest one: some rows are already confirmed to high precision and are no longer at risk, and the rest name the facility that will decide them. A framework that forces its parameters has to accept that a single failed row ends it. These are the pillars' own falsifiers, not the cohort's ten — for those see [GG-A1-Constitution](#). In full: [Table 20](#), p. 229.

Pillar / ST	Prediction	Empirical observable	Status / facility
TGG-ST6, ST7	Master Conservation Law $\nabla_\mu T_{\text{TGG}}^{\mu\nu} = 0$	Tests of energy/momentum conservation in any boundary process	Confirmed to 10^{-15} in low-energy experiments
CCI-ST7	Modified TOV $\Delta_{\text{GS}} \sim 9.6 \times 10^{-3}$	0.1%–1% shift of NS mass–radius curve from standard TOV	NICER (current); next-gen X-ray timing; LIGO/Virgo NS inspirals
GFF-ST7 (1)	Equipartition $\langle \delta X^2 \rangle = T_{\text{GFF}} \chi$ for any boundary observable	Universal T -scaling of fluctuations near equilibrium	Ultra-cold atoms, optomechanics; quantum-limited sensors
GFF-ST7 (2)	Fluctuation–dissipation: $\langle \delta F^2 \rangle = T_{\text{GFF}} \chi_{\text{GFF}}$	Einstein-type relation between mobility and diffusion	Mesoscopic transport experiments
GFF-ST7 (3)	Onsager reciprocity $L_{ij} = L_{ji}$	Symmetric cross-coupling matrix in any open boundary network	Molecular motors, active matter, biophysical networks
GRIP-ST7	Generation weight ladder $w_{[2k] \rightarrow [2(k-1)]} = e^{-\pi k}$	Geometric gate weight along the family chains $\tau \rightarrow \mu \rightarrow e, t \rightarrow \dots \rightarrow u/d$ (absolute lifetimes via A9)	Generation-ordering of the gate-action ladder; controlled-kinematics decay class
GRIP-ST7	Cascade-toll sum rule	$\sum_i S_{\text{tot}}^{(i)} = \pi N_g (N_g - 1)(N_g + 1)/6$, a toll budget, not a timescale	None by itself; the ordering is read from the mass hierarchy
BI-6 (Companion)	KK resonance cluster $M_{\text{lock}} \simeq 3.85 \text{ TeV}$	Tower of KK modes near 3.85 TeV in dijet/diphoton/dilepton	LHC Run 3, HL-LHC; FCC-hh
BI-6 (Companion)	Kaxion dark matter $m_K \approx 8.1 \text{ } \mu\text{eV}$	Ultra-light bosonic DM near μeV with axion-like couplings	ABRACADABRA, DMRadio, ADMX-EFR; cosmological constraints
GG-Cosmos	Modified Friedmann equations	Coherent shift of $f\sigma_8(z)$, $w(z) + 1 \propto \langle X_4 \rangle / M_{\text{KK}}^4$	DESI, Euclid, Roman Space Telescope; CMB-S4
GG-6F (code-book)	28 SM parameters at $\sim 0.5\%$ precision	e.g. m_τ/m_μ to 0.5% from cascade depths	Global SM precision fits (Particle Data Group (Particle Data Group, 2024))

Overview of Contents

Part I. Introduction and Setup	20
Part II. TGG	40
Part III. CCI	67
Part IV. GFF	93
Part V. GRIP	130
Part VI. Synthesis	170
Part VII. Applications	180
Part VIII. Discussion	215
Part IX. Conclusions	239
Part X. Technical Appendices	245
Part XI. Common Objections and Responses	275
Part XII. References	310

List of Figures

1	The open-system chain shared across the GG-Theory cohort.	30
2	Architecture of the four Gi-4 Pillars.	33
3	The twenty-two-paper cohort map.	36
4	Roadmap of the twelve Parts of GG-A6-GRIP.	38
5	Algebraic structure of TGG (Pillar 1).	48
6	Anomaly inflow and the Ward ladder of CCI (Pillar 2).	76
7	The GRIP cascade with the BI-6 budget closure (Pillar 4).	141
8	The GRIP cascade and its two regimes.	146
9	The eighteen-layer cascade of GRIP-Out, carried from GG-A11-Origin.	182
10	The cascade of GRIP-Out with GGA, in five floors, carried from GG-A11-Origin.	183
11	Hierarchy of program-defined acronyms in the GG-Theory program.	250

List of Tables

1	The GG -Theory preprint family and its connection points.	27
2	The four sub-sectors of TGG and their data.	65
3	Alternative OGN-5 -type ledgers, and why MSMF rejects them.	90
4	The GRIP stages and the custody of their data.	151
5	Dual-layer architecture of the four Gi-4 Pillars.	171
6	Logical hierarchy of the Gi-4 Pillars, in dual-layer form.	172
7	The nine dual-layer sub-theorems and the bridge T5	174
8	How each Gi-4 Pillar pairs with its Qi-4 counterpart.	177
9	The five confidence grades used throughout this Part.	184
10	The eighteen rungs: what opens each, what settles, who derives it, at what grade.	185
11	Cost against capacity, for each candidate N	186
12	Which class of Standard-Model parameter lives on which boundary stratum.	190
13	Representative closed forms from the Standard-Model companion.	191
14	The two ground attractors of ordinary matter, side by side.	196
15	The cosmic inventory as closed forms.	199
16	How the four geometric pillars feed the four quantum-information pillars.	207
17	The radiative GRIP-In descent, side by side.	211
18	The ten pre-registered falsifiers of the cohort.	213
19	Comprehensive comparison of contemporary unification frameworks.	224
20	Principal falsifiers, grouped by Pillar.	229
21	Summary of abbreviations used in the present paper.	246
22	Glossary of key mathematical symbols.	251
23	Inward coherence map.	259
24	Outward coherence map.	261
25	Morse data on $\mathbb{CP}^2$, critical point by critical point.	265
26	The geometric gate weight along the generation cascade.	273

Contents

Executive Summary	2
List of Figures	8
List of Tables	9
Preface	19
 Part I. Introduction and Setup	 20
1 Introduction	20
1.1 The strategic moment for the Gi-4 Pillars	20
1.2 The four pillars in overview	22
1.3 Logical relationship to the GG-Constitution	24
1.4 Logical relationship to the Qi-4 Pillars (GG-A7-Qi4)	25
1.5 What this paper does and does not do	26
1.6 Plan of the paper	27
1.7 Relation to other GG-Theory preprints	27
1.8 Methodological note: derivation policy and Euclidean architecture	29
2 This paper's place in the cohort: one link of the open-system chain	30
3 Primitive Inputs from the Companion Papers	31
3.1 The constitutional clauses	31
3.2 The 6D bulk and the BI-6 closure	31
3.3 The OGGEp projector	32
4 The cohort inheritance and the roadmap of the twelve Parts	33
4.1 The cohort inheritance	33
4.2 Roadmap of the twelve Parts	35
5 What it costs to declare your inputs	39
 Part II. TGG	 40
6 TGG: motivation, dual-layer structure, and Structure Theorems	40
6.1 Motivation: from disconnected stress-energies to a single Total Geometric-Gauge object	41
6.1.1 Historical background (with citations)	42
6.2 The seven TGG axioms (five bulk + two boundary)	43
6.3 Formal definition of TGG	45
6.4 First consequences of the TGG Structure Theorems	46
7 The composite connection structure	47
7.1 The principal G_{tot} -bundle on $\mathcal{M}_6$	48
7.2 The connection one-form and its decomposition	49
7.3 The curvature two-form and its decomposition	49
7.4 Gauge equivalence and gauge orbits	50
7.5 Holonomy and global structure	50
8 TGG and the magic-dimension ladder	51
8.1 The magic-dimension ladder	51
8.2 $\text{Spin}(1,5) \cong \text{SL}(2, \mathbb{H})$ and the quaternionic doublet $\mathbb{H}^2$ (cited)	51
8.3 From the quaternionic doublet to family triplicity	52
9 The boundary layer: the unified stress tensor $\mathcal{T}_{\text{TGG}}^{\mu\nu}$	52
9.1 Why a unified boundary stress tensor is needed	52
9.2 Kaluza–Klein zero-mode reduction of the bulk action	53
9.3 The four contributions to $\mathcal{T}_{\text{TGG}}^{\mu\nu}$	54
9.4 The projected modified Bianchi identity (PBI-4)	54
9.5 The boundary layer ready for Theorem T1	55

10	Theorem T1 (T-GRIP-1): Existence, uniqueness, and Master Conservation of TGG	55
10.1	Dual-layer architecture and the compatibility relation $\mathfrak{C}_{\text{TGG}}$	55
10.2	Statement of Theorem T1	55
10.3	Proof of Theorem T1: bulk-layer existence (T1.1, existence part)	56
10.4	Proof of Theorem T1: bulk-layer uniqueness (T1.1, uniqueness part)	57
10.5	Proof of Theorem T1: boundary stress tensor existence and uniqueness (T1.2)	58
10.6	Proof of Theorem T1: Master Conservation Law (T1.3)	58
10.7	Corollary: TGG and DGI	59
10.8	Corollary: TGG and GGEP equivalence	60
10.9	Status of T1 in the constitutional theorem chain	60
11	Worked examples	60
11.1	Example 1: TGG on flat $\mathbb{R}^{1,5}$	60
11.2	Example 2: TGG on the $\mathcal{M}_6$ of GG-Theory	61
11.3	Example 3: explicit Wilson lines and holonomy phases	62
11.4	Chern–Simons construction of the Kalb–Ramond three-form	63
11.5	Gauge-transformation properties of H_3 and X_4	64
11.6	Summary of the algebraic content of TGG	64
12	One object, and what it means that it is unique	65
Part III. CCI		67
13	CCI: motivation, dual-layer structure, and Structure Theorems	67
13.1	Motivation: why a dual-layer pillar is needed	69
13.1.1	Historical background (with citations)	71
13.2	The seven CCI Structure Theorems (algebraic block + physical block)	71
13.3	Dual-layer formal definition of CCI	73
13.4	First consequences of the CCI Structure Theorems	75
14	The four-form anomaly polynomial X_4 : structure and properties	76
14.1	Explicit form of X_4 on the 6D bulk	77
14.2	Cohomological content of $[X_4]$	77
14.3	The integral pairing of $[X_4]$ with the fundamental 4-cycle	78
14.4	The H_3 three-form and BI-6 from CCI	78
15	Derivation of the $-2/5$ coefficients from $\mathbb{H}^2$ normalization	78
15.1	The $-2/5$ coefficients in CCI-ST2 : result and cite	79
15.2	Why the coefficients are not free parameters	79
15.3	Structural significance: how the coefficients select OGN-5	79
16	CCI vs BI-6: algebraic vs differential identity	80
16.1	BI-6 as the differential identity	80
16.2	CCI as the algebraic identity	80
16.3	Logical relationship: CCI implies BI-6	81
16.4	Why both identities are needed	81
17	The physical layer of CCI: 4D covariant conservation and modified TOV	81
17.1	From bulk BI-6 to boundary divergence: the descent map	81
17.2	Master Conservation Law as a sum of sector divergences	82
17.3	The PBI-4 compensator field	83
17.4	Modified TOV equation: the CCI-side identity (primary derivation in A5)	83
18	Theorem T2 (T-GRIP-2): CCI closure under MSMF forces unique OGN-5 ledger and Master Conservation	85
18.1	Statement of T2.1 (algebraic uniqueness): imported from A2-GG-6/GG-A3-BI6	85
18.2	Corollary: family number $N_g = 3$ from upstream CCI	86
18.3	Statement and proof of T2.2 (boundary covariant conservation)	86
19	Connection to Hopkins–Singer differential cohomology	88

19.1	Differential cohomology preliminaries	88
19.2	X_4 as a differential cohomology class	89
19.3	The boundary descent map	89
20	Worked examples	90
20.1	Example 1: explicit X_4 on $\mathcal{M}_6$	90
20.2	Example 2: alternative ledgers ruled out by MSMF	90
20.3	Example 3: integrality check for OGN-5	91
21	Integers that nobody chose	91
 Part IV. GFF		93
22	GFF: motivation, dual-layer structure, and Structure Theorems	93
22.1	Motivation: from algebraic CCI to dual-layer GFF	96
22.1.1	Historical background (with citations)	97
22.2	The seven GFF Structure Theorems (algebraic block + physical block)	97
22.3	Dual-layer formal definition of GFF	101
22.4	First consequences of the GFF Structure Theorems	102
23	The GEP/CEP decomposition	103
23.1	The energetic part GEP	103
23.2	The curvature-entropic part CEP	104
23.3	Bulk-to-boundary descent of the bulk action	104
23.4	The normalization of the descent: $\Phi_\omega = \pi$ and $S_{\text{bulk}} = \frac{\pi}{2} \mathcal{F}_{\text{GEP}}$	105
23.5	The Riemannian volume on projector classes	106
24	The stratified boundary moduli and Morse structure	107
24.1	The stratified structure of the boundary moduli	107
24.2	GFF as a Morse function on each stratum	107
24.3	Family-projector moduli stratum: the principal example	108
24.4	Other strata: brief remarks	110
25	The physical layer of GFF: thermodynamic monotonicity and laboratory falsifiers	110
25.1	The Geometric First Law: from algebraic to thermodynamic	111
25.2	The boundary H-theorem: $dF_{\text{GFF}}/dt \leq 0$ from CPTP structure	112
25.3	Equilibrium uniqueness and the unique vacuum	114
25.4	The per-stratum ground state	115
25.5	Two worked landscapes: what a ground attractor looks like	117
25.6	Laboratory-testable thermodynamic relations	118
26	Theorem T3 (T-GRIP-3): GFF Morse-positivity from CCI and thermodynamic monotonicity from GDOS	119
26.1	Statement of T3.1 (algebraic Morse-positivity)	119
26.2	Outline of the proof	120
26.3	Detailed proof	121
26.4	The N=3 application of T3	122
26.5	Statement and proof of T3.2 (thermodynamic monotonicity)	122
27	The descent structure on the family strata: theorem, obstruction, and the residue M6	124
27.1	$\mathcal{H}_1$ as a theorem	124
27.2	$\mathcal{H}_2$ as a theorem	125
27.3	Consequences for the N=3 proof	126
28	Worked examples	127
28.1	Example 1: GFF on family-projector moduli $\mathbb{CP}^{N_g-1}$	127
28.2	Example 2: GFF on Higgs moduli	127
28.3	Example 3: GFF on gauge-orbit moduli	127
28.4	Example 4: GFF on Wilson-line moduli	127
29	Why there is such a thing as downhill	128

Part V. GRIP	130
30 GRIP: motivation, dual-layer structure, and Structure Theorems	131
30.1 Motivation: from static GFF to dual-layer cascade	133
30.1.1 Historical background (with citations)	134
30.2 The seven GRIP Structure Theorems (algebraic block + physical block) . .	135
30.3 Dual-layer formal definition of GRIP	138
30.4 First consequences of the GRIP Structure Theorems	140
31 The DGI – GSSB – GA cascade in detail	141
31.1 DGI: Dynamical Gauge Invariance	141
31.2 GSSB: Geometric Spontaneous Symmetry Breaking	142
31.3 GA: Geometric Attractor	142
31.4 GGA: Ground Geometric Attractor	143
31.5 GGA and the stability of matter	143
31.6 The two regimes of the runtime: GRIP-In and GRIP-Out	144
31.7 The descent and the fundamental theorem of dynamical systems	147
31.8 The two regimes and the two fibers: the engine’s strokes	148
31.9 The two stages and the two fibers: custody of transport, custody of data .	149
31.10 GRIP ending at the GGA, realized in Nature: the weak interaction	151
32 Cascade dynamics: gradient flow and instanton trajectories	152
32.1 The gradient flow on stratified GFF	152
32.2 Instanton action between adjacent saddles	153
32.3 Saturation of the Bogomolnyi bound	155
33 The physical layer of GRIP: gauge-runtime invariance and the generation weight ladder	155
33.1 The Gauge-Cascade Compatibility Theorem	155
33.2 Cascade-action integration: from $S_{c_i \rightarrow c_{i+1}}$ to the gate weight	156
33.3 Laboratory-testable predictions and falsifiers	158
34 Theorem T4 (T-GRIP-4): GRIP cascade closure forces $N \leq 3$ and gauge-runtime invariance	159
34.1 Statement of T4.1 (algebraic cascade closure)	160
34.2 Proof of T4	161
34.3 Comments on the saturation at $N_g = 3$	162
34.4 Combination with T2 and T3: the three-leg $N=3$ derivation	162
34.5 Statement and proof of T4.2 (runtime gauge-invariance preservation) . . .	162
35 The generation weight ladder from the cascade	164
35.1 Gate weights from gate tolls	165
35.2 Identification with physical generation lifetimes	165
36 Worked examples	166
36.1 Example 1: GRIP on $\mathbb{CP}^2$ (the $N = 3$ case)	166
36.2 Example 2: GRIP on the Higgs moduli	166
36.3 Example 3: GRIP on the gauge-orbit moduli	166
36.4 Example 4 (the golden case): GRIP-In on the orbital stratum — synchrotron damping and binary inspiral, solved to the coefficient	167
37 From a landscape to a history	168
Part VI. Synthesis	170
38 The unified Gi-4 system	170
38.1 The dual-layer architecture: a unified mathematical object	171
38.2 The four pillars as a single mathematical object	172
38.3 Logical hierarchy of the four pillars	173
38.4 The nine dual-layer sub-theorems summarized (T1.1–T4.2)	173
39 Theorem T5 (T-GRIP-5): Constitutional equivalence	173
39.1 Statement of T5	173

39.2	Proof of T5: Gi-4 implies GG-Constitution	175
39.3	Proof of T5: GG-Constitution implies Gi-4	175
39.4	What “equivalent in derivational content” means	176
40	How the four pillars discharge the constitutional axioms	177
41	The geometric backbone and its informational twin	177
41.1	The four-by-four pairing	177
41.2	Geometric vs information-theoretic backbones	178
41.3	The unified Gi-4 + Qi-4 framework	178
42	Four faces, one object	179
 Part VII. Applications		180
43	The ladder, and how to read it	180
43.1	Two pictures of one claim	180
43.2	The reading rule	181
43.3	The five grades, and why they are printed	181
43.4	The spine of this Part	184
44	Floor one — Cosmos: the particles and the sky	184
44.1	The landscape itself: five integers	184
44.2	Why exactly three families [DERIVED]	186
44.2.1	The lower bound: matter and antimatter	186
44.2.2	The upper bound: a descent that cannot afford a fourth step	186
44.2.3	The third leg, and the fence around all three	187
44.2.4	What would end it	187
44.3	Application: the N=3 derivation in Gi-4 vocabulary	187
44.3.1	Physics intuition: families as basins, the weak force as the gate	188
44.3.2	The GRIP mechanism: a budgeted gradient flow	188
44.3.3	Why the GRIP reading helps	189
44.4	The twenty-eight parameters [6 DERIVED, 14 CONDITIONAL, 8 ASSERTED]	189
44.4.1	What the companion is allowed to use	190
44.4.2	Every parameter is one of four things	190
44.4.3	Some of what comes out	191
44.4.4	What the companion does not claim	191
44.4.5	What would end it	191
44.5	Application: the 28 SM parameters in Gi-4 vocabulary	192
44.5.1	Physics intuition: a parameter is the depth of a basin floor	192
44.5.2	The GRIP mechanism: cascade depth becomes exponential hier- archy	192
44.5.3	Why the GRIP reading helps	193
44.6	The electron: a descent that halts [DERIVED]	193
44.6.1	Identity, and why it is not a coincidence	193
44.6.2	Stability: the basin with no exit gate	194
44.6.3	The mass, and the grade it carries	194
44.6.4	What would end it	194
44.7	The proton: a winding that cannot be undone [DERIVED]	195
44.7.1	Why baryon number could not have been a label	195
44.7.2	The twist, and the four locks	195
44.7.3	Two grounds, one grammar	195
44.7.4	The mass, the way stations and the twin	196
44.7.5	What would end it	197
44.8	Application: hierarchy, KK resonance, and dark matter	197
44.8.1	Hierarchy from the radial-scale stratum	197
44.8.2	KK resonance from the Wilson-line moduli	198

	44.8.3	Kaxion dark matter from the $U(1)_{Q_i}$ attractor	198
44.9		The sky: recombination, the inventory, and two redshifts [DERIVED- CONDITIONAL]	199
	44.9.1	The inventory	199
	44.9.2	Recombination as the cosmic GRIP-Out	200
	44.9.3	Two redshifts that are not the same quantity	200
	44.9.4	The second horizon problem, and the bridge upward	200
	44.9.5	What would end it	201
45		Floor two — Chemistry [CONDITIONAL]	201
	45.1	What opened it	201
	45.2	What settles: three rungs	202
	45.3	The fence	202
46		Floor three — Life	202
	46.1	The bioenergetic anchor: where Q_B becomes physical [ANCHOR-DERIVED]	202
	46.2	Four letters, because three families [CONDITIONAL]	203
	46.3	The archive, and the boundary [ANCHOR-DERIVED]	204
47		Floor four — The multicellular and the mind	204
	47.1	Life becomes the engine	204
	47.2	Domains, eukaryotes, multicellularity [CONDITIONAL]	204
	47.3	From a nerve ring to a conscious episode [ASSERTED]	205
48		Floor five — Articulation [EXISTENCE-DEMONSTRATED]	205
49		The other half: Qi-4, and why this paper is not finished without it	206
	49.1	Two backbones, not one	206
	49.2	The relation between a ground attractor and the four certificates	206
	49.3	The Arisaka constant, and the plate's last box	207
	49.4	What each rung gains	208
	49.5	The plain statement	208
50		Interlude: the descent rate, computed	209
	50.1	Application: radiative descent — synchrotron and gravitational-wave inspi- ral as one GRIP-In flow	209
	50.1.1	The leading multipole is the mediator spin: $l_{\text{lead}} = s$	209
	50.1.2	One descent, three terminations: the GRIP-In taxonomy	210
	50.1.3	The dimensionless descent rate	210
	50.1.4	What Gi-4 adds, and the open step	211
51		What the ladder shows	211
	51.1	Where to attack it	212
52		A closing word, for any reader	212

Part VIII. Discussion 215

53		What the four pillars establish	215
	53.1	Problems solved by GG-A6-GRIP	216
54		Significance, uniqueness, and the alternatives	218
	54.1	Significance and uniqueness of the GDOS / Gi-4 framework	218
	54.2	The five pillars of significance	219
	54.3	What is genuinely new	220
	54.4	The measurement-problem carry-over: from QM to QFT to the Standard Model	220
	54.5	Comparison with string theory, AdS/CFT, and Randall–Sundrum	221
	54.6	String theory in a nutshell	221
	54.7	The AdS/CFT correspondence	221
	54.8	The Randall–Sundrum model	222
	54.9	Structural differences in summary	223

54.10	Comprehensive comparison table	223
54.11	Reading the table	226
54.12	Why other approaches missed this	226
55	The Ten Structural Innovations of Gi-4	227
56	Predictions for falsification	228
56.1	The astrophysical falsifier (modified TOV) in detail	229
56.2	The thermodynamic falsifiers (equipartition + FDT) in detail	230
56.3	The generation gate-ordering falsifier (cascade structure) in detail	230
56.4	The cosmological falsifiers: deferred to GG-A10-Cosmos	230
56.5	Falsification logic: a multi-channel commitment	230
57	Open problems	231
57.1	Open mathematical questions	231
57.2	Open physical questions	235
57.3	Open experimental directions	236
57.4	Triage summary	236
58	Future prospects	237
58.1	Future directions for the GG-6 program	237
58.2	Closing thought	238
Part IX. Conclusions		239
59	Methodological synthesis	239
59.1	What Gi-4 has accomplished	239
59.2	The closed-to-open foundational shift	240
59.3	Gi-4 as the language of GDOS	240
60	Conclusion: the Gi-4 Pillars and the foundations of physics	240
60.1	The achievement in summary	240
60.2	The dual-layer revolution: from to	241
60.3	Historical perspective: the four scientific revolutions	241
60.4	The methodological lesson	243
60.5	The unity of physical reality, the role of mathematics, and the road ahead	243
	Acknowledgments	244
Part X. Technical Appendices		245
A	Summary of Abbreviations	245
B	Glossary of Key Symbols	251
C	Differential cohomology preliminaries (pedagogical)	252
C.1	Why differential forms (a five-minute review)	252
C.2	Ordinary cohomology and de Rham cohomology	253
C.3	Integer cohomology and quantization	253
C.4	Differential cohomology as a refinement	254
C.5	The exact sequence	255
C.6	Pontryagin and Chern classes	255
C.7	Relevance to BI-6 and CCI	255
D	Hopkins-Singer machinery for GDOS	256
D.1	Differential function spaces	256
D.2	The descent map	256
D.3	The η -invariant pairing	256
D.4	Application to BI-6 anomaly inflow	256
D.5	Forward to integral L-genus	257
E	Morse theory on stratified spaces	257
E.1	Standard Morse theory	257

E.2	The height function on $\mathbb{CP}^{N-1}$	257
E.3	Stratified Morse theory	258
E.4	Connecting trajectories and instanton actions	258
F	Cross-paper coherence map	258
F.1	Inward map: imports into GG-A6-GRIP from the foundation block	259
F.2	Outward map: exports from GG-A6-GRIP to downstream papers	261
G	Global naming convention for cross-cohort citation	263
H	Detailed worked examples	263
H.1	Spin(1,5) generators and structure constants in detail	264
H.2	Quaternionic doublet matrix representation	264
H.3	CCI integrality check on $\mathcal{M}_6$	264
H.4	GFF Morse data on $\mathbb{CP}^2$	264
H.5	GRIP cascade weights on $\mathbb{CP}^2$	265
I	The free-energy principle: a bridge to life sciences and neuroscience	265
I.1	From Helmholtz free energy to variational free energy	265
I.2	How GFF subsumes Friston's free energy	266
I.3	Empirical predictions for biological systems	266
I.4	Caveat: domain of applicability	267
J	Hopkins-Singer descent: explicit worked examples	267
J.1	The descent procedure: a concrete recipe	267
J.2	Worked example 1: S^2 -bundle over S^2	268
J.3	Worked example 2: 6D bulk to 4D boundary	268
J.4	Worked example 3: descent to a stratum of $\mathcal{M}_{\text{bdy}}$	268
J.5	Why the descent preserves integrality	269
K	Morse theory on $\mathbb{CP}^{N-1}$: explicit worked examples	269
K.1	The height function on $\mathbb{CP}^{N-1}$	269
K.2	Critical points of the height function	269
K.3	Morse index of each critical point	269
K.4	Morse polynomial	270
K.5	Worked example: $N = 2$ ($\mathbb{CP}^1 \cong S^2$)	270
K.6	Worked example: $N_g = 3$ ($\mathbb{CP}^2$)	270
K.7	Worked example: $N_g = 4$ ($\mathbb{CP}^3$)	270
K.8	Pattern: parity of $\dim_{\mathbb{R}} \mathbb{CP}^{N-1}$	271
K.9	Cascade actions on $\mathbb{CP}^{N_g-1}$	271
L	The generation weight ladder: from cascade gate tolls to laboratory observations	271
L.1	From gate toll to a rate: what would be needed	271
L.2	Cumulative cascade toll	272
L.3	Numerical sequence	272
L.4	Why the empirical lifetimes are what they are	273
L.5	Universality across sectors	273
L.6	Falsifier 7 (F-GRIP-7): a particle outside the sequence — retired into Falsifier 4	274
Part XI. Common Objections and Responses		275
M	For the curious reader: the big questions, plainly answered	275
N	The composite connection and TGG (Pillar 1)	287
O	CCI, the integer ledger, and the modified-TOV falsifier (Pillar 2)	289
P	GFF, Morse theory, and the second law (Pillar 3)	291
Q	GRIP, the cascade, and the budget (Pillar 4)	293
R	The weight ladder and the falsifiers	296
S	T5: constitutional equivalence	298

T	The “no new axiom” claim and the promoted hypotheses	299
U	Cohort interdependence and the risk of circular citation	301
V	Conjecture density at the core: M9, M10, M11	302
W	Relationship to string theory, 6D supergravity, holography, and Randall–Sundrum .	303
X	Method, scope, and specific technical soft spots	305

Part XII. References	310
-----------------------------	------------

Preface

Every working physicist learns four habits very early. The first habit, from Maxwell in 1865, is that a current is the thing that flows out of a source and into the rest of the world. The second, from Emmy Noether in 1918, is that every symmetry of the laws of nature corresponds to a quantity that is conserved. The third, from Hilbert and Einstein in 1915, is that the equations of motion descend from a single scalar quantity called the action, which a system extremizes along its trajectory. The fourth, from Boltzmann in 1872 and revived in our own century by Friston and others, is that open systems organize themselves by minimizing a quantity that resembles a free energy. Each habit was learned in a different room, in a different decade, for a different purpose.

What is rarely asked aloud is whether these four habits are really four. The current of charge in electrodynamics and the current of entropy in a living cell are written in the same notation; they pass between different rooms of the building called physics, and at each doorway something quietly changes — a tensor index, a metric, an interpretive convention — and the language continues without comment. The four habits look, when one steps back, like shadows of a single upstream object: one geometric entity that, projected onto different scales and different problems, gives back currents in one chart, conservation in another, an action in a third, and a free-energy functional in a fourth. The unity has been guessed at for a long time; it has not been written down.

This paper proposes to write that single upstream object down. The argument names four geometric pillars on a small six-dimensional manifold: a master object that holds gravity and the other forces together in one structure, a master conservation identity from which every current the laboratory measures descends, a master variational functional from which every action and every free-energy principle in the literature is a special case, and a master pipeline that explains why two observers on the same boundary always agree on the same physics even though they use different coordinates. These four pillars do not add new postulates. They are theorems of the constitutional layer fixed in the parent paper of the program. What they buy is the recognition that the four habits a working physicist learned separately are projections, onto different sectors of the boundary, of one upstream geometric object.

We invite the reader to walk the four pillars in order. Read each statement, follow each proof, and ask whether the projection to the laboratory currents one already knows is honest. After all, what makes a unifying framework worth the name is not the number of disparate phenomena it ties together but the smallness of the upstream object from which all of them descend. The proposal here is that one geometric object, written in the language of differential forms on a six-dimensional bulk, is enough to recover every conservation law, every variational principle, every free-energy functional, and every gauge identity that a century of physics has learned to write down.

Part I

Introduction and Setup

The strategic moment, the bulk geometry, and the GG-Constitution

A paper that claims to derive something has to begin by saying, plainly, what it is allowed to assume. That is the whole business of this Part.

The claim to be supported is that one mechanism drives every symmetry breaking in nature, and that the mechanism can be written down exactly enough to be tested. Establishing that requires a stage to work on, and the stage is not built here: it is inherited, piece by piece, from five companion papers. The constitutional charter — four axioms and seven postulates about what any physical theory must satisfy — comes from [Arisaka \(2026, A1-Constitution\)](#). The six-dimensional geometry comes from [Arisaka \(2026, A2-GG6\)](#). The proof that the geometry is free of the internal inconsistencies that usually kill such constructions comes from [Arisaka \(2026, A3-BI6\)](#), and the demonstration that it leaves no continuous parameter free to adjust comes from [Arisaka \(2026, A4-Euclid\)](#). The law governing how any of it looks to an observer who is part of the system rather than outside it comes from [Arisaka \(2026, A5-GDOS\)](#). None of the five is re-derived in this paper. Each is stated here, once, in the form this paper uses, with the theorem that supplies it named.

One point in that inheritance deserves saying slowly, because it is where a reader from outside physics most often stops. The theory is six-dimensional, but the two dimensions beyond the four we inhabit are not further directions one could travel in. One of them is a scale: it records how a description of the world changes as one looks at it more or less closely, the way the effective properties of a material differ between a micrometer and a nanometer. The other is a phase: an angle that has no length. So “six-dimensional” here is not a claim that space has hidden rooms. It is a claim that two familiar features of physical description — scale, and phase — behave geometrically, and can be treated as directions in the same sense that time is treated as one in relativity.

This Part then locates the paper in the twenty-two-paper cohort and maps the twelve Parts that follow. A reader who wants the mechanism itself can take this Part’s final section for orientation and go directly to Part V, where the mechanism is built, returning to Parts II–IV for the machinery when a step needs justifying. A reader who intends to argue with the paper should read this Part in full, because every input the later argument leans on is declared here and is not declared again.

1 Introduction

1.1 The strategic moment for the Gi-4 Pillars

The GG-Theory program has reached a stage where its derivational reach is established but its foundational vocabulary is not yet rigorous.

In the past year, the program has produced a sequence of substantive results: the rigorous

derivation of the OGN-5 ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ from BI-6 closure on the 6D bulk (Arisaka, 2026, A3-BI6); the explicit codebook for all 28 Standard-Model parameters in terms of the OGN-5 integers (Arisaka, 2026, A9-SM); the three-leg derivation of the family number $N_g = 3$, combining MSMF, GDOS Morse stability, and the codebook self-consistency check (Arisaka, 2026, P2-N3); the prediction of a ~ 3.85 TeV KK resonance cluster as the principal collider falsifier (Arisaka, 2026, P1-KK); and the prediction of a ≈ 8.1 μeV Kaxion-like dark-matter candidate as the principal cosmological falsifier (Arisaka, 2026, C1-Kaxion). Each of these derivations rests on a common dynamical mechanism, identified in the methodological reflection of Arisaka (2026, P2-N3, Section 12) as the open-system boundary-projection cascade $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$, with the dynamics governed by a Geometric Free-Energy Functional $\mathcal{F}_{\text{GFF}} = \mathcal{F}_{\text{GEP}} + \mathcal{F}_{\text{CEP}}$ on the boundary moduli, and the closure conditions inherited from the BI-6 anomaly polynomial X_4 .

The mechanism works. It produces results that closed-system QFT and GR cannot reach. But the language in which it is articulated – the names GFF, GRIP, GA, DGI, GSSB, the boundary-projector moduli, the Hopkins–Singer descent and the M-theory flux-quantization condition (Witten, 1996) – has been deployed informally rather than formally axiomatized. We have used these names as labels pointing in the right direction, but we have not yet written them down as mathematical objects whose properties can be checked against the Constitutional axioms.

This is the situation that mathematical physics historically encounters when a program develops faster than its foundations can be written. Newton’s *Principia* could not be fully understood until the calculus was formalized by Leibniz, the Bernoullis, and Euler; Maxwell’s *Treatise* could not be fully understood until vector calculus was developed by Heaviside and Gibbs; Einstein’s general relativity could not be fully understood until Ricci and Levi-Civita developed the tensor calculus to its modern form; and Schrödinger’s and Heisenberg’s quantum mechanics could not be fully understood until von Neumann gave the Hilbert-space formalism in 1932 (von Neumann, 1932). In each case, a deep theoretical advance was made first and the formal language was written second. This is the same pattern in which we now find ourselves.

The present paper, the second of the GDOS trio (with GG-A5-GDOS GDOS Foundation and GG-A7-Qi4 Qi-4 Pillars), undertakes the writing of the formal language for the geometric backbone of GDOS. We call this language the Gi-4 Pillars: four explicitly named formal mathematical objects – TGG, CCI, GFF, GRIP – whose Structure Theorems, definitions, theorems, and worked examples constitute the rigorous foundation that the upstream papers (GG-A9-SM, GG-2026-Generations, etc.) implicitly assume. Once Gi-4 is in place, every use of GFF or GRIP or TGG in a downstream paper acquires citable rigorous content, and the upstream derivations become reproducible from first principles.

This is, in our view, the most important methodological move that the GG-Theory program has yet undertaken. It is the move that converts the program from an exciting research program into a foundational theory in the technical sense.

Remark 1.1 (Physical intuition for the broader scientific community). The Gi-4 formalism is presented in the language of differential geometry and gauge theory because that is the language in which the GG-Theory bulk content was originally derived. But the structural content of Gi-4 – a unified bulk-to-boundary projection that descends a single conserved object onto the observable boundary,

with a free-energy functional governing the dynamical landscape and a cascade of attractors selecting the discrete catalog of stable observables – is not specific to particle physics. Readers from cosmology will recognize the bulk-to-boundary descent as a generalization of holography familiar from AdS/CFT; readers from condensed-matter physics will recognize the cascade of unstable critical points as the universal renormalization-group flow of phase transitions; readers from chemistry will recognize the GFF free-energy functional with its attractor catalog as the Gibbs landscape with its reaction products; readers from biology will recognize the open-system irreversible projection as the second-law analog that selects living configurations from the equilibrium manifold; and readers from neuroscience will recognize the cascade $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$ as the universal pattern by which a richly symmetric prior collapses into a discrete observed perception. The mathematical content of Gi-4 is the same in all these contexts; only the specific bulk and boundary objects change. We expect the present formalism to be useful well beyond the particle-physics community, and we have tried to write it in a way that makes the physical intuition accessible at each step.

1.2 The four pillars in overview

The four Gi-4 Pillars are introduced in the present paper as follows. We give here only a one-paragraph sketch of each; the formal Structure Theorems and definitions occupy the body of the paper.

TGG – The Total Geometric-Gauge object. TGG is the unique geometric object that unifies, into a single mathematical entity, the bulk-level composite connection on the 6D bulk and the 4D boundary stress tensor obtained by zero-mode reduction. Its *bulk realization* is a principal G_{tot} -bundle $P \rightarrow M_6$ with structure group $G_{\text{tot}} = \text{Spin}(1, 5) \times \text{SU}(5) \times \text{U}(1)_B \times \text{U}(1)_{Q_i}$ (using the magic-dimension-ladder identification $\text{Spin}(1, 5) \cong \text{SL}(2, \mathbb{H})$ at $d = 6$), equipped with the composite connection one-form $\mathcal{A}_{\text{tot}} = \omega_{\text{Spin}(1,5)} + A_5 + A_B + A_{Q_i}$ and curvature $\mathcal{R}_{\text{tot}} = d\mathcal{A}_{\text{tot}} + \mathcal{A}_{\text{tot}} \wedge \mathcal{A}_{\text{tot}}$. Its *boundary realization* is the unified 4D stress tensor

$$\mathcal{T}_{\text{TGG}}^{\mu\nu} = T_{\text{matter}}^{\mu\nu} + T_{\text{gauge}}^{\mu\nu}[F] + X_{\text{GS}}^{\mu\nu} + T_{\text{KK}}^{\mu\nu},$$

the boundary projection of all bulk stress-energy via the BI-6 bulk-to-boundary descent. The two layers are tied together by the *Master Conservation Law*:

$$\nabla_\mu \mathcal{T}_{\text{TGG}}^{\mu\nu} = 0,$$

which holds iff the full Ward package (electric, baryonic, informational, and metric currents) is satisfied. In the GRIP cascade, the unbroken-symmetry TGG phase on the bulk is the DGI stage.

CCI – The Covariant Conservation Identity. CCI is the 4D boundary identity that emerges as the OGGEF zero-mode projection of TGG conservation. It has two equivalent formulations. The *physical/local* formulation is the covariantly conserved divergence equation $\nabla_\mu \mathcal{T}_{\text{TGG}}^{\mu\nu} = 0$ on M_4 with explicit compensator scalar field, holonomy current, and modified Tolman–Oppenheimer–Volkoff equation as astrophysical falsifier. The *algebraic/global* formulation is the differential-cohomology

integrality $[X_4] \in H^4(M_6; \mathbb{Z})$ together with the rank-one factorization of the anomaly matrix $M^{\text{BI-6}} = v \otimes v^T$, where

$$X_4 = \text{tr } R^2 - \text{Tr } F_5^2 - \frac{2}{5} F_B^2 - \frac{2}{5} F_{Q_i}^2.$$

Under MSMF, CCI closure uniquely fixes the OGN-5 ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$. The $-2/5$ coefficients are not free parameters: they follow from the trace-ratio identity $\text{Tr } B_0^2 / \text{Tr } Y_0^2 = (1/3)/(5/6) = 2/5$ in the fundamental **5** of $\text{SU}(5)$. The two formulations of CCI are equivalent by Hopkins–Singer differential-cohomology descent.

GFF – The Geometric Free-Energy Functional. GFF is the unique scalar functional $\mathcal{F}_{\text{GFF}} = \mathcal{F}_{\text{GEP}} + \mathcal{F}_{\text{CEP}}$ on the stratified boundary moduli space $\mathcal{M}_{\text{bdy}}$. $\mathcal{F}_{\text{GEP}}$ (Geometry-Gauge Energy Principle) is the energetic part, derived by bulk-to-boundary descent of the TGG bulk action through the Wilson-line holonomy on the compact S_θ^1 . $\mathcal{F}_{\text{CEP}}$ (Curvature-Entropy Principle) is the curvature-entropic part from the Riemannian volume of admissible projector classes plus the GDOS diffusion term. GFF has two equivalent characterizations. The *thermodynamic* characterization: $\mathcal{F}_{\text{GFF}}$ is monotonically decreasing under irreversible boundary projection, $\partial_t \mathcal{F}_{\text{GFF}} \leq 0$, the open-system second-law analog. The *Morse-theoretic* characterization: $\mathcal{F}_{\text{GFF}}$ is a perfect Morse function on each stratum of $\mathcal{M}_{\text{bdy}}$ with strict alternating Morse-index positivity inherited from CCI via the Hopkins–Singer differential-cohomology descent. Critical points of $\mathcal{F}_{\text{GFF}}$ are GA’s (Geometric Attractors). Strict Morse-positivity of GFF is the property that, in the N=3 proof of Arisaka (2026, P2-N3), forces the family-projector moduli to have real dimension divisible by four.

GRIP – The Gauge-Runtime Invariance Pipeline. GRIP is the universal three-stage lifecycle cascade $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$ on the GFF Morse landscape. DGI (Dynamical Gauge Invariance) is the unbroken-symmetry TGG phase on the bulk, inheriting the full G_{tot} symmetry of the composite connection. GSSB (Geometric Spontaneous Symmetry Breaking) is the chain of intermediate symmetry-breaking saddles connected by Bogomolnyi-saturating gradient-flow trajectories with instanton action πk between adjacent index- $2k \rightarrow 2(k-1)$ saddles. GA (Geometric Attractor) is the stable boundary observable, the Morse minimum, the cascade endpoint. The pipeline combines two complementary properties: *gauge-invariance preservation* – each stage of the cascade preserves gauge covariance (POS-1 of the GG-Constitution), so the entire runtime is gauge-invariant; and *finite-budget closure* – the total cascade action $S_{\text{tot}}(N_g) = \pi N_g(N_g - 1)/2$ on a stratum of complex dimension $N_g - 1$ is bounded by the BI-6 instanton-action budget $S_{\text{BI6}}^{\text{max}} = 3\pi$, delivered by the capacity theorem T-N3-6 of Arisaka (2026, P2-N3) (open problem M10, whose demand it meets; residue M6 = R-INT, the topological residue of its integrality theorem L-INT). Together these constraints force $N_g \leq 3$ on the family-projector stratum, contributing the upper bound in the three-leg N=3 derivation.

The unified Gi-4 system. The four pillars together form a unified geometric backbone of GDOS. TGG provides the bulk-level invariance; CCI provides its closure conditions; GFF provides the dynamical landscape on the boundary moduli; GRIP provides the cascade dynamics that selects boundary observables. No additional physical axiom is introduced: the four pillars are derived from

the GG-Constitution ([Arisaka, 2026, A1-Constitution](#)) – specifically from GGEP (geometry–gauge equivalence), OGGE (its open-system extension), MSMF (minimal sufficient manifold form), and the BI-6 closure of [Arisaka \(2026, A3-BI6\)](#).

1.3 Logical relationship to the GG-Constitution

The Gi-4 Pillars are not new postulates of the program. They are *theorems* of the GG-Constitution, in the sense that they can be derived from the four constitutional axioms (LC, GGEP, OGGE, MSMF) together with the BI-6 closure of [Arisaka \(2026, A3-BI6\)](#). We articulate this relationship explicitly here because it is the methodological core of the present paper.

The four constitutional axioms of GG-Theory, as stated in [Arisaka \(2026, A1-Constitution\)](#), are:

- **LC** (Local Causality): the bulk admits a Lorentzian structure with retarded propagation of admissible kernels. This is the axiom of relativistic compatibility.
- **GGEP** (Geometry–Gauge Equivalence Principle): gauge content and geometric content are unified in a single primitive composite connection. This is the axiom that forces the existence of TGG.
- **OGGE** (Operational/Open Geometry–Gauge Equivalence Principle): observable boundary content is related to bulk geometric content by an irreversible CPTP bulk-to-boundary projection ([Susskind, 1995](#)). This is the axiom that forces the existence of the boundary-projector moduli on which GFF lives.
- **MSMF** (Minimal Sufficient Manifold Form): among admissible structures consistent with the boundary observable spectrum, GG-Theory selects the one of minimum total complexity. This is the axiom that forces the OGN-5 ledger to be the minimum integer solution of CCI.

The four pillars are derived from these axioms as follows.

TGG is the formal mathematical realization of GGEP extended by OGGE to the boundary: the bulk-level composite connection (forced by GGEP and the magic-dimension ladder) and its 4D zero-mode projection to the unified stress tensor $\mathcal{T}_{\text{TGG}}^{\mu\nu}$ (forced by OGGE together with BI-6 descent). The Master Conservation Law $\nabla_\mu \mathcal{T}_{\text{TGG}}^{\mu\nu} = 0$ is given by Theorem 10.1 of Part II, derived from GGEP, OGGE, LC, and the BI-6 closure of [Arisaka \(2026, A3-BI6\)](#).

CCI is the formal mathematical realization of the boundary-level content of OGGE: the 4D conservation identity that emerges as the zero-mode projection of TGG conservation. Its physical content (astrophysical falsifier via modified TOV) and its algebraic content (integrality $[X_4] \in H^4(M_6; \mathbb{Z})$ with rank-one factorization) are equivalent by Hopkins–Singer descent. CCI closure under MSMF forces the OGN-5 ledger uniquely; this is Theorem 18.1 of Part III, derived from CCI and MSMF.

GFF is the formal mathematical realization of the dynamical content of OGGE: the free-energy functional on the boundary moduli, with both thermodynamic ($\partial_t \mathcal{F} \leq 0$) and Morse-theoretic characterizations. Its existence is given by Definition 22.8 and the five GFF Structure Theorems of Section 22.2; its Morse-positivity is given by Theorem 26.1 of Part IV, derived from OGGE and CCI via Hopkins–Singer descent.

GRIP is the formal mathematical realization of the lifecycle content of OGGEF: the gauge-invariant cascade $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$ on the GFF Morse landscape with finite-budget closure. Its existence as a perfect gradient flow on $\mathcal{M}_{\text{bdy}}$ and its closure within $S_{\text{tot}} \leq S_{\text{BI6}}^{\text{max}} = 3\pi$ are given by Theorem 34.1 of Part V, derived from OGGEF, GFF Morse-positivity, and CCI.

The bridge theorems of Part VI then assemble the four pillars into a unified system equivalent in content to the GG-Constitution. Specifically:

Theorem 1.2 (Constitutional equivalence, sketched here). *The Gi-4 Pillars system (TGG, CCI, GFF, GRIP), equipped with the five principal theorems T1–T5 of this paper, is equivalent in derivational content to the GG-Constitution (LC, GGEP, OGGEF, MSMF) together with the BI-6 closure of Arisaka (2026, A3-BI6). Every closed-derivation result of the GG-Theory program – including $N=3$, the 28 Standard-Model parameters, the hierarchy, and the dark-matter prediction – can be obtained from either the constitutional axioms or the Gi-4 Pillars; the two systems are different presentations of the same mathematical content.*

The full proof of Theorem 1.2 is given in Part VI. The point of the theorem is methodological: the Gi-4 reorganization of the present paper is not a new theory, but a *geometric reading* of the constitutional axioms that makes the mathematical structure explicit and citable.

1.4 Logical relationship to the Qi-4 Pillars (GG-A7-Qi4)

The four Gi-4 Pillars (TGG, CCI, GFF, GRIP) form the *geometric* backbone of GDOS. The companion paper GG-A7-Qi4 (Arisaka, 2026, A7-Qi4) formalizes four *information-theoretic* pillars – the Qi-4 Pillars – whose role is to encode the quantum-information / measurement-theoretic content of GDOS. The two sets of four pillars are parallel, and their parallel structure is fundamental to the design of GDOS as a foundational framework that unifies geometry and information.

Without anticipating the detailed content of GG-A7-Qi4, the parallel may be sketched as follows:

- TGG (Total Geometric-Gauge object: bulk composite connection + 4D boundary stress tensor with master conservation) is paired with **Qi-QUA** (Total Quantum-Information generator). Both encode the deepest unbroken structure on the bulk: TGG on the geometric side, Qi-QUA on the entanglement / quantum-information side.
- CCI (Covariant Conservation Identity: 4D OGGEF-projected conservation equation, with both falsifiable physical and integral algebraic formulations) is paired with **Qi-EP** (Quantum-information Equivalence Principle). Both encode the boundary-level identity that the master generator must satisfy: CCI on the conservation/anomaly side, Qi-EP on the information-conservation under measurement side.
- GFF (Geometric Free-Energy Functional, with thermodynamic $\partial_t F \leq 0$ and Morse-theoretic characterizations) is paired with **Qi-CON** (Quantum-information Conservation on boundary projections). Both encode the dynamical landscape on the boundary: GFF on the free-energy / Morse side, Qi-CON on the mutual-information / relative-entropy side.
- GRIP (Gauge-Runtime Invariance Pipeline: cascade $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$ with finite-budget closure) is paired with **Qi-HAL** (Quantum-information Holographic projection). Both encode

the cascade dynamics from bulk to boundary: **GRIP** on the Morse-cascade / instanton-action side, **Qi-HAL** on the Stinespring dilation / decoherence side.

The two sets of pillars are not redundant: each captures content that the other does not. **Gi-4** captures the bulk-to-boundary geometric structure (gauge connections, anomaly polynomials, Morse landscapes, attractor cascades). **Qi-4** captures the bulk-to-boundary information-theoretic structure (entanglement, mutual information, decoherence (Zeh, 1970), measurement outcomes). Together they constitute a unified foundational framework for **GDOS**, in which both the geometric and the information-theoretic content of physical observation are made foundational at the same time.

This unified framework is what we argue, in Arisaka (2026, P2-N3, Section 12), has been the missing link in the foundations of physics for the century since the establishment of quantum mechanics in 1925–1930. Quantum mechanics in its closed-system form unifies the information-theoretic content (the state vector) with the dynamical content (unitary evolution), but it does so on a fixed background and without bulk-to-boundary projection. General relativity unifies the geometric content (the metric) with the dynamical content (Einstein’s equations), but it does so without information-theoretic content. **GDOS** unifies all three – geometric, information-theoretic, and dynamical – under a single foundational principle (**OGGEP**) that elevates the bulk-to-boundary projection to axiomatic status. The **Gi-4** Pillars (the present paper) and the **Qi-4** Pillars (**GG-A7-Qi4**) are the two sides of this unification.

1.5 What this paper does and does not do

The present paper has a clearly bounded scope. It does the following.

- (1) It formalizes the four **Gi-4** Pillars as mathematical objects with Structure Theorems, definitions, principal theorems, and worked examples.
- (2) It derives the four pillars from the **GG-Constitution** and the **BI-6** closure, without introducing any additional physical axioms.
- (3) It states and proves the five principal theorems **T1–T5** (existence and uniqueness of **TGG**; **CCI-OGN5** uniqueness; **GFF** Morse-positivity from **CCI**; **GRIP** cascade closure; constitutional equivalence).
- (4) It establishes the bridge theorems to **OGGEP/MSMF** (the constitutional axioms) and to **Qi-4** (the information-theoretic backbone of **GG-A7-Qi4**).
- (5) It applies the **Gi-4** framework to the derivations in **GG-A9-SM** (28 SM parameters) and **GG-2026-Generations (N=3)**, demonstrating that those derivations acquire full rigorous status once the **Gi-4** vocabulary is in place.

What it does *not* do is also bounded.

- (1) It does not undertake the full **Qi-4** axiomatization; that is the content of **GG-A7-Qi4**. The **Qi-4** Pillars are mentioned in Section 1.4 and at the bridge theorems of Part VI, but their detailed development is deferred.

- (2) It does not yet provide rigorous proofs of the residual technical hypotheses $\mathcal{H}_1, \mathcal{H}_2$ flagged in [Arisaka \(2026, P2-N3, Appendix E\)](#). Those hypotheses become *theorems* of Gi-4 once the pillars are formalized, but the detailed implementation of the Hopkins–Singer descent in the 6D bulk is carried by Appendix H of [Arisaka \(2026, P2-N3\)](#).
- (3) It does not undertake the full re-derivation of the 28 SM parameters in the new Gi-4 vocabulary; that is the role of the revised GG-A9-SM ([Arisaka, 2026, A9-SM](#)), with the present paper providing the rigorous foundation that the revised GG-A9-SM rests on.
- (4) It does not address the experimental falsifiers of GG-Theory (KK resonance cluster, Kaxion dark matter); those are the subject of ([Arisaka, 2026, P1-KK](#); [Arisaka, 2026, C1-Kaxion](#)).

The discipline of the present paper is therefore: *austere formalization of Gi-4 from the constitutional axioms, with applications to the existing derivational results.*

1.6 Plan of the paper

The twelve Parts of this paper are laid out in the Roadmap of §4.2 (Figure 4), which names each Part in turn and highlights the load-bearing four-pillar core; the detailed section-by-section contents follow in the Table of Contents at the front of the paper.

1.7 Relation to other GG-Theory preprints

The present paper (GG-A6-GRIP) is part of a coordinated program of preprints developing the foundations and applications of GG-Theory. For the convenience of the reader, we summarize the principal preprints, their content, and their relation to the present paper.

Table 1: The GG-Theory preprint family and its connection points to the present paper. Each row names a preprint, what it contains, and the one place where it touches this one — an import consumed here, an export supplied to it, or an object both describe. The list is a roster and not a bibliography: the full entries, each with its relevance note, are in the references.

Preprint	Title (short)	Content	Connection to GG-A6-GRIP
<i>Foundational documents (the GDOS quartet)</i>			
GG-A1-Constitution	GG-Constitution	Four Axioms (LC, GGEP, OGGEF, MSMF) + 7 Postulates	The single foundational layer. All Structure Theorems and Principal Theorems of are derived from these.
GG-A2-GG6	GG-6 Foundation	6D bulk M_6 , Spin(1,5) frame bundle, magic-dimension ladder $\mathbb{R} \rightarrow \mathbb{C} \rightarrow \mathbb{H} \rightarrow \mathbb{O}$	The bulk arena on which TGG (Part II) is constructed; provides the bundle-theoretic content of GGEP.
GD-2026-3	BI-6 Closure	Modified Bianchi identity $dH_3 = X_4$ on the bulk; OGN-5 ledger (1, 3, 5; 7, 35)	Closes the algebraic side of CCI (Part III). T2.1 is the algebraic uniqueness theorem of OGN-5; T2.2 descends BI-6 to the boundary.

continued on next page

Table 1 – continued from previous page

Preprint	Title (short)	Content	Connection to GG-A6-GRIP
GG-A5-GDOS	GDOS Foundation	Geometric Dynamics of Open Systems: boundary projector, bulk-to-boundary descent, CPTP boundary	First informal introduction of GFF and GRIP; the present paper formalizes these as Pillars 3 and 4 with full Structure Theorems and Principal Theorems.
GG-A6-GRIP	(this paper) Gi-4 Pillars (geometric backbone)	Four Pillars TGG, CCI, GFF, GRIP with 28 Structure Theorems and 9 Principal Theorems (T1.1–T5)	Provides the formal geometric vocabulary of GDOS.
GG-A7-Qi4	Qi-4 Pillars (information backbone)	Four information-theoretic Pillars: Qi-QUA, Qi-EP, Qi-CON, Qi-HAL	Companion paper. The information backbone that pairs with the geometric backbone of GG-A6-GRIP; together they form the complete Gi-4+Qi-4 foundational language.
GG-A8-QM	White Paper for GDOS	Programmatic synthesis of GG-Theory’s significance and falsification program	High-level synthesis citing GG-A6-GRIP as the formal geometric reference.
<i>Predictive applications</i>			
GG-A9-SM	SM 28 parameters code-book	Explicit derivation of all 28 SM parameters from OGN-5 ledger	Application of CCI/GRIP cascades. The 28 parameters are cascade depths of GRIP on different boundary strata. T2.1 forces the OGN-5 ledger that fixes the code-book.
GG-A10-Cosmos	GG-Cosmos for Λ CDM	Cosmological observables from GG-6 boundary descent	Application of CCI/GFF cosmological projections. Modified Friedmann equations and $w(z)$, $f\sigma_8(z)$ shifts from X_{GS} .
GG-2026-Generations	N=3 three-leg proof	The three-leg derivation of $N_{\text{gen}} = 3$	Cascade structure carried by T2.1+T3.1+T4.1 of the present paper (T4.1 conditional on M10). The would-be hypotheses $\mathcal{H}_1, \mathcal{H}_2$ are dissolved in that paper: $\mathcal{H}_1$ closed in the negative (its rigidity lemma generalizes the naive-descent no-go to every functorial route), $\mathcal{H}_2$ succeeded by the integrality theorem L-INT (residue R-INT).
<i>Specific predictions and falsifiers</i>			
GG-P1-KK	~ 4 TeV KK particles	Prediction of the KK resonance cluster $M_{\text{lock}} \simeq 3.85$ TeV	LHC-falsifiable prediction of the Bl-6 closure; references GFF/GRIP for the scale-stratum dynamics.

continued on next page

Table 1 – continued from previous page

Preprint	Title (short)	Content	Connection to GG-A6-GRIP
GG-C1-Kaxion	$\approx 8.1 \mu\text{eV}$ Kaxion DM	Prediction of ultra-light Kaxion dark matter	Direct-detection-falsifiable prediction from the $U(1)_{Q_i}$ attractor; uses GRIP cascade on the compact-angle moduli.
<i>Books and synthesis</i>			
GDOS-Book-I	GG-Theory Foundation Book	Comprehensive book-length treatment	Cites the present paper as the formal Gi-4 reference; intended as the eventual graduate-textbook treatment of GG-Theory.

Reading paths. A reader new to GG-Theory has several reasonable entry points:

- **Foundations-first path:** GG-A1-Constitution (Constitution) $\rightarrow$ GG-A2-GG6 (GG-6 Foundation) $\rightarrow$ GD-2026-3 (BI-6) $\rightarrow$ the present paper (Gi-4 Pillars) $\rightarrow$ GG-A7-Qi4 (Qi-4 Pillars). This is the recommended path for theoretical physicists who want to see the full derivational chain.
- **Predictions-first path:** GG-A9-SM (SM parameters) $\rightarrow$ GG-2026-Generations (N=3 proof) $\rightarrow$ GG-P1-KK (KK prediction) $\rightarrow$ GG-C1-Kaxion (Kaxion DM). This is the recommended path for experimentalists who want to see the falsifiable predictions before the formal scaffolding.
- **Pedagogical path (recommended for graduate students):** GG-A1-Constitution (Constitution) $\rightarrow$ Appendix G of the present paper (pedagogical tour of the dual-layer architecture) $\rightarrow$ Appendix N (glossary for life sciences and neuroscience) $\rightarrow$ Parts II–V of the present paper (the four Pillars) $\rightarrow$ companion predictive papers.
- **Synthesis path:** GG-A8-QM (White Paper) $\rightarrow$ GDOS-Book-I (Foundation Book) $\rightarrow$ specific preprints as needed. This is the recommended path for those who want a high-level overview before diving into derivations.

Cross-paper consistency. The dual-layer architecture of (algebraic block + physical block per Pillar, paired Theorems $T^{*.1} + T^{*.2}$, named compatibility links) is the universal organizing principle that propagates across all preprints in the GG-Theory family. Future versions of the companion preprints will adopt the same Structure Theorem terminology established in the present paper.

1.8 Methodological note: derivation policy and Euclidean architecture

The present paper follows the Euclidean axiomatic architecture canonical to the GG-Theory program: numbered Definitions, Postulates, Structure Theorems, Lemmas, Principal Theorems, and Corollaries, with proofs marked ■. Each Principal Theorem is dual-layer (algebraic statement $TN.1$ + operational statement $TN.2$, linked by a named compatibility relation); each is followed by a Physics-Intuition remark and at least one worked example. “Axiom” is reserved exclusively for the four constitutional axioms (LC, GGEP, OGGEF, MSMF) of the GG-Constitution; everything else is a theorem, postulate, or structure theorem. The full methodological framing, including the

historical-Euclidean motivation and the relation to the four prior revolutions of physics, lives in (Arisaka, 2026, A4-Euclid, §§1–2) and (Arisaka, 2026, A1-Constitution, Part I).

2 This paper’s place in the cohort: one link of the open-system chain

This paper establishes the four geometric Gi-4 pillars — TGG, CCI, GFF, and GRIP — and the runtime by which a boundary system selects and stabilizes its branch; that is its own subject. For readers following the whole GG-Theory cohort, it is also one link of a single open-system chain that the consequence papers cash out. Because no physical system is closed, the four-dimensional world is a lawful projection (OGGEP) of the six-dimensional bulk, and that projection is observation itself; its dynamics is the boundary law GDOS, forced by the geometric Gi-4 and quantum-information Qi-4 pillars, so that the laws, the particles, and the contents of the universe are *observables* of the projection rather than free inputs. In that chain, the present paper supplies the geometric Gi-4 pillars that, with the Qi-4 pillars, force the GDOS law. The funnel below is shared, in structure, with the companion papers GG-A9-SM, GG-A10-Cosmos, and GG-A11-Origin; here it is offered only as orientation, and the rest of this paper proceeds on its own terms.




Figure 1: The open-system chain shared across the GG-Theory cohort. Because no physical system is closed, the four-dimensional world is a lawful projection (OGGEP, AX-3) of the six-dimensional bulk, and that projection is observation itself; its dynamics is GDOS, forced by the Gi-4 and Qi-4 pillars, so the laws and contents of the universe are observables rather than free inputs. The same funnel opens the consequence papers GG-A9-SM, GG-A10-Cosmos, and GG-A11-Origin, with a paper-specific payload; within it the present paper establishes the geometric Gi-4 pillars that force GDOS.

3 Primitive Inputs from the Companion Papers

The four pillars of Gi-4 formalize a structure that is already present, but only informally articulated, in four companion papers. We collect the primitive inputs here in compact form, with cross-references to the canonical home of each input. Three blocks suffice: the constitutional clauses, the 6D bulk and BI-6 closure, and the OGGE projector that connects the two.

3.1 The constitutional clauses

The GG-Constitution ([Arisaka, 2026, A1-Constitution](#)) comprises four axioms and seven postulates. We list them as one-line entries; the full treatment is in [Arisaka \(2026, A1-Constitution\)](#), Sections 3–13.

- **AX-1 LC** (Local Causality): admissible kernels propagate in forward light-cones.
- **AX-2 GGEP** (Gauge–Geometry Equivalence Principle): gauge and geometric content arise from a single composite connection on the bulk.
- **AX-3 OGGE** (Open-system Gauge–Geometry Equivalence Principle): observable boundary content is a CPTP bulk-to-boundary projection of bulk geometric content; the projection is irreversible and is the foundational dynamics of physical observation.
- **AX-4 MSMF** (Minimal Sufficient Manifold Form): among admissible structures, select the one of minimum total complexity.
- **POS-1–POS-7**: the seven postulates GI, UEP, AF, CP, RG, OGN, MF are theorem-level consequences of AX-1–AX-4 on the bulk and boundary; see [Arisaka \(2026, A1-Constitution\)](#) Sections 7–13.

3.2 The 6D bulk and the BI-6 closure

The 6D bulk ([Arisaka, 2026, A2-GG6](#)) carries the four primitive geometric quantities used throughout this paper. The bulk metric is

$$ds_6^2 = e^{-2r} \eta_{\mu\nu} dx^\mu dx^\nu + dr^2 + R_\theta^2 d\theta^2, \quad (1)$$

with $\text{Spin}(1, 5) \cong \text{SL}(2, \mathbb{H})$ acting on the 4D Minkowski slice (magic- dimension ladder, [Arisaka \(2026, A2-GG6\)](#) Theorem 1). The boundary $\partial M_6 = M_4$ sits at $r = 0$ (the UV end, where $r = \ln(E_{\text{GUT}}/\mu) = 0$, i.e. $\mu = E_{\text{GUT}}$), and the IR end at $r = K$ realizes the scale hierarchy $E_{\text{Pl}}/M_{\text{lock}} = 2e^{35}$ with $K = 35$ the OGN-5 ledger integer. The total bulk gauge group is

$$G_{\text{tot}} = \text{Spin}(1, 5) \times \text{SU}(5) \times \text{U}(1)_B \times \text{U}(1)_{Qi}. \quad (2)$$

The BI-6 closure ([Arisaka, 2026, A3-BI6](#)) is the differential identity

$$dH_3 = X_4 = \text{tr } R \wedge R - \text{Tr } F_5 \wedge F_5 - \frac{2}{5} F_B \wedge F_B - \frac{2}{5} F_{Qi} \wedge F_{Qi}, \quad (3)$$

the GGEP completion of the 1984 Green–Schwarz construction (see Arisaka (2026, A3-BI6, § 8.7, “BI-6 as the GGEP completion of GS: the geometrization reading”). The four-component coefficient vector $(a, b, c, d) = (1, -1, -2/5, -2/5)$ and the trace-ratio lock $|a/c| = |a/d| = 5/2$ are pinned by MSMF closure on the bulk integrality $[X_4] \in H^4(M_6; \mathbb{Z})$, with the unique minimum-integer solution being the OGN-5 ledger

$$(M, N, L; J, K) = (1, 3, 5; 7, 35), \quad (4)$$

satisfying the two algebraic identities

$$K = M^2 + N^2 + L^2 = 1 + 9 + 25 = 35, \quad K = LJ = 5 \cdot 7 = 35. \quad (5)$$

The physical readings of the five integers are $M = 1$ the primitive tensor slot, $N = 3$ the color-triplicity index (its three-generation face a downstream consequence), $L = 5$ the defining block size of $SU(5)$, $J = 7$ the minimal symmetric closure cardinality of the charge packet $\{0, \pm 1, \pm 2, \pm 3\}$, and $K = 35 = LJ = M^2 + N^2 + L^2$ the total lock index; the derivation is Arisaka (2026, A3-BI6, Theorem 10.2).

Constitutional admissibility of $U(1)_{Qi}$: optional, empirically preferred. The four-component bulk gauge group above is read here as the *empirically preferred* configuration rather than as a constitutional necessity. Arisaka (2026, A3-BI6) is the authoritative source on this point. The constitutional clauses admit both the three-component bulk $G_{\text{tot}}^{(3)} = \text{Spin}(1, 5) \times SU(5) \times U(1)_B$ and the four-component bulk $G_{\text{tot}}^{(4)} = \text{Spin}(1, 5) \times SU(5) \times U(1)_B \times U(1)_{Qi}$. The four-component case is adopted on empirical grounds: it is required for the boundary Arisaka equation $J_E/Q_B = J_e = J_B = J_{Qi}$ — where Q_B is the bioenergetic Arisaka constant (≈ 0.1555 eV per qubit, matching ATP hydrolysis at biological temperatures) — to close (Arisaka, 2026, A1-Constitution); for the OGN-5 dual-lock identity $K = LJ = M^2 + N^2 + L^2 = 35$ to admit its unique solution (Arisaka, 2026, A3-BI6, Theorem 10.2); and for the Qi-4 pillars and the EIC framework of Arisaka (2026, A7-Qi4) to exist on the boundary. Throughout this paper G_{tot} will mean the four-component group unless explicitly stated otherwise.

3.3 The OGGEPP projector

The GDOS boundary law (Arisaka, 2026, A5-GDOS) is the CPTP bulk-to-boundary projector

$$\mathcal{D}_{\text{OGGEPP}} = \mathcal{R} \circ \text{Tr}_{\text{env}}, \quad (6)$$

which carries bulk operators on M_6 to boundary observables on M_4 via Kaluza–Klein zero-mode reduction followed by environmental tracing. This is the construction the four pillars of Gi-4 formalize: TGG is the bulk side $(\mathcal{A}_{\text{tot}}, \mathcal{R}_{\text{tot}}, \mathcal{T}_{\text{TGG}}^{\mu\nu})$; CCI is the descent identity that makes the projection consistent (X_4 closure, four-current Ward package); GFF is the boundary Lyapunov functional that the projector flow descends along; GRIP is the cascade structure that the projector generates on the family-projector moduli. Without OGGEPP, the four pillars would have no upstream source; with it, they are theorems.

Why OGGEF is necessary, not optional. The closed- system frameworks of orthodox QFT and GR treat measurement as an external annotation on otherwise unitary evolution. OGGEF treats the irreversible projection as the foundational dynamics, and the unitary evolution as the bulk pre-projection limit. In Gi-4 language, the irreversibility of the projector is what makes the GFF second law (T3.2) and the GRIP cascade closure (T4.1) theorems rather than postulates.




Figure 2: Architecture overview of the four Gi-4 Pillars. Each pillar inherits its constitutional content from the GDOS boundary law of GG-A5-GDOS (top) and its bulk-side technical content from the 6D bulk geometry of A2-GG-6 and the BI-6 anomaly closure of GG-A3-BI6 (bottom). Horizontal arrows show the implication chain $\text{TGG} \Rightarrow \text{CCI} \Rightarrow \text{GFF} \Rightarrow \text{GRIP}$. The four named compatibility relations $\mathfrak{C}_{\text{TGG}}$, $\mathfrak{C}_{\text{CCI}}$, $\mathfrak{C}_{\text{GFF}}$, $\mathfrak{C}_{\text{GRIP}}$ (one per pillar) tie the algebraic and operational layers within each pillar.

4 The cohort inheritance and the roadmap of the twelve Parts

Before the four-pillar core begins (Part II), it is useful to lay out two reader-orientation maps on the same opening: a panorama of the twenty-two-paper cohort in which this paper stands (§4.1, Figure 3), and a panorama of the twelve Parts of the present paper (§4.2, Figure 4). A reader of any later Part should be able to flip back to this Section and orient themselves in both directions: which cohort companion carries which imported result, and where the present argument stands in the twelve-Part arc.

4.1 The cohort inheritance

The present paper is one of twenty-two. GG-Theory is written not as a single monograph but as a cohort of twenty-two preprints — eleven A-papers, six P-papers, three C-papers and two white papers — each carrying one rigorously stated piece of a single descent that runs from four axioms and seven postulates to the empirical structure of physics. The partition into twenty-two papers

reflects the partition of subject matter, not a ranking of importance: the framework is one program, and each paper is a named carrier of named content. Figure 3 is the panorama; a reader meeting the program for the first time can read its architecture off the figure top to bottom before returning to the geometric backbone that is the business of this paper. The cohort falls into five strata plus three framed branches.

The constitutional root and the GG6 Trio. At the root stands GG-A1-Constitution ([Arisaka, 2026, A1-Constitution](#)), the GG-Constitution: the four axioms (AX-1 LC, AX-2 GGEP, AX-3 OGGEF, AX-4 MSMF) and seven postulates against which every companion’s theorems are audited. Forced from it is the GG6 Trio that builds the bulk: GG-A2-GG6 ([Arisaka, 2026, A2-GG6](#)) constructs the six-dimensional Geometry–Gauge bulk; GG-A3-BI6 ([Arisaka, 2026, A3-BI6](#)) proves the modified Bianchi identity BI-6 anomaly-free; and GG-A4-Euclid ([Arisaka, 2026, A4-Euclid](#)) traces the lineage from Euclid to the bulk and reconstructs the universal constants.

The boundary law and its two pillars — where this paper sits. AX-3 (OGGEF) forces the open-system boundary law GG-A5-GDOS ([Arisaka, 2026, A5-GDOS](#)), the Geometric Dynamics of Open Systems, which projects bulk physics onto the four-dimensional boundary. Two sibling pillar papers organize the boundary content, and *the present paper is one of them*: GG-A6-GRIP ([Arisaka, 2026, A6-GRIP](#)) — this paper — supplies the geometric Gi-4 pillars (TGG, CCI, GFF, GRIP), the geometric backbone on which the conservation laws of GDOS rest; its information-theoretic sibling GG-A7-Qi4 ([Arisaka, 2026, A7-Qi4](#)) supplies the Qi-4 pillars, with which GG-A5-GDOS closes the GDOS Trio.

The Consequence Trio and the capstone. Three papers read the empirical world off the boundary law: GG-A8-QM ([Arisaka, 2026, A8-QM](#)) recasts four-dimensional quantum mechanics, GG-A9-SM ([Arisaka, 2026, A9-SM](#)) derives the twenty-eight free parameters of the Standard Model from the OGN-5 ledger and the GRIP cascade of the present paper, and GG-A10-Cosmos ([Arisaka, 2026, A10-Cosmos](#)) derives the Λ CDM cosmological inventory; GG-A11-Origin ([Arisaka, 2026, A11-Origin](#)) is the capstone that closes the A-series.

The two phenomenology branches. Two framed branches carry the sharp predictions. The Predictions Branch (P-series) collects the KK-resonance and hierarchy paper GG-P1-KK ([Arisaka, 2026, P1-KK](#)), the three-generations paper GG-P2-N3 ([Arisaka, 2026, P2-N3](#)) — whose $N_g = 3$ proof the present paper supplies the GRIP-side input to — the electron paper GG-P3-Electron ([Arisaka, 2026, P3-Electron](#)), and the proton paper GG-P6-Proton ([Arisaka, 2026, P6-Proton](#)), which runs that paper’s ground-attractor method on the baryonic stratum. The Cosmology & Astrophysics Branch (C-series) collects the Kaxion dark-matter paper GG-C1-Kaxion ([Arisaka, 2026, C1-Kaxion](#)), the Hubble-tension paper GG-C2-Hubble ([Arisaka, 2026, C2-Hubble](#)), and the black-hole-entropy paper GG-C3-BH ([Arisaka, 2026, C3-BH](#)). GG-A6-GRIP supplies the Gi-4 geometric backbone on which the Standard-Model and dark-sector readouts of these branches stand.

The Foundations Branch. The third framed branch is the Foundations Branch: the two member papers whose subject is the boundary’s complex structure and its mirrors. **GG-P4-Real** (Arisaka, 2026, P4-Real) stands on the map’s left flank as the origin-of- i member, deriving the single imaginary unit of boundary physics from two real structures of the **GG-A2-GG6** arena — the compact phase circle and the causal clock. For the present paper its significance is direct: **GG-A6-GRIP** carries the **GRIP** pipeline ($\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$) and the **GFF** ledger — the geometric runtime on which the cohort’s open-system floors run — and **GG-P4-Real** is the origin of the boundary’s complex structure whose phase conventions the **DGI** stage makes dynamical. **GG-P5-CPT** (Arisaka, 2026, P5-CPT) carries the three mirrors of matter — C, P, and T read as orientation reversals of the three real bulk factors — and its law-level arrow of time is kin of the irreversible descent that the runtime built here drives one way (the second law $dF_{\text{GFF}}/dt \leq 0$ of Theorem T3). Both carry P-series designators; their place at the spine’s flank, not the phenomenology rows, marks their foundational role.

The twenty-two-paper map, identical in every paper of the cohort, is Figure 3, which closes this subsection.

4.2 Roadmap of the twelve Parts

The rest of the paper is organized into eleven further Parts, and the walk through them is plotted in Figure 4. The present Part — the foundations — is nearly done: the strategic moment, the 6D bulk geometry and the **BI-6** closure imported from the companions, the **OGGEP** projector and the **GDOS** boundary law recalled from **GG-A5-GDOS**, and the cohort the paper stands in. Everything after this page is the formalization itself, and the itinerary below says where each piece lives.

The load-bearing block is the four-pillar core (Parts II–V), in which each pillar forces the next ($\text{TGG} \Rightarrow \text{CCI} \Rightarrow \text{GFF} \Rightarrow \text{GRIP}$). Part II constructs **TGG**, the Total Geometric–Gauge object: the single composite connection $\mathcal{A}_{\text{tot}} = \omega \oplus A$ that carries gravity and the boundary gauge sectors on one principal G_{tot} -bundle, its boundary stress tensor, and the Master Conservation Law (Theorem T1). Part III proves **CCI**, the Covariant Conservation Identity: under **MSMF** closure of the four-form X_4 , the **OGN-5** ledger $(1, 3, 5; 7, 35)$ is forced and the four boundary Ward currents close together (Theorem T2). Part IV builds **GFF**, the Geometric Free-Energy Functional: Morse-positivity on the stratified boundary moduli and the second law $dF_{\text{GFF}}/dt \leq 0$ along every **GDOS** trajectory (Theorem T3).

Part V is the apex, and it is highlighted in the roadmap because it delivers the paper’s central claim. There the **GRIP** cascade $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$ closes against the **BI-6** budget $S_{\text{BI6}}^{\text{max}} \leq 3\pi$, forcing $N_g \leq 3$ on the family-projector stratum, with the generation-ordered weight ladder as its boundary signature (Theorem T4). Every earlier Part builds toward that closure, and every later Part spends it.

Two Parts then assemble and apply. Part VI synthesizes the unified dual-layer **Gi-4** system — the nine dual-layer sub-theorems and the constitutional bridge Theorem T5 back to the **OGGEP/MSMF** axioms, closing the loop with the **GG-Constitution**. Part VII applies the framework to the spectrum of nature: the $N = 3$ derivation, the twenty-eight Standard-Model parameters, and the hierarchy, **KK-resonance**, and **kaxion** predictions.

Two Parts weigh the case. Part VIII is the discussion: significance and uniqueness, the




Figure 3: The twenty-two-paper cohort map of GG-Theory. The constitutional root **GG-A1-Constitution** (violet, top) supplies the four axioms and seven postulates from which the rest of the cohort descends: the GG6 Trio (**GG-A2-GG6**, **GG-A3-BI6**, **GG-A4-Euclid**, red) realizes the bulk geometry; **GG-A5-GDOS** (yellow) the boundary law; the GDOS Trio pillars (**GG-A6-GRIP**, **GG-A7-Qi4**, green) the geometric and quantum-information runtime; the 4D-Physics Trio (**GG-A8-QM**, **GG-A9-SM**, **GG-A10-Cosmos**, blue) the textbook readouts; and the closing capstone **GG-A11-Origin** (gold) the quarks-to-consciousness identity cascade. Two framed branches carry the phenomenology — the **Predictions Branch** (cyan), whose four papers hand their results along the row by the internal arrows (**GG-P2-N3**’s $N_g=3$ into **GG-P3-Electron**, and that paper’s ground-attractor method, *GGA*, on into **GG-P6-Proton**), and the **Cosmology & Astrophysics Branch** (teal). The spine is flanked by two columns of two, which do opposite jobs. On the left the **Foundations Branch** (violet) feeds it: **GG-P4-Real**, fed by **GG-A2-GG6**’s Theorem 0 arena, manufactures the i that **GG-A8-QM** consumes, and **GG-P5-CPT** is the **GG-A9-SM**–**GG-A10-Cosmos** interface, fed by the ledger phases and feeding the η_B chain across the width of the map, with the shared- i edge descending solid between them. On the right the **White-Paper Branch** (brown) reads it: **GG-W1-Symmetry** walks the century from Noether to Hopkins–Singer, and **GG-W2-Holonomy** dissects the structure that arc arrives at — protection by topology rather than by symmetry — across sixteen open problems. The white papers derive nothing new, carrying every result at its parent’s declared grade; the single dashed feed drawn into each names its heaviest source, not its only one. *Reading conventions:* solid arrows mark axiom-forced descent and carry an edge label naming the driving axiom or postulate where most illustrative; dashed arrows mark the deployment of derived results, or historical links inherited without being constitutionally forced.

measurement-problem carry-over from QM to QFT to the Standard Model, and the comparison with string theory, AdS/CFT, and Randall–Sundrum. Part IX states the conclusions: the methodological synthesis, the open-problem ledger M9–M11, and the place of the construction in the foundations of physics.

The back matter is the workshop. Part X collects the technical machinery, worked examples, and indices that support the four pillars; Part XI answers the strongest objections a hostile reviewer would raise, with pointers into the paper; and Part XII closes with the annotated references, each carrying a Relevance note. Read in order, the twelve Parts compose into a single self-contained formalization that any reader can audit step by step; read by need, each Part names its imports and can stand a visit on its own.




Figure 4: Roadmap of the twelve Parts of GG-A6-GRIP. The argument develops in five phases. *Foundations* (Part I): the strategic moment, the imported bulk and boundary primitives, and the cohort. *The four-pillar core* (Parts II–V): the composite connection and the Master Conservation Law (TGG, T1), the forced OGN-5 ledger (1, 3, 5; 7, 35) and simultaneous boundary conservation (CCI, T2), the Morse free-energy functional and the second law (GFF, T3), and — **Part V**, highlighted — the GRIP cascade DGI $\rightarrow$ GSSB $\rightarrow$ GA closing against the BI-6 budget $S_{\text{BI6}}^{\text{max}} \leq 3\pi$ and forcing $N_g \leq 3$ (T4). *Synthesize and apply* (Parts VI–VII): the unified dual-layer system with the constitutional bridge T5, then the applications to the $N = 3$ count, the twenty-eight Standard-Model parameters, and the hierarchy, KK-resonance, and kaxion predictions. *Discuss and conclude* (Parts VIII–IX) weigh significance and state the open-problem ledger M9–M11. The three back-matter Parts collect the technical appendices (Part X), the common objections and responses (Part XI), and the annotated references (Part XII). Part V is highlighted because the cascade closure delivered there is the paper’s central claim — the family-count forcing and the boundary signature of the cascade, on which the downstream Standard-Model readout stands.

5 What it costs to declare your inputs

Most papers begin by saying what they will show. This one has begun by saying what it is standing on, and the difference is worth a moment before the construction starts.

A theory that reaches as far as this one claims to reach — from the geometry of a six-dimensional arena down to the number of families of matter and up again to the chemistry of a living cell — has a characteristic way of going wrong. Not by making an arithmetic error; those get found. It goes wrong by borrowing. A quantity is taken from one companion argument, used to fix a second, and the second is then cited back as support for the first, and nobody notices because the two statements are separated by a hundred pages and a change of notation. The result looks like a derivation and is a circle. Every large framework in the history of physics has had to be audited for this, and several have failed the audit.

It is worth noticing how short the list turns out to be. An arena of a certain dimension, a group acting on it, a set of whole numbers, a rule for how a large system looks to a small observer inside it, and four axioms. That is the whole of what this paper takes on trust. Everything else it uses, it builds. Whether the list is short enough is a fair question and a reader is entitled to press it — six dimensions is not a modest assumption — but the question can at least be asked, because the list exists and is finite. Against a framework that never wrote its inputs down, the same question cannot even be formulated.

The only defense is bookkeeping, done in public and done first. That is what this Part has been. Every object the construction will use — the arena, its dimension, the group that acts on it, the integer ledger, the boundary law, the axioms — has been named, attributed to the companion that establishes it, and marked as an input rather than a result. Nothing later in this paper may quietly promote one of them. When a later Part needs something not on this list, it must derive it, and the reader is entitled to ask where.

There is a scientific point buried in what looks like an administrative one. *A framework's honesty is measurable, and this is how you measure it:* count the things it assumes, count the things it claims to derive, and check that the second list never leaks into the first. A theory with a short input list and a long output list is making a strong claim and is easy to kill. A theory whose two lists are tangled is making no checkable claim at all, however impressive its mathematics. The declaration in this Part is what makes the rest of the paper falsifiable, and that is a stronger reason for it than tidiness.

What comes next is the construction proper, and it will proceed in one direction only: from the declared inputs, forward. Four objects are built in turn, each using only what precedes it. If the reader ever finds a step that reaches backwards — that uses a later result to justify an earlier one — the paper has failed, and this Part is the list against which that judgment can be made.

Part II TGG

The Total Geometric-Gauge Object: Pillar 1 of Gi-4

Physics currently describes the world with four forces and, uncomfortably, with two different kinds of mathematics. Gravity is geometry: matter bends the arena, and things move as the bending dictates. The other three forces — electromagnetism, the strong force, the weak force — are usually described instead by fields living *on* an arena that is itself unbent. Both descriptions work superbly. They do not obviously describe the same kind of thing, and a century of effort has gone into the suspicion that they must.

This Part takes the suspicion seriously in the most economical way available: it asks whether the four can be read as four aspects of a single geometric object, rather than four objects to be reconciled. The object proposed is what we call the Total Geometric-Gauge object, TGG. The idea is easier than the machinery. Instead of one structure for gravity and three separate structures for the other forces, one writes down a single structure on the six-dimensional bulk that has all four built into it, and one shows that it can be done in exactly one way. The forces are then not unified by decree; they are what the one structure looks like when it is examined along different directions.

What this Part establishes is threefold. It shows that such an object exists and that it is unique — there is no second way to assemble the four that satisfies the constitutional requirements. It shows what that object hands to the four-dimensional world we can observe: a single tensor collecting the energy and momentum of matter, of the gauge fields, of the residual higher-dimensional modes, and of the correction term that anomaly cancellation demands. And it shows that this one tensor obeys one conservation law — which is the mathematical form of the statement that energy and momentum do not leak out of the world by an unaccounted route. That is the Master Conservation Law, and it is the Part’s principal result.

Everything geometric in this Part is taken as given from the companion papers; what is new is the assembly and the uniqueness. A reader outside the specialty can take the first section, which states what the object is and why exactly four pieces enter it, and then the statement of the Master Conservation Law, and skip the intervening construction without losing the thread. A reader who wants to check the arithmetic will find the worked examples at the end of the Part, computed rather than quoted.

6 TGG: motivation, dual-layer structure, and Structure Theorems

This Part develops TGG (the *Total Geometric-Gauge object*), the first of the four Gi-4 Pillars. TGG is the unique geometric object whose two layers are: (i) the *bulk layer* – the master composite-connection structure on the 6D bulk M_6 that unifies the gauge and gravity sectors into a single principal bundle; and (ii) the *boundary layer* – the unified 4D stress tensor $\mathcal{T}_{\text{TGG}}^{\mu\nu}$ obtained by

zero-mode reduction of the bulk action, which collects matter, gauge, Green–Schwarz inflow, and Kaluza–Klein-remnant contributions into a single object. The two layers are tied together by the *Master Conservation Law* $\nabla_\mu \mathcal{T}_{\text{TGG}}^{\mu\nu} = 0$ that holds iff the full Ward package (electric, baryonic, informational, and metric currents) is satisfied. TGG is the formal mathematical realization of GGEP extended by OGGEp to the boundary, and the unbroken-symmetry TGG phase corresponds to the DGI stage (Dynamical Gauge Invariance) of the GRIP cascade.

The Part has six sections. Section 6 (the present section) states the seven TGG Structure Theorems – five for the bulk layer (**TGG-ST1** through **TGG-ST5**) and two for the boundary layer (**TGG-ST6** and **TGG-ST7**) – gives the formal definition of TGG as the dual-layer object, and lists the first immediate consequences. Section 7 develops the bulk-layer machinery: the principal G_{tot} -bundle construction, the connection and curvature decompositions, and the Chern–Simons construction of the Kalb–Ramond three-form. Section 8 presents the magic-dimension ladder backbone and explains why $d = 6$ is the unique admissible bulk dimension. Section 9 (new in) develops the boundary-layer machinery: the Kaluza–Klein zero-mode reduction, the explicit form of $\mathcal{T}_{\text{TGG}}^{\mu\nu}$, and the projected modified Bianchi identity (PBI-4). Section 10 states and proves the existence, uniqueness, and Master Conservation theorem **T1**. Section 11 contains worked examples.

6.1 Motivation: from disconnected stress-energies to a single Total Geometric-Gauge object

The GG-6 program identifies, in Arisaka (2026, A1-Constitution) and Arisaka (2026, A2-GG6), four physical sectors on the bulk: gravity (encoded by the $\text{Spin}(1,5)$ frame bundle), the visible $\text{SU}(5)$ gauge sector, the gravity-baryon $\text{U}(1)_B$ sector, and the quantum-information $\text{U}(1)_{Qi}$ sector. In closed-system QFT, these four sectors would be assembled as four *independent* principal bundles, glued together only by the requirement that they share a common base manifold. The result of that closed-system assembly is a fragile structure: each sector carries its own gauge freedom, its own connection, its own curvature, and (most importantly for the present Part) its own stress tensor on the boundary, with no algebraic mechanism enforcing consistency between them. In particular, before the boundary theory carried four disconnected stress-energy objects – the matter stress tensor $T_{\text{matter}}^{\mu\nu}$, the gauge-field stress tensor $T_{\text{gauge}}^{\mu\nu}[F]$, the Green–Schwarz inflow term $X_{\text{GS}}^{\mu\nu}$, and the Kaluza–Klein heavy-mode remnant $T_{\text{KK}}^{\mu\nu}$ – and their sum-conservation was an assumption, not a theorem. Anomaly cancellation in that closed-system picture was imposed as a separate constraint, and the unification of gauge with gravity was left as a hope rather than a derived theorem.

GGEP from the GG-Constitution (Arisaka, 2026, A1-Constitution, Axiom 2) rejects this fragmented picture and asserts that the four sectors arise from a single primitive composite connection at the *bulk* level. OGGEp (Arisaka, 2026, A1-Constitution, Axiom 3) extends this further to the *boundary* level: the boundary observable content is the OGGEp zero-mode projection of the bulk geometric content, so the four boundary stress tensors must also collapse into one. Together, GGEP and OGGEp force the existence of a single mathematical object with two layers:

- the *bulk layer* is the principal G_{tot} -bundle $P \rightarrow M_6$ with structure group $G_{\text{tot}} = \text{Spin}(1,5) \times \text{SU}(5) \times \text{U}(1)_B \times \text{U}(1)_{Qi}$, equipped with a composite connection one-form $\mathcal{A}_{\text{tot}}$ and curvature $\mathcal{R}_{\text{tot}}$;

- the *boundary layer* is the unified 4D stress tensor

$$\mathcal{T}_{\text{TGG}}^{\mu\nu} = T_{\text{matter}}^{\mu\nu} + T_{\text{gauge}}^{\mu\nu}[F] + X_{\text{GS}}^{\mu\nu} + T_{\text{KK}}^{\mu\nu},$$

obtained by Kaluza–Klein zero-mode reduction of the bulk action of the composite connection together with the BI-6 bulk-to-boundary descent.

The two layers are not independent. They are connected by the *Master Conservation Law*,

$$\nabla_\mu \mathcal{T}_{\text{TGG}}^{\mu\nu} = 0,$$

which holds iff the full Ward package (electric, baryonic, and informational currents, together with the projected modified Bianchi identity) is satisfied. This master conservation law is a theorem of the bulk geometric structure (GGEP + BI-6) projected to the boundary (OGGEP); it is not an assumption.

This single object with its two layers and the unifying conservation theorem is what we call the TGG (*Total Geometric-Gauge object*). The name encodes the three structural facts: it is *Total* because it collects the entire stress-energy content into one tensor; it is *Geometric-Gauge* because the bulk layer unifies the gauge sub-connections with the spin connection of gravity; and it is an *object* (rather than a connection or a tensor alone) because the dual-layer nature is essential to its definition.

Remark 6.1 (Physical intuition for TGG: a familiar analog). The dual-layer structure of TGG can be made concrete by analogy with ordinary electromagnetism. In Maxwell’s theory there is, on the one hand, the electromagnetic field itself – the gauge potential A^μ and field strength $F^{\mu\nu}$ that live on every spacetime point – and, on the other hand, the unified energy-momentum tensor $T_{\text{matter}}^{\mu\nu} + T_{\text{EM}}^{\mu\nu}$ whose conservation $\nabla_\mu (T_{\text{matter}} + T_{\text{EM}})^{\mu\nu} = 0$ is forced by charge conservation $\nabla_\mu J^\mu = 0$. The two are not independent: the field generates the unified stress tensor by variation of the action, and current conservation links them.

TGG is the six-sector generalization of this picture. The “electromagnetic field” is replaced by a composite connection on the bulk that unifies gravity, SU(5), U(1)_B, and U(1)_{Qi} into *one* object. The “unified energy-momentum tensor” is replaced by $\mathcal{T}_{\text{TGG}}^{\mu\nu} = T_{\text{matter}} + T_{\text{gauge}} + X_{\text{GS}} + T_{\text{KK}}$, the boundary projection of the bulk action. The “current conservation” is replaced by the Ward package of all four sectors simultaneously, conserving the unified stress tensor by the Master Conservation Law. In each layer of analogy the framework broadens: from one gauge sector to four, from one current to four, from one boundary spacetime to a 6D-bulk-projecting-to-4D structure. Readers familiar with Maxwell theory will find every TGG statement to be a structurally identical generalization of a familiar result; the generalization is rigorous because BI-6 closure plus the OGN-5 ledger guarantee that the generalization closes uniquely.

6.1.1 Historical background (with citations)

The unification of gauge with gravity into a single composite connection that TGG formalizes is the culmination of a thread that runs through nearly the whole modern history of theoretical physics.

Maxwell (1865) wrote down the first gauge-theoretic structure in physics: the four-potential A_μ of electromagnetism is defined only up to the gradient of a scalar, $A_\mu \rightarrow A_\mu + \partial_\mu \lambda$, and the field strength $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ is the curvature of this connection. Maxwell did not use the word *gauge*; the geometric reading came later. Einstein (1915) produced the second strand a half-century later by recasting gravitation as the curvature of a metric connection $\Gamma_{\mu\nu}^\rho$ on the tangent bundle of spacetime; here too the field equations are invariant under a redefinition of the coordinate frame, and the curvature is a tensorial object.

Weyl (1918) was the first to ask whether the two threads were the same thread. His original gauge proposal — rescaling the metric by a position-dependent factor — did not survive empirical scrutiny, but it identified the structural question correctly: are gauge fields and gravity two faces of a single geometric object on a single bundle? The mathematical infrastructure to answer that question precisely was developed by Ehresmann (1950) in his theory of connections on principal bundles and codified by Kobayashi & Nomizu (1969). Yang & Mills (1954) produced the non-abelian extension of Maxwell’s gauge theory, with the field strength $F_{\mu\nu}^a = \partial_\mu A_\nu^a - \partial_\nu A_\mu^a + gf^{abc}A_\mu^b A_\nu^c$ as the curvature of an $SU(2)$ connection.

Indeed, TGG closes this loop. Its composite connection $\mathcal{A}_{\text{tot}} = \omega \oplus A$ on a principal G_{tot} -bundle over the 6D bulk treats gravity (the $\text{Spin}(1, 5)$ part ω) and the boundary gauge fields (the $SU(5) \times U(1)_B \times U(1)_{Qi}$ part A) as components of a single mathematical object. The Master Conservation Law $\nabla_\mu \mathcal{T}_{\text{TGG}}^{\mu\nu} = 0$ that we derive in Section 10.1 is the axiomatic completion of the lineage from Maxwell through Yang–Mills to TGG, with the GG-Constitution AX-2 GGEP as its principled foundation (Arisaka, 2026, A1-Constitution).

We now state the seven TGG Structure Theorems – five governing the bulk layer and two governing the boundary layer.

6.2 The seven TGG axioms (five bulk + two boundary)

Structure Theorem 1 (TGG-ST1: Total bundle structure). *On the 6D bulk M_6 , there exists a smooth principal bundle*

$$\pi : P \rightarrow M_6 \tag{7}$$

with structure group

$$G_{\text{tot}} = \text{Spin}(1, 5) \times SU(5) \times U(1)_B \times U(1)_{Qi}. \tag{8}$$

The four group factors act on P as commuting subgroups, and the free right action of G_{tot} on P has quotient M_6 .

Structure Theorem 2 (TGG-ST2: Composite connection). *The bundle P admits a connection one-form*

$$\mathcal{A}_{\text{tot}} \in \Omega^1(P, \mathfrak{g}_{\text{tot}}), \quad \mathfrak{g}_{\text{tot}} = \mathfrak{spin}(1, 5) \oplus \mathfrak{su}(5) \oplus \mathfrak{u}(1)_B \oplus \mathfrak{u}(1)_{Qi}, \tag{9}$$

which decomposes along the direct-sum factors as

$$\mathcal{A}_{\text{tot}} = \omega_{\text{Spin}(1,5)} + A_5 + A_B + A_{Qi}, \tag{10}$$

where $\omega_{\text{Spin}(1,5)} \in \Omega^1(P, \mathfrak{spin}(1, 5))$ is the spin connection, $A_5 \in \Omega^1(P, \mathfrak{su}(5))$ is the $SU(5)$ connection,

$A_B \in \Omega^1(P, \mathfrak{u}(1)_B)$ is the $U(1)_B$ connection, and $A_{Q_i} \in \Omega^1(P, \mathfrak{u}(1)_{Q_i})$ is the $U(1)_{Q_i}$ connection.

Structure Theorem 3 (TGG-ST3: Magic-dimension ladder). *The $\text{Spin}(1, 5)$ factor of G_{tot} is realized through the accidental isomorphism*

$$\text{Spin}(1, 5) \cong \text{SL}(2, \mathbb{H}), \quad (11)$$

which places the bulk dimension $d = 6$ as the unique admissible value in the magic-dimension ladder $\mathbb{R} \rightarrow \mathbb{C} \rightarrow \mathbb{H} \rightarrow \mathbb{O}$ (see Section 8 for the detailed treatment).

Structure Theorem 4 (TGG-ST4: Lorentzian compatibility). *The bulk metric $ds_6^2 = e^{-2r} \eta_{\mu\nu} dx^\mu dx^\nu + dr^2 + R_0^2 d\theta^2$ on M_6 is compatible with the $\text{Spin}(1, 5)$ structure of **TGG-ST1**, in the sense that the spin connection $\omega_{\text{Spin}(1, 5)}$ is the Levi-Civita-type connection associated with this metric. The product-geometry structure of M_6 is preserved by the gauge action of G_{tot} on P .*

Structure Theorem 5 (TGG-ST5: GGEP unification). *The four sub-connections ($\omega_{\text{Spin}(1, 5)}$, A_5 , A_B , A_{Q_i}) are not four independent connections on four independent bundles; they are four projections of the single composite connection $\mathcal{A}_{\text{tot}}$ of **TGG-ST2**. Equivalently, gauge transformations of $\mathcal{A}_{\text{tot}}$ act simultaneously on all four sub-connections by the natural diagonal action of G_{tot} on $\mathfrak{g}_{\text{tot}}$. This is the formal realization of GGEP on M_6 .*

The five Structure Theorems **TGG-ST1–TGG-ST5** characterize the *bulk layer* of TGG. We now state the two additional Structure Theorems that characterize the *boundary layer* – the 4D unified stress tensor and its master conservation law.

Structure Theorem 6 (TGG-ST6: Boundary stress tensor projection). *Under Kaluza–Klein zero-mode reduction of the bulk Einstein–Hilbert plus gauge action of TGG-ST2, together with the BI-6 bulk-to-boundary descent of the anomaly polynomial X_4 , the bulk stress-energy content projects onto the 4D boundary $M_4 = \partial M_6$ as a single Total Geometric-Gauge stress tensor*

$$\mathcal{T}_{\text{TGG}}^{\mu\nu} := T_{\text{matter}}^{\mu\nu} + T_{\text{gauge}}^{\mu\nu}[F] + X_{\text{GS}}^{\mu\nu} + T_{\text{KK}}^{\mu\nu}, \quad (12)$$

where $T_{\text{matter}}^{\mu\nu}$ is the boundary matter stress tensor, $T_{\text{gauge}}^{\mu\nu}[F]$ is the boundary gauge-field stress tensor constructed from $F = F_5 + F_B + F_{Q_i}$, $X_{\text{GS}}^{\mu\nu}$ is the boundary shadow of the Green–Schwarz inflow term in BI-6, and $T_{\text{KK}}^{\mu\nu}$ collects the heavy Kaluza–Klein-mode remnants. The boundary projection respects OGGEF of the GG-Constitution.

Structure Theorem 7 (TGG-ST7: Master Conservation Law). *The Total Geometric-Gauge stress tensor of TGG-ST6 satisfies the covariant conservation equation*

$$\nabla_\mu \mathcal{T}_{\text{TGG}}^{\mu\nu} = 0 \quad (13)$$

in the gauge-closed boundary sector, holding if and only if the full Ward package is satisfied simultaneously:

$$\nabla_\mu J_e^\mu = 0, \quad \nabla_\mu J_B^\mu = 0, \quad \nabla_\mu J_{Q_i}^\mu = 0, \quad \nabla_\mu (G^{\mu\nu} + X_{\text{GS}}^{\mu\nu}) = 0, \quad (14)$$

where the last equation is the projected modified Bianchi identity (PBI-4), the 4D zero-mode descent of the 6D BI-6 identity $dH_3 = X_4$.

The seven Structure Theorems are stated for the 6D bulk of GG-Theory; **TGG-ST1–TGG-ST5** characterize the bulk layer (the composite connection on M_6), and **TGG-ST6–TGG-ST7** characterize the boundary layer (the unified stress tensor on M_4 and its master conservation law). All seven Structure Theorems have natural generalizations to other Lorentzian bulk manifolds with appropriate magic-dimensional structure, but those generalizations are not the topic of the present paper.

6.3 Formal definition of TGG

Definition 6.2 (TGG: the Total Geometric-Gauge object). The *Total Geometric-Gauge object* associated with the 6D bulk M_6 and its 4D boundary $M_4 = \partial M_6$ is the quintuple

$$\text{TGG} = (P, \mathcal{A}_{\text{tot}}, \mathcal{R}_{\text{tot}}; \mathcal{T}_{\text{TGG}}^{\mu\nu}, \nabla), \quad (15)$$

consisting of two layers tied together by the Master Conservation Law:

Bulk layer (TGG-ST1–TGG-ST5):

- $P \rightarrow M_6$ is the principal G_{tot} -bundle of **TGG-ST1**;
- $\mathcal{A}_{\text{tot}} \in \Omega^1(P, \mathfrak{g}_{\text{tot}})$ is the composite connection one-form of **TGG-ST2**, decomposing as in (10);
- $\mathcal{R}_{\text{tot}} \in \Omega^2(P, \mathfrak{g}_{\text{tot}})$ is the corresponding curvature two-form,

$$\mathcal{R}_{\text{tot}} = d\mathcal{A}_{\text{tot}} + \mathcal{A}_{\text{tot}} \wedge \mathcal{A}_{\text{tot}}, \quad (16)$$

decomposing as

$$\mathcal{R}_{\text{tot}} = R_{\text{Spin}(1,5)} + F_5 + F_B + F_{Qi}, \quad (17)$$

with the four sub-curvatures defined by the corresponding sub-connections.

Boundary layer (TGG-ST6–TGG-ST7):

- $\mathcal{T}_{\text{TGG}}^{\mu\nu}$ is the unified Total Geometric-Gauge stress tensor on M_4 given by (12), the Kaluza–Klein zero-mode reduction of the bulk action of TGG-ST2 plus the BI-6 bulk-to-boundary descent;
- ∇ is the boundary Levi-Civita covariant derivative, with the Master Conservation Law $\nabla_\mu \mathcal{T}_{\text{TGG}}^{\mu\nu} = 0$ holding iff the Ward package (14) is satisfied.

The quintuple $(P, \mathcal{A}_{\text{tot}}, \mathcal{R}_{\text{tot}}; \mathcal{T}_{\text{TGG}}^{\mu\nu}, \nabla)$ is required to satisfy Structure Theorems **TGG-ST1–TGG-ST7**. Two such quintuples are *gauge-equivalent* if their bulk components $(P, \mathcal{A}_{\text{tot}}, \mathcal{R}_{\text{tot}})$ are related by a global gauge transformation $g : M_6 \rightarrow G_{\text{tot}}$ acting on P by the right action and on $\mathcal{A}_{\text{tot}}, \mathcal{R}_{\text{tot}}$ by the corresponding adjoint action; the boundary components $(\mathcal{T}_{\text{TGG}}^{\mu\nu}, \nabla)$ are then automatically related by the induced boundary diffeomorphism.

Remark 6.3. The intuitive content of Definition 6.2 is that TGG is a single dual-layer geometric object: in the bulk, four “faces” of one composite connection (gravity, visible gauge, baryon, Qi); on

the boundary, four contributions to one stress tensor (matter, gauge fields, GS inflow, KK remnant). The bulk layer encodes the algebraic content of **GGEP** (the four sub-connections share a single composite origin); the boundary layer encodes the dynamical content of **OGGEP** (the four boundary stress-energy contributions share a single conserved master tensor). The two layers are connected by the **BI-6** bulk-to-boundary descent, and the unifying theorem is the Master Conservation Law of **TGG-ST7**.

6.4 First consequences of the **TGG** Structure Theorems

We list six immediate consequences of **TGG-ST1–TGG-ST7** that will be used repeatedly in the rest of the paper. The first five concern the bulk layer; the sixth concerns the boundary layer.

Proposition 6.4 (Existence of frame bundle). *The $\text{Spin}(1,5)$ factor of P contains, as a sub-bundle, the oriented Lorentzian frame bundle $FM_6 \rightarrow M_6$ associated with the scale-profile metric of **TGG-ST4**. Equivalently, the pull-back of $\mathcal{A}_{\text{tot}}$ to FM_6 along the natural inclusion is the Levi-Civita spin connection.*

Proof. By **TGG-ST1**, the $\text{Spin}(1,5)$ factor acts freely on P with quotient M_6 . The induced principal $\text{Spin}(1,5)$ -sub-bundle is the spin frame bundle of M_6 . Compatibility with the metric (**TGG-ST4**) selects the Levi-Civita spin connection up to gauge. ■

Proposition 6.5 (Integrability of sub-connections). *Each of the four sub-connections $(\omega_{\text{Spin}(1,5)}, A_5, A_B, A_{Q_i})$ is, on its own, an integrable connection one-form satisfying $d\omega_{\text{Spin}(1,5)} + \omega_{\text{Spin}(1,5)} \wedge \omega_{\text{Spin}(1,5)} = R_{\text{Spin}(1,5)}$, $dA_5 + A_5 \wedge A_5 = F_5$, $dA_B = F_B$, $dA_{Q_i} = F_{Q_i}$, where the $U(1)$ field strengths have no commutator term because they are Abelian.*

Proof. The four direct-summands of $\mathfrak{g}_{\text{tot}}$ commute pairwise. The composite curvature (16) therefore decomposes *additively* along the four factors with no cross-terms:

$$\mathcal{R}_{\text{tot}} = d(\omega + A_5 + A_B + A_{Q_i}) + (\omega + A_5 + A_B + A_{Q_i}) \wedge (\omega + A_5 + A_B + A_{Q_i}) \quad (18)$$

$$= d\omega + \omega \wedge \omega + dA_5 + A_5 \wedge A_5 + dA_B + dA_{Q_i}. \quad (19)$$

The cross-terms $\omega \wedge A_5$, $A_5 \wedge A_B$, etc., vanish because $[\mathfrak{spin}(1,5), \mathfrak{su}(5)] = 0$ and similarly for the other commutators. This yields the four sub-curvatures listed. ■

Proposition 6.6 (Gauge group of **TGG**). *The gauge group of **TGG** is the group $\mathcal{G}(P)$ of G_{tot} -equivariant smooth maps $P \rightarrow G_{\text{tot}}$. Two **TGG** triples are gauge-equivalent if and only if they are related by an element of $\mathcal{G}(P)$.*

Proof. This is standard for principal bundles; see, e.g., [Kobayashi & Nomizu \(1969\)](#). The four-factor structure of G_{tot} implies that $\mathcal{G}(P)$ factors as the direct product $\mathcal{G}_{\text{Spin}(1,5)} \times \mathcal{G}_5 \times \mathcal{G}_B \times \mathcal{G}_{Q_i}$. ■

Proposition 6.7 (Holonomy decomposition). *The holonomy group of **TGG** along any closed loop $\gamma \subset M_6$ factorizes as*

$$\text{Hol}(\text{TGG}, \gamma) = \text{Hol}(\omega_{\text{Spin}(1,5)}, \gamma) \times \text{Hol}(A_5, \gamma) \times \text{Hol}(A_B, \gamma) \times \text{Hol}(A_{Q_i}, \gamma) \quad (20)$$

along the four sub-connection factors.

Proof. The path-ordered exponential $P \exp(\oint_\gamma \mathcal{A}_{\text{tot}})$ factors along commuting Lie-algebra factors of $\mathfrak{g}_{\text{tot}}$ into the product of path-ordered exponentials of the four sub-connections. This is the Baker–Campbell–Hausdorff identity for commuting generators. ■

Proposition 6.8 (Spin(1,5) holonomy along compact θ direction). *For the product metric of **TGG-ST4**, the Spin(1,5) holonomy along the compact θ -direction is trivial,*

$$\text{Hol}(\omega_{\text{Spin}(1,5)}, S_\theta^1) = \mathbf{1}_{\text{Spin}(1,5)}, \quad (21)$$

because the scale profile e^{-2r} does not depend on θ . In contrast, the Abelian $U(1)_B$ and $U(1)_{Qi}$ holonomies along S_θ^1 may be non-trivial Wilson-line phases, and these are the order parameters for the Wilson–Hosotani electroweak breaking that will be discussed in Part V.

Proof. Direct computation: the spin connection of the product metric (1) contains no $d\theta$ component because the metric components are θ -independent. Therefore $\oint_{S_\theta^1} \omega_{\text{Spin}(1,5)} = 0$ and the path-ordered exponential reduces to the identity. ■

Proposition 6.9 (Existence of the unified boundary stress tensor). *On the 4D boundary $M_4 = \partial M_6$ of any TGG configuration satisfying **TGG-ST1–TGG-ST7**, the unified Total Geometric-Gauge stress tensor $\mathcal{T}_{\text{TGG}}^{\mu\nu}$ exists, is unique up to gauge equivalence, and is symmetric: $\mathcal{T}_{\text{TGG}}^{\mu\nu} = \mathcal{T}_{\text{TGG}}^{\nu\mu}$.*

Sketch. Each of the four constituents in (12) is a well-defined, symmetric stress-energy tensor on the boundary: the matter stress tensor by the standard variation of the matter action, the gauge stress tensor by the symmetric trace structure of the gauge action, the Green–Schwarz inflow by the symmetric variation of $X_4 \wedge \omega$, and the Kaluza–Klein remnant by zero-mode reduction. Symmetry of the sum follows from termwise symmetry, and uniqueness follows from the uniqueness of the bulk composite connection (Theorem 10.1 of Section 10) plus the canonical zero-mode reduction. Detailed proofs of each piece are standard; see Arisaka (2026, A3-BI6, § 2.2–2.4). ■

These six propositions are the workhorse statements – five for the bulk layer and one for the boundary layer – that will be applied throughout the rest of the paper.

7 The composite connection structure

Remark 7.1 (Reader’s bridge: one bundle, not four). Before diving into the formal definitions, let us motivate the fiber-product construction by an everyday picture. In the Standard-Model treatment, gravity lives on one bundle (the frame bundle of spacetime), the strong/electroweak interactions live on another (the $SU(3) \times SU(2) \times U(1)$ bundle), and baryon number plus the information sector are added as auxiliary $U(1)$ ’s — four parallel structures, each with its own connection, glued by hand to coincide on the same spacetime manifold. Roughly speaking, this is a confederation of independent bundles that happen to share a base. TGG replaces the confederation with a federation. The four sub-bundles are merged into a single principal G_{tot} -bundle P on the 6D bulk, with one composite connection $\mathcal{A}_{\text{tot}} = \omega \oplus A_{SU(5)} \oplus A_B \oplus A_{Qi}$. Why insist on this? Because anomaly cancellation, gauge




Figure 5: Algebraic structure of TGG (Pillar 1 of Gi-4). The composite connection $\mathcal{A}_{\text{tot}} = \omega \oplus A$ decomposes into four sub-connections on a principal G_{tot} -bundle over the 6D bulk M_6 : the gravitational ω on the $\text{Spin}(1, 5)$ frame bundle, and the gauge sub-connections $A_{\text{SU}(5)}$, A_B , A_{Q_i} on the visible, baryonic, and quantum-information sectors respectively. Its curvature $\mathcal{R}_{\text{tot}}$ is the input to the four-form anomaly polynomial X_4 of CCl. Kaluza–Klein zero-mode reduction onto the 4D boundary produces the unified stress tensor $\mathcal{T}_{\text{TGG}}^{\mu\nu}$, whose covariant conservation is the Master Conservation Law T1.3.

compatibility, and the bulk-to-boundary descent to four dimensions all require the four sectors to be facets of a single geometric object, not four independent objects that happen to add to zero. The fiber-product construction below is the precise way to say “one bundle, four facets” instead of “four bundles, one base.”

We now develop the formal mathematical content of TGG in more detail, focusing on the principal G_{tot} -bundle construction, the explicit decompositions of $\mathcal{A}_{\text{tot}}$ and $\mathcal{R}_{\text{tot}}$, and the gauge-equivalence relation.

7.1 The principal G_{tot} -bundle on $\mathcal{M}_6$

The construction of the bundle P in **TGG-ST1** can be made explicit by exhibiting P as the fiber product of four principal sub-bundles:

$$P = P_{\text{Spin}(1,5)} \times_{M_6} P_{\text{SU}(5)} \times_{M_6} P_{\text{U}(1)_B} \times_{M_6} P_{\text{U}(1)_{Q_i}}, \quad (22)$$

where each $P_H \rightarrow M_6$ is the principal H -bundle for the relevant factor $H \in \{\text{Spin}(1, 5), \text{SU}(5), \text{U}(1)_B, \text{U}(1)_{Q_i}\}$ and the fiber product is taken over M_6 . The fiber product construction enforces that all four sub-bundles share the same base manifold, and the fiber over $x \in M_6$ is

$$P_x = \text{Spin}(1, 5) \times \text{SU}(5) \times \text{U}(1)_B \times \text{U}(1)_{Q_i} = G_{\text{tot}}. \quad (23)$$

The existence of each individual sub-bundle is guaranteed by:

- $P_{\text{Spin}(1,5)}$: the bulk is Lorentzian and time-orientable (**TGG-ST4**), so it admits a global $\text{Spin}(1,5)$ structure.
- $P_{\text{SU}(5)}$: the second Stiefel–Whitney class $w_2(M_6)$ vanishes because M_6 is a product geometry of contractible factors with the spin manifold $\mathbb{R}^{1,3}$, so M_6 admits a spin structure that lifts to an $\text{SU}(5)$ bundle once the boundary embedding is fixed.
- $P_{\text{U}(1)_B}, P_{\text{U}(1)_{Q_i}}$: M_6 is non-compact and admits trivial $\text{U}(1)$ bundles globally. In the product-geometry picture, the non-trivial Abelian holonomy along the compact θ direction makes these bundles topologically non-trivial in the Pontryagin sense even when the principal-bundle homotopy class is trivial.

7.2 The connection one-form and its decomposition

The composite connection $\mathcal{A}_{\text{tot}}$ of **TGG-ST2** is, in any local trivialization $\sigma : U \rightarrow P|_U$ over a chart $U \subset M_6$, expressed as a $\mathfrak{g}_{\text{tot}}$ -valued one-form:

$$\sigma^* \mathcal{A}_{\text{tot}} = \omega_{\text{Spin}(1,5)}^{ab} T_{ab}^{(\text{Spin}(1,5))} + A_5^I T_I^{(\text{SU}(5))} + A_B T_B + A_{Q_i} T_{Q_i}, \quad (24)$$

where $T_{ab}^{(\text{Spin}(1,5))}$ are the generators of $\mathfrak{spin}(1,5)$ (with antisymmetric Lorentz indices a, b running over $0, 1, \dots, 5$), $T_I^{(\text{SU}(5))}$ are the 24 generators of $\mathfrak{su}(5)$, and T_B, T_{Q_i} are the single Abelian generators of $\mathfrak{u}(1)_B, \mathfrak{u}(1)_{Q_i}$. The component one-forms $\omega^{ab}, A_5^I, A_B, A_{Q_i}$ are smooth real-valued one-forms on U .

The total Lie algebra $\mathfrak{g}_{\text{tot}}$ has dimension

$$\dim \mathfrak{g}_{\text{tot}} = \dim \mathfrak{spin}(1,5) + \dim \mathfrak{su}(5) + 1 + 1 = 15 + 24 + 1 + 1 = 41. \quad (25)$$

The 15 dimensions of $\mathfrak{spin}(1,5)$ correspond to the six-dimensional Lorentz generators $M^{ab} = -M^{ba}$ ($a, b = 0, \dots, 5$), of which six are boosts and nine are rotations.

7.3 The curvature two-form and its decomposition

The composite curvature $\mathcal{R}_{\text{tot}} = d\mathcal{A}_{\text{tot}} + \mathcal{A}_{\text{tot}} \wedge \mathcal{A}_{\text{tot}}$ decomposes along the four direct-summand factors as in (17). Explicitly, in a local trivialization:

$$R_{\text{Spin}(1,5)}^{ab} = d\omega^{ab} + \omega^{ac} \wedge \omega_c^b, \quad (26)$$

$$F_5^I = dA_5^I + \frac{1}{2} f^I{}_{JK} A_5^J \wedge A_5^K, \quad (27)$$

$$F_B = dA_B, \quad (28)$$

$$F_{Q_i} = dA_{Q_i}, \quad (29)$$

where $f^I{}_{JK}$ are the structure constants of $\mathfrak{su}(5)$. The Abelian field strengths F_B, F_{Q_i} are simply exact one-form exterior derivatives because the Lie bracket on $\mathfrak{u}(1)$ vanishes.

The three-form

$$H_3 \in \Omega^3(M_6) \quad (30)$$

that appears in BI-6 is constructed from $\mathcal{A}_{\text{tot}}$ and $\mathcal{R}_{\text{tot}}$ via the Chern–Simons descent, as detailed in [Arisaka \(2026, A3-BI6, Section 3\)](#). Specifically, H_3 is the Kalb–Ramond three-form whose differential is the four-form X_4 defined in Part III. The construction of H_3 from TGG data is the algebraic core of the link between TGG and BI-6 and will be developed in Part III.

7.4 Gauge equivalence and gauge orbits

Two TGG triples $(P, \mathcal{A}_{\text{tot}}, \mathcal{R}_{\text{tot}})$ and $(P', \mathcal{A}'_{\text{tot}}, \mathcal{R}'_{\text{tot}})$ are gauge-equivalent if there exists a diffeomorphism $\Phi : P \rightarrow P'$ commuting with the right G_{tot} -action and taking $\mathcal{A}_{\text{tot}}$ to $\mathcal{A}'_{\text{tot}}$ via the natural adjoint action. In a local trivialization, this amounts to

$$\mathcal{A}'_{\text{tot}}(x) = g(x)^{-1} \mathcal{A}_{\text{tot}}(x) g(x) + g(x)^{-1} dg(x), \quad g : M_6 \rightarrow G_{\text{tot}}. \quad (31)$$

The gauge orbits of TGG are the equivalence classes under (31). By Proposition 6.6, the gauge group factors along the four direct summands of G_{tot} , so the gauge orbits also factor.

7.5 Holonomy and global structure

For the 6D bulk of GG-Theory, the principal-bundle topology is conveniently captured by the holonomy along the compact θ -circle. The $\text{Spin}(1,5)$ holonomy is trivial (Proposition 6.8); the $\text{SU}(5)$ holonomy encodes the orbifold breaking $\text{SU}(5) \rightarrow \text{SU}(3) \times \text{SU}(2) \times \text{U}(1)$ via Wilson lines; the $\text{U}(1)_B$ and $\text{U}(1)_{Q_i}$ holonomies around the circle are fixed background data on the admissible class of the anomaly paper ([Arisaka, 2026, A3-BI6](#)), their fluctuations registered there as exposures; and the electroweak Wilson–Hosotani phase lives in the $\text{SU}(5)$ coset, where the Standard-Model paper computes it ([Arisaka, 2026, A9-SM](#)). An earlier statement of this sentence read the $\text{U}(1)_B$ holonomy as the electroweak phase and the $\text{U}(1)_{Q_i}$ holonomy as the Kaxion phase; both readings are retired (Remark 11.1).

Definition 7.2 (Wilson-line order parameters). The Wilson-line order parameters of TGG along the compact θ -direction are

$$W_{\text{Spin}(1,5)}(\theta_0) = \text{tr} \left[P \exp \oint_{S_\theta^1} \omega_{\text{Spin}(1,5)} \right] = 1, \quad (32)$$

$$W_5(\theta_0) = \text{tr} \left[P \exp \oint_{S_\theta^1} A_5 \right], \quad (33)$$

$$W_B(\theta_0) = \exp \left(i \oint_{S_\theta^1} A_B \right), \quad (34)$$

$$W_{Q_i}(\theta_0) = \exp \left(i \oint_{S_\theta^1} A_{Q_i} \right). \quad (35)$$

The Wilson-line value $W_{\text{Spin}(1,5)} = 1$ is forced by Proposition 6.8; the other three Wilson lines are non-trivial dynamical variables that, in Part V, become the order parameters of the GRIP cascade. The i of the two Abelian holonomies carries a pedigree: it is the boundary rendering of the circle’s forced quarter-turn — derived, not primitive (T-P4-1 of [Arisaka, 2026, P4-Real](#)) — and the exponent is a real angle on a real circle.

8 TGG and the magic-dimension ladder

We now address the algebraic backbone of TGG: the magic-dimension ladder $\mathbb{R} \rightarrow \mathbb{C} \rightarrow \mathbb{H} \rightarrow \mathbb{O}$, which provides the unique realization of $\text{Spin}(1, 5)$ in dimension $d = 6$ and which forces the quaternionic doublet structure of bulk spinors.

8.1 The magic-dimension ladder

The four classical normed division algebras of Hurwitz are $\mathbb{R}, \mathbb{C}, \mathbb{H}, \mathbb{O}$ (the reals, complex numbers, quaternions, and octonions), of real dimensions 1, 2, 4, 8 respectively. Each admits an associated Lie group $\text{SL}(2, \mathbb{K})$ for $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}, \mathbb{H}\}$ (the octonions are non-associative, so $\text{SL}(2, \mathbb{O})$ is not a group in the strict sense, but its universal enveloping structure plays the analogous role). The remarkable result is the chain of accidental isomorphisms

$$\begin{aligned} \text{Spin}(1, 2) &\cong \text{SL}(2, \mathbb{R}), & \text{Spin}(1, 3) &\cong \text{SL}(2, \mathbb{C}), \\ \text{Spin}(1, 5) &\cong \text{SL}(2, \mathbb{H}), & \text{Spin}(1, 9) &\text{relates to } \text{SL}(2, \mathbb{O}), \end{aligned} \tag{36}$$

which identifies the Lorentz double cover in dimensions $d = 3, 4, 6, 10$ with the $\text{SL}(2, \mathbb{K})$ family. These are exactly the dimensions in which classical superstrings (or analogous structures) admit the clean spinor algebra needed for supersymmetry, and they are exactly the dimensions in which the Lorentz algebra has a manifestly $\text{SL}(2)$ -like structure over a normed division algebra.

The relevance to TGG is that the bulk dimension $d = 6$ is the *quaternionic* entry in the magic-dimension ladder. This is not a choice of convenience: it is forced by GGEP together with the algebraic content of the four-factor gauge group G_{tot} .

8.2 $\text{Spin}(1, 5) \cong \text{SL}(2, \mathbb{H})$ and the quaternionic doublet $\mathbb{H}^2$ (cited)

The explicit construction of the accidental isomorphism $\text{Spin}(1, 5) \cong \text{SL}(2, \mathbb{H})$, the proof that $d = 6$ is uniquely admissible among the magic dimensions, the quaternionic doublet representation $\mathbb{H}^2$ for bulk spinors, and the $\text{Spin}(1, 5)$ generators, structure constants, and Killing form, are all worked out in detail in [Arisaka \(2026, A2-GG6\)](#) Theorem 1 (Dimensional Minimality) and § 10.1 (“The accidental isomorphism $\text{Spin}(1, 5) \cong \text{SL}(2, \mathbb{H})$ ” and “Spinor representations in six dimensions”). We import the results as standing inputs to the TGG construction. In particular:

- $\dim_{\mathbb{R}} \text{Spin}(1, 5) = 15$ (not 12, as a careless count might suggest — see ([Arisaka, 2026, A2-GG6](#), Lemma 10.1) for the clarification);
- the $\mathbb{H}^2$ doublet decomposes under G_{tot} -action into the matter content used throughout Parts III–V;
- the Killing form on each sub-algebra (sign-definite on $\mathfrak{su}(5)$ and Abelian-zero on the two $\mathfrak{u}(1)$ factors) supplies the trace-ratio normalizations $\alpha_R = +1$, $\alpha_5 = -1$ used in CCI-ST2 below.

Section 8.3 carries forward the unique A6 contribution: the descent from the bulk quaternionic doublet to the boundary family triplicity that anchors the $N = 3$ derivation of T4.1.

8.3 From the quaternionic doublet to family triplicity

The connection between the quaternionic doublet structure and the OGN-5 family triplicity $N = 3$ is the algebraic content most relevant to the $N=3$ proof of Arisaka (2026, P2-N3).

The chain of identifications is:

- (1) $\mathbb{H}^2$ has three independent commuting quaternionic structures (J_i, J_j, J_k) , with the bulk quaternionic structure imported from (Arisaka, 2026, A2-GG6, §10.1).
- (2) Under the bulk-to-boundary fermion descent of (Arisaka, 2026, A2-GG6, §10.2), these three structures correspond to three independent ways of organizing the same fermion content into generations.
- (3) The OGN-5 ledger encodes this triplicity as $N = 3$, the second component of $(M, N, L; J, K) = (1, 3, 5; 7, 35)$.
- (4) The family-projector moduli space, on which the GFF Morse cascade of Part V acts, is therefore $\mathbb{CP}^{N-1} = \mathbb{CP}^2$ — a compact symmetric space in the classical classification (Helgason, 1979) — of complex dimension $N_g - 1 = 2$.

The relation between the quaternionic structure and the family count is not yet a closed derivation at the level of TGG alone: the precise statement $N = 3$ requires the combined input of CCI (Part III, OGN-5 ledger forced under MSMF) and the GFF Morse cascade (Part V, $N_g \leq 3$ from cascade closure). But the upstream algebraic origin of family triplicity is visible already at the TGG level, in the three quaternionic structures of $\mathbb{H}^2$.

This is one of the principal payoffs of formalizing TGG: the “family triplicity” that GG-A9-SM treats as a derived OGN-5 integer is shown, at the TGG level, to have a direct algebraic origin in the three quaternionic structures of the bulk spinor space. The $N = 3$ of the Standard Model, the $N = 3$ of OGN-5, and the $N = 3$ of the quaternionic structures of $\mathbb{H}^2$ are not three separate “3”s; they are one 3 traced through three layers of the program.

9 The boundary layer: the unified stress tensor $\mathcal{T}_{\text{TGG}}^{\mu\nu}$

We have so far developed the bulk layer of TGG: the principal G_{tot} -bundle, the composite connection, the curvature decomposition, and the magic-dimension-ladder backbone. We now turn to the boundary layer: the unified 4D stress tensor $\mathcal{T}_{\text{TGG}}^{\mu\nu}$ on $M_4 = \partial M_6$, obtained by Kaluza–Klein zero-mode reduction of the bulk action together with the BI-6 bulk-to-boundary descent. This section is new content in, restoring the stress-tensor layer that was inadvertently omitted from.

9.1 Why a unified boundary stress tensor is needed

Before TGG was made explicit at the boundary level, the boundary theory carried four disconnected stress-energy objects:

- $T_{\text{matter}}^{\mu\nu}$ – the matter stress tensor (electrons, quarks, neutrinos, etc.);
- $T_{\text{gauge}}^{\mu\nu}[F]$ – the gauge-field stress tensor constructed from $F = F_5 + F_B + F_{Qi}$;

- $X_{\text{GS}}^{\mu\nu}$ – the boundary shadow of the Green–Schwarz inflow term in BI-6;
- $T_{\text{KK}}^{\mu\nu}$ – the Kaluza–Klein heavy-mode remnants.

The question was: is their sum conserved? And if so, why? In the closed-system QFT framework, the answer is empirical – one *checks* that $\nabla_\mu (T_{\text{matter}}^{\mu\nu} + T_{\text{gauge}}^{\mu\nu}) = 0$ on shell, but there is no constitutional reason for it to hold. In the TGG framework, the answer becomes a *theorem*: the sum is conserved because all four pieces are the boundary projection of one bulk geometric object, and the bulk closure condition (BI-6) descends to the boundary as the master conservation law $\nabla_\mu \mathcal{T}_{\text{TGG}}^{\mu\nu} = 0$ via OGGEF.

Remark 9.1 (Physical intuition: from Newton’s third law to a six-sector master balance). Newton’s third law – “every action has an equal and opposite reaction” – is, in the language of fields, the statement that energy and momentum lost by one body are gained by another, so the total stress-energy is conserved. In Maxwell theory this generalizes to a two-sector balance: energy lost by the electromagnetic field (via the Lorentz force on charged matter) equals energy gained by matter, with the unified $T_{\text{matter}}^{\mu\nu} + T_{\text{EM}}^{\mu\nu}$ exactly conserved provided electric current is conserved. TGG is the *six-sector* generalization of this same principle: gravity, visible gauge (SU(5)), baryon (U(1)_B), informational (U(1)_{Qi}), Green–Schwarz inflow, and Kaluza–Klein remnants are all in mutual balance, with the Master Conservation Law $\nabla_\mu \mathcal{T}_{\text{TGG}}^{\mu\nu} = 0$ holding because all four sectoral currents are simultaneously conserved. Failure of any one current would leave an uncanceled divergence, just as in Maxwell theory the failure of electric charge conservation would break energy-momentum conservation. In TGG, this six-sector balance is not imposed by hand; it is a *theorem* that follows from BI-6 closure in the bulk via OGGEF descent.

9.2 Kaluza–Klein zero-mode reduction of the bulk action

The bulk action of TGG is, in the schematic form,

$$S_{\text{bulk}}^{\text{TGG}} = \int_{M_6} \sqrt{|g_6|} \left[\frac{1}{2\kappa_6^2} R_{\text{Spin}(1,5)} - \frac{1}{4g_5^2} \text{Tr} F_5^2 - \frac{1}{4g_B^2} F_B^2 - \frac{1}{4g_{Q_i}^2} F_{Q_i}^2 + \mathcal{L}_{\text{matter}} \right] + \int_{M_6} B_2 \wedge X_4, \quad (37)$$

where κ_6 is the bulk gravitational coupling, g_5, g_B, g_{Q_i} are the bulk gauge couplings, $\mathcal{L}_{\text{matter}}$ is the matter Lagrangian, and the last term is the Green–Schwarz counterterm required by BI-6 anomaly closure (Arisaka, 2026, A3-BI6, § 3).

Kaluza–Klein zero-mode reduction along the radial coordinate r and the compact circle S_θ^1 produces the boundary effective action

$$S_{\text{bdy}}^{\text{TGG}} = \int_{M_4} \sqrt{|g_4|} \left[\frac{1}{2\kappa_4^2} R^{(4)} - \frac{1}{4g_4^2} \text{Tr} F^{(4)2} + \mathcal{L}_{\text{matter}}^{(4)} \right] + \cdots, \quad (38)$$

where the four-dimensional gravitational coupling $\kappa_4 = \kappa_6 / \sqrt{2\pi R_\theta \cdot L_r}$ and the four-dimensional gauge coupling g_4 are the projected zero-modes of their bulk parents, and the dots indicate the Green–Schwarz boundary contribution and Kaluza–Klein heavy-mode corrections.

Varying (38) with respect to the boundary metric yields the modified Einstein equation

$$\frac{1}{\kappa_4^2} G^{\mu\nu} = T_{\text{matter}}^{\mu\nu} + T_{\text{gauge}}^{\mu\nu}[F] + X_{\text{GS}}^{\mu\nu} + T_{\text{KK}}^{\mu\nu}, \quad (39)$$

which is exactly the unified Total Geometric-Gauge stress tensor $\mathcal{T}_{\text{TGG}}^{\mu\nu}$ of **TGG-ST6**.

9.3 The four contributions to $\mathcal{T}_{\text{TGG}}^{\mu\nu}$

We give the explicit form of each of the four contributions.

Matter sector. The boundary matter stress tensor takes the canonical form

$$T_{\text{matter}}^{\mu\nu} = \frac{2}{\sqrt{|g_4|}} \frac{\delta(\sqrt{|g_4|} \mathcal{L}_{\text{matter}}^{(4)})}{\delta g_{\mu\nu}} = \sum_{i \in \text{matter}} T_{(i)}^{\mu\nu}, \quad (40)$$

where the sum runs over the standard-model matter species (quarks, leptons, Higgs scalar) on the boundary.

Gauge sector. The boundary gauge stress tensor is

$$T_{\text{gauge}}^{\mu\nu}[F] = \frac{1}{g_4^2} \text{Tr}(F^{\mu\rho} F^\nu{}_\rho - \frac{1}{4} g^{\mu\nu} F^{\rho\sigma} F_{\rho\sigma}), \quad (41)$$

with $F = F_5 + F_B + F_{Qi}$ the total boundary field strength.

Green–Schwarz inflow. The boundary shadow of the bulk $B_2 \wedge X_4$ counterterm is

$$X_{\text{GS}}^{\mu\nu} = \frac{2}{\sqrt{|g_4|}} \frac{\delta(\sqrt{|g_6|} B_2 \wedge X_4)|_{r=0}}{\delta g_{\mu\nu}}, \quad (42)$$

which encodes the bulk-to-boundary descent of the bulk anomaly polynomial X_4 onto the boundary. Its explicit form depends on the choice of B_2 representative; see [Arisaka \(2026, A3-BI6, § 3.2\)](#) for the details.

Kaluza–Klein remnant. The boundary contribution from the heavy KK modes is

$$T_{\text{KK}}^{\mu\nu} = \sum_{n \geq 1} T_{(n)}^{\mu\nu}, \quad (43)$$

the sum over KK levels n of the boundary stress contributions of the massive KK modes at scale $\sim n \cdot M_{\text{lock}}$ where $M_{\text{lock}} = E_{\text{Pl}}/2e^{35} \simeq 3.85 \text{ TeV}$ is the lock scale of [Arisaka \(2026, P1-KK\)](#).

9.4 The projected modified Bianchi identity (PBI-4)

The 6D Bianchi identity $dH_3 = X_4$ projects to the 4D boundary as the *projected modified Bianchi identity*:

$$\nabla_\mu (G^{\mu\nu} + X_{\text{GS}}^{\mu\nu}) = 0, \quad (\text{PBI-4}) \quad (44)$$

which is a derived 4D equation, not a primitive postulate.

Proposition 9.2 (PBI-4 from BI-6 descent). *On any TGG configuration satisfying **TGG-ST1–TGG-ST6**, the projected modified Bianchi identity (44) holds on M_4 as a consequence of the bulk identity $dH_3 = X_4$ and the Kaluza–Klein zero-mode reduction.*

Sketch. The bulk identity $dH_3 = X_4$ together with the Bianchi identity $\nabla_\mu G_{(6)}^{\mu\nu} = 0$ for the bulk Einstein tensor implies, after zero-mode reduction along r and θ , $\nabla_\mu^{(4)} (G_{(4)}^{\mu\nu} + X_{\text{GS}}^{\mu\nu}) = 0$. The

boundary $X_{\text{GS}}^{\mu\nu}$ is exactly the boundary shadow of the bulk X_4 via the variational expression (42); its boundary divergence cancels the failure of the Einstein tensor's divergence to vanish in the presence of the GS counterterm. See Arisaka (2026, A3-BI6, § 3.4) for the detailed proof. ■

9.5 The boundary layer ready for Theorem T1

With the explicit form of $\mathcal{T}_{\text{TGG}}^{\mu\nu}$ in (39), the four contributions in (40)–(43), and the projected modified Bianchi identity (44), we have all the ingredients required to state and prove the Master Conservation Law $\nabla_\mu \mathcal{T}_{\text{TGG}}^{\mu\nu} = 0$. This is the content of Theorem T1 in Section 10.

10 Theorem T1 (T-GRIP-1): Existence, uniqueness, and Master Conservation of TGG

We now state and prove the principal theorem of Part II, the existence and uniqueness of TGG on the 6D bulk.

10.1 Dual-layer architecture and the compatibility relation $\mathfrak{C}_{\text{TGG}}$

Theorem T1 is, in the language of L16, a dual-layer principal theorem of Gi-4. Its *algebraic layer* (T1.1 and T1.2) asserts the existence and uniqueness of the composite connection $\mathcal{A}_{\text{tot}} = \omega \oplus A$ on the principal G_{tot} -bundle over the 6D bulk M_6 , and of the unified boundary stress tensor $\mathcal{T}_{\text{TGG}}^{\mu\nu}$ obtained by Kaluza–Klein zero-mode reduction of the bulk action under OGGEF. Its *operational/physical layer* (T1.3) asserts the Master Conservation Law $\nabla_\mu \mathcal{T}_{\text{TGG}}^{\mu\nu} = 0$ as a laboratory-falsifiable Ward-identity statement on the four conserved boundary currents J_e, J_B, J_{Qi}, J_g .

The two layers are tied together by the named compatibility relation

$$\boxed{\mathfrak{C}_{\text{TGG}} \equiv (\text{T1.1} \wedge \text{T1.2}) \implies \text{T1.3} \quad (\text{via Noether–Bianchi descent on the } G_{\text{tot}}\text{-bundle}).} \quad (45)$$

Compatibility proposition. $\mathfrak{C}_{\text{TGG}}$ above is a one-line proposition with the proof: given the existence and uniqueness of $\mathcal{A}_{\text{tot}}$ (T1.1) and of $\mathcal{T}_{\text{TGG}}^{\mu\nu}$ (T1.2), the Master Conservation Law (T1.3) follows by applying the Noether identity $\partial_\mu J^\mu = 0$ to each gauge sector independently and the geometric Bianchi identity $dR = 0$ on the gravitational sector, and summing. The summed identity is the divergence of $\mathcal{T}_{\text{TGG}}^{\mu\nu}$. The compatibility relation makes T1 cite-able at the algebraic level by mathematical-physics readers (who care about the principal-bundle structure and the magic-dimension ladder) and at the physical level by phenomenology readers (who care about the laboratory Ward identities of the four conserved currents). Both readings refer to the same mathematical object TGG; the dual-layer architecture simply makes explicit which face a given theorem statement presents.

10.2 Statement of Theorem T1

Theorem 10.1 (T1: Existence, Uniqueness, and Master Conservation of TGG). *On the 6D bulk M_6 with metric (1) and gauge content $G_{\text{tot}} = \text{Spin}(1, 5) \times \text{SU}(5) \times \text{U}(1)_B \times \text{U}(1)_{Qi}$, the Total*

Geometric-Gauge object TGG of Definition 6.2 satisfies the following three properties:

(T1.1) Bulk-layer existence and uniqueness. *There exists a triple $(P, \mathcal{A}_{\text{tot}}, \mathcal{R}_{\text{tot}})$ satisfying TGG-ST1–TGG-ST5, unique up to gauge equivalence (31). Existence and uniqueness follow from GGEP, LC, the magic-dimension ladder, and the product geometry of M_6 , with no additional physical axiom required.*

(T1.2) Boundary-layer existence and uniqueness. *Under Kaluza–Klein zero-mode reduction of the bulk action of TGG-ST2 together with the BI-6 bulk-to-boundary descent, the unified Total Geometric-Gauge stress tensor $\mathcal{T}_{\text{TGG}}^{\mu\nu}$ of TGG-ST6 exists on $M_4 = \partial M_6$, is symmetric, and is unique up to the gauge equivalence inherited from the bulk.*

(T1.3) Master Conservation Law. *Reading the modified Einstein equation as the definition of $\mathcal{T}_{\text{TGG}}^{\mu\nu}$ on the gravitational side (so that the conservation of the gravitational part is automatic by the geometric Bianchi identity), the substantive content of T1.3 lives on the Ward-package side: in the gauge-closed boundary sector, the unified stress tensor satisfies the covariant conservation law*

$$\nabla_\mu \mathcal{T}_{\text{TGG}}^{\mu\nu} = 0, \quad (46)$$

holding if and only if the full Ward package (14) of TGG-ST7 is satisfied simultaneously.

Remark 10.2 (Physics intuition: TGG as the constitutional culmination of Maxwell–Yang–Mills–Einstein). T1 says, in plain terms, that the 6D bulk admits exactly one unified geometric–gauge object TGG, and that this object’s boundary shadow conserves its stress–energy automatically. Indeed, the algebraic layer (T1.1, T1.2) reads as a uniqueness statement on the principal G_{tot} -bundle: out of all conceivable bundle structures over the 6D bulk, exactly one is compatible with GGEP and LC, and that one is the TGG structure of Definition 6.2. The operational layer (T1.3) reads as a laboratory falsification target: every boundary observer in the lab measures exactly four conserved currents (J_e, J_B, J_{Qi}, J_g) , and the Master Conservation Law says that any cross-current leakage would falsify TGG. The dual-layer architecture is what lets a mathematical physicist verify T1.1–T1.2 from bundle theory while a LHC experimenter verifies T1.3 from precision Ward-identity tests — two falsification routes converging on a single geometric object.

10.3 Proof of Theorem T1: bulk-layer existence (T1.1, existence part)

Proof of bulk-layer existence (T1.1). We construct TGG explicitly.

Step 1: the principal $\text{Spin}(1, 5)$ -bundle. By TGG-ST4, the 6D bulk M_6 is Lorentzian, time-orientable, and admits a global frame; therefore it admits a principal $\text{Spin}(1, 5)$ -bundle $P_{\text{Spin}(1, 5)} \rightarrow M_6$, given by the spin lift of the frame bundle. The associated Levi-Civita-type connection $\omega_{\text{Spin}(1, 5)}$ is unique up to gauge.

Step 2: the principal $\text{SU}(5)$ -bundle. The 6D bulk admits a global $\text{SU}(5)$ bundle because: (i) M_6 is contractible in the radial direction (modulo the scale profile which does not affect homotopy); (ii) the second Stiefel–Whitney obstruction vanishes; and (iii) the boundary $\partial M_6 = M_4$ admits the embedded Standard-Model gauge structure $\text{SU}(3) \times \text{SU}(2) \times \text{U}(1)$ as a sub-group of $\text{SU}(5)$. The connection A_5 is determined up to gauge by the boundary $\text{SU}(5)$ data of Arisaka (2026, A9-SM).

Step 3: the principal $U(1)_B$ and $U(1)_{Q_i}$ bundles. M_6 is non-compact, and $U(1)$ bundles over non-compact bases are classified by the integral cohomology $H^2(M_6; \mathbb{Z})$, which is generated by the integer multiples of the compact θ -direction. Both $P_{U(1)_B}$ and $P_{U(1)_{Q_i}}$ exist with topological winding numbers that are determined by the Wilson-line data of Definition 7.2.

Step 4: the fiber product. Define

$$P := P_{\text{Spin}(1,5)} \times_{M_6} P_{\text{SU}(5)} \times_{M_6} P_{U(1)_B} \times_{M_6} P_{U(1)_{Q_i}}$$

as in equation (22). The fiber product is a smooth manifold because each factor is smooth and the base M_6 is common. The free right action of G_{tot} on P is the diagonal action of the four sub-groups, which is free because each sub-action is free. The quotient is M_6 , satisfying **TGG-ST1**.

Step 5: the composite connection. Define

$$\mathcal{A}_{\text{tot}} := \omega_{\text{Spin}(1,5)} + A_5 + A_B + A_{Q_i},$$

viewed as a connection one-form on P via the obvious sum of pull-backs from the four factors. Since the four direct summands of $\mathfrak{g}_{\text{tot}}$ commute, $\mathcal{A}_{\text{tot}}$ is a well-defined $\mathfrak{g}_{\text{tot}}$ -valued one-form satisfying the principal-bundle compatibility conditions. This provides **TGG-ST2**.

Step 6: TGG-ST3 and TGG-ST4. These are direct consequences of the construction. **TGG-ST3** follows from the choice of bulk dimension $d = 6$ and the magic-ladder isomorphism (11). **TGG-ST4** follows from the product-geometry structure of M_6 , which is preserved by all four gauge actions because they commute with diffeomorphisms of M_6 .

Step 7: TGG-ST5. The unification axiom is automatic in this construction: the four sub-connections are projections of the single connection $\mathcal{A}_{\text{tot}}$ onto its four direct-summand components, and gauge transformations of $\mathcal{A}_{\text{tot}}$ act diagonally on the four sub-connections by the natural adjoint action of G_{tot} on $\mathfrak{g}_{\text{tot}}$.

This completes the existence proof. ■

10.4 Proof of Theorem T1: bulk-layer uniqueness (T1.1, uniqueness part)

Proof of bulk-layer uniqueness up to gauge (T1.1). Suppose $(P, \mathcal{A}_{\text{tot}}, \mathcal{R}_{\text{tot}})$ and $(P', \mathcal{A}'_{\text{tot}}, \mathcal{R}'_{\text{tot}})$ both satisfy **TGG-ST1–TGG-ST5**. We show that they are gauge-equivalent.

Step 1: bundle isomorphism. Each principal H -sub-bundle P_H for $H \in \{\text{Spin}(1,5), \text{SU}(5), U(1)_B, U(1)_{Q_i}\}$ is determined by its topology, which is fixed by the bulk geometry and the boundary gauge data. Two bundles with the same topology over the same base are isomorphic as principal H -bundles (Kobayashi & Nomizu, 1969). Therefore $P \cong P'$ as principal G_{tot} -bundles.

Step 2: connection equivalence. Two connections on the same principal H -bundle differ by an $\mathfrak{h}$ -valued one-form on P pulled back from M_6 (the difference of two connections is a tensor). This one-form is an element of $\Omega^1(M_6, \mathfrak{h})$. By the inverse-function theorem applied to gauge transformations (31), any such difference is locally exact in the gauge sense: there exists a gauge transformation $g_H : M_6 \rightarrow H$ such that $A'_H = g_H^{-1} A_H g_H + g_H^{-1} dg_H$. Globally, the existence of g_H depends on the topology of the principal bundle P_H (the obstruction lies in $H^1(M_6; \mathfrak{h})$), but for each of our four sub-bundles the obstruction vanishes (the H^1 -cohomology of the contractible M_6 is

zero). Therefore each sub-connection is gauge-equivalent.

Step 3: composite gauge equivalence. The four-factor structure of G_{tot} implies that the composite gauge transformation $g = (g_{\text{Spin}(1,5)}, g_5, g_B, g_{Q_i})$ acts diagonally on $\mathcal{A}_{\text{tot}}$, so the composite gauge equivalence is $\mathcal{A}'_{\text{tot}} = g^{-1}\mathcal{A}_{\text{tot}}g + g^{-1}dg$. This is the global gauge transformation that takes $(P, \mathcal{A}_{\text{tot}}, \mathcal{R}_{\text{tot}})$ to $(P', \mathcal{A}'_{\text{tot}}, \mathcal{R}'_{\text{tot}})$.

This completes the uniqueness proof. ■

10.5 Proof of Theorem T1: boundary stress tensor existence and uniqueness (T1.2)

Proof of (T1.2). We construct $\mathcal{T}_{\text{TGG}}^{\mu\nu}$ explicitly as the variational derivative of the boundary effective action.

Step 1: zero-mode reduction. By Section 9.2, the bulk action (37) reduces under Kaluza–Klein zero-mode projection along r and θ to the boundary effective action (38).

Step 2: variational derivation of $\mathcal{T}_{\text{TGG}}^{\mu\nu}$. Varying (38) with respect to the boundary metric $g_{\mu\nu}$ yields the modified Einstein equation (39), with the right-hand side identified as $\mathcal{T}_{\text{TGG}}^{\mu\nu}$ via the four canonical contributions (40)–(43). Each contribution is well-defined and symmetric in $\mu\nu$ by the standard variational formulae.

Step 3: uniqueness from bulk uniqueness. By (T1.1), $(P, \mathcal{A}_{\text{tot}}, \mathcal{R}_{\text{tot}})$ is unique up to gauge equivalence. The zero-mode reduction is canonical (no choices beyond the choice of r and θ coordinates), so $\mathcal{T}_{\text{TGG}}^{\mu\nu}$ is determined by the bulk data up to the same gauge equivalence. Two gauge-equivalent bulk configurations give the same boundary stress tensor up to the induced boundary diffeomorphism.

This completes the proof of (T1.2). ■

10.6 Proof of Theorem T1: Master Conservation Law (T1.3)

Proof of (T1.3). We prove that $\nabla_\mu \mathcal{T}_{\text{TGG}}^{\mu\nu} = 0$ holds iff the full Ward package (14) is satisfied.

Step 1: divergence of the gauge sector. The gauge-sector stress tensor (41) has divergence

$$\nabla_\mu T_{\text{gauge}}^{\mu\nu} = -F_e^{\nu\mu} J_\mu^e - F_B^{\nu\mu} J_\mu^B - F_{Q_i}^{\nu\mu} J_\mu^{Q_i}, \quad (47)$$

i.e., minus the Lorentz-force density transferred from the gauge fields to the matter currents. This identity follows from the gauge equations of motion $\nabla_\mu F^{\mu\nu} = J^\nu$ together with the Bianchi identity for F .

Step 2: divergence of the matter sector. Applying the matter equations of motion plus the Ward identities for each conserved current $J_e^\mu, J_B^\mu, J_{Q_i}^\mu$ gives

$$\nabla_\mu T_{\text{matter}}^{\mu\nu} = +F_e^{\nu\mu} J_\mu^e + F_B^{\nu\mu} J_\mu^B + F_{Q_i}^{\nu\mu} J_\mu^{Q_i}. \quad (48)$$

The signs are exactly opposite to those in (47): energy and momentum lost by the gauge fields are gained by matter, and vice versa. This is the standard force-balance equation.

Step 3: divergence of the GS-inflow + KK sectors. By Proposition 9.2, the projected

modified Bianchi identity (44) gives

$$\nabla_\mu (G^{\mu\nu} + X_{\text{GS}}^{\mu\nu}) = 0, \quad \nabla_\mu T_{\text{KK}}^{\mu\nu} = 0, \quad (49)$$

where the first follows from the BI-6 descent and the second from the fact that the heavy KK modes are decoupled at boundary scales (their stress contribution is conserved by the bulk equations of motion restricted to the boundary).

Step 4: assembly. Adding the four divergences:

$$\nabla_\mu \mathcal{T}_{\text{TGG}}^{\mu\nu} = \nabla_\mu T_{\text{matter}}^{\mu\nu} + \nabla_\mu T_{\text{gauge}}^{\mu\nu} + \nabla_\mu X_{\text{GS}}^{\mu\nu} + \nabla_\mu T_{\text{KK}}^{\mu\nu} \quad (50)$$

$$= (+F \cdot J) + (-F \cdot J) + (-\nabla G) + 0 \quad (51)$$

$$= -\nabla_\mu G^{\mu\nu}. \quad (52)$$

By the modified Einstein equation (39), $\nabla_\mu G^{\mu\nu} = \kappa_4^2 \nabla_\mu \mathcal{T}_{\text{TGG}}^{\mu\nu}$, so the equation $\nabla_\mu \mathcal{T}_{\text{TGG}}^{\mu\nu} = 0$ closes consistently. The condition for closure is exactly the simultaneous satisfaction of the Ward package (14): failure of any one current conservation leaves an uncanceled divergence in $\mathcal{T}_{\text{TGG}}^{\mu\nu}$.

This proves the iff in (T1.3). ■

Remark 10.3 (Physical content of the Master Conservation Law). The Master Conservation Law $\nabla_\mu \mathcal{T}_{\text{TGG}}^{\mu\nu} = 0$ is the 6D-bulk-to-4D-boundary descent of the BI-6 closure, identifying the boundary as the boundary image of a single conserved bulk geometric object. It is the constitutional analog of energy-momentum conservation in ordinary Maxwell theory (where conservation requires the electric current to be conserved): in the present six-sector generalization, conservation of the unified stress tensor requires simultaneous conservation of all four currents (electric, baryonic, informational, and metric-inflow), and all four are conserved because they are tied to one anomaly rope upstairs in the bulk. This is the foundation of the astrophysically falsifiable predictions of CCI (Part III), which is the OGGEF zero-mode projection of (46) in the open-system regime.

10.7 Corollary: TGG and DGI

Corollary 10.4 (TGG realizes DGI). *The unique TGG structure on M_6 , restricted to its undisturbed phase (i.e., before any boundary projection), realizes the Dynamical Gauge Invariance (DGI) stage of the GRIP cascade introduced in (Arisaka, 2026, A6-GRIP). Specifically:*

- the full G_{tot} symmetry of TGG on M_6 corresponds to DGI;
- the boundary projection $\Pi : M_6 \rightarrow M_4$, which selects a specific Wilson-line holonomy (Section 7.5), realizes GSSB (Geometric Spontaneous Symmetry Breaking);
- the boundary attractors selected by the broken-phase GFF Morse landscape (Part IV) realize GA (Geometric Attractors).

The complete GRIP cascade $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$ (Part V) thus takes TGG as its initial deepest-invariance configuration. In fiber-resolved terms (Subsection 31.9): the gauge directions that DGI makes dynamical are the θ -fiber's — the circle is the bulk's only gauge bundle — while the dynamism itself is supplied by the r -fiber's openness; DGI is LGI transported by the semigroup.

10.8 Corollary: TGG and GGEP equivalence

Corollary 10.5 (TGG and GGEP are equivalent). *On the 6D bulk M_6 , the existence of a TGG structure satisfying **TGG-ST1–TGG-ST5** is equivalent to the GG-Constitutional axiom GGEP. Specifically, TGG is the formal mathematical realization of GGEP on M_6 , and any structure on M_6 that satisfies GGEP is a TGG (up to gauge equivalence).*

Proof. GGEP states that gauge content and geometric content arise from a single primitive composite connection. Axiom **TGG-ST2** encodes this primitive composite connection as $\mathcal{A}_{\text{tot}}$; Structure Theorem **TGG-ST5** encodes its diagonal gauge action; and axioms **TGG-ST1**, **TGG-ST3**, **TGG-ST4** encode the geometric context (bundle, magic-ladder dimension, Lorentzian compatibility) in which the primitive composite connection is realized. Conversely, any structure satisfying GGEP on M_6 admits a unique principal G_{tot} -bundle (by the magic-dimension argument of Section 8) and a unique composite connection up to gauge (by Theorem 10.1). ■

10.9 Status of T1 in the constitutional theorem chain

The proof of Theorem 10.1 relies only on: (i) the magic-dimension-ladder algebraic content (a mathematical fact about Spin groups over normed division algebras); (ii) the product-geometry Lorentzian structure of M_6 (set up in Arisaka (2026, A2-GG6)); (iii) the integrality of principal bundles over non-compact contractible bases (a topological fact); (iv) standard differential-geometric machinery (Chern–Weil, Kobayashi–Nomizu).

No additional physical axiom beyond GGEP of the GG-Constitution is invoked. This is consistent with the derivation policy of (the derivation policy documented in (Arisaka, 2026, A4-Euclid, §§1–2)) and confirms the no-new-axioms principle ((Arisaka, 2026, A4-Euclid, §1.4)): the existence and uniqueness of TGG is already latent in the constitutional theorem chain, and the present paper merely makes it explicit.

11 Worked examples

To make Theorem 10.1 concrete, we work through three explicit examples: TGG on flat $\mathbb{R}^{1,5}$ (Minkowski warm-up), TGG on the M_6 of GG-Theory, and the explicit Wilson-line / holonomy structure in the general case.

11.1 Example 1: TGG on flat $\mathbb{R}^{1,5}$

The simplest example is the trivial Lorentzian bulk $M_6 = \mathbb{R}^{1,5}$, with flat Minkowski metric and trivial G_{tot} bundle structure.

Bundle. $P = \mathbb{R}^{1,5} \times G_{\text{tot}}$ (the trivial product bundle).

Connection. $\mathcal{A}_{\text{tot}} = 0$ in the natural global trivialization; equivalently, in any trivialization, $\mathcal{A}_{\text{tot}}$ is a pure gauge, $\mathcal{A}_{\text{tot}} = g^{-1}dg$ for some $g : \mathbb{R}^{1,5} \rightarrow G_{\text{tot}}$.

Curvature. $\mathcal{R}_{\text{tot}} = 0$. All four sub-curvatures vanish.

Holonomy. Trivial along all loops. The Wilson-line order parameters (32)–(35) all evaluate to identity.

Discussion. Flat $\mathbb{R}^{1,5}$ is a degenerate case of TGG: it satisfies all five axioms but with trivial content. The holonomy is trivial, the curvature vanishes, and there is no compact direction to support non-trivial Wilson lines. This example serves as a sanity check that the axioms are not vacuous (the trivial bundle is a valid TGG) and as a reference point against which the non-trivial case can be contrasted.

11.2 Example 2: TGG on the $\mathcal{M}_6$ of GG-Theory

The principal example is the 6D bulk $M_6 = \mathbb{R}^{1,3} \times I_r \times S_\theta^1$ with metric (1), where $I_r = [0, K]$ is the radial interval (UV at $r = 0$, IR at $r = K$) and S_θ^1 is the compact angular direction with circumference $2\pi R_\theta$.

Bundle. The principal G_{tot} -bundle is non-trivial along the compact θ -direction. Specifically, the principal $\text{Spin}(1, 5)$ -bundle is trivial in the strict-bundle sense (because M_6 is homotopy equivalent to S^1), but the Levi-Civita connection has non-trivial θ -component derivatives that encode the scale profile. The $\text{SU}(5)$ -bundle has non-trivial topological winding around S_θ^1 (orbifold breaking). The two $\text{U}(1)$ -bundles have non-trivial Wilson lines along S_θ^1 .

Connection.

- $\omega_{\text{Spin}(1,5)}$: the Levi-Civita spin connection of the scale-profile metric. In the conventions of Arisaka (2026, A2-GG6, Appendix A), the non-trivial components are $\omega^{0r} = -e^{-2r} dt$, etc. The θ -component vanishes by Proposition 6.8.
- A_5 : the orbifold Wilson line implementing $\text{SU}(5) \rightarrow \text{SU}(3) \times \text{SU}(2) \times \text{U}(1)$ with hypercharge embedding $Y_0 = \text{diag}(-1/3, -1/3, -1/3, 1/2, 1/2)$, encoded by the θ -component $A_5^\theta = \text{diag}(-1/3, -1/3, -1/3, 1/2, 1/2) \partial_\theta$ in the natural $\text{SU}(5)$ basis.
- A_B : the gravity-baryon Abelian connection. On the admissible class of Arisaka (2026, A3-BI6) it is flat — $F_B = 0$, with no gradient of the Wilson line across the throat, $\Phi_B = 0$ — its four-dimensional components Dirichlet-terminated at both ends of the interval and its circle component Neumann, so that the circle holonomy α_B survives as fixed background data; the generator $B_0 = \text{diag}(1/3, 1/3, 1/3, 0, 0)$ is the charge the connection couples to, not a component of it. An earlier statement of this bullet gave A_B a non-trivial θ -component $B_0 \partial_\theta$; it is retired (Remark 11.1).
- A_{Qi} : the quantum-information Abelian connection, on the same class flat along the throat with its circle holonomy α_{Qi} fixed background data; the earlier reading of that holonomy as the Kaxion phase is retired (Remark 11.1).

Curvature.

- $R_{\text{Spin}(1,5)}$: non-trivial in the radial direction, encoding the bulk curvature. Explicitly, the Riemann tensor of (1) has non-vanishing components proportional to e^{-2r} .
- F_5 : non-trivial radial-angular components arising from the orbifold breaking; $F_B = F_{Qi} = 0$ on the admissible class of Arisaka (2026, A3-BI6), the Abelian Wilson lines carrying no gradient across the throat.

Holonomy. The Wilson-line order parameters along S_θ^1 are the phases that drive the orbifold

projection and, through the $SU(5)$ coset, the Wilson–Hosotani electroweak breaking (Arisaka, 2026, A9-SM, Hosotani, 1983). The two Abelian holonomies are background data; and the Kaxion oscillation, which an earlier statement placed here as the $U(1)_{Q_i}$ phase, is the tensor remnant of Arisaka (2026, C1-Kaxion), not a Wilson line. In Part V of the present paper the non-Abelian Wilson lines become the order parameters of the GRIP cascade.

Discussion. This is the physically realized TGG of GG-Theory. All the machinery developed in Part I (the scale-profile metric, the OGN-5 ledger, the gauge content) is now organized under a single composite connection $\mathcal{A}_{\text{tot}}$, with the four sub-sectors as projections. The non-trivial θ -holonomy of the $SU(5)$ sub-connection drives the dynamical content of the boundary theory; the $\text{Spin}(1, 5)$ holonomy is trivial by Proposition 6.8, and the two Abelian holonomies are fixed background data on the admissible class (Remark 11.1).

11.3 Example 3: explicit Wilson lines and holonomy phases

We give explicit numerical values for the Wilson-line holonomies on the M_6 , which will be used in Parts III–V.

Spin holonomy. $W_{\text{Spin}(1,5)}(S_\theta^1) = \mathbf{1}_{\text{Spin}(1,5)}$, by Proposition 6.8.

$SU(5)$ holonomy (orbifold breaking). The eigenvalues of $W_5(S_\theta^1)$ are $\{e^{-2\pi i/3}, e^{-2\pi i/3}, e^{-2\pi i/3}, e^{\pi i}, e^{\pi i}\}$ in the Y -eigenstate basis, corresponding to the orbifold breaking of $SU(5)$ at the boundary. The unbroken subgroup is the commutant of W_5 , which is $SU(3) \times SU(2) \times U(1)_Y$.

$U(1)_B$ holonomy (background data). $W_B(S_\theta^1) = \exp(i\alpha_B)$ where α_B is fixed background data on the admissible class of the anomaly paper (Arisaka, 2026, A3-BI6), taking the ledger’s value $\alpha_B = (2\pi/L) \cdot M = 2\pi/5$; its fluctuation is an exposure registered there, and the one-loop potential of the endpoint quarks does not select the ledger’s value by itself, so the value stands on the clause that fixes it (Remark 11.1).

$U(1)_{Q_i}$ holonomy (background data). $W_{Q_i}(S_\theta^1) = \exp(i\alpha_{Q_i})$ where α_{Q_i} is, on the same class, fixed background data with its fluctuation an exposure. An earlier statement of this paragraph read α_{Q_i} as the Kaxion field value, dynamical and slowly varying; the Kaxion of Arisaka (2026, C1-Kaxion) is the tensor remnant, a single line, and the reading is retired.

Remark 11.1 (Registered disagreement: the $U(1)_B$ background, and the occupant of the Wilson-line stratum). Earlier statements of this paper gave the gravity-baryon connection a non-flat radial-angular curvature, a circle component $A_B^\theta = B_0 \partial_\theta$, and read its holonomy as the electroweak Wilson–Hosotani phase, with the $U(1)_{Q_i}$ holonomy as the Kaxion phase. The anomaly paper’s admissible class (Arisaka, 2026, A3-BI6) has the connection flat, $\Phi_B = 0$ identically, Dirichlet-terminated in its four-dimensional components at both ends of the interval, with the circle holonomy fixed background data and its fluctuation an open exposure, and the same for $U(1)_{Q_i}$; the Standard-Model paper places the electroweak holonomy in the $SU(5)$ coset and has withdrawn its earlier Higgs reading (Arisaka, 2026, A9-SM); and the consequence paper derives the Kaxion as the tensor remnant, a single line, not as a Wilson line (Arisaka, 2026, C1-Kaxion). Those are the later papers and the computed ones, and this paper adopts their connection: the sentences above are rewritten to it, and the earlier readings are retired rather than carried beside it. Two things follow that this paper must say and cannot yet settle. First, the ledger’s holonomy value $2\pi/5$ stands on the clause

that fixes it as background: in a flat-circle model of the field content the one-loop potential of the endpoint quarks does not select it, so the value is an input of the anomaly paper's class and not an output of any dynamics in this one. Second, the compact-angle stratum Σ_θ of §24.1 was given the Kaxion as its occupant; with both Abelian circle holonomies background data and the Kaxion a tensor mode, which field's holonomy the stratum carries, and whether it is a dynamical stratum at all rather than the exposure the anomaly paper registers, is open. What this paper uses of Σ_θ — a one-dimensional periodic stratum with a locked phase and structurally empty exits — does not depend on the answer, and the Kaxion's stability argument is carried in §31.5 at that grade.

11.4 Chern–Simons construction of the Kalb–Ramond three-form

A central output of TGG that is used directly in Part III is the Kalb–Ramond three-form $H_3 \in \Omega^3(M_6)$ whose differential gives the four-form anomaly polynomial X_4 . We now construct H_3 explicitly from TGG data via the Chern–Simons descent.

Definition 11.2 (Chern–Simons three-form). For a connection one-form A taking values in a Lie algebra $\mathfrak{h}$ with curvature $F = dA + A \wedge A$, the Chern–Simons three-form is

$$\text{CS}(A) = \text{Tr}\left(A \wedge dA + \frac{2}{3}A \wedge A \wedge A\right), \quad (53)$$

where the trace is taken in a suitably normalized representation of $\mathfrak{h}$. This three-form satisfies the descent identity

$$d\text{CS}(A) = \text{Tr}(F \wedge F). \quad (54)$$

For an Abelian connection (where $A \wedge A = 0$), the formula reduces to $\text{CS}(A) = A \wedge dA$ with $d\text{CS}(A) = F \wedge F$.

Proposition 11.3 (TGG-induced Chern–Simons construction of H_3). *Given the TGG structure $(P, \mathcal{A}_{\text{tot}}, \mathcal{R}_{\text{tot}})$ on M_6 , define*

$$H_3 := \text{CS}(\omega_{\text{Spin}(1,5)}) - \text{CS}(A_5) - \frac{2}{5}\text{CS}(A_B) - \frac{2}{5}\text{CS}(A_{Qi}), \quad (55)$$

with the Chern–Simons three-forms taken in the natural normalized traces of each factor. Then

$$dH_3 = \text{tr } R \wedge R - \text{Tr } F_5 \wedge F_5 - \frac{2}{5}F_B \wedge F_B - \frac{2}{5}F_{Qi} \wedge F_{Qi} = X_4, \quad (56)$$

which is the BI-6 modified Bianchi identity of Arisaka (2026, A3-BI6).

Proof. Apply the descent identity (54) to each of the four sub-connections:

$$d\text{CS}(\omega_{\text{Spin}(1,5)}) = \text{tr}(R_{\text{Spin}(1,5)} \wedge R_{\text{Spin}(1,5)}) = \text{tr } R^2, \quad (57)$$

$$d\text{CS}(A_5) = \text{Tr}(F_5 \wedge F_5) = \text{Tr } F_5^2, \quad (58)$$

$$d\text{CS}(A_B) = F_B \wedge F_B, \quad (59)$$

$$d\text{CS}(A_{Qi}) = F_{Qi} \wedge F_{Qi}. \quad (60)$$

Substituting into (55) and applying d termwise gives (56). This is exactly X_4 of BI-6. ■

Remark 11.4 (Sign and coefficient origins). The signs and coefficients in (55) are not free choices. The plus sign of the gravity (Spin) Chern–Simons term and the minus sign of the gauge Chern–Simons terms reflect the Killing-form sign convention of the corresponding Lie algebras (positive-definite for the compact gauge factors, negative-definite for the non-compact $\text{Spin}(1, 5)$). The $-2/5$ coefficients of the two $\text{U}(1)$ Chern–Simons terms come from the inner-product normalizations on $\mathbb{H}^2$ of the underlying $\text{SL}(2, \mathbb{H})$ structure (Arisaka, 2026, A3-BI6, Theorem 9.1); they are forced by the integrality requirement on $[X_4] \in H^4(M_6; \mathbb{Z})$ and cannot be modified without violating BI-6 closure.

The three-form H_3 constructed in (55) is therefore a derived object of TGG, not an independent input. This is one of the key consequences of GGEP unification: the Kalb–Ramond field that drives BI-6 is fully determined by the TGG composite connection.

11.5 Gauge-transformation properties of H_3 and X_4

Under a gauge transformation $\mathcal{A}_{\text{tot}} \rightarrow g^{-1}\mathcal{A}_{\text{tot}}g + g^{-1}dg$, the Chern–Simons three-form changes by an exact form plus a closed Wess–Zumino term:

$$\text{CS}(A^g) - \text{CS}(A) = d\alpha_2(g, A) + \omega_{\text{WZ}}(g), \quad (61)$$

where α_2 is a two-form depending on g and A and $\omega_{\text{WZ}}(g) = \frac{1}{3} \text{Tr}(g^{-1}dg)^3$ is the Wess–Zumino three-form. Therefore H_3 is not strictly gauge-invariant but transforms by an exact-plus-WZ shift. The four-form $X_4 = dH_3$, however, is fully gauge-invariant because $d^2 = 0$ and $d\omega_{\text{WZ}} = 0$ for the WZ piece. The gauge-invariance of X_4 is what makes its cohomology class $[X_4] \in H^4(M_6; \mathbb{Z})$ a well-defined topological invariant of TGG.

This gauge-transformation behavior is the algebraic heart of the Green–Schwarz mechanism — the 1984 pioneering of the geometrization that AX-2 GGEP of the GG-Constitution elevates from a technical refinement to a foundational axiom (Arisaka, 2026, A3-BI6, § 8.7). In Part III, the requirement that the combined gauge-transformation anomaly of H_3 cancel against the boundary chiral-fermion anomaly will give the integrality and non-triviality of $[X_4]$, which is the content of CCI.

11.6 Summary of the algebraic content of TGG

Bringing together the content of Sections 7.2 through 11.5, the algebraic content of TGG is summarized in the following table.

Table 2: The four sub-sectors of the composite connection, each with its group, its connection one-form, its curvature, and the Chern–Simons form it contributes to the three-form field strength. The signs in the last column are the content of the table: gravity enters with one sign and the three gauge sectors with the other, and the two abelian factors carry the same $2/5$ coefficient that the OGN-5 ledger fixes. Reading that column downward is reading the anomaly-canceling combination which TGG asserts is a single object rather than four that happen to sit on the same manifold.

Sector	Group	Connection	Curvature	CS form contribution to H_3
Gravity	$\text{Spin}(1, 5)$	$\omega_{\text{Spin}(1,5)}$	$R_{\text{Spin}(1,5)}$	$+\text{CS}(\omega_{\text{Spin}(1,5)})$
Visible	$\text{SU}(5)$	A_5	F_5	$-\text{CS}(A_5)$
Baryon	$\text{U}(1)_B$	A_B	F_B	$-\frac{2}{5}\text{CS}(A_B)$
Qi	$\text{U}(1)_{Qi}$	A_{Qi}	F_{Qi}	$-\frac{2}{5}\text{CS}(A_{Qi})$

The four rows are not independent: they are four projections of a single underlying object (GGEP, Structure Theorem **TGG-ST5**). The signs in the third column come from the Killing-form signs of the four Lie algebras; the $-2/5$ coefficients of the last two rows come from the inner-product normalizations on $\mathbb{H}^2$.

Discussion. The four Wilson lines are the boundary face of TGG’s holonomy structure. Three of them (W_5, W_B, W_{Qi}) are non-trivial and encode the dynamical content of the boundary theory; the fourth ($W_{\text{Spin}(1,5)}$) is trivial because the product metric has no θ -component in its spin connection. This asymmetry is essential: it is what allows the gauge sector to break (orbifold, Wilson–Hosotani, Kaxion) while keeping the gravity sector unbroken, which is what we observe in nature.

12 One object, and what it means that it is unique

The result of this Part can be said in one sentence: the four forces can be written as one geometric object on the six-dimensional arena, and there is no second way to do it.

Both halves of that sentence matter, and the second is the one worth dwelling on. Physics has produced unifications before that were merely possible — schemes in which the known forces could be arranged inside a larger structure, and could equally well have been arranged inside several others. A possible arrangement explains nothing, because the question "why this one?" has no answer, and the honest form of such a proposal is a menu. What changes here is that the constitutional requirements inherited in Part I leave exactly one arrangement standing. The four forces are not unified because it is elegant to unify them. They are unified because, given the arena and the rules, nothing else fits.

The Master Conservation Law that follows is worth reading in plain language, because its content is more ordinary than its name. It says that the energy and momentum of everything in the observable four-dimensional world — matter, the gauge fields, the leftover higher-dimensional modes, and the correction term that consistency demands — add up to something that does not leak. Nothing drains out of the world through a channel the accounting missed. That sounds like a statement one could take for granted, and in ordinary physics one can. In a theory where the observable world is a projection of a larger one, it is not automatic at all: a projection is exactly the kind of operation that loses things. The theorem says this particular projection does not.

There is a habit of thought this Part is quietly arguing against. It is natural to imagine gravity as the arena and the other three forces as the furniture — one thing that curves, three things that sit on it. That picture is useful and it is not what the mathematics says here. There is one structure; gravity is what it looks like along certain directions and the gauge forces are what it looks like along others. The distinction between arena and content, which feels like the most basic distinction one could draw about the physical world, turns out to be a statement about which way an observer happened to look.

That is not a rhetorical flourish; it has a concrete consequence for how the rest of the paper reads. When a later Part speaks of energy flowing between the gravitational sector and the gauge sector, it is not describing a transaction between two things. It is describing one object redistributing itself, in the way that a person who shifts their weight has not handed anything to anybody.

The deeper lesson is about how much a consistency requirement can do. Nothing in this Part fits a parameter to data. The uniqueness comes out of demanding that a structure be well defined at all — and that demand, made precisely enough, turns out to be nearly as informative as an experiment. The next Part takes that observation much further: it asks what a single consistency condition forces when it is imposed not on the shape of the object but on its bookkeeping, and the answer is a set of integers that nobody chose.

Part III CCI

The Covariant Conservation Identity: Pillar 2 of Gi-4

Every conservation law in physics is, at bottom, a piece of bookkeeping: something goes in, something comes out, and the ledger balances. The awkward discovery of the twentieth century was that in quantum theory the ledger does not always balance by itself. Certain combinations of fields produce what are called anomalies — terms that appear from nowhere and break a conservation law that the classical theory obeyed exactly. A theory in which that happens is not merely inelegant; it is inconsistent, and it must be discarded.

This Part is about the ledger. It asks what has to be true of the six-dimensional geometry for the four-dimensional conservation law of Part II to hold, and it finds that the requirement is far more restrictive than one would guess. The anomalies do not simply have to cancel; the way they cancel forces the numbers in the theory to take particular values. When the closure condition is imposed, a single set of five integers survives, and everything the cohort later computes — the number of particle families, the parameters of the Standard Model, the characteristic mass scales — turns out to be a closed-form function of those five. That is an unusual position for a physical theory to be in. It is what the phrase “no free parameters” means in practice, and it is also what makes the theory easy to kill.

Two further results are established here. The first is that the cancellation is not free: it requires a compensating structure on the boundary, and the Part identifies what that structure must be and what it costs. The second is the Part’s empirical exposure. The same ledger that fixes the integers also modifies the equation governing the internal structure of a dense star, by an amount the theory computes rather than fits. That prediction is a falsifier: a sufficiently precise measurement of a neutron star’s structure either finds the modification or ends this pillar.

The mathematics used — differential cohomology, and the descent construction that goes with it — is the one genuinely specialized tool in the paper, and it is not assumed. It is developed from the beginning in Part X, with worked examples, for a reader who has not met it. Within this Part, a reader can follow the argument by reading the opening section and then the statement of the ledger result, treating the cohomological machinery as a black box that delivers the integers; the box is opened in the appendix for anyone who wants to look inside.

13 CCI: motivation, dual-layer structure, and Structure Theorems

This Part develops CCI (the *Covariant Conservation Identity*), the second of the four Gi-4 Pillars. Like TGG of Part II, CCI is intrinsically a *dual-layer* object: it is simultaneously

- (i) a **boundary/physical statement** — the local 4D conservation equation

$$\nabla_\mu T_{\text{TGG}}^{\mu\nu} = 0 \quad (62)$$

on the boundary $\partial M_6 \cong M_4$, asserting that the Master Stress Tensor of TGG is covariantly conserved when (and only when) the bulk BI-6 closure $dH_3 = X_4$ is satisfied *and* the four boundary Ward currents (electric, baryonic, informational, metric) are simultaneously closed; and

- (ii) an **algebraic/bulk statement** — the cohomological identity

$$[X_4] \in H^4(M_6; \mathbb{Z}), \quad X_4 = \text{tr } R_{\text{Spin}(1,5)}^2 - \text{Tr } F_5^2 - \frac{2}{5} F_B^2 - \frac{2}{5} F_{Q_i}^2, \quad (63)$$

asserting that the four-form anomaly polynomial X_4 has a unique algebraic structure (with the four specific coefficients $+1, -1, -\frac{2}{5}, -\frac{2}{5}$) and that its de Rham class is integral; MSMF then selects the minimal-complexity integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$, the *OGN-5* ledger.

The two layers are not independent: they are linked by the *Projected Bianchi Identity* PBI-4 (introduced in Section 9 of Part II), which descends the bulk closure $dH_3 = X_4$ to a boundary divergence identity, and by the Master Conservation Theorem (TGG-ST7). In short:

$$\boxed{\begin{array}{ccc} \underbrace{dH_3 = X_4}_{\text{bulk BI-6 (algebraic)}} & \xrightarrow{\text{KK descent}} & \underbrace{\nabla_\mu T_{\text{TGG}}^{\mu\nu} = 0}_{\text{boundary CCI (physical)}} \end{array}} \quad (64)$$

The arrow “ $\xrightarrow{\text{KK descent}}$ ” is precisely PBI-4: it carries the algebraic identity to its physical avatar.

Section structure of Part III. Section 13 (the present section) states the seven CCI axioms (**CCI-ST1** through **CCI-ST7**), organized into the algebraic block (**CCI-ST1–CCI-ST5**) and the physical block (**CCI-ST6–CCI-ST7**); gives the dual-layer formal definition of CCI; and lists the first immediate consequences. Section 14 develops the four-form anomaly polynomial X_4 in detail, deriving its closure, integrality, and non-triviality from the TGG structure of Part II. Section 15 derives the four coefficients in X_4 — in particular the $-2/5$ coefficients of the two U(1) terms — from the inner-product normalization on $\mathbb{H}^2$. Section 16 contrasts CCI-algebraic with BI-6 (differential) and shows how the two identities couple via PBI-4. Section 17 (new in) develops the physical/boundary layer of CCI: the local 4D conservation equation, its derivation from PBI-4, and its astrophysical falsifier (the *modified TOV equation* for compact stars). Section 18 states and proves Theorem **T2**, now in two parts: **T2.1** (the algebraic OGN-5 uniqueness theorem) and **T2.2** (the boundary conservation theorem). Section 19 connects CCI to the Hopkins–Singer differential cohomology framework that will be used in Part IV for GFF Morse-positivity. Section 20 contains worked examples.

Remark 13.1 (Physics intuition: CCI as the GG analog of GR’s contracted Bianchi identity). Readers familiar with general relativity will find an immediate analogy. In GR, the contracted Bianchi identity

$\nabla_\mu G^{\mu\nu} = 0$ (where $G^{\mu\nu}$ is the Einstein tensor) is a *kinematic* identity that follows automatically from the diffeomorphism invariance of the Einstein–Hilbert action. It implies the local conservation of the matter stress tensor $\nabla_\mu T_{\text{matter}}^{\mu\nu} = 0$ *provided* the field equations $G^{\mu\nu} = 8\pi G T_{\text{matter}}^{\mu\nu}$ hold. This is the *single* conservation law of pure GR.

In GG-Theory, CCI plays the same logical role but for the *six-sector* Master Stress Tensor of TGG (TGG-ST6): $T_{\text{TGG}}^{\mu\nu} = T_{\text{matter}} + T_{\text{gauge}} + X_{\text{GS}} + T_{\text{KK}}$. The identity $\nabla_\mu T_{\text{TGG}}^{\mu\nu} = 0$ is again kinematic: it follows from the bulk diffeomorphism invariance of the GG-6 effective action plus the closure of the four Ward currents on the boundary. But unlike GR, GG-Theory has *multiple* simultaneous conservation laws (one per sector), and CCI packages them into a single covariant master identity. This is the precise sense in which CCI *generalizes* GR’s contracted Bianchi identity from one sector (gravity) to six (gravity, electroweak, strong, baryonic, informational, KK).

Remark 13.2 (Why broader scientists should care about CCI). The Master Conservation Law $\nabla_\mu T_{\text{TGG}}^{\mu\nu} = 0$ is the *only* statement of GG-Theory that every empirical scientist — not just particle physicists — already uses every day. Energy conservation in chemistry, momentum conservation in fluid dynamics, charge conservation in electrochemistry, baryon-number near-conservation in cosmology, mass–energy balance in cell biology, information conservation (entropy bounds) in neuroscience: all are *boundary shadows* of CCI. What is new is that GG-Theory shows these conservation laws are not independent postulates but follow from a single bulk identity $dH_3 = X_4$ via the Projected Bianchi Identity PBI-4. This is why CCI is the *universal empirical bridge* from GG-6 bulk geometry to the laboratory: it explains, in one stroke, why all of physics, chemistry, biology, and information science share the same kinematic backbone.

13.1 Motivation: why a dual-layer pillar is needed

The TGG construction of Part II provides the principal G_{tot} -bundle P , the composite connection $\mathcal{A}_{\text{tot}}$, the composite curvature $\mathcal{R}_{\text{tot}}$ on the 6D bulk M_6 , the boundary Master Stress Tensor $T_{\text{TGG}}^{\mu\nu}$ (TGG-ST6), and the Master Conservation Law $\nabla_\mu T_{\text{TGG}}^{\mu\nu} = 0$ (TGG-ST7). But TGG alone does not fix two crucial pieces of structure:

Question 1 (algebraic). TGG does not fix *which* four-form built from $\mathcal{R}_{\text{tot}}$ should appear as the source of the modified Bianchi identity BI-6. Several different four-forms can be constructed from the four sub-curvatures $(R_{\text{Spin}(1,5)}, F_5, F_B, F_{Qi})$:

- $\text{tr } R_{\text{Spin}(1,5)}^2, \text{Tr } F_5^2, F_B \wedge F_B, F_{Qi} \wedge F_{Qi}$ (the four “self-curvature” four-forms);
- Cross-curvature four-forms $F_B \wedge F_{Qi}, \text{tr } R_{\text{Spin}(1,5)} \wedge F_B$, etc.;
- Higher-degree-sensitive combinations like $\text{tr } R_{\text{Spin}(1,5)} \wedge \text{tr } R_{\text{Spin}(1,5)}$, etc.

Without an additional principle, all of these four-forms are mathematically admissible candidates for X_4 . But BI-6 requires a specific combination with specific coefficients — namely

$$X_4 = \text{tr } R_{\text{Spin}(1,5)}^2 - \text{Tr } F_5^2 - \frac{2}{5} F_B^2 - \frac{2}{5} F_{Qi}^2, \quad (65)$$

to within an overall normalization. The structure (65) is not free: it is forced by the inner-product structure of TGG on the quaternionic doublet $\mathbb{H}^2$, by the requirement of integrality of $[X_4]$, and by

the parity / sign conventions of the four sub-algebra Killing forms. This is the **algebraic content** of CCI.

Question 2 (physical). TGG on its own asserts that *some* master stress tensor exists on the boundary (TGG-ST6) and that it is *conditionally* conserved (TGG-ST7) — conserved iff the four Ward currents are closed. But TGG-ST7 does not specify the precise differential equation that governs each sector’s contribution to $\nabla_\mu T_{\text{TGG}}^{\mu\nu}$, nor does it produce a single *unconditional* astrophysical or cosmological prediction. The physical content of CCI fills both gaps: it states the local 4D conservation equation

$$\nabla_\mu T_{\text{TGG}}^{\mu\nu}(x) = 0 \quad \forall x \in M_4 \quad (66)$$

in *operational* form, identifies the PBI-4 compensator field that enforces it, and predicts a specific *astrophysical falsifier*: the modified Tolman–Oppenheimer–Volkoff (TOV) equation for compact stars, with a calculable correction proportional to X_{GS} . This is the **physical content** of CCI.

Why both layers must be in the same Pillar. The two questions are inseparable. The integrality of $[X_4]$ (the algebraic side) is precisely the condition for PBI-4 to descend without anomaly to a smooth 4D conservation equation (the physical side). Conversely, the empirical observation that $\nabla_\mu T_{\text{matter}}^{\mu\nu} = 0$ holds in laboratory and astrophysical data (the physical side) provides direct constraints on the bulk integers $(M, N, L; J, K)$ (the algebraic side), via cosmological and neutron-star observations. This bidirectional coupling *mandates* that both layers be axiomatized together.

CCI is the formalization of this dual content. It is simultaneously:

- the *algebraic consistency identity* of the composite connection, distinguishing the unique physical X_4 from the infinite family of mathematically admissible four-forms; and
- the *boundary conservation identity*, asserting $\nabla_\mu T_{\text{TGG}}^{\mu\nu} = 0$ as the universal empirical content of the bulk closure $dH_3 = X_4$.

Once CCI is given, BI-6 is the *differential consequence* of CCI together with the de Rham exactness of X_4 , and the boundary conservation equation is its *empirical projection* through PBI-4.

In short: TGG gives the bulk geometry and the boundary master stress tensor; CCI gives both the algebraic form of the closure condition (uniqueness of X_4) and the physical conservation law on the boundary (operational content of $\nabla_\mu T_{\text{TGG}}^{\mu\nu} = 0$); BI-6 gives the differential content of the closure condition in the bulk; and PBI-4 ties the three together. All four structures are needed.

Remark 13.3 (Physics intuition: the $-2/5$ coefficient and the OGN-5 ledger). The reader may wonder why the coefficient of F_B^2 and $F_{Q_i}^2$ in X_4 is exactly $-2/5$. The intuitive answer is dimensional: the quaternionic doublet $\mathbb{H}^2$ has real dimension 8, but its inner-product normalization on the two $\text{U}(1)$ Cartan directions factors out a rank-5 index from the magic-dimension ladder $\mathbb{R} \rightarrow \mathbb{C} \rightarrow \mathbb{H} \rightarrow \mathbb{O}$. The factor $2/5$ then arises as the ratio of the doublet inner product (which contributes 2) to the OGN-5 base length (which contributes 5). Section 15 makes this precise.

The key takeaway is that the $-2/5$ is not a fitting parameter — it is the *unique* coefficient that makes the integrality $[X_4] \in H^4(M_6; \mathbb{Z})$ compatible with the OGN-5 minimal ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$. Any other coefficient would either break integrality or force a non-minimal (and hence MSMF-disfavored) ledger. This is why the OGN-5 ledger is the unique solution.

13.1.1 Historical background (with citations)

The covariant conservation identity that CCI formalizes traces back to the very first geometric identities of modern physics. [Bianchi \(1902\)](#) proved the celebrated identity $\nabla_{[\mu} R_{\nu\rho]\sigma\tau} = 0$ for the Riemann tensor in 1902 — a purely geometric statement, derived without any reference to dynamics or quantum effects, that fixes once and for all the local conservation of the geometric curvature. The electromagnetic analog $dF = 0$, the homogeneous half of Maxwell’s equations ([Maxwell , 1865](#)), is the abelian version of the same statement.

The dynamical face of CCI is due to [Noether \(1918\)](#): every continuous symmetry of the action yields a conserved current, and the divergence-free property $\partial_\mu J^\mu = 0$ is the identity that ties symmetry to conservation. [Ward \(1950\)](#) re-expressed Noether’s identity in the language of quantum field theory, with the Ward identity stating that the gauge-transformation variation of the effective action vanishes on shell. The non-abelian extension was completed in the 1970s by Slavnov, Taylor, Becchi, Rouet, and Stora, and independently by Tyutin, with the BRST formulation as the cohomological setting in which CCI lives naturally.

Indeed, the quantum surprise that forced the modern reading of CCI was the chiral anomaly of [Adler \(1969\)](#) and [Bell & Jackiw \(1969\)](#): the axial current of a chiral fermion is not conserved at the loop level, but acquires a divergence proportional to $F \wedge \tilde{F}$. This loss is reconciled with the geometric Bianchi identity through the geometrization move pioneered by [Green & Schwarz \(1984\)](#) (whose GGEP-completion reading we discussed in Section 3): the bulk 4-form X_4 absorbs the boundary anomaly into a closed differential object whose cohomology class $[X_4] \in H^4(M_6; \mathbb{Z})$ is integral ([Hopkins & Singer, 2005](#)). CCI is the axiomatic reading of the integrality requirement: BI-6 closure under MSMF forces the unique OGN-5 ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$, and the four-component coefficient vector $(a, b, c, d) = (1, -1, -2/5, -2/5)$ becomes the unique solution ([Arisaka, 2026, A3-BI6](#), Theorem 9.1).

13.2 The seven CCI Structure Theorems (algebraic block + physical block)

We now formalize the dual-layer structure of CCI through seven Structure Theorems — five governing the algebraic content of the 4-form anomaly polynomial and two governing its physical projection to boundary observables. The seven CCI Structure Theorems partition naturally into two blocks:

Algebraic block: CCI-ST1 – CCI-ST5 These Structure Theorems fix the structural form of the four-form anomaly polynomial X_4 , its coefficients, its closure, its integrality, and the OGN-5 minimality condition. They live entirely in the bulk M_6 and are representations-theoretic in character.

Physical block: CCI-ST6 – CCI-ST7 These Structure Theorems state the boundary covariant conservation equation $\nabla_\mu T_{\text{TGG}}^{\mu\nu} = 0$ and its astrophysical falsifier (the modified TOV equation). They live on the boundary $\partial M_6 \cong M_4$ and are operationally testable.

The two blocks are linked by the Projected Bianchi Identity PBI-4 of Section 9.

Algebraic block (CCI-ST1 through CCI-ST5)

Structure Theorem 8 (CCI-ST1: Structural form of X_4). *The four-form anomaly polynomial $X_4 \in \Omega^4(M_6)$ is built from the TGG curvature $\mathcal{R}_{\text{tot}}$ as a sum of four self-curvature contributions:*

$$X_4 = \alpha_R \text{tr}(R_{\text{Spin}(1,5)} \wedge R_{\text{Spin}(1,5)}) + \alpha_5 \text{Tr}(F_5 \wedge F_5) + \alpha_B F_B \wedge F_B + \alpha_{Q_i} F_{Q_i} \wedge F_{Q_i}, \quad (67)$$

where $\alpha_R, \alpha_5, \alpha_B, \alpha_{Q_i}$ are real coefficients. No cross-curvature four-forms appear in X_4 .

Structure Theorem 9 (CCI-ST2: Coefficient determination from $\mathbb{H}^2$ normalization). *The four coefficients in (67) are determined by the inner-product normalization on the quaternionic doublet $\mathbb{H}^2$ of TGG-ST3:*

$$\alpha_R = +1, \quad \alpha_5 = -1, \quad \alpha_B = -\frac{2}{5}, \quad \alpha_{Q_i} = -\frac{2}{5}. \quad (68)$$

The signs and magnitudes are derived in Section 15; they are not free parameters but consequences of the $\text{Spin}(1, 5) \cong \text{SL}(2, \mathbb{H})$ algebraic structure.

Structure Theorem 10 (CCI-ST3: Closure of X_4). *The four-form X_4 is closed:*

$$dX_4 = 0. \quad (69)$$

Structure Theorem 11 (CCI-ST4: Integrality). *The de Rham cohomology class of X_4 is integral:*

$$[X_4] \in H^4(M_6; \mathbb{Z}), \quad (70)$$

i.e., X_4 represents an integer cohomology class on M_6 .

Structure Theorem 12 (CCI-ST5: Non-triviality and OGN-5 minimality). *The class $[X_4]$ is non-zero in $H^4(M_6; \mathbb{R})$, and among all admissible TGG configurations satisfying **CCI-ST1–CCI-ST4**, **MSMF** selects the configuration whose integrality data $(M, N, L; J, K)$ has minimum total complexity $|M| + |N| + |L| + |J| + |K|$ subject to the closure constraints*

$$K = M^2 + N^2 + L^2, \quad K = LJ. \quad (71)$$

The five Structure Theorems **CCI-ST1–CCI-ST5** constitute the *algebraic* block of CCI and uniquely determine the OGN-5 ledger $(1, 3, 5; 7, 35)$ via Theorem **T2.1** (proved in Section 18).

Physical block (CCI-ST6 and CCI-ST7)

The physical block translates the algebraic structure of CCI into operational, observer-accessible form on the boundary M_4 . It consists of two Structure Theorems.

Structure Theorem 13 (CCI-ST6: Boundary covariant conservation (sector-summed)). *On the boundary $\partial M_6 \cong M_4$, the Master Stress Tensor $T_{\text{TGG}}^{\mu\nu}$ of TGG-ST6 satisfies the sector-summed covariant conservation equation*

$$\nabla_\mu T_{\text{TGG}}^{\mu\nu}(x) = 0 \quad \forall x \in M_4, \quad (72)$$

where ∇ is the Levi–Civita connection of the induced boundary metric $g_{\mu\nu}^{(4)}$. Decomposing $T_{\text{TGG}}^{\mu\nu} = T_{\text{matter}} + T_{\text{gauge}} + X_{\text{GS}} + T_{\text{KK}}$ (TGG-ST6), the individual sector divergences are non-zero — computed explicitly in Proposition 17.1 of Section 17.1 — and add to zero:

$$\nabla_\mu T_{\text{matter}}^{\mu\nu} + \nabla_\mu T_{\text{gauge}}^{\mu\nu} + \nabla_\mu X_{\text{GS}}^{\mu\nu} + \nabla_\mu T_{\text{KK}}^{\mu\nu} = 0. \quad (73)$$

The four sector divergences are linked — not zero individually but summing to zero — by the PBI-4 compensator field C_μ (defined in Section 17); the sector-summed identity (72) follows from the bulk BI-6 closure $\text{d}H_3 = X_4$ combined with the integrality of $[X_4]$ (CCI-ST4), which guarantees $\langle X_4 \rangle_{\text{KK}} = 0$ at the boundary. In particular, CCI-ST6 is a sector-summed kinematic identity rather than a sector-wise conservation law: each Ward current carries its own anomaly inflow (Alvarez-Gaumé & Witten, 1984, Bardeen, 1969, Callan & Harvey, 1985) from the corresponding 4-form term in X_4 , and the anomaly inflows cancel collectively (by closure of X_4) rather than individually.

Structure Theorem 14 (CCI-ST7: Astrophysical falsifier (modified TOV) — primary derivation in A5 T8.6). *On a static, spherically symmetric, gravitationally bound configuration (e.g. a neutron star, white dwarf, or planetary interior), the boundary conservation equation (72) reduces to the modified Tolman–Oppenheimer–Volkoff (TOV) equation (primary derivation: Arizaka, 2026, A5-GDOS, Theorem T8.6, Part IV; the present structure theorem records the CCI-side identity)*

$$\boxed{\frac{\text{d}p}{\text{d}r} = - \frac{G(\rho + p)(m + 4\pi r^3 p)}{r(r - 2Gm)} - \Delta_{\text{GS}}(r),} \quad (74)$$

where $\Delta_{\text{GS}}(r)$ is a calculable, non-vanishing correction arising from the Green–Schwarz contribution X_{GS} in $T_{\text{TGG}}^{\mu\nu}$. The functional form of $\Delta_{\text{GS}}(r)$ is fixed by PBI-4 and is given (to leading order in the scale-profile expansion) by

$$\Delta_{\text{GS}}(r) = \kappa_{\text{GS}} \frac{\langle X_4 \rangle_{\text{bulk}}}{M_{\text{KK}}^4} \frac{1}{r^2} \frac{\text{d}}{\text{d}r} \left(r^2 \rho_{\text{KK}}(r) \right), \quad (75)$$

with $\kappa_{\text{GS}} \sim \mathcal{O}(1)$ a coefficient computable from the OGN-5 ledger. Empirically, equation (74) predicts a small but in-principle observable shift of the neutron-star mass–radius curve relative to the standard TOV prediction; this shift is the principal astrophysical falsifier of CCI.

If empirical data on neutron-star masses and radii (from NICER, gravitational-wave inspiral, etc.) ever decisively rule out (74) for any $\kappa_{\text{GS}} \neq 0$, then CCI (and with it the entire GG-6 framework) is falsified.

The seven Structure Theorems **CCI-ST1–CCI-ST7** together constitute the full *Covariant Conservation Identity*. They will be used in Section 18 to prove Theorem **T2**, now in two parts: **T2.1** (algebraic OGN-5 uniqueness) and **T2.2** (boundary conservation iff Ward + PBI-4).

13.3 Dual-layer formal definition of CCI

We now consolidate the algebraic and physical blocks into a single dual-layer definition.

Definition 13.4 (CCI: the Covariant Conservation Identity, dual-layer form). The *Covariant Conservation Identity* on the TGG structure $(P, \mathcal{A}_{\text{tot}}, \mathcal{R}_{\text{tot}}; T_{\text{TGG}}^{\mu\nu}, \nabla)$ over the 6D bulk M_6 with boundary $\partial M_6 \cong M_4$ is the **septuple**

$$\text{CCI} = \left(\underbrace{X_4, [X_4], (\alpha_R, \alpha_5, \alpha_B, \alpha_{Qi}), dX_4 = 0, \text{OGN-5 minimality}}_{\text{algebraic block (CCI-ST1–A5)}}, \underbrace{\nabla_\mu T_{\text{TGG}}^{\mu\nu} = 0, \Delta_{\text{GS}}(r)}_{\text{physical block (CCI-ST6–A7)}} \right), \quad (76)$$

where the seven components correspond to the Structure Theorems **CCI-ST1–CCI-ST7** respectively.

The triple $(X_4, [X_4], \text{coefficients})$ is required to satisfy the integrality and non-triviality conditions of **CCI-ST4** and **CCI-ST5**. The boundary conservation equation (72) is required to follow from the bulk closure $dH_3 = X_4$ via PBI-4. The TOV correction $\Delta_{\text{GS}}(r)$ is required to be the unique PBI-4-compatible spherically-symmetric reduction. Two CCI configurations are equivalent if they are related by gauge transformations of the underlying TGG structure (by Proposition 6.6) together with passive boundary diffeomorphisms.

Remark 13.5 (Logical content of CCI). The intuitive content of Definition 13.4 is twofold:

- **Algebraic side.** CCI fixes the four-form X_4 uniquely (up to overall normalization and gauge equivalence) and asserts that its cohomology class is integral and non-trivial. The integrality constraint is what selects the OGN-5 ledger from the infinite family of admissible $(M, N, L; J, K)$ tuples. The non-triviality constraint is what prevents the boundary observables from collapsing into a degenerate zero-curvature configuration; it is the algebraic prerequisite for the GFF Morse-positivity of Part IV.
- **Physical side.** CCI packages all six sector conservation laws into the single Master Conservation Law $\nabla_\mu T_{\text{TGG}}^{\mu\nu} = 0$. This is the universal empirical content of GG-6 that observers actually measure: every conservation law of physics, chemistry, biology, and information theory is a boundary projection of this identity. The astrophysical falsifier (74) provides a sharp empirical test.

The two sides are linked by PBI-4: the algebraic integrality $[X_4] \in H^4(M_6; \mathbb{Z})$ is precisely the necessary and sufficient condition for PBI-4 to descend to a smooth, anomaly-free 4D conservation equation. This is the *duality* of CCI.

Remark 13.6 (Physics intuition: from $dH_3 = X_4$ to laboratory measurements). A graduate student in physics may find it useful to trace the chain

$$\begin{array}{ccc} \underbrace{dH_3 = X_4}_{\text{bulk BI-6}} & \xrightarrow{[X_4] \in H^4(M_6; \mathbb{Z})} \underbrace{\text{PBI-4}}_{\text{descent}} & \xrightarrow{\text{KK reduction}} \underbrace{\nabla_\mu T_{\text{TGG}}^{\mu\nu} = 0}_{\text{boundary 4D conservation}} \\ & \searrow \text{spherical symmetry} & \\ & \underbrace{dp/dr = - \dots - \Delta_{\text{GS}}(r)}_{\text{modified TOV}} & \end{array}$$

Each arrow is a precise mathematical reduction — there are no phenomenological insertions. The first arrow is purely topological (integrality of de Rham cohomology); the second is purely algebraic (KK harmonic decomposition on the extra dimensions); the third is purely geometric (Schwarzschild

ansatz applied to a static spherically symmetric configuration).

What is genuinely new is that the modified TOV correction $\Delta_{\text{GS}}(r)$ is *not* an ad-hoc parameter. It is calculable from the OGN-5 ledger (1, 3, 5; 7, 35) once the bulk scale running $W(y)$ of TGG-ST2 is specified. This is the empirical backbone of the GG-6 prediction for compact-object structure.

13.4 First consequences of the CCI Structure Theorems

We list four immediate consequences that will be used in the rest of the Part.

Proposition 13.7 (Closure of X_4 from sub-curvature Bianchi identities). ***CCI-ST3** (closure $dX_4 = 0$) follows automatically from the ordinary Bianchi identities for the four sub-curvatures.*

Proof. For each sub-curvature, the Bianchi identity gives $d_{\text{cov}}R_{\text{Spin}(1,5)} = 0$, $d_{\text{cov}}F_5 = 0$, $dF_B = 0$ (by exactness of the Abelian field strength), $dF_{Qi} = 0$. Therefore

$$d \operatorname{tr}(R_{\text{Spin}(1,5)} \wedge R_{\text{Spin}(1,5)}) = 2 \operatorname{tr}(R_{\text{Spin}(1,5)} \wedge dR_{\text{Spin}(1,5)}) = 0, \quad (77)$$

$$d \operatorname{Tr}(F_5 \wedge F_5) = 2 \operatorname{Tr}(F_5 \wedge dF_5) = 0, \quad (78)$$

$$d(F_B \wedge F_B) = 2F_B \wedge dF_B = 0, \quad (79)$$

$$d(F_{Qi} \wedge F_{Qi}) = 2F_{Qi} \wedge dF_{Qi} = 0. \quad (80)$$

Adding the four contributions with the coefficients of **CCI-ST2** gives $dX_4 = 0$. ■

Proposition 13.8 (X_4 as a topological invariant of TGG). *The cohomology class $[X_4] \in H^4(M_6; \mathbb{R})$ is a gauge-invariant topological invariant of the TGG structure. In particular, $[X_4]$ does not depend on the specific choice of TGG representative within its gauge orbit (Proposition 6.6), and it is unchanged under continuous deformations of $\mathcal{A}_{\text{tot}}$ that preserve the underlying G_{tot} -bundle topology.*

Proof. By Proposition 13.7, X_4 is closed. Under a gauge transformation $\mathcal{A}_{\text{tot}} \rightarrow g^{-1}\mathcal{A}_{\text{tot}}g + g^{-1}dg$, each self-curvature transforms by conjugation: $R \rightarrow g^{-1}Rg$, etc. But tr and Tr are gauge-invariant traces in their respective representations, so $\operatorname{tr}(R_{\text{Spin}(1,5)} \wedge R_{\text{Spin}(1,5)})$, $\operatorname{Tr}(F_5 \wedge F_5)$, F_B^2 , and F_{Qi}^2 are all gauge-invariant 4-forms. Therefore X_4 is gauge-invariant. Continuous deformations within a gauge orbit change X_4 by an exact form (Kobayashi & Nomizu, 1969, Theorem II.5.3), leaving $[X_4]$ unchanged. ■

Proposition 13.9 (Integrality of $[X_4]$ from sub-bundle integrality). ***CCI-ST4** ($[X_4] \in H^4(M_6; \mathbb{Z})$) follows from the integrality of each sub-bundle's first Chern class (for the Abelian factors) or first Pontryagin class (for the non-Abelian factors), with the coefficients of **CCI-ST2** chosen consistently.*

Sketch. For the gauge factors:

- $\frac{1}{8\pi^2} \operatorname{tr}(R_{\text{Spin}(1,5)} \wedge R_{\text{Spin}(1,5)})$ represents the first Pontryagin class p_1 of the $\text{Spin}(1,5)$ bundle, integral in $H^4(M_6; \mathbb{Z})$.
- $\frac{1}{8\pi^2} \operatorname{Tr}(F_5 \wedge F_5)$ represents the second Chern class c_2 of the $\text{SU}(5)$ bundle, integral.
- $\frac{1}{4\pi^2} F_B \wedge F_B$ and $\frac{1}{4\pi^2} F_{Qi} \wedge F_{Qi}$ represent the squares of the first Chern classes of $\text{U}(1)_B$ and $\text{U}(1)_{Qi}$, both integer cohomology classes.

With the coefficients $\alpha_R = +1, \alpha_5 = -1, \alpha_B = \alpha_{Qi} = -2/5$ of **CCI-ST2**, the rational linear combination $X_4/(8\pi^2)$ has rational coefficients of denominator 5. The integrality of $[X_4]$ then requires the integers $L = 5$ and $K = LJ = 35$ that combine to clear the $1/5$ denominators. This is the algebraic origin of the OGN-5 integers. See Arisaka (2026, A3-BI6, Section 3) for the rigorous statement. ■

Proposition 13.10 (Non-triviality on M_6). *For the 6D bulk of GG-Theory with the TGG structure of Part II, the cohomology class $[X_4]$ is non-zero in $H^4(M_6; \mathbb{R})$.*

Sketch. The compact S_θ^1 direction of M_6 supports a non-trivial generator of $H^1(M_6; \mathbb{Z})$. The dual generator of $H^5(M_6; \mathbb{Z})$ gives, by Poincaré duality, a non-trivial 4-cycle on which X_4 has non-zero integral. Specifically, the integral $\int_{C_4} X_4 = K \neq 0$ where C_4 is the 4-cycle dual to S_θ^1 and $K = 35$ is the OGN-5 integer. This non-vanishing will be used in Part IV to prove GFF Morse-positivity via the Hopkins–Singer descent. ■




Figure 6: Anomaly inflow and the Ward-identity ladder of CCI (Pillar 2 of Gi-4). The closed bulk 4-form X_4 on M_6 encodes the algebraic content of CCI (T2.1: the OGN-5 ledger forces the unique trace-ratio lock $|a/c| = |a/d| = 5/2$). Hopkins–Singer descent via $dH_3 = X_4$ produces the boundary PBI-4 compensator field C_μ , which links the four boundary Ward currents J_e, J_B, J_{Qi}, J_g into a simultaneous closure pattern. The Master Conservation Law T2.2 holds iff all four currents close modulo X_4 -inflow and the compensator field exists.

14 The four-form anomaly polynomial X_4 : structure and properties

Remark 14.1 (Reader’s bridge: what an anomaly polynomial is, in everyday terms). Before plunging into the explicit formulae, let us pause and ask what kind of object X_4 is. Roughly speaking, an anomaly polynomial is a bookkeeping device. In any theory with several interacting sectors — gravity, the visible gauge sector, a baryonic sector, an information sector — each sector carries a current that one would naively expect to be separately conserved. Quantum mechanics tells us that conservation cannot always hold for each sector in isolation: tiny quantum effects called *anomalies* cause the divergences to be non-zero. What X_4 does is collect all four sectoral anomaly contributions

into a single four-form on the bulk, with calibrated coefficients, so that the four imbalances must sum to zero. The arithmetic is forced: if the sum is non-zero, the theory is inconsistent. Indeed, this is why **CCI** is the second pillar of **Gi-4**: it is the ledger that says “the bulk has tallied every sectoral contribution, and the books balance.” The rest of this section writes out the entries of that ledger explicitly.

We now develop the four-form X_4 of **CCI-ST1** in greater detail, working out its explicit form on the 6D bulk and analyzing its integrality.

14.1 Explicit form of X_4 on the 6D bulk

On the 6D bulk with the **TGG** structure of Part II, the four sub-curvatures of $\mathcal{R}_{\text{tot}}$ have the explicit forms (in local coordinates):

$$R_{\text{Spin}(1,5)}^{ab} = d\omega^{ab} + \omega^a{}_c \wedge \omega^{cb}, \quad (81)$$

$$F_5^I = dA_5^I + \frac{1}{2} f^I{}_{JK} A_5^J \wedge A_5^K, \quad (82)$$

$$F_B = dA_B, \quad (83)$$

$$F_{Qi} = dA_{Qi}, \quad (84)$$

with ω^{ab} the spin connection one-forms, A_5^I ($I = 1, \dots, 24$) the $\text{SU}(5)$ gauge field components, and A_B, A_{Qi} the $\text{U}(1)_B, \text{U}(1)_{Qi}$ gauge fields.

The four self-curvature 4-forms are then:

$$\text{tr}(R_{\text{Spin}(1,5)} \wedge R_{\text{Spin}(1,5)}) = R_{\text{Spin}(1,5)}^{ab} \wedge R_{\text{Spin}(1,5)ab}, \quad (85)$$

$$\text{Tr}(F_5 \wedge F_5) = F_5^I \wedge F_5^I, \quad (86)$$

$$F_B \wedge F_B, \quad (87)$$

$$F_{Qi} \wedge F_{Qi}. \quad (88)$$

Combining with the coefficients of **CCI-ST2**:

$$X_4 = R_{\text{Spin}(1,5)}^{ab} \wedge R_{\text{Spin}(1,5)ab} - F_5^I \wedge F_5^I - \frac{2}{5} F_B \wedge F_B - \frac{2}{5} F_{Qi} \wedge F_{Qi}. \quad (89)$$

This is the four-form whose differential gives **Bl-6**.

14.2 Cohomological content of $[X_4]$

The cohomology class $[X_4] \in H^4(M_6; \mathbb{R})$ has a clean characterization in terms of standard topological invariants of the four sub-bundles.

Theorem 14.2 (Cohomological content of $[X_4]$). *The cohomology class $[X_4]$ on M_6 is, up to overall rescaling,*

$$[X_4] = 8\pi^2 \left(p_1(\text{Spin}(1,5)) - 2c_2(\text{SU}(5)) - \frac{2}{5} c_1(\text{U}(1)_B)^2 - \frac{2}{5} c_1(\text{U}(1)_{Qi})^2 \right), \quad (90)$$

where p_1 is the first Pontryagin class, c_2 is the second Chern class, and c_1 is the first Chern class. All four classes are integer cohomology classes on M_6 .

Proof. Apply the Chern–Weil identifications: $\frac{1}{8\pi^2} \text{tr } R_{\text{Spin}(1,5)}^2 \equiv p_1(\text{Spin}(1,5))$, $\frac{1}{8\pi^2} \text{Tr } F_5^2 \equiv 2c_2(\text{SU}(5))$ (the factor of 2 is the standard normalization), $\frac{1}{4\pi^2} F_B^2 \equiv c_1(\text{U}(1)_B)^2$, similarly for $\text{U}(1)_{Q_i}$. Substituting into (89) and extracting the overall factor of $8\pi^2$ gives (90). ■

14.3 The integral pairing of $[X_4]$ with the fundamental 4-cycle

For the $M_6 = \mathbb{R}^{1,3} \times I_r \times S_\theta^1$, the relevant 4-cycle for the integral pairing with X_4 is the 4-dimensional slice $C_4 = \{x^0 = \text{const}, x^1 = \text{const}\} \times I_r \times S_\theta^1$ (suitably compactified). The integral

$$\int_{C_4} X_4 = K \quad (91)$$

defines the integer OGN-5 ledger entry K via the integrality of $[X_4]$. In the OGN-5 ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$, the value $K = 35$ is the integral pairing of $[X_4]$ with the fundamental 4-cycle of the bulk.

14.4 The H_3 three-form and BI-6 from CCI

Given CCI, the modified Bianchi identity BI-6 of Arisaka (2026, A3-BI6) is the differential consequence:

$$dH_3 = X_4, \quad (92)$$

where H_3 is constructed from the TGG data via the Chern–Simons descent (Proposition 11.3 of Part II).

Proposition 14.3 (BI-6 as a derived consequence of CCI). *Given the TGG structure of Part II and the CCI axioms CCI-ST1–CCI-ST4, the modified Bianchi identity BI-6 of (92) holds globally on M_6 , with H_3 given by the Chern–Simons construction of Proposition 11.3.*

Proof. By Proposition 11.3, $dH_3 = X_4$ where H_3 is constructed from the four sub-Chern–Simons forms with the coefficients of CCI-ST2. This is exactly X_4 of CCI-ST1 with the prescribed coefficients. Hence BI-6 holds. ■

The hierarchy is therefore: TGG $\rightarrow$ CCI $\rightarrow$ BI-6. TGG provides the bundle and connection structure (Part II); CCI provides the algebraic structure of X_4 and its integrality (Part III); BI-6 is the differential identity that relates them.

15 Derivation of the $-2/5$ coefficients from $\mathbb{H}^2$ normalization

We now derive the four coefficients of CCI-ST2, with particular attention to the $-2/5$ coefficients of the two $\text{U}(1)$ terms. This is the algebraic content that makes the OGN-5 integer $L = 5$ inevitable.

15.1 The $-2/5$ coefficients in CCI-ST2: result and cite

The four coefficients $(\alpha_R, \alpha_5, \alpha_B, \alpha_{Qi}) = (+1, -1, -\frac{2}{5}, -\frac{2}{5})$ in X_4 are not free parameters but algebraic consequences of the inner-product normalization on the bulk fermion content. The derivation chain is:

- $\alpha_R = +1$ from the positive sub-block of the $\mathfrak{spin}(1, 5)$ Killing form (sign-definite on the rotation generators, mixed-sign overall);
- $\alpha_5 = -1$ from the standard $\mathfrak{su}(5)$ Killing form $B(T_I, T_J) = 10\delta_{IJ}$ in the conventional $\text{Tr}(T_I T_J) = \frac{1}{2}\delta_{IJ}$ normalization;
- $\alpha_B = \alpha_{Qi} = -\frac{2}{5}$ from the inner-product normalization $\langle T_B, T_B \rangle_{\mathbb{H}^2} = \frac{2}{5}$ induced by the quaternionic doublet structure of TGG and locked by the BI-6 closure equation on the 6D bulk.

The full derivation — combining the trace-squared identity on the fundamental **5** of $\text{SU}(5)$, the BI-6 anomaly-inflow coefficients, and the OGN-5 integer constraints $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ — is given in Arisaka (2026, A3-BI6, Theorem 9.1) (“Abelian trace ratio is $2/5$ ”) and Arisaka (2026, A2-GG6, Theorem 4, Lemma 4.3). In the present paper we import the result; readers who wish to retrace the derivation step-by-step are referred to those theorems.

Remark 15.1 (Physics intuition: why $-2/5$, not a free parameter). The coefficient $-2/5$ is the algebraic content that makes the OGN-5 integer $L = 5$ inevitable. In the doublet inner product $\langle T_B, T_B \rangle_{\mathbb{H}^2}$, the numerator is the trace-squared of the $\text{U}(1)_B$ generator on the color-electroweak fundamental **5**, and the denominator is the dimension of the representation ($= 5$). “Numerator $\div$ denominator $= \frac{2}{5}$ ” is, in this sense, the cleanest algebraic expression of the five-fold structural rigidity of $\text{SU}(5)$ unification. No adjustable parameter; no choice of normalization freedom.

15.2 Why the coefficients are not free parameters

The coefficients $\alpha_R, \alpha_5, \alpha_B, \alpha_{Qi}$ in **CCI-ST2** are not free fitting parameters. They are forced by:

- (1) the Killing-form sign conventions of the four sub-algebras ((Arisaka, 2026, A3-BI6, Theorem 9.1));
- (2) the inner-product normalization on $\mathbb{H}^2$ ((Arisaka, 2026, A3-BI6, Theorem 9.1); see also §15.1);
- (3) the integrality condition **CCI-ST4** on the resulting cohomology class;
- (4) the requirement that the descent to BI-6 produces a well-defined modified Bianchi identity.

Any other choice of coefficients would either violate integrality of $[X_4]$ (failing **CCI-ST4**), or produce inconsistent gauge-anomaly cancellation, or break the magic-dimension-ladder structure $\text{Spin}(1, 5) \cong \text{SL}(2, \mathbb{H})$. This is the algebraic content of GGEP: the four sub-sectors are tied together so tightly that no relative coefficient remains free.

15.3 Structural significance: how the coefficients select OGN-5

The $-2/5$ coefficients of **CCI-ST2** are the algebraic origin of the OGN-5 integer $L = 5$. Specifically, integrality of $[X_4]$ requires that the $1/5$ -denominator contributions from F_B^2 and F_{Qi}^2

clear, which forces the $U(1)_B$ and $U(1)_{Q_i}$ Chern classes to have integer multiples of 5 in their pairings with the fundamental 4-cycle C_4 .

This is the rigidity that selects **OGN-5** from the infinite family of candidate ledgers. Without the $-2/5$ coefficients, the integrality constraint would admit larger or smaller solutions; with the $-2/5$ coefficients fixed by $\mathbb{H}^2$ normalization, the unique minimum is $L = 5$, hence the rest of the ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ as we will prove in Section 18.

16 CCI vs BI-6: algebraic vs differential identity

We now contrast CCI (algebraic) with BI-6 (differential) to clarify their distinct roles in the Gi-4 framework.

16.1 BI-6 as the differential identity

BI-6, as developed in Arisaka (2026, A3-BI6), is the *differential* identity

$$dH_3 = X_4 \tag{93}$$

on the 6D bulk. It is a statement about the exterior derivative of the Kalb–Ramond three-form H_3 . It involves only differential forms and the de Rham operator, with no algebraic structure beyond the inner-product on the bulk Lie algebra.

BI-6 has the following content:

- it is an equation between differential forms on M_6 ;
- it is a local statement (it holds at every point);
- it is gauge-covariant under bulk gauge transformations;
- it determines X_4 from H_3 , but does not, on its own, say which four-form is X_4 (any closed four-form would, in principle, satisfy BI-6 for some H_3).

16.2 CCI as the algebraic identity

CCI, by contrast, is an *algebraic* identity that specifies which four-form is X_4 . It has the following content:

- it specifies the structural form of X_4 as a sum of four self-curvature contributions (**CCI-ST1**);
- it determines the four coefficients in that sum from the $\mathbb{H}^2$ inner-product normalization (**CCI-ST2**);
- it asserts integrality and non-triviality of the resulting cohomology class (**CCI-ST4**, **CCI-ST5**);
- it is global rather than local (the cohomology class is a global topological invariant);
- it determines X_4 uniquely up to gauge, leaving BI-6 as a derived consequence.

16.3 Logical relationship: CCI implies BI-6

Proposition 16.1 (CCI implies BI-6). *Given the TGG structure of Part II and the CCI axioms CCI-ST1–CCI-ST4, the modified Bianchi identity BI-6 ($dH_3 = X_4$) follows automatically, with H_3 given by the Chern–Simons construction of Proposition 11.3.*

Proof. This is Proposition 14.3. ■

16.4 Why both identities are needed

The reader may ask: if CCI implies BI-6, why state both as separate identities? The answer has three parts.

First (logical separation). CCI and BI-6 are logically distinct identities of distinct mathematical types (algebraic versus differential). They formalize different aspects of the bulk closure: CCI formalizes the algebraic structure of the four-form, BI-6 formalizes its differential descent to a three-form.

Second (modularity). Stating them separately makes the Gi-4 framework modular. TGG (Part II) provides the bundle and connection. CCI (Part III) provides the algebraic closure. BI-6 provides the differential closure. These three layers can be modified independently in extensions of the framework. For instance, one can imagine variants of CCI with different sub-algebra content that would lead to different BI-6-type identities; the present GG-Theory commits to a specific CCI that gives the OGN-5 ledger.

Third (descent to the boundary). In the descent to the boundary moduli (Part IV), CCI is the relevant identity for the anomaly inflow (Hopkins–Singer (Hopkins & Singer, 2005)), whereas BI-6 is the relevant identity for the bulk-level geometric structure. The GFF Morse-positivity theorem of Part IV uses CCI directly; BI-6 enters only through its consequence that X_4 is exact on M_6 (so the descent map is well-defined).

17 The physical layer of CCI: 4D covariant conservation and modified TOV

This section develops the *physical* content of CCI (Structure Theorems CCI-ST6 and CCI-ST7) in operational detail: the local 4D covariant conservation equation $\nabla_\mu T_{\text{TGG}}^{\mu\nu} = 0$, the PBI-4 compensator field that enforces it, and the explicit derivation of the modified TOV equation (74) for a static spherically symmetric configuration. The mathematics here is straightforward; the conceptual novelty is the realization that the *single* bulk identity $dH_3 = X_4$ produces, by bulk-to-boundary descent, the *universal* empirical conservation laws that observers measure in every laboratory.

17.1 From bulk BI-6 to boundary divergence: the descent map

We recall from Part II (Section 9) that the Projected Bianchi Identity PBI-4 arises as the boundary projection of the bulk closure $dH_3 = X_4$. Specifically, integrating X_4 over the two-cycle of the extra dimensions (y^4, y^5) gives a 2-form $\tilde{X}_2$ on the boundary M_4 , and the boundary divergence

of any sector current J^μ is constrained by

$$\nabla_\mu J^\mu = \star \tilde{X}_2 \cdot \langle \cdot \rangle_{\text{KK}}, \quad (94)$$

where $\star$ is the Hodge dual on M_4 and $\langle \cdot \rangle_{\text{KK}}$ denotes the Kaluza–Klein average over the extra-dimensional profile. This is the precise content of PBI-4.

Proposition 17.1 (Sector divergences from PBI-4). *For each of the four boundary Ward currents J_e, J_B, J_{Qi}, J_g (electric, baryonic, informational, metric), the boundary divergence is*

$$\nabla_\mu J_e^\mu = -\frac{2}{5} \langle F_B \wedge F_B \rangle_{\text{KK}}, \quad (95)$$

$$\nabla_\mu J_B^\mu = -\frac{2}{5} \langle F_B \wedge F_B \rangle_{\text{KK}} \cdot \delta_{\text{baryon}}, \quad (96)$$

$$\nabla_\mu J_{Qi}^\mu = -\frac{2}{5} \langle F_{Qi} \wedge F_{Qi} \rangle_{\text{KK}}, \quad (97)$$

$$\nabla_\mu J_g^\mu = + \langle \text{tr } R_{\text{Spin}(1,5)}^2 \rangle_{\text{KK}} - \langle \text{Tr } F_5^2 \rangle_{\text{KK}}, \quad (98)$$

where δ_{baryon} is the relative coupling of the baryonic to the electromagnetic sector (set to zero in the Standard-Model limit). The four divergences sum to zero by closure of X_4 : $\sum_a \nabla_\mu J_a^\mu = -\langle X_4 \rangle_{\text{KK}} = 0$ identically when $[X_4]$ is integral.

Proof. Each Ward current J_a ($a \in \{e, B, Qi, g\}$) is associated with a boundary symmetry of the corresponding sector connection in the TGG composite. By the standard Ward identity for that sector, $\nabla_\mu J_a^\mu$ equals the boundary anomaly inflow from the corresponding 4-form term in X_4 . Summing over the four sectors and applying the closure of X_4 (Structure Theorem **CCI-ST3**) gives the total cancellation. ■

17.2 Master Conservation Law as a sum of sector divergences

Theorem 17.2 (**T2.2 prelude**: Master Conservation from PBI-4). *Given the four sector divergences of Proposition 17.1, the Master Stress Tensor of TGG-ST6 satisfies the boundary covariant conservation equation*

$$\nabla_\mu T_{\text{TGG}}^{\mu\nu} = 0 \quad (\text{on } M_4, \text{ identically}) \quad (99)$$

provided each sector stress tensor satisfies its own sector conservation modulo the corresponding Ward divergence:

$$\nabla_\mu T_a^{\mu\nu} = -J_a^\nu \cdot \nabla_\mu J_a^\mu \cdot g_{(4)}^{\nu\sigma} / |J_a|^2, \quad a \in \{\text{matter, gauge, GS, KK}\}. \quad (100)$$

Proof sketch. Sum the four sector conservation equations (100) and apply the divergence cancellation $\sum_a \nabla_\mu J_a^\mu = 0$ from Proposition 17.1. The cross terms cancel by the closure of X_4 (Structure Theorem **CCI-ST3**) and the integrality of $[X_4]$ (Structure Theorem **CCI-ST4**). Equation (99) follows. ■

This establishes that the boundary 4D conservation equation $\nabla_\mu T_{\text{TGG}}^{\mu\nu} = 0$ is a *kinematic* consequence of the bulk identity $dH_3 = X_4$ together with the closure of the four Ward currents. No additional dynamical equation is needed.

17.3 The PBI-4 compensator field

The Projected Bianchi Identity PBI-4 can be packaged into a single boundary 1-form C_μ , called the PBI-4 *compensator field*, defined as follows.

Definition 17.3 (PBI-4 compensator field). The PBI-4 compensator field C_μ on the boundary M_4 is the unique 1-form such that

$$dC = \tilde{X}_2, \quad \nabla^\mu C_\mu = 0, \quad (101)$$

where $\tilde{X}_2 = \int_{(y^4, y^5)} X_4$ is the KK-reduced 2-form on M_4 (a closed boundary 2-form by **CCI-ST3**). The compensator C_μ exists and is unique up to gauge transformations $C \rightarrow C + d\chi$, by the Poincaré lemma applied to the simply connected boundary M_4 .

Proposition 17.4 (Operational form of CCI-ST6 via C_μ). *The boundary covariant conservation equation (72) is equivalent to the statement that the four sector divergences are linked by C_μ :*

$$\nabla_\mu T_a^{\mu\nu}(x) = \alpha_a \cdot C^\nu(x) \cdot \rho_a(x), \quad a \in \{\text{matter, gauge, GS, KK}\}, \quad (102)$$

where α_a are the four **CCI-ST2** coefficients and $\rho_a(x)$ are the local densities of each sector. The sum $\sum_a \alpha_a \rho_a(x) C^\nu(x) = 0$ by closure of C .

Proof. This follows from Proposition 17.1 by pulling out the common factor C_μ from each Ward divergence and applying the algebraic relation $\sum_a \alpha_a = 1 - 1 - 2/5 - 2/5 = -4/5 \neq 0$, but the *density-weighted* sum vanishes identically because $\sum_a \alpha_a \rho_a(x) = 0$ on the physical boundary by the constraint $\langle X_4 \rangle_{\text{KK}} = 0$ (integrality). ■

Remark 17.5 (Physics intuition: the compensator field as a “glue” across sectors). The PBI-4 compensator C_μ has no analog in the Standard Model. It is a new boundary field that “glues together” the four sector conservation laws so that they cancel pairwise, producing a single global Master Conservation Law. In the limit where one sector dominates (e.g., in pure gravity), C_μ becomes degenerate and the master equation reduces to that sector’s local conservation. The non-trivial content of CCI-ST6 is that, in the full multi-sector configuration, C_μ is non-zero and links the sectors.

In phenomenological terms, C_μ is a candidate for a new ultra-weak boundary field with calculable couplings. It is too weak to detect at LHC energies but contributes at the level of 10^{-7} to 10^{-4} corrections in compact astrophysical configurations, where the sector densities ρ_a are extreme. This is the source of the modified TOV correction $\Delta_{\text{GS}}(r)$ of **CCI-ST7**.

17.4 Modified TOV equation: the CCI-side identity (primary derivation in A5)

The boundary covariant conservation law (72) of **CCI-ST6**, evaluated on a static, spherically symmetric configuration with the PBI-4 compensator C_μ of Definition 17.3, produces the *modified Tolman–Oppenheimer–Volkoff* equation (74) of **CCI-ST7**.

The full step-by-step derivation — setup of the Schwarzschild ansatz, sector-by-sector contributions to $T_{\text{TGG}}^{\mu\nu}$ (matter / gauge / Green–Schwarz / Kaluza–Klein), application of the Master

Conservation Law, identification of $\Delta_{\text{GS}}(r)$ with the bracket $[\nabla_\mu X_{\text{GS}}^{\mu r} + \nabla_\mu T_{\text{KK}}^{\mu r}]$, and evaluation of the leading-order OGN-5 coefficient $\kappa_{\text{GS}} = (2/5)(N/K)(J/L) = 42/4375 \approx 9.6 \times 10^{-3}$ — is the boundary-law projection of GDOS Theorem T8.6 (modified TOV from OGGEp descent) in the companion foundation paper GG-A5-GDOS (Arisaka, 2026, A5-GDOS, § T8.6, Part IV). The present paper imports the derivation result. The CCI-side reading is that the modified TOV equation is the spherically-symmetric reduction of the PBI-4-linked sector divergences of Proposition 17.1; the $\Delta_{\text{GS}}(r)$ correction is the explicit C_μ -mediated gluing across the four sector stress tensors, as expressed in (75).

Empirical scale of the prediction. For a typical neutron star ($M_{\text{NS}} \sim 1.4 M_\odot$, $R \sim 10$ km, $\rho_{\text{central}} \sim 10^{15}$ g/cm³), the ratio $\Delta_{\text{GS}}/(\text{dp/dr})_{\text{TOV}} \sim 10^{-3}$ to 10^{-2} , translating to a 0.1%–1% shift of the neutron-star mass–radius relation — at the threshold of NICER measurements (Miller et al., 2021, Riley et al., 2021) and well within the reach of next-generation observatories. Numerical details and the full OGN-5 coefficient chain are recorded in A5 Part IV (Theorem T8.6 and the worked-example calculation following it).

Remark 17.6 (Physics intuition: why the modified TOV is the right falsifier). The modified TOV equation is the cleanest astrophysical falsifier of CCI for three reasons. First, it is a *static* prediction, free of dynamical assumptions about stellar evolution. Second, the correction $\Delta_{\text{GS}}(r)$ is *calculable* from the OGN-5 ledger with no free parameters — the only input is the assumed value of the bulk integral $\langle X_4 \rangle_{\text{bulk}}$, which is bounded by the integers (1, 3, 5; 7, 35). Third, the prediction is *universal* across all compact stars (neutron stars, white dwarfs, magnetars) up to a scaling by the central density — so a correlated population study (rather than a single-star measurement) provides the strongest test.

If multiple compact stars across a wide mass range *all* agree with the standard TOV prediction to better than 0.1%, then CCI (and the entire GG-6 framework) is in serious trouble. Conversely, if the next-generation NICER and gravitational-wave inspiral measurements show a coherent 10^{-3} to 10^{-2} deviation matching the OGN-5 prediction, this would be the first direct empirical confirmation of GG-Theory.

Remark 17.7 (Companion physical predictions: cosmological observables). The same Master Conservation Law that produces the modified TOV equation also predicts:

- A small modification of the Friedmann equations at high redshift (sourced by the same X_{GS} contribution), discussed in detail in GG-C1-Kaxion;
- A correlated shift of the cosmological growth-of-structure function $f\sigma_8(z)$ at the 10^{-3} level;
- A specific relation between the dark-energy equation-of-state parameter $w(z)$ and the bulk integral $\langle X_4 \rangle_{\text{bulk}}$, of the form $w(z) + 1 \propto \langle X_4 \rangle / M_{\text{KK}}^4$.

These are all kinematic projections of the single bulk identity $\text{d}H_3 = X_4$ via PBI-4. The fact that one bulk identity produces a coherent set of falsifiable predictions across both astrophysics and cosmology is the strongest empirical signature of GG-Theory.

18 Theorem T2 (T-GRIP-2): CCI closure under MSMF forces unique OGN-5 ledger and Master Conservation

We now state and prove Theorem **T2**, the principal theorem of Part III. The theorem is presented in two parts: **T2.1** (algebraic OGN-5 uniqueness, the original statement) and **T2.2** (boundary covariant conservation, new in to capture the physical content of CCI).

Dual-layer architecture and the compatibility relation $\mathfrak{C}_{\text{CCI}}$. Per the L16 architecture, T2.1 is the *algebraic layer* (the OGN-5 integer ledger as the unique MSMF-admissible solution of the CCI closure equations), and T2.2 is the *operational/physical layer* (the boundary Master Conservation Law as a Ward-identity-falsifiable iff statement). The two layers are tied together by the named compatibility relation

$$\mathfrak{C}_{\text{CCI}} \equiv \text{T2.1} \implies \text{T2.2} \quad (\text{via the PBI-4 compensator and the Ward package of } \mathbf{CCI-ST6}).$$

(103)

$\mathfrak{C}_{\text{CCI}}$ above is a one-line proposition with the proof: the OGN-5 uniqueness at the algebraic level (T2.1) forces the four-component coefficient vector $(a, b, c, d) = (1, -1, -2/5, -2/5)$ in the anomaly polynomial X_4 ; the PBI-4 compensator construction of **CCI-ST6** then carries this trace-ratio lock to a divergence-balance identity on the four boundary Ward currents, which is the iff content of T2.2.

18.1 Statement of T2.1 (algebraic uniqueness): imported from A2-GG-6/GG-A3-BI6

The OGN-5 integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ is the unique MSMF-admissible solution of the integer closure constraints $K = M^2 + N^2 + L^2$ and $K = LJ$ together with the constitutional inputs $L = 5$ (SU(5) rank), $M \geq 1$ (one Higgs scalar) and $N \geq 3$ (lower bound from observed CP violation). We import this uniqueness as the algebraic layer of T2.

Theorem 18.1 (T2.1: OGN-5 characterization of CCI closure (imported)). *Among integer 5-tuples $(M, N, L; J, K) \in \mathbb{Z}_{>0}^5$ satisfying the CCI-closure constraints*

$$K = M^2 + N^2 + L^2, \quad K = LJ, \tag{104}$$

*together with the constitutional inputs $L = 5$, $M \geq 1$, $N \geq 3$, the MSMF criterion of **CCI-ST5** selects the unique solution*

$$(M, N, L; J, K) = (1, 3, 5; 7, 35). \tag{105}$$

Proof (imported). The full Diophantine proof — combine the two closure constraints, specialize to $L = 5$, work modulo 5 to identify admissible (M, N) pairs, enumerate small solutions, and apply the ℓ^1 MSMF criterion — is owned by the bulk-side companion papers in the program. Specifically, the Diophantine uniqueness is (Arisaka, 2026, A2-GG6, Theorem 6) (the OGN-5 ledger as the MSMF minimum of the BI-6 anomaly-closure equation) and (Arisaka, 2026, A3-BI6, Theorem 7.2) (the boundary trace-ratio lock $|a/c| = |a/d| = 5/2$ characterized by the same ledger). The present paper

imports this uniqueness as Theorem T2.1 and records the CCI-side reading below; the reader who wishes to retrace the Diophantine analysis step-by-step is referred to those two companion theorems. One presentational difference is worth a line: the cascade of Arisaka (2026, A2-GG6, Theorem 6) carries the additional closure-index clause $J = 2N + 1$ (its Step 4), which the hypotheses of T2.1 above do not assume; T2.1 therefore states the uniqueness over the larger admissible set, and the selection is unaffected — indeed independent of the complexity norm (see the M9 entry, §57). ■

Remark 18.2 (Physics intuition: why $(1, 3, 5; 7, 35)$ is the only ledger that survives). The ledger $(1, 3, 5; 7, 35)$ packages every empirically known integer count of the Standard Model: $M = 1$ Higgs scalar, $N = 3$ fermion families, $L = 5$ SU(5) gauge rank, $J = 7$ topological winding count, $K = 35$ total degrees of freedom. Indeed, the MSMF criterion explains why we do not see $M = 2$ or $N_g = 4$: any such extension would force a non-minimal ledger that MSMF disfavors. The constitutional input $L = 5$ encodes the visible SU(5) unifying gauge group; a hypothetical $L = 4$ Pati–Salam or $L = 10$ E_8 GG-Theory would yield a different ledger, which is why the $L = 5$ choice is read as a constitutional commitment rather than a derivation output. In honest summary: T2.1 is a uniqueness statement *within the SU(5) unification family*, with the Diophantine machinery delivered by A2-GG-6 Theorem 6 and the boundary trace-ratio reading by GG-A3-BI6 Theorem 7.2. What T2.1 contributes uniquely in the present paper is the CCI-side reading: the same ledger that closes the BI-6 anomaly polynomial in the bulk also characterizes the MSMF-minimal solution of the CCI integrality constraints on the boundary. This dual-side characterization is the content of (105).

18.2 Corollary: family number $N_g = 3$ from upstream CCI

Corollary 18.3 ($N_g = 3$ from upstream CCI). *The family number N_g — which is not itself a ledger entry — is forced to equal exactly 3 by the combination of:*

- (1) *the lower bound $N \geq 3$ from CP violation (the empirical constraint on the Jarlskog invariant);*
- (2) *the upper bound $N \leq 3$ from the cascade closure of GFF/GRIP (Theorem T4 of Part V, conditional on M10 per Remark 34.2);*
- (3) *the MSMF minimality criterion applied to the integrality of $[X_4]$ as in Theorem 18.1 above.*

The combined effect is that $N_g = 3$ is the unique admissible family number, and the OGN-5 ledger is fixed at $(1, 3, 5; 7, 35)$.

Proof. By Theorem 18.1, with $L = 5$ fixed by constitutional input, the family number N_g is determined by the Diophantine + MSMF selection to be $N_g = 3$ (the minimum-complexity admissible value, established in (Arisaka, 2026, A2-GG6, Theorem 6)). Combined with the empirical lower bound $N \geq 3$ and the cascade-closure upper bound $N \leq 3$ (Theorem T4 of Part V), $N = 3$ is forced. ■

18.3 Statement and proof of T2.2 (boundary covariant conservation)

Indeed, the second half of Theorem T2 packages CCI-ST6 into a precise iff statement — the *boundary covariant conservation theorem* that links the bulk BI-6 closure to the boundary 4D conservation equation.

Theorem 18.4 (T2.2: Master Conservation iff Ward + PBI-4). *Let $(P, \mathcal{A}_{\text{tot}}, \mathcal{R}_{\text{tot}}; T_{\text{TGG}}^{\mu\nu}, \nabla)$ be the TGG structure of **TGG-ST6**, with $T_{\text{TGG}}^{\mu\nu} = T_{\text{matter}} + T_{\text{gauge}} + X_{\text{GS}} + T_{\text{KK}}$ on the boundary $\partial M_6 \cong M_4$. Let $X_4 = \text{tr } R_{\text{Spin}(1,5)}^2 - \text{Tr } F_5^2 - \frac{2}{5}F_B^2 - \frac{2}{5}F_{Q_i}^2$ satisfy Structure Theorems **CCI-ST1–CCI-ST5** on the bulk M_6 . Then*

$$\nabla_\mu T_{\text{TGG}}^{\mu\nu} = 0 \quad (\text{on } M_4, \text{ identically}) \quad (106)$$

if and only if the following two conditions both hold:

- (1) **Ward closure.** *The four boundary Ward currents J_e, J_B, J_{Q_i}, J_g are individually closed: $d\star J_a = 0$ for all $a \in \{e, B, Q_i, g\}$, modulo the anomaly inflow from X_4 given by Proposition 17.1;*
- (2) **PBI-4 compensator.** *The PBI-4 compensator field C_μ of Definition 17.3 exists, satisfies $dC = \tilde{X}_2$, and links the four sector divergences as in Proposition 17.4.*

Proof. ($\Leftarrow$) Assume conditions (1) and (2) hold. By Proposition 17.1, the four sector divergences sum to $-\langle X_4 \rangle_{\text{KK}}$, which vanishes by integrality **CCI-ST4** together with closure **CCI-ST3**. By Theorem 17.2, the Master Stress Tensor of TGG-ST6 then satisfies $\nabla_\mu T_{\text{TGG}}^{\mu\nu} = 0$.

($\Rightarrow$) Conversely, assume $\nabla_\mu T_{\text{TGG}}^{\mu\nu} = 0$ on M_4 . Decomposing into the four sector contributions of TGG-ST6 and applying the linearity of ∇_μ , we obtain $\sum_a \nabla_\mu T_a^{\mu\nu} = 0$. By the duality between sector stress tensors and sector Ward currents (the standard Noether construction), each $\nabla_\mu T_a^{\mu\nu}$ is proportional to a combination of the Ward divergences $\nabla_\mu J_a^\mu$. The vanishing of the total then forces the Ward divergences to satisfy the closure relations of Proposition 17.1, which is condition (1). The existence of a single 1-form C_μ linking these divergences (condition (2)) then follows *locally* from the Poincaré lemma applied to the closed boundary 2-form $\tilde{X}_2$; its global, single-valued (monodromy-free) extension across the codimension-2 instanton stratum is open problem M11 (Remark 18.5). ■

Remark 18.5 (Status: the reverse direction of T2.2 is conditional on M11). The ($\Rightarrow$) direction of T2.2 constructs the PBI-4 compensator C_μ *locally*, by the Poincaré lemma. Its global, single-valued existence — vanishing of the monodromy of C_μ around the codimension-2 instanton stratum — is the same topological non-obstruction that conditions T4.2, and is recorded as open problem M11 — now sharpened to a spectral-flow form with a verified benchmark instance (see the M11 entry in §57). The iff of T2.2 therefore holds pointwise (verified case-by-case on $\mathbb{CP}^{N-1}$ for $N_g \leq 3$, sufficient for the Standard-Model application) and globally subject to M11; the forward ($\Leftarrow$) direction is unconditional. This matches the honest status carried by T4.2 (Remark 34.6).

Remark 18.6 (Logical structure of T2: T2.1 + T2.2). Theorem **T2** as stated in is the *conjunction* of **T2.1** and **T2.2**. T2.1 fixes the algebraic side: the unique OGN-5 ledger is forced by the Diophantine closure constraints together with MSMF. T2.2 fixes the physical side: the unique boundary conservation law is forced by Ward closure together with the PBI-4 compensator. The two halves are linked by integrality of $[X_4]$: T2.1 gives the integers $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ that make $\langle X_4 \rangle_{\text{KK}} = 0$ on the boundary, which is precisely the anomaly-cancellation condition needed for T2.2 (cf. the extended-model conditions of Kanno, 2010).

This is the precise sense in which the algebraic side of CCI *determines* the physical side: the OGN-5 ledger of T2.1 is exactly the integer data that makes the Master Conservation Law of T2.2

anomaly-free. No other ledger gives a consistent boundary 4D conservation equation.

Remark 18.7 (Physics intuition: why T2.2 is the deeper half). A graduate student in physics may at first think that T2.1 is the “deeper” theorem because it produces specific integers (1, 3, 5; 7, 35). But T2.2 is in fact the deeper statement: it is the assertion that those integers *must* be the unique data that make the boundary world consistent. T2.1 alone only says that (1, 3, 5; 7, 35) has minimum complexity — one could imagine philosophical objections to MSMF. T2.2 says that no other ledger *works*: any other choice would produce a 4D world with broken conservation laws (anomalies), which is empirically excluded by the laboratory observation that energy and momentum are conserved to better than 10^{-15} precision in low-energy experiments.

In this sense, T2.2 is the *empirical* arm of T2: the integers (1, 3, 5; 7, 35) are forced not by aesthetics (MSMF) but by the empirical reality of conservation laws. The convergence of T2.1 (aesthetic) and T2.2 (empirical) on the same answer is one of the strongest internal consistency checks of GG-Theory.

Corollary 18.8 (Astrophysical falsifier from T2.2). *Combining Theorem 18.4 with the spherically symmetric reduction of Section 17.4, the modified TOV equation (74) with the PBI-4-determined correction $\Delta_{\text{GS}}(r)$ (the OGN-5 coefficient is $\kappa_{\text{GS}} = (2/5)(N/K)(J/L) = 42/4375 \approx 9.6 \times 10^{-3}$, with the full derivation in (Arisaka, 2026, A5-GDOS, Theorem T8.6, Part IV)) is a kinematic prediction of GG-Theory. An empirical refutation of this equation across a coherent population of compact stars would constitute a falsification of CCI and hence of the entire GG-6 framework.*

Proof. This is the spherically symmetric specialization of T2.2 with the OGN-5 coefficient $\kappa_{\text{GS}} = 42/4375 \approx 9.6 \times 10^{-3}$; the full derivation is GG-A5-GDOS Theorem T8.6, summarized in Section 17.4. ■

19 Connection to Hopkins-Singer differential cohomology

We now articulate the connection between CCI and the Hopkins–Singer framework of differential cohomology (Hopkins & Singer, 2005), which will be the technical machinery used in Part IV to prove GFF Morse-positivity.

19.1 Differential cohomology preliminaries

Differential cohomology refines ordinary cohomology by combining topological data (cohomology classes) and differential-form data (curvature representatives). For a manifold X and integer $k \geq 0$, the differential cohomology group $\check{H}^k(X; \mathbb{Z})$ fits into exact sequences relating it to ordinary cohomology and to closed differential forms. Specifically:

$$0 \rightarrow H^{k-1}(X; \mathbb{R}/\mathbb{Z}) \rightarrow \check{H}^k(X; \mathbb{Z}) \xrightarrow{\text{curv}} \Omega_{\text{closed}}^k(X) \rightarrow 0, \quad (107)$$

where $\text{curv} : \check{H}^k \rightarrow \Omega_{\text{closed}}^k$ extracts the curvature representative. An element of $\check{H}^k(X; \mathbb{Z})$ is, roughly speaking, a cohomology class with a chosen differential-form representative.

19.2 X_4 as a differential cohomology class

Definition 19.1 (Differential cohomology class of X_4). Given the CCI structure of **CCI-ST1–CCI-ST5**, the four-form X_4 together with its integrality condition $[X_4] \in H^4(M_6; \mathbb{Z})$ (**CCI-ST4**) lifts to a differential cohomology class

$$\check{x}_4 \in \check{H}^4(M_6; \mathbb{Z}), \quad (108)$$

called the *differential cohomology class* of X_4 . Its curvature representative is X_4 itself: $\text{curv}(\check{x}_4) = X_4$.

The lift to differential cohomology is what enables the bulk-to-boundary descent of X_4 to the boundary moduli (Part IV). Without the lift, one would have only the ordinary cohomology class $[X_4]$, and the descent would lose the differential-form structure needed to define the Geometric Free-Energy Functional rigorously.

19.3 The boundary descent map

The Hopkins–Singer descent map for differential cohomology, adapted to the 6D bulk, is a homomorphism

$$\check{\rho}_{\text{bdy}}: \check{H}^4(M_6; \mathbb{Z}) \rightarrow \check{H}^4(\mathcal{M}_{\text{bdy}}; \mathbb{Z}), \quad (109)$$

where $\mathcal{M}_{\text{bdy}}$ is the boundary moduli space of OGGEF (the stratified space of admissible boundary projector configurations). The descent map restricts a bulk differential cohomology class to its boundary counterpart via integration along the projection fibers.

Proposition 19.2 (Descent of $\check{x}_4$ to family-projector moduli). *The descent of $\check{x}_4 \in \check{H}^4(M_6; \mathbb{Z})$ to the family-projector moduli stratum $\mathbb{CP}^{N-1} \subset \mathcal{M}_{\text{bdy}}$ is the differential cohomology class*

$$\check{\rho}_{\text{bdy}}(\check{x}_4)|_{\mathbb{CP}^{N-1}} = c_R \check{p}_1(\mathbb{CP}^{N-1}) \in \check{H}^4(\mathbb{CP}^{N-1}; \mathbb{Z}), \quad (110)$$

where $\check{p}_1(\mathbb{CP}^{N-1})$ is the first Pontryagin class of $\mathbb{CP}^{N-1}$ (lifted to differential cohomology) and $c_R \in \mathbb{Q}_{>0}$ is the rational coefficient set by the $\text{Spin}(1, 5)$ Killing-form sign.

Sketch. This is essentially Lemma 3.1 of the N=3 proof of Arisaka (2026, P2-N3). The descent integrates X_4 along the fibers of the boundary projector $\Pi: M_6 \rightarrow M_4$ restricted to the family-projector slice; the result, by the Chern–Weil identification of $\text{tr } R_{\text{Spin}(1,5)}^2$ with p_1 , is the first Pontryagin class on $\mathbb{CP}^{N-1}$ with coefficient c_R . The vanishing of the gauge-curvature contributions on $\mathbb{CP}^{N-1}$ (Lemma 3.2 of Arisaka (2026, P2-N3)) removes the $\text{SU}(5)$, $\text{U}(1)_B$, $\text{U}(1)_{Q_i}$ terms. ■

The chain into Part IV. The descent class (110) is the input to the GFF Morse-positivity theorem of Part IV. Specifically, the strict positivity of the alternating Morse-index sum on $\mathbb{CP}^{N-1}$ is equivalent, via Hirzebruch’s signature theorem, to the non-vanishing of the integrated L-genus, which by (110) is non-zero exactly when $[X_4] \neq 0$ in $H^4(M_6; \mathbb{R})$. This is the technical chain by which CCI forces GFF Morse-positivity.

The detailed proof is given in Part IV (Theorem 26.1).

20 Worked examples

To make the CCI content concrete, we work through three explicit examples.

20.1 Example 1: explicit X_4 on $\mathcal{M}_6$

We compute X_4 in the explicit 6D bulk of GG-Theory.

Spin curvature contribution. The Riemann tensor of the scale-profile metric (1) has non-vanishing components in the radial sector. The leading 4-form $\text{tr } R_{\text{Spin}(1,5)}^2$ has support on the radial slice and integrates to a non-trivial 4-cycle pairing (Arisaka, 2026, A3-BI6, Section 3).

SU(5) curvature contribution. The orbifold breaking $\text{SU}(5) \rightarrow \text{SU}(3) \times \text{SU}(2) \times \text{U}(1)_Y$ along S_θ^1 gives a non-trivial $\text{Tr } F_5^2$ contribution centered on the boundary.

U(1)_B, U(1)_{Qi} contributions. On the admissible class of Arisaka (2026, A3-BI6) the Abelian connections are flat along the throat, so F_B^2 and $F_{Q_i}^2$ contribute nothing here; the Abelian slot of X_4 is carried by the unit flux class of the index background (Remark 23.5, premise (P-ii)), not by a curvature of the Wilson lines.

Total integral. $\int_{C_4} X_4 = K = 35$, the OGN-5 integer.

20.2 Example 2: alternative ledgers ruled out by MSMF

We explicitly enumerate the smallest few alternative integer 5-tuples satisfying the closure constraints (104) and constraint (iii) ($L = 5$), then show how MSMF eliminates them in favor of $(1, 3, 5; 7, 35)$.

Table 3: Every alternative integer ledger of the OGN-5 type that clears the prior constraints, with its complexity and its fate. The selected ledger $(1, 3, 5; 7, 35)$ scores 51; every competitor scores 65 or more, and MSMF rejects all of them. Listing the rejected ledgers rather than merely asserting minimality is the point of the table: the reader can check the arithmetic, and the margin between 51 and the nearest alternative is wide enough that the selection does not turn on a fine judgment about how complexity is counted.

M	N	L	J	K	Complexity	Status
1	3	5	7	35	51	Selected by MSMF
2	4	5	9	45	65	MSMF reject
3	4	5	10	50	72	MSMF reject
4	3	5	10	50	72	MSMF reject
2	6	5	13	65	91	MSMF reject
3	6	5	14	70	98	MSMF reject
6	3	5	14	70	98	MSMF reject
1	7	5	15	75	103	MSMF reject
5	5	5	15	75	105	MSMF reject

The minimum complexity is achieved uniquely by $(1, 3, 5; 7, 35)$ with complexity 51. All other admissible ledgers have complexity ≥ 65 , and they all become accessible only at higher integer values that MSMF rejects. In fact more is true: every admissible competitor dominates the selected

ledger coordinatewise after sorting, so the selection is independent of the choice of complexity norm altogether (the dominance argument is recorded in the M9 entry, §57).

20.3 Example 3: integrality check for OGN-5

We verify the integrality $[X_4] \in H^4(M_6; \mathbb{Z})$ for the OGN-5 ledger.

Sub-bundle Chern classes.

- $p_1(\text{Spin}(1, 5))$ pairs with C_4 to give the integer $\langle p_1, C_4 \rangle = K = 35$ (the gravity Pontryagin class).
- $c_2(\text{SU}(5))$ pairs with C_4 to give the integer $\langle 2c_2, C_4 \rangle = 2L = 10$ (twice the $\text{SU}(5)$ rank, with the standard normalization factor of 2).
- $c_1(\text{U}(1)_B)$ pairs with S_θ^1 to give the integer $\langle c_1, S_\theta^1 \rangle = M = 1$. Squared: 1.
- $c_1(\text{U}(1)_{Q_i})$ pairs with S_θ^1 to give the integer $N = 3$. Squared: 9.

Combined integrality check.

$$\langle [X_4], C_4 \rangle = K - 2L - \frac{2}{5}M^2 - \frac{2}{5}N^2 \cdot (\text{geom. factor}) \quad (111)$$

$$= 35 - 10 - \frac{2}{5}(1) - \frac{2}{5}(9) \cdot (\text{geom. factor}). \quad (112)$$

With the geometric factors set by the bulk volume normalization, the result is an integer (specifically, the algebraic combination of M, N, L, J, K that gives a clean integer is the OGN-5 closure identity itself). The detailed verification is in [Arisaka \(2026, A3-BI6, Section 4\)](#).

Discussion. The integrality of $[X_4]$ is non-trivial: the $1/5$ coefficients of the two $\text{U}(1)$ contributions threaten to produce non-integer cohomology values, but the OGN-5 ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ exactly clears the denominators in the admissible combinations. This is the algebraic content of **CCI-ST4**: the integrality requirement, combined with **CCI-ST2**'s coefficient determination, forces the OGN-5 integers.

21 Integers that nobody chose

Something happens in this Part that is worth stopping to notice, because it is the kind of thing that does not happen often and is easy to read past.

A consistency condition was imposed — roughly, the requirement that a certain quantity which must cancel does in fact cancel, everywhere and exactly. No data entered. No parameter was tuned. And out of that requirement came a short list of whole numbers, which then turn out to control how many families of matter there are, how the charges are assigned, and what a compact star should weigh.

Whole numbers behave differently from other quantities in a theory, and the difference is what makes this Part load-bearing. A continuous parameter can be adjusted; if a measurement disagrees, the parameter moves and the theory survives, a little weaker each time. An integer cannot move. It is three or it is four, and if the world says four when the ledger says three, nothing can be done — there is no small adjustment available, because there is no such thing as a small change to a whole

number. *A theory that has been forced onto integers has given up its capacity to retreat.* That is a loss of flexibility, and it is the whole point.

It is worth being clear about what has and has not been shown. The integers are forced *given* the arena and the constitutional requirements, both of which were declared in Part I and are inputs. This Part does not derive six dimensions; it shows that once you have them, the bookkeeping has exactly one solution. That is a conditional result, and stating it as conditional costs nothing and buys the right to be believed about the rest.

There is a general point here about what makes a prediction cheap or expensive. A theory that predicts a number to three decimal places has said something, but it has usually paid for the precision with a parameter, and a parameter can absorb a disagreement. A theory that predicts a whole number has paid with nothing and can absorb nothing. This is why counting predictions — how many families, how many charges, how many dimensions — are the most dangerous kind a framework can make, and why frameworks able to make them are rare.

The physical consequence carried out of this Part is small, specific, and measurable: a correction to the relation between a neutron star's mass and its radius, of a size fixed entirely by the integers, with no freedom to adjust it. A neutron star is a very long way from an anomaly polynomial, and the fact that the two are connected at all is the sort of thing this framework exists to claim. If the shift is not there at the predicted size, the ledger is wrong and this Part fails — not approximately, not in some regime, but as a whole.

A last thought, which the rest of the paper will lean on. The integers that come out of this Part were not measured, and they were not chosen for their elegance either. They fell out of a demand for consistency imposed on an object built for entirely different reasons. If that is what is happening — if the world's small whole numbers are the residue of self-consistency rather than facts about it — then a great many questions of the form "why this value?" have been asked in the wrong grammar all along.

What the next Part adds is motion. So far the paper has a structure and a set of numbers describing a world that does not yet do anything. The question becomes what shape the possibilities form when laid out as a landscape, and why a system placed on that landscape has anywhere to go.

Part IV

GFF

The Geometric Free-Energy Functional: Pillar 3 of Gi-4

Nature does not roll downhill toward the lowest energy. It rolls downhill toward the lowest *free* energy — energy discounted by what the system gains in disorder — which is why ice melts above zero degrees and why a protein folds into one shape out of astronomically many. Anyone who works on phase transitions, on self-assembly, or on the thermodynamics of living systems already knows this in their bones. The remark opening the Part’s first section develops the point with the familiar example, and it is worth reading even by a physicist, because the whole of the following two Parts is an argument about a landscape.

The question here is whether the six-dimensional geometry supplies such a landscape for the universe as a whole. Not an analogy — an actual function, defined on the space of configurations the boundary world could be in, whose value the world decreases as it evolves. This Part constructs that function and shows it has the two properties a landscape needs. It has to be shaped like a landscape rather than a plateau: there must be genuine basins, with walls, so that a system arriving in one stays there. And motion on it has to run one way, so that a descent is a descent and not an aimless wander.

Those two properties are established as the Part’s main results, and the second is worth naming carefully, because it is where the arrow of time enters this paper. The function decreases along every physically admissible evolution of an open system. That statement is the second law of thermodynamics in the form this framework produces it — not assumed, not argued from statistics, but a consequence of the geometry and of the fact that no system is closed. A companion paper carries the law-level treatment of time’s direction ([Arisaka, 2026, P5-CPT](#)); what is proved here is the runtime version, the one that makes the descent of the next Part irreversible.

The landscape has one further feature, and it is the feature that makes the rest of the paper possible: it is not smooth. It comes in layers — strata — of different dimension, and a system can sit inside a layer or fall between layers, which is what will later distinguish a slow settling from a sudden reorganization. Also established here is a test: if the landscape is real, then any observable on the boundary must satisfy two classical relations from statistical physics, equipartition and the fluctuation–dissipation relation, and their failure would end this pillar. A reader from the life sciences may find Part X’s section on the free-energy principle a useful entry point, since it connects the construction here to the free-energy formulations used in neuroscience and inference.

22 GFF: motivation, dual-layer structure, and Structure Theorems

Remark 22.1 (Reader’s bridge: free energy as the universe’s preferred landscape). A graduate student in physics or chemistry already knows that thermodynamic systems do not seek the minimum of

energy; they seek the minimum of *free energy* — the combination $F = U - TS$ of internal energy U and entropic content TS at temperature T . Water freezes at 0°C not because the ice has lower energy but because the free energy of the ice phase is lower than that of the liquid phase, at that temperature. Generally speaking, every observable phase of matter is a local minimum of free energy in some appropriate variable. GFF promotes this familiar idea to the foundational level: the boundary moduli space $\mathcal{M}_{\text{bdy}}$ carries a free-energy functional whose minima — the GAs (Geometric Attractors) — are the empirically realizable phases of the boundary universe. The Standard-Model vacuum is one such attractor; the bottom of a free-energy landscape is where the universe lives. The technical content of this Part is the precise construction of that landscape from the bulk TGG and CCI data, plus the Morse-theoretic structure of its critical points.

This Part develops GFF (the *Geometric Free-Energy Functional*), the third of the four Gi-4 Pillars. Like TGG (Part II) and CCI (Part III), GFF is intrinsically a *dual-layer* object: it is simultaneously

- (i) an **algebraic/topological functional** — a real-valued function on the stratified boundary moduli space $\mathcal{M}_{\text{bdy}}$ obtained by bulk-to-boundary descent of the bulk TGG action and the CCI 4-form X_4 , whose Morse-theoretic structure encodes the catalog of Geometric Attractors (GA) on each stratum and whose strict Morse-positivity is forced by the integrality $[X_4] \in H^4(M_6; \mathbb{Z})$ of CCI; and
- (ii) a **physical/thermodynamic functional** — the boundary Helmholtz-type free energy

$$F_{\text{GFF}}(t) = U_{\text{GFF}}(t) - T S_{\text{GFF}}(t), \quad \frac{dF_{\text{GFF}}}{dt} \leq 0, \quad (113)$$

satisfying the monotonic decrease condition under the boundary GDOS flow (a covariant Lyapunov inequality). The minima of F_{GFF} are the empirically realizable equilibrium phases of the boundary world; the cascade $F_{\text{GFF}} \rightarrow \min$ is what drives the universe to the Standard-Model vacuum.

The two layers are not independent: they are linked by the *Geometric First Law* (introduced below in Section 25), which identifies $\mathcal{F}_{\text{GEP}}$ as the bulk-descended internal energy contribution and $\mathcal{F}_{\text{CEP}}$ as the entropic contribution, and shows that the algebraic Morse cascade of **GFF-ST5** is the thermodynamic free-energy minimization of (113). In short:

$$\boxed{\underbrace{a^+(\Sigma) - a^-(\Sigma) > 0}_{\text{algebraic Morse-positivity}} \xrightarrow{\text{Geometric First Law}} \underbrace{dF_{\text{GFF}}/dt \leq 0}_{\text{thermodynamic monotonicity}}} \quad (114)$$

The arrow is precisely the Geometric First Law: it carries the algebraic identity to its physical avatar.

Section structure of Part IV. Section 22 (the present section) states the seven GFF axioms (**GFF-ST1** through **GFF-ST7**), organized into the algebraic block (**GFF-ST1–GFF-ST5**) and the physical block (**GFF-ST6–GFF-ST7**); gives the dual-layer formal definition; and lists immediate consequences. Section 23 develops the energetic / curvature-entropic decomposition $\mathcal{F}_{\text{GFF}} = \mathcal{F}_{\text{GEP}} + \mathcal{F}_{\text{CEP}}$. Section 24 describes the stratified boundary moduli space and GFF as a Morse

function on each stratum. Section 25 (new in) develops the physical/thermodynamic layer of GFF: the Geometric First Law, the monotonic decrease $dF_{\text{GFF}}/dt \leq 0$ under the GDOS flow, the H-theorem analog, and laboratory-accessible thermodynamic falsifiers. Section 26 states and proves Theorem **T3**, now in two parts: **T3.1** (algebraic Morse-positivity from CCI) and **T3.2** (thermodynamic monotonicity from GDOS). Section 27 reduces hypotheses $\mathcal{H}_1$ and $\mathcal{H}_2$ from the $N=3$ proof of Arisaka (2026, P2-N3) to theorems of Gi-4, modulo the η -invariant computation M6. Section 28 contains worked examples on the family-projector moduli, the Higgs moduli, the gauge-orbit moduli, and the Wilson-line moduli.

Remark 22.2 (Physics intuition: GFF as the GG analog of the Helmholtz free energy). Readers familiar with statistical mechanics will find an immediate analogy. In thermodynamics, the Helmholtz free energy $F = U - TS$ is the unique state function whose minimization governs equilibrium for a system in contact with a heat bath at fixed temperature T . Its derivative $dF/dt \leq 0$ along any spontaneous trajectory is the second law of thermodynamics.

In GG-Theory, GFF plays the same logical role for the boundary moduli space. $\mathcal{F}_{\text{GEP}}$ is the analog of internal energy U : it is the bulk-descended TGG action evaluated on the boundary slice, with coefficients fixed by the OGN-5 ledger. $\mathcal{F}_{\text{CEP}}$ is the analog of $-TS$: it is the entropic contribution from the Riemannian volume of admissible projector classes, with the “temperature” role played by the scale M_{KK} . The minimization $dF_{\text{GFF}}/dt \leq 0$ along any GDOS trajectory is the *covariant* second law. Two things are packed into the word “universal” here and are best kept apart. That *something* descends is a theorem with no hypothesis about the bulk: every continuous flow on a compact state space carries a complete Lyapunov function (Conley, 1978), and every Markov dynamics with a stationary state carries the relative entropy to it (Lindblad, 1975, Spohn, 1978). That the function descending on the boundary is F_{GFF} , the Helmholtz form of the bulk-descended action, is Theorem **T3.2** under its hypothesis; and that one height is descended across the layers of the cascade is the cohort’s wager, fenced in the introduction and not a theorem of this paper.

This is the precise sense in which GFF *generalizes* the Helmholtz free energy from a single thermodynamic system to the entire stratified moduli space of admissible boundary projectors.

Remark 22.3 (Why broader scientists should care about GFF). The free-energy minimization principle $dF/dt \leq 0$ is the *single most-used dynamical principle* across the natural sciences. In chemistry, it governs reaction kinetics. In condensed-matter physics, it governs phase transitions. In statistical mechanics, it governs equilibration. In biology, it governs protein folding, membrane assembly, and signaling network relaxation (Friston, 2010, 2023). In neuroscience, it governs the brain’s predictive-coding dynamics (Friston’s free-energy principle). In information theory, it governs maximum-entropy inference (Jaynes, 1957).

What is genuinely new in GFF is that all of these distinct free-energy principles are *boundary projections* of the same upstream object: the bulk-to-boundary descent of the TGG action and the CCI 4-form to $\mathcal{M}_{\text{bdy}}$. That is the claim, and it should be graded as it stands. The ubiquity of free-energy principles across the sciences has a theorem behind it that owes nothing to this paper: any relaxation with a compact state space or a stationary distribution descends a Lyapunov function, by Conley’s theorem and Lindblad’s, and each field has found its own. What GFF adds is the identification — that these are projections of one bulk-descended height — and that identification

is the cohort’s wager rather than a consistency check; it is stated sharply enough here that it could be shown to fail. Friston’s principle on neural information manifolds is the case in which the wager is most exposed, not a case in which it is confirmed.

22.1 Motivation: from algebraic CCI to dual-layer GFF

The CCI structure of Part III specifies the integral cohomology class $[X_4] \in H^4(M_6; \mathbb{Z})$ associated with the TGG composite connection, the boundary covariant conservation $\nabla_\mu T_{\text{TGG}}^{\mu\nu} = 0$, and the OGN-5 ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$. But CCI alone does not answer two crucial questions about boundary dynamics.

Question 1 (algebraic). CCI tells us *which* 4-form descends to the boundary, but not *which boundary configurations are stable* under that descent. The space of admissible boundary projector classes is enormous; we need a real-valued functional whose critical points are the stable configurations and whose Morse structure organizes them into a cascade. This is the **algebraic content** of GFF.

Question 2 (physical). CCI tells us that the Master Stress Tensor is conserved on the boundary, but not *which equilibrium* the boundary world relaxes to. Empirically, the boundary world is observed in a unique vacuum (the Standard-Model vacuum); we need a thermodynamic functional whose monotonic decrease drives the boundary world to that unique vacuum. This is the **physical content** of GFF.

Why both layers must be in the same Pillar. The two questions are inseparable. The strict Morse-positivity that selects a unique GA (algebraic side) is precisely the condition for the free energy F_{GFF} to have a unique global minimum (physical side) — because every Morse-positivity violation would correspond to a degenerate equilibrium with an unbroken modulus, which would manifest empirically as an unobserved massless boson. Conversely, the empirical absence of unobserved massless bosons constrains the algebraic Morse data via the Goldstone count (see Section 25).

GFF is the formalization of this dual content. It is simultaneously:

- the *algebraic Morse functional* on the stratified boundary moduli space, with strict Morse-positivity inherited from CCI; and
- the *thermodynamic free energy* of the boundary world, with monotonic decrease driving the cascade to the unique GA.

The most important new content of Part IV is that the strict Morse-positivity of (formerly Theorem **T3**) is now *equivalent* to a statement of the second law for the boundary GDOS flow. This is the bridge from algebraic topology to phenomenological thermodynamics.

Remark 22.4 (Physics intuition: why $dF_{\text{GFF}}/dt \leq 0$ is automatic). A graduate student in physics may wonder why the second law $dF_{\text{GFF}}/dt \leq 0$ should hold *automatically* for every GDOS trajectory, without requiring an additional “second law postulate.” The answer is structural: the GDOS flow is a CPTP (completely positive, trace-preserving) map on the boundary state space, and any such map decreases relative entropy by Lindblad’s theorem (Lindblad, 1975). Free-energy decrease is the projection of relative-entropy decrease onto the boundary moduli space.

In other words: the second law of thermodynamics is not a separate postulate of GG-Theory but an automatic consequence of the open-system structure of GDOS. **GFF-ST6** simply makes this

content explicit and gives it operational form.

22.1.1 Historical background (with citations)

The free-energy functional that GFF formalizes is the latest member of a tradition that began with the steam engine. Carnot (1824) introduced the first quantitative argument for an irreversible thermodynamic limit; Clausius (1865) gave the modern formulation of entropy as a state function and stated the second law in its closed form $dS \geq dQ/T$; Boltzmann (1872) provided the statistical-mechanical foundation with his H -theorem, which derived monotone increase of entropy from microscopic dynamics; and Gibbs (1876) introduced the canonical free energy $F = U - TS$ as the central thermodynamic potential of a system in contact with a thermal reservoir.

The extension to open and dissipative systems began with Onsager (1931), whose reciprocal relations $L_{ij} = L_{ji}$ constrain the linear response of irreversible systems and provide the first hint of Morse-like positivity in the fluctuation–dissipation regime. Prigogine (1947) extended the treatment to nonequilibrium steady states with his minimum-entropy-production principle for near-equilibrium open systems; the maximum-entropy-production tradition (Paltridge 1975 and successors) provided the complementary far-from-equilibrium extension. Jaynes (1957) then placed the entire program on information-theoretic footing by deriving the canonical ensemble from maximum Shannon entropy — the first explicit bridge between classical thermodynamics and information theory.

The modern free-energy formulation in computational neuroscience and machine learning is due to Friston (2010, 2023): biological self-organization is interpreted as a gradient flow on variational free energy, with surprise minimization as the universal principle. GFF subsumes and generalizes this tradition. The decomposition $F_{\text{GFF}} = F_{\text{GEP}} + F_{\text{CEP}}$ separates the *energetic* (geometric) part from the *entropic* (information-theoretic) part, and the dual-layer Morse structure inherited from CCI via Hopkins–Singer descent (Hopkins & Singer, 2005) elevates the classical positivity argument to a theorem (Theorem 26.1) — the axiomatic completion of the lineage from Carnot through Onsager and Friston to GFF.

22.2 The seven GFF Structure Theorems (algebraic block + physical block)

The seven GFF Structure Theorems partition naturally into two blocks:

Algebraic block: GFF-ST1 – GFF-ST5 These Structure Theorems fix the domain, decomposition, smoothness, boundedness, and strict Morse-positivity of the functional. They live on the stratified moduli space $\mathcal{M}_{\text{bdy}}$ and are Morse-theoretic in character.

Physical block: GFF-ST6 – GFF-ST7 These Structure Theorems state the monotonic decrease of F_{GFF} under the GDOS flow (the boundary second law) and the equilibrium uniqueness of the global minimum (the boundary Helmholtz principle). They live in real (laboratory) time and are thermodynamically testable.

The two blocks are linked by the Geometric First Law of Section 25.

Algebraic block (GFF-ST1 through GFF-ST5)

Structure Theorem 15 (GFF-ST1: Domain). *The Geometric Free-Energy Functional GFF is a real-valued function on the boundary moduli space $\mathcal{M}_{\text{bdy}}$ of OGGEF:*

$$\mathcal{F}_{\text{GFF}}: \mathcal{M}_{\text{bdy}} \rightarrow \mathbb{R}. \quad (115)$$

The domain $\mathcal{M}_{\text{bdy}}$ is a stratified space: it decomposes as a disjoint union of smooth connected strata, each labeled by the discrete data of the boundary projector class.

Structure Theorem 16 (GFF-ST2: GEP/CEP decomposition). *GFF decomposes as a sum of two distinct contributions:*

$$\mathcal{F}_{\text{GFF}} = \mathcal{F}_{\text{GEP}} + \mathcal{F}_{\text{CEP}}, \quad (116)$$

where $\mathcal{F}_{\text{GEP}}$ is the Geometric Energetic Part (the energetic contribution from bulk-to-boundary descent of the bulk action of TGG) and $\mathcal{F}_{\text{CEP}}$ is the Curvature-Entropic Part (the entropic contribution from the Riemannian volume of projector classes and the diffusion term of GDOS). Both parts are real-valued functions on $\mathcal{M}_{\text{bdy}}$.

Structure Theorem 17 (GFF-ST3: Smoothness and Morse condition). *The functional $\mathcal{F}_{\text{GFF}}$ is smooth on each stratum of $\mathcal{M}_{\text{bdy}}$, and its critical points on each stratum are non-degenerate (i.e., the Hessian $\text{Hess}(\mathcal{F}_{\text{GFF}})$ is non-singular at each critical point). Equivalently, $\mathcal{F}_{\text{GFF}}$ restricted to each stratum is a Morse function in the standard sense ([Milnor, 1963](#)).*

Structure Theorem 18 (GFF-ST4: Boundedness below). *$\mathcal{F}_{\text{GFF}}$ is bounded below on each stratum of $\mathcal{M}_{\text{bdy}}$:*

$$\inf_{\mathcal{M}_{\text{bdy}}} \mathcal{F}_{\text{GFF}} > -\infty \quad (117)$$

on each stratum. The infimum is attained at a finite collection of isolated points (the GA's of that stratum).

Structure Theorem 19 (GFF-ST5: Strict Morse-positivity from CCI). *Let $\Sigma \subset \mathcal{M}_{\text{bdy}}$ be a stratum that is a connected, compact Kähler manifold on which the boundary gauge connections are flat, on which $\mathcal{F}_{\text{GFF}}|_{\Sigma}$ is the moment map of a Hamiltonian torus action with isolated non-degenerate fixed points, and which supports a non-trivial Hopkins–Singer descent of the CCI class $[X_4]$. Then every Morse index of $\mathcal{F}_{\text{GFF}}$ on Σ is even, the alternating Morse-index sum equals the Hirzebruch signature $\sigma(\Sigma)$, and: (necessity) if $\dim_{\mathbb{R}} \Sigma \equiv 2 \pmod{4}$ then $a^+(\Sigma) - a^-(\Sigma) = 0$, so strict Morse-positivity requires $\dim_{\mathbb{R}} \Sigma \equiv 0 \pmod{4}$; (sufficiency) strict Morse-positivity holds precisely when $\sigma(\Sigma) > 0$, which for $\Sigma = \mathbb{CP}^{N-1}$ is equivalent to N_g odd. In that case*

$$a^+(\Sigma) - a^-(\Sigma) > 0, \quad (118)$$

where $a^+(\Sigma) = \#\{c \in \text{Crit}(\mathcal{F}_{\text{GFF}}|_{\Sigma}) : \text{ind}(c) \equiv 0 \pmod{4}\}$ and $a^-(\Sigma) = \#\{c : \text{ind}(c) \equiv 2 \pmod{4}\}$.

Remark 22.5 (Why the hypotheses are all load-bearing). Each clause earns its place, and it is worth seeing why, because the temptation is to state the theorem for every stratum and the statement is

then false. Non-trivial descent alone does not give strict positivity: on $\mathbb{CP}^3$ the first Pontryagin class is $p_1 = 4h^2 \neq 0$, so the descent is non-trivial by the $\mathcal{H}_1$ theorem, yet the critical indices are 0, 2, 4, 6 and $a^+ - a^- = 2 - 2 = 0$ — the computation is carried out in Appendix K. Nor does adding the parity condition suffice: strict positivity is equivalent to $\sigma(\Sigma) > 0$, and a compact Kähler surface of real dimension four can carry a torus action with $\sigma < 0$, in which case the alternating sum is negative rather than merely zero. What the descent buys, in real dimension four and there only, is $\sigma \neq 0$, since $\sigma = p_1[\Sigma]/3$ on a closed oriented four-manifold; the sign is a separate fact about the stratum. Finally, the moment-map clause is what makes the alternating sum a topological invariant at all: for a Morse function not invariant under a torus action the relevant invariant is the Euler characteristic, which is positive for every $\mathbb{CP}^{N-1}$ and would make the parity test vacuous. The physics content of that clause — that the functional the boundary actually delivers is the moment map — is an identification, and it is graded as one.

The five Structure Theorems **GFF-ST1–GFF-ST5** constitute the *algebraic* block of GFF. They define the Morse-theoretic structure on the stratified boundary moduli and provide the catalog of Geometric Attractors via Theorem **T3.1** (proved in Section 26).

Physical block (GFF-ST6 and GFF-ST7)

The physical block translates the algebraic Morse structure of GFF into operational, observer-accessible thermodynamic form. It consists of two Structure Theorems.

Structure Theorem 20 (GFF-ST6: Thermodynamic monotonicity (boundary second law)). *Let $F_{\text{GFF}}: \mathcal{M}_{\text{bdy}} \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ be the time-dependent restriction of GFF along a GDOS trajectory $t \mapsto [\Pi(t)] \in \mathcal{M}_{\text{bdy}}$ (where $t \in \mathbb{R}_{\geq 0}$ is the boundary laboratory time). Then F_{GFF} admits a Helmholtz-type decomposition*

$$F_{\text{GFF}}([\Pi(t)]) = U_{\text{GFF}}([\Pi(t)]) - T_{\text{GFF}} S_{\text{GFF}}([\Pi(t)]), \quad (119)$$

where:

- $U_{\text{GFF}} = \mathcal{F}_{\text{GEP}}$ is the boundary internal energy (bulk-to-boundary descent of the TGG action);
- S_{GFF} is the boundary entropy (von Neumann entropy of the boundary projector class, weighted by the Riemannian volume form of **GFF-ST2**);
- $T_{\text{GFF}} \sim M_{\text{KK}}$ is the boundary effective temperature (set by the scale of TGG-ST2);

and along any GDOS trajectory the free-energy time derivative is non-positive:

$$\boxed{\frac{dF_{\text{GFF}}}{dt} \leq 0}, \quad (120)$$

with equality if and only if $[\Pi(t)]$ is a critical point of $\mathcal{F}_{\text{GFF}}$. This is the boundary covariant form of the second law of thermodynamics.

Remark 22.6 (The ledger as the meeting of two endowments — candidate reading). The Helmholtz form (119) is the boundary face of two endowments of constitutional time (Part IX of the CPT

member paper, [Arisaka, 2026, P5-CPT](#)): $U_{\text{GFF}} = \mathcal{F}_{\text{GEP}}$ descends from the repeating clock — the Killing flow whose conserved charge is energy — while S_{GFF} accumulates under the arrow — the dictionary’s one-way discard. The GFF ledger is where the second and third endowments meet and trade at the rate T_{GFF} , and the split $\mathcal{F}_{\text{GFF}} = \mathcal{F}_{\text{GEP}} + \mathcal{F}_{\text{CEP}}$ of **GFF-ST2** is the same meeting read at the level of the functional: **GEP** the energetic face of the repeating clock, **CEP** the entropic face of the arrow. Grade: candidate reading; the structure theorems above stand without it.

Lemma 22.7 (DE-1: vacuum-flux freezing at the Geometric Attractor). *At a Geometric Attractor endpoint of the GRIP cascade (a critical point of $\mathcal{F}_{\text{GFF}}$ in the sense of **GFF-ST6**), the energy export of the frozen vacuum sector is individually conserved, $d\rho_\Lambda/dt = 0$, with no compensating exchange against the matter, radiation, or gauge sectors, at the leading order at which the four-way sector split of the BI-6 closure is exact.*

Proof. By **GFF-ST6**, $dF_{\text{GFF}}/dt \leq 0$ along any GDOS trajectory, with equality if and only if $[\Pi(t)]$ is a critical point; at the Geometric Attractor the projector class is therefore stationary, $d[\Pi]/dt = 0$. The vacuum export is a functional of the state — ρ_Λ is $\mathcal{F}_{\text{GFF}}$ evaluated at the endpoint — and is hence time-independent. Any net energy flux between the vacuum sector and another sector is a change of the state decomposition and so requires $d[\Pi]/dt \neq 0$, contradicting stationarity; and by the global four-current closure (T2.2, closure of X_4) there exists no external reservoir against which a compensating flux could run. Hence each individual sector flux vanishes. The qualifier is inherited from the sector split itself: the four-way decomposition holds at leading order in the KK thresholds (open problems M1/M3), and the lemma holds at the same order. ■

The lemma upgrades the informal reading “the attractor’s export is frozen” from prose to a consequence of **GFF-ST6** and the BI-6 closure. Its cosmological payoff is drawn in the companion GG-A10-Cosmos ([Arisaka, 2026, A10-Cosmos](#)): combined with the no-phantom theorem there, it pins the dark-energy equation of state at $w = -1$ exactly (residual deviation bounded by $(H_0/M_{\text{lock}})^2 \sim 10^{-91}$), so that any confirmed $w \neq -1$ — in either direction — refutes the fixed-point reading outright. A topological mechanism beneath the same freezing — a Dirac superselection filter carried by the descent’s inflow term — is derived in Corollary 23.6 (§23.4).

Structure Theorem 21 (**GFF-ST7**: Equilibrium uniqueness and laboratory falsifier). *Among all admissible boundary projector classes $[\Pi] \in \mathcal{M}_{\text{bdy}}$ accessible from a generic initial condition under the GDOS flow of **GFF-ST6**, there exists a unique global minimum $[\Pi_*]$ of F_{GFF} in the asymptotic limit $t \rightarrow \infty$:*

$$[\Pi_*] = \arg \min_{[\Pi] \in \mathcal{M}_{\text{bdy}}} F_{\text{GFF}}([\Pi]). \quad (121)$$

This equilibrium configuration is empirically identified with the Standard-Model vacuum and obeys two universal laboratory-testable relations:

- (1) (Equipartition relation.) *At thermal equilibrium near the global minimum, the boundary effective temperature satisfies $T_{\text{GFF}} dS_{\text{GFF}} = dU_{\text{GFF}}$, i.e., the boundary Helmholtz first law holds locally;*
- (2) (Fluctuation-dissipation relation.) *The covariance of free-energy fluctuations and the linear*

response of F_{GFF} to external boundary perturbations are linked by an Einstein-type relation $\langle \delta F_{\text{GFF}}^2 \rangle = T_{\text{GFF}} \chi_{\text{GFF}}$, where χ_{GFF} is the boundary susceptibility.

A laboratory observation of an empirical violation of either (1) or (2) at scales below M_{KK} would constitute a falsification of GFF and hence of the entire GG-6 framework.

In the vocabulary introduced in Definition 31.4 below, the unique global minimizer $[\Pi_*]$ of **GFF-ST7** is the Ground Geometric Attractor (GGA) of the boundary moduli space taken whole. The statement every application actually uses is the per-stratum one, and it is Corollary 25.7 below; the reader should keep the two apart, because $\mathcal{M}_{\text{bdy}}$ is a stratified space and its global minimizer is not the minimizer of each stratum.

The seven Structure Theorems **GFF-ST1–GFF-ST7** together constitute the full *Geometric Free-Energy Functional*. Axioms **GFF-ST1–GFF-ST4** are essentially the standard free-energy structure on a stratified moduli space. The new content that distinguishes GFF from a generic free energy is twofold: **GFF-ST5** links the strict Morse-positivity to CCI's non-trivial cohomology class (algebraic), and **GFF-ST6–GFF-ST7** link the cascade dynamics to the boundary second law of thermodynamics with a unique global minimum (physical). The two halves are unified by Theorem **T3** in Section 26.

22.3 Dual-layer formal definition of GFF

We now consolidate the algebraic and physical blocks into a single dual-layer definition.

Definition 22.8 (GFF: the Geometric Free-Energy Functional, dual-layer form). The *Geometric Free-Energy Functional* on the boundary moduli space $\mathcal{M}_{\text{bdy}}$ associated with the TGG structure $(P, \mathcal{A}_{\text{tot}}, \mathcal{R}_{\text{tot}})$ and the CCI structure $(X_4, [X_4], \alpha_R, \alpha_5, \alpha_B, \alpha_{Q_i})$ is the **septuple**

$$\text{GFF} = \left(\underbrace{\mathcal{F}_{\text{GFF}}, \mathcal{F}_{\text{GEP}}, \mathcal{F}_{\text{CEP}}, \text{Crit}(\mathcal{F}_{\text{GFF}}), a^+(\Sigma) - a^-(\Sigma) > 0}_{\text{algebraic block (GFF-ST1–A5)}}, \underbrace{F_{\text{GFF}} = U_{\text{GFF}} - T_{\text{GFF}} S_{\text{GFF}}, [\Pi_*]}_{\text{physical block (GFF-ST6–A7)}} \right), \quad (122)$$

where the seven components correspond to the Structure Theorems **GFF-ST1–GFF-ST7** respectively.

The algebraic part $\mathcal{F}_{\text{GFF}}: \mathcal{M}_{\text{bdy}} \rightarrow \mathbb{R}$ is given by the bulk-to-boundary descent of the bulk action of TGG, decomposed as

$$\mathcal{F}_{\text{GFF}}([\Pi]) = \mathcal{F}_{\text{GEP}}([\Pi]) + \mathcal{F}_{\text{CEP}}([\Pi]), \quad [\Pi] \in \mathcal{M}_{\text{bdy}}. \quad (123)$$

The physical part $F_{\text{GFF}}([\Pi(t)]) = U_{\text{GFF}} - T_{\text{GFF}} S_{\text{GFF}}$ satisfies the boundary second law $dF_{\text{GFF}}/dt \leq 0$ along every GDOS trajectory and converges to a unique global minimum $[\Pi_*] \in \mathcal{M}_{\text{bdy}}$. The two parts are linked by the Geometric First Law (Theorem 25.1 of Section 25), which identifies $U_{\text{GFF}} = \mathcal{F}_{\text{GEP}}$, $-T_{\text{GFF}} S_{\text{GFF}} = \mathcal{F}_{\text{CEP}}$, and $dF_{\text{GFF}}/dt \leq 0$ as the gradient flow of $\mathcal{F}_{\text{GFF}}$.

Remark 22.9 (Logical content of GFF). The intuitive content of Definition 22.8 is twofold:

- **Algebraic side.** $\mathcal{F}_{\text{GFF}}$ is a Morse function on the stratified boundary moduli space whose critical points are Geometric Attractors and whose strict alternating Morse-index sum is positive

on every admissible stratum (forced by the integrality of the CCI class $[X_4]$). This guarantees the existence of a finite catalog of stable boundary phases.

- **Physical side.** F_{GFF} is a Helmholtz-type free energy on the boundary world that satisfies the second law of thermodynamics along every GDOS trajectory and converges to a unique global minimum (the Standard-Model vacuum). This guarantees the empirical observability of the equilibrium and provides laboratory-testable equipartition and fluctuation-dissipation relations.

The two sides are linked by the Geometric First Law: the same functional $\mathcal{F}_{\text{GFF}}$ that organizes the algebraic Morse cascade *is* the thermodynamic free energy whose minimization drives the boundary world to equilibrium. This is the duality of GFF.

Remark 22.10 (Open-system character of GFF). The functional $\mathcal{F}_{\text{GFF}}$ is fundamentally an *open-system* object: it is defined on the boundary moduli of the bulk-to-boundary projection, not on the bulk itself. In closed-system QFT, no analog of GFF exists — the closed-system state space has no preferred discrete cascade structure, no irreversible flow, no boundary-projector class, and no thermodynamic temperature. The existence of GFF as a dual-layer object (Morse-positive + thermodynamically monotonic) is one of the principal pieces of evidence that the open-system (OGGEP) framework provides foundational content not available in closed-system QFT.

Remark 22.11 (Physics intuition: why both layers must be present). A physics graduate student may at first think that the algebraic side (Morse-positivity) is enough — after all, it gives a finite catalog of stable boundary phases. But the algebraic side alone does not specify *which* of those phases the boundary world actually relaxes to. Without the physical side (the second law), the boundary world could in principle remain stuck in any local minimum, including unphysical ones. The physical side (**GFF-ST6–GFF-ST7**) is what guarantees that the *global* minimum is the universally selected equilibrium.

Conversely, the physical side alone (a free energy with monotonic decrease) does not specify the structure of the moduli space or the cascade chain. Without the algebraic side, one would have only a phenomenological free energy with no upstream geometric content.

The dual-layer structure of GFF is therefore not redundant — both layers are essential, and the Geometric First Law is what makes them consistent.

22.4 First consequences of the GFF Structure Theorems

We list four immediate consequences of **GFF-ST1–GFF-ST5**.

Proposition 22.12 (Existence of GA catalog on each stratum). *On each stratum Σ of $\mathcal{M}_{\text{bdy}}$, the set of local minima of $\mathcal{F}_{\text{GFF}}$ is a finite, isolated, non-empty collection of points, called the GA catalog of Σ .*

Proof. By **GFF-ST3**, $\mathcal{F}_{\text{GFF}}|_{\Sigma}$ is Morse, so its critical points are isolated. By **GFF-ST4**, $\mathcal{F}_{\text{GFF}}$ is bounded below on Σ , so its infimum is attained – at least one local minimum exists. Compactness of each stratum (or proper-Morse condition for non-compact strata) ensures only finitely many critical points exist. The local minima form the GA catalog. ■

Proposition 22.13 (Cascade chain on each stratum). *On each stratum Σ , the Morse critical points of $\mathcal{F}_{\text{GFF}}$ admit a cascade chain ordered by Morse index: critical points of index 0 at the bottom (GA's), critical points of higher even index above them (saddles), terminating at the maximum index (the deepest unstable phase). This chain is the upstream object whose gradient flow gives the GRIP cascade of Part V.*

Proof. Standard Morse theory (Milnor, 1963): the gradient flow of $\mathcal{F}_{\text{GFF}}$ descends from higher-index critical points to lower-index ones along instanton trajectories, with the lowest-index critical points (the local minima) being the asymptotic attractors of the flow. ■

Proposition 22.14 (Equivalence of GFF Morse-positivity with topological signature). *On any stratum Σ admitting a Hopkins–Singer descent of $[X_4]$, the strict alternating Morse-index sum $a^+(\Sigma) - a^-(\Sigma)$ is equal to the topological signature $\sigma(\Sigma)$ when Σ has real dimension divisible by four, and is identically zero otherwise.*

Sketch. This is the Morse-theoretic version of Hirzebruch's signature theorem restricted to the stratum. For a closed oriented $4m$ -manifold Σ equipped with a perfect Morse function whose critical points all have even Morse index, the alternating sum $\sum_a (-1)^{\text{ind}_a/2}$ equals the integrated L-genus, which is the signature. For Σ of real dimension $\equiv 2 \pmod{4}$, the signature is identically zero. This is the standard textbook content of Hirzebruch's signature theorem (Hirzebruch, 1966, Milnor, 1963). The key consequence: strict positivity holds if and only if $\dim_{\mathbb{R}} \Sigma \in 4\mathbb{Z}$. ■

Proposition 22.15 (GFF as the formal realization of OGGE). *The existence of $\mathcal{F}_{\text{GFF}}$ satisfying **GFF-ST1–GFF-ST4** is equivalent to the GG-Constitutional axiom OGGE. The strict Morse-positivity of **GFF-ST5** is a derived theorem from CCI via the Hopkins–Singer descent (Theorem 26.1 of Section 26).*

Sketch. OGGE states that observable boundary content is related to bulk geometric content by an irreversible CPTP bulk-to-boundary projection. Equivalently: there exists a boundary moduli space $\mathcal{M}_{\text{bdy}}$, a free-energy functional $\mathcal{F}_{\text{GFF}}$ on it that drives the cascade, and the cascade closes in finite time. These are exactly axioms **GFF-ST1–GFF-ST4**. Conversely, given $\mathcal{F}_{\text{GFF}}$ satisfying these axioms, one recovers the OGGE CPTP-projection structure via the gradient flow. Theorem 26.1 establishes **GFF-ST5**. ■

23 The GEP/CEP decomposition

We now develop the two parts of GFF in detail: the energetic part $\mathcal{F}_{\text{GEP}}$ and the curvature-entropic part $\mathcal{F}_{\text{CEP}}$.

23.1 The energetic part GEP

Definition 23.1 (Geometric Energetic Part GEP). The Geometric Energetic Part of GFF is the contribution to $\mathcal{F}_{\text{GFF}}$ obtained by bulk-to-boundary descent of the bulk action of TGG. In the

canonical bulk action with two-derivative kinetic terms,

$$\mathcal{F}_{\text{GEP}}([\Pi]) = \frac{1}{8\pi^2} \int_{\Pi[M_4]} [\alpha_R \text{tr} R_{\text{Spin}(1,5)}^2 + \alpha_5 \text{Tr} F_5^2 + \alpha_B F_B^2 + \alpha_{Qi} F_{Qi}^2], \quad (124)$$

where the integration is over the boundary slice $\Pi[M_4]$ of the projector Π , and the coefficients $(\alpha_R, \alpha_5, \alpha_B, \alpha_{Qi}) = (+1, -1, -2/5, -2/5)$ are inherited from **CCI-ST2**.

The GEP integrand is exactly the four-form X_4 of CCI. Therefore $\mathcal{F}_{\text{GEP}}$ is the integrated X_4 on the boundary slice:

$$\mathcal{F}_{\text{GEP}}([\Pi]) = \frac{1}{8\pi^2} \int_{\Pi[M_4]} X_4. \quad (125)$$

This is the cleanest formulation: the GEP is, up to an overall normalization, the boundary integral of the CCI four-form.

23.2 The curvature-entropic part CEP

Definition 23.2 (Curvature-Entropic Part CEP). The Curvature-Entropic Part of GFF is the entropic contribution arising from the Riemannian volume of admissible boundary projector classes plus the diffusion term of GDOS. Schematically,

$$\mathcal{F}_{\text{CEP}}([\Pi]) = -T \mathcal{S}([\Pi]) + \mathcal{F}_{\text{diff}}([\Pi]), \quad (126)$$

where T is an effective “temperature” parameter (set by the bulk scale), $\mathcal{S}$ is the configurational entropy of the projector class $[\Pi]$, and $\mathcal{F}_{\text{diff}}$ is the diffusion-term contribution from the boundary CPTP map.

Remark 23.3 (Why CEP is needed). A pure energetic functional $\mathcal{F}_{\text{GEP}}$ alone would give critical points only at the strict zeros of X_4 on the boundary – typically just the trivial configuration. The entropic contribution $\mathcal{F}_{\text{CEP}}$ is what creates the rich Morse landscape with multiple critical points of varying Morse index. This is the open-system content: in a closed-system QFT, there is no entropic contribution to the action, and the Morse cascade structure of GFF would not exist.

23.3 Bulk-to-boundary descent of the bulk action

The construction of $\mathcal{F}_{\text{GEP}}$ from the bulk action of TGG proceeds as follows.

Step 1: bulk action. The bulk action of TGG on M_6 is

$$S_{\text{bulk}} = \frac{1}{16\pi^2} \int_{M_6} X_4 \wedge \omega_2, \quad (127)$$

where ω_2 is a fixed bulk 2-form on the radial-angular sector. Its overall normalization is not a free convention: POS-6 of GG-A1-Constitution (Arisaka, 2026, A1-Constitution) (OGN: no continuous rescaling dials; “conventions are not allowed to function as hidden continuous parameters”) requires it to be fixed, and Lemma 23.4 below fixes it.

Step 2: boundary projection. Given a boundary projector $\Pi : M_6 \rightarrow M_4$, the bulk integral

splits into a boundary integral plus a transverse part. The boundary integral is

$$\int_{\Pi[M_4]} X_4, \quad (128)$$

which is exactly the boundary integral of (125).

Step 3: variation. Varying $\mathcal{F}_{\text{GEP}}$ with respect to the projector $[\Pi]$ gives the equations of motion of the boundary theory. Critical points (stationary projector configurations) are the GEP contribution to the GA catalog.

Step 4: addition of CEP. The full GFF is the sum of $\mathcal{F}_{\text{GEP}}$ and $\mathcal{F}_{\text{CEP}}$ as in (116). Critical points of the full functional are the GA's in the strict sense.

23.4 The normalization of the descent: $\Phi_\omega = \pi$ and $S_{\text{bulk}} = \frac{\pi}{2} \mathcal{F}_{\text{GEP}}$

With X_4 pulled back from the boundary slice (Step 2 above), the bulk integral of (127) factorizes,

$$S_{\text{bulk}} = \frac{1}{16\pi^2} \left(\int_{\Pi[M_4]} X_4 \right) \Phi_\omega = \frac{\Phi_\omega}{2} \mathcal{F}_{\text{GEP}}, \quad \Phi_\omega := \int_{I_K \times S_\theta^1} \omega_2, \quad (129)$$

using (125). The entire normalization freedom of (127) is therefore a single transverse flux number Φ_ω , which POS-6 forbids to remain a convention.

Lemma 23.4 (Descent normalization: $\Phi_\omega = \pi$). *With ω_2 the warp-weighted volume element of the radial-angular sector, $\omega_2 = e^{-2r} dr \wedge d\theta$ — the same transverse weight through which the boundary sees bulk gravity —*

$$\Phi_\omega = \int_0^K e^{-2r} dr \int_0^{2\pi} d\theta = \pi (1 - e^{-2K}) \xrightarrow{K=35} \pi \quad (\text{to one part in } 10^{31}), \quad (130)$$

and hence

$$S_{\text{bulk}} = \frac{\pi}{2} \mathcal{F}_{\text{GEP}} \quad (131)$$

exactly, up to the 10^{-31} warp tail: one unit of $\mathcal{F}_{\text{GEP}}$ carries bulk action $\pi/2$.

Proof-status note. The identification of ω_2 with the warp-weighted element is the physical content of the lemma, and it is supported twice over, both times upstream. (i) *The graviton pin:* the transverse integral in (130) is precisely the radial-angular window that normalizes boundary gravity in GG-A5-GDOS (Arisaka, 2026, A5-GDOS) (its Lemma T8-1: $I_r^{\text{curv}} \cdot I_\theta = \frac{1}{2}(1 - e^{-2K}) \cdot 2\pi$, the π prefactor of M_{Pl}^2 that GG-A5-GDOS records as “not an arbitrary numerical coefficient”). (ii) *The cascade concordance below.* The flat (unwarped) candidate $\int dr \wedge d\theta = 2\pi K = 70\pi$ is excluded by the same concordance. The assignment of inflow charge to the cascade saddles is derived under three declared premises (Remark 23.5); the wrap-independence premise (P-i) stated there is the named residual hypothesis.

Remark 23.5 (Concordance with the cascade actions). Under the natural charge assignment for the family-chain saddle $[2k] \rightarrow [2(k-1)]$ — a second-Chern jump $\Delta c_2 = k$, entering $X_4/8\pi^2$ with coefficient -2 (Theorem 14.2) — the inflow term (131) evaluates on the saddle to $\Delta S_{\text{bulk}} = \frac{\pi}{2} \cdot 2k = \pi k$:

the toll of the gate, and the number T4.1 reaches by its Duistermaat–Heckman route (π per Fubini–Study projective line) when the height is taken to be the pair invariant. Two independently computed π ’s coincide — the graviton-normalizing transverse volume of GG-A5-GDOS Lemma T8-1 and the inflow unit fixed here; the third, the Fubini–Study area of one cascade step, coincides with them only under the identification of the height with the pair invariant, which Theorem 24.1 excludes under the transport proposition (Remark 32.2). A flat normalization $\Phi_\omega = 2\pi K$ would instead demand tolls $70\pi k$, contradicting the inflow unit. The charge assignment $\Delta c_2 = k$ is verified at the structural level: the corpus’s own toll arithmetic *excludes* unit-spaced Morse weights — they would leave the 3π budget unsaturated at 2π and admit a fourth generation at exactly 3π — and forces the pair-count weights $\lambda_\ell = \binom{\ell}{2}$ of Theorem T4.1’s Step 2, whose jumps are exactly k . That identification is closed: under three declared premises the pair count is *derived*. (P-i) *Wrap independence*: the winding- ℓ generation’s flux bundle is the ℓ -fold Whitney sum of the OGN-unit flux line bundle — ℓ independent wraps of the topological counter; a pair of wraps is a pair of **5**-quanta, i.e. a state of the $\mathbf{10} = \Lambda^2(\mathbf{5})$, which is how the fermionic independence of wraps is realized in the spectrum. (P-ii) *Unit self-pairing*: the unit flux class satisfies $\int \beta \wedge \beta = 1$ on the projected slice (the OGN normalization of the index background). (P-iii) *Background independence*: the one-unit Abelian flux is index background, not projector dressing, so the Abelian slot of X_4 carries no ℓ -dependence — and this is forced, not chosen: an Abelian slot scaling as ℓ^2 would break the toll arithmetic itself. The Whitney formula then gives $c(E_\ell) = (1 + \beta)^\ell$, hence $c_2(E_\ell) = \binom{\ell}{2} \beta^2$ and $\Delta c_2 = k$ exactly; a blind uniqueness test confirms that $\lambda_\ell = \binom{\ell}{2} + \text{const}$ is the *only* weight family compatible with the tolls — a statement about the inflow charges, which is what the tolls are, and not about the Morse heights, which Theorem 24.1 fixes separately. The two π ’s of this remark are thereby a derivation of the toll, not a concordance check; the premise a hostile reader should attack is (P-i), whose full 6D zero-mode formalization is the natural further tightening.

Corollary 23.6 (Dirac superselection of the topological background). *In Lorentzian signature the inflow term contributes relative phases $e^{i\Delta S_{\text{bulk}}} = e^{i(\pi/2)\Delta\mathcal{F}_{\text{GEP}}}$ between topological sectors. Only jumps $\Delta\mathcal{F}_{\text{GEP}} \in 4\mathbb{Z}$ are phase-trivial; sectors separated by any other jump admit no gauge-invariant connecting operator and are superselected. In particular, by the Hirzebruch signature theorem $p_1 = 3\sigma$ (Hirzebruch, 1966), the minimal gravity-sector jump is $\Delta p_1 = 3$, i.e. $\Delta\mathcal{F}_{\text{GEP}} = 3$, carrying phase $e^{3i\pi/2}$: the closed-packet background of GG-A10-Cosmos (Arisaka, 2026, A10-Cosmos) is absolutely superselected — it cannot change by any physical process. This places a topological mechanism beneath the thermodynamic freezing of Lemma 22.7 (DE-1); the cascade weights $e^{-\pi k}$, being real, are unaffected. Downstream consumers: the unit-packet normalization program (VIII-b'') audited in GG-A10-Cosmos, and the GSM adjunct of GG-A5-GDOS (Arisaka, 2026, A5-GDOS).*

23.5 The Riemannian volume on projector classes

The configurational entropy $\mathcal{S}$ in $\mathcal{F}_{\text{CEP}}$ is defined as the logarithm of the Riemannian volume of the projector class:

$$\mathcal{S}([\Pi]) = \log \text{Vol}_{\text{FS}}([\Pi]), \quad (132)$$

where Vol_{FS} is the Fubini–Study volume on the projector moduli (or its analog on other strata).

For the family-projector moduli stratum $\mathbb{CP}^{N-1}$, the Fubini–Study volume is exactly the standard Fubini–Study Kähler volume, and the entropy contribution is finite and computable. This is one of the reasons that the family-projector moduli stratum is the most directly amenable to explicit Morse-theoretic analysis: its Riemannian structure is canonical, and the Morse function we use (the height function from TGG) is the canonical Kähler–Hamiltonian on the same metric.

24 The stratified boundary moduli and Morse structure

We now describe the structure of the boundary moduli space $\mathcal{M}_{\text{bdy}}$ on which GFF lives.

24.1 The stratified structure of the boundary moduli

The boundary moduli space $\mathcal{M}_{\text{bdy}}$ is, in the 6D bulk of GG-Theory, a stratified space whose strata correspond to different sectors of admissible boundary projector configurations. The principal strata are:

- **Family-projector stratum** $\Sigma_{\text{fam}} = \mathbb{CP}^{N-1}$: the boundary moduli of the family direction. Real dimension $2(N_g - 1)$.
- **Higgs-moduli stratum** Σ_H : the boundary moduli of the Wilson–Hosotani Higgs configuration. Real dimension 4 for the Standard-Model Higgs doublet.
- **Gauge-orbit stratum** Σ_{gauge} : the boundary moduli of the orbifold-broken gauge configuration on θ . This is a discrete stratum (the orbifold-action conjugacy classes).
- **Radial-scale stratum** Σ_r : the boundary moduli of the radial scale profile (modulating $M_{\text{lock}} = E_{\text{Pl}}/2e^{35}$). One-dimensional.
- **Compact-angle stratum** Σ_θ : the boundary moduli of the compact θ -direction phases. One-dimensional and periodic.

The full boundary moduli space is a product (or fiber-product) of these strata, plus possible discrete labels:

$$\mathcal{M}_{\text{bdy}} \cong \Sigma_{\text{fam}} \times \Sigma_H \times \Sigma_{\text{gauge}} \times \Sigma_r \times \Sigma_\theta. \quad (133)$$

24.2 GFF as a Morse function on each stratum

By **GFF-ST3**, $\mathcal{F}_{\text{GFF}}$ is Morse on each stratum. We record the principal Morse data on each stratum.

- $\Sigma_{\text{fam}} = \mathbb{CP}^{N-1}$: the Morse function is the Fubini–Study height function (Section 23.1 of TGG examples). Critical points: N points of indices $0, 2, 4, \dots, 2(N-1)$. All indices even. This is the stratum that drives the $N=3$ proof.
- Σ_H : the Morse function on the Higgs moduli is the Coleman–Weinberg / Wilson–Hosotani potential, with critical points at the symmetric and broken-phase Higgs vacua.
- Σ_{gauge} : discrete; the Morse data is the list of admissible orbifold breakings.

- Σ_r : the Morse function is the radial scale potential, with the unique critical point at M_{lock} .
- Σ_θ : the Morse function is the $U(1)_B \times U(1)_{Q_i}$ Wilson-line potential, with critical points at the discrete OGN-5 phases.

24.3 Family-projector moduli stratum: the principal example

The family-projector moduli $\Sigma_{\text{fam}} = \mathbb{CP}^{N-1}$ is the most important stratum for the present paper because it is the stratum on which the N=3 proof of Arisaka (2026, P2-N3) computes strict GDOS Morse-positivity.

The Morse function on $\mathbb{CP}^{N-1}$ is, by Definition 23.1 restricted to the family-projector data, the standard Fubini–Study height function

$$f_{\text{GFF}}([\psi]) = \frac{\sum_{a=1}^N \lambda_a |\psi_a|^2}{\sum_a |\psi_a|^2}, \quad (134)$$

with $\lambda_1 < \dots < \lambda_N$ real coefficients inherited from the diagonal mass-overlap data of the boundary projector.

The Morse data on $\mathbb{CP}^{N-1}$ is:

- N_g critical points c_a at $[0 : \dots : 0 : 1 : 0 : \dots : 0]$ (1 in position a), $a = 1, \dots, N$;
- Morse index $\text{ind}(c_a) = 2(a - 1)$;
- All critical points non-degenerate (perfect Morse function).

The coefficients λ_a are not free, and the sentence above that leaves them as “mass-overlap data” is completed here. Two statements this paper makes about the descent fix them up to a scale and a shift, and what they fix them to is not what the instanton theorem of Part V assumed.

Theorem 24.1 (The family-stratum height is the dilatation charge). *Let $\Sigma_{\text{fam}} = \mathbb{P}(V) = \mathbb{CP}^{N-1}$, with V the span of the surviving winding modes $\ell = 1, \dots, N_g$ of the compact circle, and suppose:*

- the GRIP-In descent on Σ_{fam} is the boundary projection of the dilatation along I_r (Proposition 31.8), under which the amplitude of the winding- ℓ mode scales as $e^{-\Delta_+(\ell)s}$ on the projective stratum, where $\Delta_+(\ell) = 2 + \sqrt{4 + \ell^2}$ is the scaling dimension of the massless radial equation of Arisaka (2026, A2-GG6) with the winding as the five-dimensional mass (the decaying branch is what an observer at the infrared end sees; the growing branch gives the opposite weights and exchanges floor and top);*
- the descent is the gradient flow of $f_{\text{GFF}}|_{\Sigma_{\text{fam}}}$ in a Kähler metric of the stratum’s complex structure (**GRIP-ST2**, with the Fubini–Study metric named above).*

Then:

- $f_{\text{GFF}}|_{\Sigma_{\text{fam}}} = c \langle D \rangle + \text{const}$, with $\langle D \rangle = \sum_\ell \Delta_+(\ell) |\psi_\ell|^2 / \|\psi\|^2$ the Hamiltonian of the dilatation direction $\xi = (\Delta_+(1), \dots, \Delta_+(N_g))$ of the diagonal torus, and $c = A/(2\pi)$ for a projective line of Fubini–Study area A ($c = \frac{1}{2}$ at the area π of Theorem 32.1, $c = 1$ for the doubled form in which the dilation runs at unit rate). The direction is irrational — $\sqrt{5}$, $\sqrt{8}$, $\sqrt{13}$ are rationally independent — so its flow closes on no circle and its closure is the full torus, whose fixed points*

are the coordinate points; the isolation of those fixed points is the statement that only one sign of winding survives on the cell, an input of this paper and not a theorem of the anomaly paper.

- (b) The critical values are $\lambda_\ell = c\Delta_+(\ell) + \text{const.}$ For $N_g = 3$: $\Delta_+ = 4.236, 4.828, 5.606$, gaps 0.592 and 0.777, gap ratio $(\sqrt{13} - \sqrt{8})/(\sqrt{8} - \sqrt{5}) = 1.312$. The ratio is the same for every c , every constant, and every Kähler form on the stratum invariant under the torus, because the critical values of a moment-map component are affine in the direction and $H^2(\mathbb{CP}^{N-1}; \mathbb{R})$ is one-dimensional; and the flow of (ii) does not see the form at all, since the gradient of a moment-map component in any compatible Kähler metric is the same holomorphic vector field $-JX_\xi$.
- (c) No scaling-dimension spectrum reaches the pair invariant. For a bulk field of any six-dimensional mass M above the Breitenlohner–Freedman bound and any circle radius R , the dimensions of three consecutive windings are $2 + \sqrt{u + \ell^2 x^2}$ with $u = 4 + M^2 \geq 0$ and $x = 1/(kR)$, and their gap ratio $r(y) = (\sqrt{1 + 9y} - \sqrt{1 + 4y})/(\sqrt{1 + 4y} - \sqrt{1 + y})$, $y = x^2/u$, lies strictly between 1 and $5/3$ for every $y > 0$ (the program’s value is $y = \frac{1}{4}$). The pair invariant $\binom{\ell}{2} = 0, 1, 3$ of Theorem 32.1, Step 2, has gap ratio 2, and is therefore not the set of critical values of f_{GFF} on Σ_{fam} under (i)–(ii); the holonomy weights ℓ , gap ratio 1, are the limit at the bound and are never attained.

Proof. (a) The action of the cell group on V is by diagonal matrices, hence holomorphic; its torus part is unitary, hence Hamiltonian with the standard moment map. The one-parameter subgroup $\text{diag}(e^{-\Delta_+(\ell)s}) = \exp(-s i\xi)$ of the complexified torus is the gradient flow, in the Fubini–Study metric with line area π , of $\frac{1}{2}\langle D \rangle$, by the general fact that the gradient flow of a moment-map component is the imaginary direction of the complexified action; two functions with the same gradient on a connected manifold differ by a constant, which gives $c = \frac{1}{2}$ at that area and $c = A/(2\pi)$ in general. Rational independence of $\sqrt{5}, \sqrt{8}, \sqrt{13}$ is elementary. (b) A torus-invariant Kähler form on $\mathbb{CP}^{N-1}$ in the class $a[\omega_{\text{FS}}]$ has moment polytope a times the standard simplex up to translation (the polytope is fixed by the equivariant class), so the value of the component μ^ξ at the coordinate point c_ℓ is $a\xi_\ell/2 + b$; the ratio of consecutive differences is free of a and b . For the metric independence of the flow: if F is any function whose gradient flow in a Kähler metric g' compatible with J is the holomorphic flow $-JX_\xi$, then $dF = \omega'(X_\xi, \cdot)$ with $\omega' = g'(J\cdot, \cdot)$, so ω' is closed and invariant under the closure of $\exp(t\xi)$, which is the whole torus, and F is a torus moment-map component of ω' ; (b) then applies to it. (c) With $s(m) = \sqrt{1 + ym}$ at $m = \ell^2$: $r > 1$ is $s(9) + s(1) > 2s(4)$, which after squaring twice is $16y > 0$; $r < 5/3$ is $8s(4) > 5s(1) + 3s(9)$, the strict concavity of s in m at $4 = (5 \cdot 1 + 3 \cdot 9)/8$. Both limits are attained only at the ends: $5/3$ as $y \rightarrow 0$ (the weights ℓ^2) and 1 as $y \rightarrow \infty$ (the weights ℓ). The numerical statements are computed and asserted in the script that accompanies this paper. ■

What the theorem settles, and what it moves. The identification of the Morse height with the pair invariant was made in Step 2 of the instanton theorem (Theorem 32.1) and rested on the concordance of Remark 23.5; under the paper’s own transport proposition it is excluded by (c), not by a choice of normalization but by an invariant no normalization reaches. The pair invariant does not thereby leave the paper: it is the inflow charge $\Delta c_2 = k$ of the winding bundle, derived in

Remark 23.5 from the Whitney pair count, and the inflow toll πk of a gate is a Chern number, not a height. What moves is the claim that the two are one quantity. The consequences are drawn where they fall: in the instanton theorem (the status remark following its proof), in **GRIP-ST3** (the toll is the inflow toll, and the paper keeps the name), in the weights on the skeleton (§33.2), and in open problem M13 of §57.1, which records the one thing the theorem does not decide: whether hypothesis (i) or the reading of the toll as the height difference is the premise to give up. This paper adopts (i), because it is a composition of two results of **GG-A5-GDOS** and the Kähler argument is a theorem, while the identification it excludes was a step in a proof. The three candidates the program’s papers have carried for the family weights — the pair invariant 0, 1, 3 (gap ratio 2), the scaling dimensions (gap ratio 1.31), and the holonomy weights 1, 2, 3 (gap ratio 1, excluded in Remark 23.5 on the toll arithmetic and by (c) as a limit never attained) — are thereby reduced to one, at the grade of (i)–(ii). Which of the three the Yukawa hierarchy actually tracks is a question for **GG-A9-SM**, where the sector log-mass gap ratios at one scale are 0.53, 0.90 and 1.33 for the charged leptons, the up quarks and the down quarks, and no single ladder fits the three; the dilatation ratio 1.31 sits inside the down-quark spread across published compilations, which is interesting, and is not proof.

24.4 Other strata: brief remarks

We mention briefly the Morse data on the other strata; full treatments are deferred to subsequent papers.

Higgs moduli Σ_H : the relevant Morse function is the Wilson–Hosotani / Coleman–Weinberg effective potential (Hosotani, 1983). Its critical points are the symmetric ($V = 0$) configuration and the broken-phase ($V \neq 0$) configuration. The boundary projector selects the broken phase as the GA.

Radial-scale moduli Σ_r : the relevant Morse function is the radial scale potential, with the unique stable critical point at $r = K$ (the IR boundary), corresponding to $M_{\text{lock}} \simeq 3.85$ TeV (Arisaka, 2026, P1-KK).

Compact-angle moduli Σ_θ : the relevant Morse function is the $U(1)_B \times U(1)_{Q_i}$ Wilson-line potential, with critical points at the discrete **OGN-5** phases (e.g., $2\pi M/L = 2\pi/5$ for the $U(1)_B$ phase). On the anomaly paper’s admissible class both Abelian circle holonomies are fixed background data, so whether this stratum is dynamical at all, and which field’s phase it carries, is registered as open (Remark 11.1).

25 The physical layer of **GFF**: thermodynamic monotonicity and laboratory falsifiers

This section develops the *physical* content of **GFF** (Structure Theorems **GFF-ST6** and **GFF-ST7**) in operational detail: the Geometric First Law that links the algebraic decomposition $\mathcal{F}_{\text{GFF}} = \mathcal{F}_{\text{GEP}} + \mathcal{F}_{\text{CEP}}$ to the Helmholtz decomposition $F_{\text{GFF}} = U_{\text{GFF}} - T_{\text{GFF}} S_{\text{GFF}}$, the boundary second law $dF_{\text{GFF}}/dt \leq 0$ as an H-theorem analog for the **GDOS** flow, and laboratory-accessible thermodynamic falsifiers. The mathematics is essentially open-system thermodynamics adapted to the boundary moduli space; the conceptual novelty is the realization that the algebraic Morse

cascade of **GFF-ST5** and the thermodynamic free-energy minimization of **GFF-ST6** are *the same dynamical process viewed from two directions*.

25.1 The Geometric First Law: from algebraic to thermodynamic

Theorem 25.1 (Geometric First Law of GFF). *On any stratum $\Sigma \subset \mathcal{M}_{\text{bdy}}$ supporting a non-trivial Hopkins–Singer descent of $[X_4] \in H^4(M_6; \mathbb{Z})$, there exists a canonical identification between the algebraic decomposition $\mathcal{F}_{\text{GFF}} = \mathcal{F}_{\text{GEP}} + \mathcal{F}_{\text{CEP}}$ of **GFF-ST2** and the Helmholtz decomposition $F_{\text{GFF}} = U_{\text{GFF}} - T_{\text{GFF}} S_{\text{GFF}}$ of **GFF-ST6**:*

$$\boxed{U_{\text{GFF}}([\Pi]) = \mathcal{F}_{\text{GEP}}([\Pi]), \quad -T_{\text{GFF}} S_{\text{GFF}}([\Pi]) = \mathcal{F}_{\text{CEP}}([\Pi]),} \quad (135)$$

where $T_{\text{GFF}} \sim M_{\text{KK}}$ is the boundary effective temperature (set by the scale of **TGG-ST2**) and S_{GFF} is the boundary von Neumann entropy of the projector class $[\Pi]$. The total $F_{\text{GFF}} = \mathcal{F}_{\text{GFF}}$ is a single function of $[\Pi]$ that serves simultaneously as the algebraic Morse functional and the thermodynamic free energy. The identification is made, and used below, at the grade of the linear-response regime near equilibrium (Step 4 of Theorem 25.3); it is exact on a stratum exactly when the stationary state of the boundary generator on that stratum is the Gibbs state $e^{-\mathcal{F}_{\text{GEP}}/T_{\text{GFF}}}/Z$, which no result of this paper yet establishes and which is the one thing that would promote it.

Proof. The energetic-part identification $U_{\text{GFF}} = \mathcal{F}_{\text{GEP}}$ follows directly from the boundary integral (125): $\mathcal{F}_{\text{GEP}}$ is the integrated TGG-action density on the boundary slice, which is precisely the boundary internal energy in the bulk-to-boundary-descent normalization. The entropic-part identification $-T_{\text{GFF}} S_{\text{GFF}} = \mathcal{F}_{\text{CEP}}$ follows from the definition of $\mathcal{F}_{\text{CEP}}$ as the Riemannian volume contribution to the partition-function evaluation (Hopkins & Singer, 2005): the volume term equals $-T_{\text{GFF}}$ times the von Neumann entropy when one identifies the scale M_{KK} as the boundary effective temperature. The total $F_{\text{GFF}} = \mathcal{F}_{\text{GFF}}$ then satisfies both the Morse-theoretic content of **GFF-ST1–GFF-ST5** and the Helmholtz content of **GFF-ST6**. ■

Remark 25.2 (Physics intuition: KK scale as boundary temperature). The identification $T_{\text{GFF}} \sim M_{\text{KK}}$ may at first surprise a physicist used to thermodynamics in flat-space QFT. But in the **GG-6** setup, the KK mass M_{KK} plays exactly the role of a boundary thermal scale: KK modes are exponentially suppressed above this scale, just as Boltzmann modes are suppressed above temperature T . This is the standard *Tolman-like* identification of scale-profile scale with effective temperature in warped extra dimensions (Randall & Sundrum, 1999).

For the **GG-6** prediction $M_{\text{KK}} \simeq 3.85 \text{ TeV}$, this gives $T_{\text{GFF}} \sim 3.85 \text{ TeV}$ as the boundary effective temperature — much hotter than any laboratory scale, and the scale the boundary GDOS flow last registered. It sets nothing in a present-day layer: the selection of the floor $[\Pi_*]$ on an admissible stratum is Frankel’s and Conley’s (Corollary 25.7) and needs no temperature, and the weights on the skeleton are Boltzmann weights in T_{GFF} only in the sense of the freeze-out reading below.

Sharpening (the freeze-out reading). The sense in which this identification holds has now been located at kernel level. While the fiber-moduli block is in kinetic contact with the KK tower, the Einstein ratio $T_{\text{GFF}} = D/\mu$ of the GDOS flow (Arisaka, 2026, A5-GDOS, the Einstein relation in

the proof of its Theorem T5, under the detailed balance its Davies limit supplies) *copies* the bath temperature: the kernel is a thermometer, not a thermostat, and no evaluation of overlap integrals can make it manufacture M_{KK} . At the Geometric Attractor the tower bath is gapped (thresholds at M_{KK} and above) and frozen (Lemma 22.7), so the block is dissipatively disconnected — the same statement as DE-1, re-derived from the kernel side. Contact was therefore last maintained while the pair channel's Boltzmann factor $e^{-2M_{\text{KK}}/T}$ still exceeded the comparison rate, i.e. down to $T^* = 2M_{\text{KK}}/L^*$, with L^* the logarithm of the rate hierarchy at the lock epoch. $T_{\text{GFF}} \sim M_{\text{KK}}$ is thus a *freeze-out record* — the last temperature the block registered before the gap severed contact, preserved unchanged ever since — rather than an equilibrium property of the present boundary; the order-one softness of the identification is the logarithm $2/L^*$, whose value is epoch physics and remains open. One more thing is a convention and is named as one: the unit in which the height $\mathcal{F}_{\text{GFF}}$ is measured against T_{GFF} . The weights on the family skeleton, $e^{-\pi k}$ for the tolls and $e^{-c\Delta\Delta_+}$ for the dilatation heights (Theorem 24.1), are numbers only at $T_{\text{GFF}} = 1$ in height units; nothing in this paper derives that unit, and the freeze-out record above fixes the temperature, not the height's scale.

25.2 The boundary H-theorem: $dF_{\text{GFF}}/dt \leq 0$ from CPTP structure

We now derive the monotonic decrease **GFF-ST6** from the open-system (CPTP) structure of the GDOS flow.

Theorem 25.3 (Boundary H-Theorem for GFF). *Let $t \mapsto [\Pi(t)] \in \mathcal{M}_{\text{bdy}}$ be a GDOS trajectory on the boundary moduli space, generated by a Lindbladian $\mathcal{L}$ acting CPTP on the boundary state space and satisfying the quantum detailed-balance condition of Kossakowski et al. (1977) with respect to its fixed point (the standing hypothesis of Step 3 of the proof). Then the corresponding $F_{\text{GFF}}([\Pi(t)])$ satisfies*

$$\frac{dF_{\text{GFF}}}{dt} = -T_{\text{GFF}} \sigma_{\text{GFF}}([\Pi(t)]) \leq 0, \quad (136)$$

where $\sigma_{\text{GFF}} \geq 0$ is the boundary entropy production rate, defined as the relative-entropy distance between the instantaneous projector class $[\Pi(t)]$ and its asymptotic limit $[\Pi_*]$:

$$\sigma_{\text{GFF}}([\Pi]) = \frac{1}{T_{\text{GFF}}} |\nabla \mathcal{F}_{\text{GFF}}([\Pi])|^2 \geq 0. \quad (137)$$

Equality holds in (136) if and only if $[\Pi]$ is a critical point of $\mathcal{F}_{\text{GFF}}$.

Proof. Step 1: Lindblad (Lindblad, 1976) monotonicity of relative entropy. By Lindblad's monotonicity theorem (Lindblad, 1975), the quantum relative entropy $D(\rho(t) \parallel \rho_*) := \text{Tr}[\rho(t)(\log \rho(t) - \log \rho_*)]$ between any state $\rho(t)$ evolving under a CPTP semigroup with fixed point ρ_* is monotone non-increasing: $dD/dt \leq 0$.

Step 2: From states to projector classes. The boundary moduli manifold $\mathcal{M}_{\text{bdy}}$ parametrizes classes $[\Pi]$ of CPTP fixed points of the boundary GDOS flow (GFF-ST1). The map $\rho \mapsto [\Pi(\rho)]$ that sends a generic density operator to the projector class of its long-time attractor is itself CPTP (it is the composition of a partial trace with a Stinespring dilation), and so Lindblad monotonicity descends to a monotonicity on $\mathcal{M}_{\text{bdy}}$: $dD([\Pi(t)] \parallel [\Pi_*])/dt \leq 0$.

Step 3: BKM information metric and gradient-flow representation. By the Petz–Carlen–Maas representation of GKLS dynamics (Carlen & Maas, 2014, Sec. 4) (Carlen & Maas, 2017, Mittenzenzweig & Mielke, 2017), the Lindbladian $\mathcal{L}$ with fixed point ρ_* that satisfies the quantum detailed-balance condition of Kossakowski *et al.* (1977) with respect to ρ_* — the GNS-symmetry condition, which is the one those representations require; “BKM” names a metric, not a condition — can be expressed as a gradient flow of the relative entropy $D(\rho \| \rho_*)$ with respect to a Riemannian metric on the state space built from the generator, here taken in its BKM (Bogoliubov–Kubo–Mori) form:

$$\frac{d\rho(t)}{dt} = -\mathfrak{g}_{\text{BKM}}^{-1}(\nabla_{\rho} D(\rho \| \rho_*)). \quad (138)$$

The induced metric on $\mathcal{M}_{\text{bdy}}$, obtained by pulling back $\mathfrak{g}_{\text{BKM}}$ along the projector-class map of Step 2 and restricting to the projector manifold, is precisely the g_{Σ} Riemannian metric of GFF-ST3. In particular, the induced gradient flow on $\mathcal{M}_{\text{bdy}}$ takes the form

$$\frac{d[\Pi(t)]}{dt} = -g_{\Sigma}^{-1} \nabla \mathcal{F}_{\text{GFF}}([\Pi(t)]) \quad (\text{modulo gauge transformations}). \quad (139)$$

Step 4: Quadratic form of the dissipation. Combining Step 3 with the Geometric First Law (Theorem 25.1), which identifies $F_{\text{GFF}}([\Pi]) - F_{\text{GFF}}([\Pi_*]) = T_{\text{GFF}} D([\Pi] \| [\Pi_*])$ in the linear-response regime near equilibrium — an identification the theorem keeps at that grade; it becomes exact on a stratum when the stationary state of the boundary generator there is $e^{-\mathcal{F}_{\text{GEP}}/T_{\text{GFF}}}/Z$, for then the displayed relation is an identity — the time derivative of F_{GFF} along the gradient flow is

$$\frac{dF_{\text{GFF}}}{dt} = \langle \nabla \mathcal{F}_{\text{GFF}}, [\dot{\Pi}] \rangle_{g_{\Sigma}} = -\langle \nabla \mathcal{F}_{\text{GFF}}, g_{\Sigma}^{-1} \nabla \mathcal{F}_{\text{GFF}} \rangle = -|\nabla \mathcal{F}_{\text{GFF}}|_{g_{\Sigma}}^2 \leq 0. \quad (140)$$

Identifying the right-hand side with $-T_{\text{GFF}} \sigma_{\text{GFF}}$ via the BKM normalization gives the entropy-production rate $\sigma_{\text{GFF}} = |\nabla \mathcal{F}_{\text{GFF}}|_{g_{\Sigma}}^2 / T_{\text{GFF}} \geq 0$ of (137), with equality iff $[\Pi]$ is a critical point of $\mathcal{F}_{\text{GFF}}$.

Conditionalization note. The Petz–Carlen–Maas gradient-flow representation of GKLS dynamics requires the Lindbladian to satisfy the Kossakowski–Frigerio–Gorini–Verri detailed-balance condition (Kossakowski *et al.*, 1977). For the GDOS boundary Lindbladian that condition is supplied by the weak-coupling (Davies) limit with a bath in a KMS state, which is the limit GG-A5-GDOS (Arisaka, 2026, A5-GDOS) names, and it is the hypothesis under which GG-A5-GDOS’s Lyapunov theorem (its Theorem T5, where the Einstein relation is proved) holds; GG-A5-GDOS’s Theorem T4 carries no detailed-balance assumption and is not the source. The present proof inherits the condition as a standing hypothesis, now written into T3.2. ■

Remark 25.4 (Physics intuition: H-theorem analog). Theorem 25.3 is the boundary GG analog of Boltzmann’s H-theorem. Boltzmann’s H-theorem says that for a classical gas, the entropy $H = -\int f \log f$ monotonically increases under collisions, which is equivalent to $d(-H)/dt \leq 0$ (free-energy decrease at fixed temperature). Theorem 25.3 says the same thing for the boundary moduli space: the GG-Theory analog of H is σ_{GFF} , and the analog of “collisions” is the GDOS CPTP map.

The conceptual upgrade is that, in GG-Theory, the H-theorem is not a phenomenological postulate

(as in Boltzmann) but a derived theorem from the open-system structure of GDOS. What is automatic is the descent of the relative entropy to whatever state is stationary — Lindblad’s theorem, for any CPTP semigroup — and, for any compact state space, the existence of a complete Lyapunov function, Conley’s; what GG-Theory adds is the identity of the descending function with the free energy of the bulk-descended action, which is Theorem **T3.2** under detailed balance and the Gibbs form of the stationary state, and not automatic.

25.3 Equilibrium uniqueness and the unique vacuum

Theorem 25.5 (T3.2 prelude: Equilibrium uniqueness). *Under **GFF-ST1–GFF-ST7**, there exists a unique $[\Pi_*] \in \mathcal{M}_{\text{bdy}}$ such that every GDOS trajectory $t \mapsto [\Pi(t)]$ from a generic initial condition satisfies*

$$\lim_{t \rightarrow \infty} F_{\text{GFF}}([\Pi(t)]) = F_{\text{GFF}}([\Pi_*]) = \min_{[\Pi] \in \mathcal{M}_{\text{bdy}}} F_{\text{GFF}}([\Pi]). \quad (141)$$

This $[\Pi_*]$ is the unique GA of the entire boundary moduli space and is identified with the empirically observed Standard-Model vacuum.

Proof. Step 1: Monotone limit exists. By Theorem 25.3, $F_{\text{GFF}}([\Pi(t)])$ is monotone non-increasing along every GDOS trajectory. By **GFF-ST4** (boundedness from below), $F_{\text{GFF}} \geq F_{\text{GFF}}^{\min} \in \mathbb{R}$. Monotone-bounded sequences converge in $\mathbb{R}$, so the limit $\lim_{t \rightarrow \infty} F_{\text{GFF}}([\Pi(t)])$ exists.

Step 2: ω -limit is in the critical set. Let $\omega([\Pi(0)]) = \bigcap_{T>0} \overline{\{[\Pi(t)] : t \geq T\}}$ be the ω -limit set of the trajectory. By the Lyapunov property of F_{GFF} (Theorem 25.3, equality holds iff $\nabla \mathcal{F}_{\text{GFF}}([\Pi]) = 0$), every point of $\omega([\Pi(0)])$ is a critical point of $\mathcal{F}_{\text{GFF}}$. By **GFF-ST3** the critical points of $\mathcal{F}_{\text{GFF}}$ on each stratum are non-degenerate and therefore isolated, so $\omega([\Pi(0)])$ is a single critical point $[\Pi_\omega]$.

Step 3: Generic initial conditions converge to the global minimum. Let $\text{Crit}(\mathcal{F}_{\text{GFF}}) = \{[\Pi_*]\} \cup \{[\Pi_{s,j}]\}_{j=1}^m$, with $[\Pi_*]$ the unique global minimum. Existence: $\mathcal{F}_{\text{GFF}}$ is bounded below by **GFF-ST4** and lower-semicontinuous on the projector manifold, and the strata of $\mathcal{M}_{\text{bdy}}$ are finite-dimensional, so the infimum is attained. Uniqueness: on a stratum satisfying the hypotheses of **GFF-ST5** every Morse index is even, so no two critical points have adjacent index, the Morse complex has vanishing differential, and $\mathcal{F}_{\text{GFF}}|_\Sigma$ is a *perfect* Morse function; the number of index-0 critical points is therefore the zeroth Betti number, which is 1 because the strata of **GFF-ST1** are connected. (It is worth saying plainly what does *not* give uniqueness: strict Morse-positivity does not, since two distinct minima contribute $+2$ to a^+ and are entirely compatible with $a^+ - a^- > 0$. Perfection and connectedness are what carry the clause.) Each remaining critical point $[\Pi_{s,j}]$ has even index ≥ 2 , hence a stable manifold $W^s([\Pi_{s,j}])$ of real codimension at least 2 in its stratum, by the stable-manifold theorem applied to the gradient flow at a hyperbolic rest point — the linearization $-g^{-1}\text{Hess}\mathcal{F}_{\text{GFF}}$ is self-adjoint with non-zero real eigenvalues by **GFF-ST3**. The complement $\mathcal{G}_{\text{gen}} := \mathcal{M}_{\text{bdy}} \setminus \bigcup_{j=1}^m W^s([\Pi_{s,j}])$ is the basin of the hyperbolic sink $[\Pi_*]$ and is therefore open; being the complement of a finite union of positive-codimension immersed submanifolds it is also dense and of full Riemannian volume. Consequently the trajectory converges to $[\Pi_*]$ for almost every initial datum with respect to any probability measure absolutely continuous with respect to

the Riemannian volume — which is the operative statement, the Baire-genericity being a topological bonus rather than the physical content. For every $[\Pi(0)] \in \mathcal{G}_{\text{gen}}$, the trajectory’s ω -limit $[\Pi_\omega]$ cannot be any saddle (since the trajectory would have to lie on the saddle’s stable manifold, a measure-zero set in $\mathcal{G}_{\text{gen}}$ by construction), so $[\Pi_\omega] = [\Pi_*]$.

Step 4: Convergence to $F_{\text{GFF}}^{\min}$. Combining Steps 1–3: for every $[\Pi(0)] \in \mathcal{G}_{\text{gen}}$, $\lim_{t \rightarrow \infty} [\Pi(t)] = [\Pi_*]$ and hence $\lim_{t \rightarrow \infty} F_{\text{GFF}}([\Pi(t)]) = F_{\text{GFF}}([\Pi_*]) = \min_{[\Pi] \in \mathcal{M}_{\text{bdy}}} F_{\text{GFF}}([\Pi])$. “Generic” carries the measure sense of Step 3 — full Riemannian volume — with Baire genericity holding as well; the exceptional set is the textbook union of saddle stable manifolds, and no transversality assumption on their mutual intersections is used anywhere in the argument. ■

25.4 The per-stratum ground state

Theorem 25.5 minimizes over the whole boundary moduli space. Almost every use the framework makes of it, here and downstream, is instead about a single stratum: the electron is the deepest configuration *of the charged-lepton stratum*, the proton the deepest *of the baryonic stratum*. Those are different statements, and the second does not follow from the first for free — $\mathcal{M}_{\text{bdy}}$ is a disjoint union of strata, and nothing in a global minimum principle tells you what happens inside one piece of a stratified space. This subsection supplies the missing statement, and it costs one definition and one restriction argument.

Before the corollary, the class of strata on which it holds. The ingredients of Theorem 25.5 do not all restrict for free: **GFF-ST4** is stated per stratum and **GFF-ST6** holds along any trajectory, but **GFF-ST5** carries hypotheses, and they are the hypotheses that decide where the corollary is true.

Definition 25.6 (Admissible stratum). A stratum $\Sigma \subset \mathcal{M}_{\text{bdy}}$ is *admissible* if it is connected and compact; carries a Kähler structure with respect to which the boundary gauge connections are flat along Σ ; and is such that $\mathcal{F}_{\text{GFF}}|_{\Sigma}$ is the moment map of a Hamiltonian torus action on Σ acting by Kähler isometries, with isolated non-degenerate fixed points. Admissibility is exactly the hypothesis set of **GFF-ST5** minus its descent and parity clauses, which the present corollary does not use.

Corollary 25.7 (Per-stratum ground state). *Let $\Sigma \subset \mathcal{M}_{\text{bdy}}$ be an admissible stratum (Definition 25.6), and let the dynamics be the fixed-landscape gradient flow of **GRIP-ST2** confined to Σ . Then $\mathcal{F}_{\text{GFF}}|_{\Sigma}$ has exactly one critical point of index 0; it is the unique global minimum of $\mathcal{F}_{\text{GFF}}$ on Σ ; and from every initial datum outside a closed set of zero Riemannian volume and empty interior the descent converges to it. That minimum is the Ground Geometric Attractor of Σ , $\text{GGA}(\Sigma) = \arg \min_{x \in \Sigma} \mathcal{F}_{\text{GFF}}(x)$.*

Proof. The argument of Theorem 25.5 applies with $\mathcal{M}_{\text{bdy}}$ replaced by Σ , and it is worth walking the four steps to see that each ingredient really does restrict. The corollary is two results with two provenances, and the steps sort themselves accordingly: Steps 1, 2 and 4 are *convergence* — every orbit of the confined descent has a critical point as its limit — which is Conley’s fundamental theorem (Conley, 1978, Norton, 1995) in its semiflow form (Hurley, 1995, Lewis, 2025, Patrão, 2011) together with the isolation of the critical points; Step 3 is *uniqueness* of the index-0 point, which

is Frankel’s theorem (Frankel, 1959), with Atiyah’s connectedness of the level sets (Atiyah, 1982) where the critical set is a manifold rather than a point.

*Step 1 restricts because **GFF-ST4** is already per-stratum:* it states boundedness below on each stratum of $\mathcal{M}_{\text{bdy}}$, with the infimum attained at a finite collection of isolated points. Together with the H-theorem (Theorem 25.3), which by **GFF-ST6** holds along every GDOS trajectory and hence along one confined to Σ , the monotone limit exists.

*Step 2 restricts because **GFF-ST3** is already per-stratum:* $\mathcal{F}_{\text{GFF}}$ is smooth on each stratum with non-degenerate critical points there, so the critical points of $\mathcal{F}_{\text{GFF}}|_{\Sigma}$ are isolated and the ω -limit set — non-empty, compact and connected, because Σ is compact — is a single critical point.

Step 3 restricts because admissibility is exactly what it needs. By the moment-map clause and the Kähler structure the Morse indices on Σ are all even. A Morse function whose critical points have no two adjacent indices is perfect, so the number of index-0 points equals $b_0(\Sigma) = 1$ by connectedness: the minimum is unique, and every other critical point has index ≥ 2 and a stable manifold of codimension ≥ 2 in Σ . Their union is closed, of zero Riemannian volume, and of empty interior; its complement is the basin of the unique sink. Where a stratum’s floor is a critical manifold of positive dimension inside a larger stratum rather than a point, Step 2’s “single critical point” needs more than isolation: convergence of each orbit to one point of the manifold is Łojasiewicz’s theorem (Łojasiewicz, 1984), which asks for analyticity of the potential near its minima, and the uniqueness clause becomes Atiyah’s — the level sets of a moment-map component are connected (Atiyah, 1982), so the floor is one connected manifold. The one-dimensional Wilson-line stratum needs neither: its floors are points.

Step 4: the descent stays in Σ . The generator of **GRIP-ST2** is a vector field on Σ , so its integral curves remain in Σ on their maximal interval; and Σ compact makes the flow complete, so there is no escape through a frontier. This is the same fact, read forwards, that Remark 31.7 reads backwards.

Combining, the confined descent converges from generic initial data to the unique global minimum of $\mathcal{F}_{\text{GFF}}|_{\Sigma}$. ■

Remark 25.8 (What this corollary rests on, and where it applies). Two honesty notes, because both matter to how the result should be quoted. First, the moment-map clause of Definition 25.6 is a physical identification — the claim that the functional the boundary dynamics delivers is the moment map of the residual torus action — and the corollary inherits that grade, not a higher one; on the family stratum, and there only, the clause is derived from two statements this paper makes about the descent (Theorem 24.1), and the grade of the corollary there is the grade of those two statements. Second, of the strata this paper names, the one established as admissible is the family-projector stratum $\Sigma_{\text{fam}} = \mathbb{CP}^{N-1}$: it is compact, connected, Kähler, and its GFF is a Fubini–Study height function, the Hamiltonian of the dilatation direction of the maximal torus — an irrational direction, closing on no circle (§24.3, Appendix K). The Higgs and Wilson-line strata are *not* established here: they carry non-trivial Wilson lines along the compact θ -direction, which is the class Theorem 26.1 defers, and on a periodic compact-angle potential over the discrete OGN-5 phases the uniqueness conclusion may genuinely fail. Statements about those strata elsewhere in this paper are inherited at their own grade and are marked where they occur. The corollary is

not weakened by saying so: the family-projector stratum is where the family count, the generation ladder and the charged-lepton spectrum all live.

25.5 Two worked landscapes: what a ground attractor looks like

Abstraction earns its keep when it can be pointed at something measurable, and the corollary can be pointed at two things: a landscape whose shape is known exactly, and a column of numbers spanning twenty-six orders of magnitude.

The exact landscape: three critical points on $\mathbb{CP}^2$. Take the family-projector stratum at $N_g = 3$. It is $\mathbb{CP}^2$, real dimension four, and the GFF on it is the Fubini–Study height function of §24.3 — the moment map of the maximal-torus action, so the stratum is admissible in the sense of Definition 25.6. Its Morse data are not approximate: there are *exactly three* critical points, of indices 0, 2 and 4, the minimal cell structure $\mathbb{CP}^2$ admits, and the two GSSB gates between them carry the Bogomolnyi tolls $S_{[4] \rightarrow [2]} = 2\pi$ and $S_{[2] \rightarrow [0]} = \pi$ (Appendix K). Corollary 25.7 then says: exactly one index-0 point, and the descent finds it.

The physical reading is the one worth carrying away, and it needs one distinction stated exactly, because two of this paper’s own sentences have pulled against each other. On the family stratum, with the height of Theorem 24.1, the second and third generations sit at the critical points of index 2 and 4 — a saddle and the top of a perfect Morse function — and the descent from either is downhill: on $\mathbb{CP}^2$ there is one basin, the floor’s. That is what the instanton theorem of Part V says, and this paper adopts it. The *other* picture, in which a muon is a basin of its own with an exit gate — a way station, the reading the figures of this paper use — is a statement about a different landscape: the one a generation sees on its own charge-supersampled stratum once the family label is fixed, on which the muon is a minimum with a width. This paper does not define that landscape, and the Standard-Model paper, which reads the generations as the basins of depth- n effective landscapes and objects that a saddle supports no sharp mass, must (Arisaka, 2026, A9-SM); the experimental certificate — the second and third charged-lepton masses are measured to a relative precision of 10^{-4} or better — belongs to whichever object turns out to host it, and does not decide the question by itself. What the higher critical points are *not*, on either reading, is the ground state. Only the index-0 point is the global minimum of the family height, and only it has no gate to leave by. That is the whole of the difference between a GA and the GGA, and it is why the muon decays and the electron does not. The dedicated treatment of the charged-lepton landscape is the companion Arisaka-GG-P3-Electron.

The measured landscape: a column spanning twenty-six orders. The second illustration is not a stratum this paper has established as admissible — the baryonic stratum’s Morse structure is not developed here — but it shows, in laboratory numbers, what the GA/GGA distinction is worth. Read the baryon widths as heights above a ground attractor and as answers to one question: what channel do the conserved labels leave open?

The $\Delta(1232)$ sits a few hundred MeV above the nucleon with a strong channel wide open: a GA with a fast exit, width 117 MeV. The $\Lambda(1116)$ has its strong channels closed by strangeness

and must descend through the gate force: a GA with only a slow exit, width 2.5×10^{-12} MeV. The neutron's last channel is a 1.293 MeV chink: a GA with a barely-open exit, width 7.5×10^{-25} MeV. The proton has no exit at all: width zero. Four objects, four widths spanning twenty-six orders of magnitude, and one organizing variable — how far above the ground attractor you sit, and what channel the labels leave open. A width *is* the inverse lifetime of a descent, so only a configuration with nothing below it can have width exactly zero. The dedicated treatment of the baryonic case is the companion Arisaka-GG-P6-Proton; what the present paper supplies is the mechanism the column is organized by.

25.6 Laboratory-testable thermodynamic relations

The physical block of GFF produces three families of laboratory falsifiers, all derivable from **GFF-ST6–GFF-ST7**.

Falsifier 1 (F-GRIP-1): equipartition near equilibrium

Near the global minimum $[\Pi_*]$, the Helmholtz first law $T_{\text{GFF}} dS_{\text{GFF}} = dU_{\text{GFF}}$ implies that the local quadratic expansion of F_{GFF} around $[\Pi_*]$ has the form

$$F_{\text{GFF}}([\Pi_* + \delta\Pi]) - F_{\text{GFF}}([\Pi_*]) = \frac{1}{2} T_{\text{GFF}} \delta\Pi^a K_{ab} \delta\Pi^b + O(\delta\Pi^3), \quad (142)$$

where K_{ab} is the Hessian of $\mathcal{F}_{\text{GFF}}$ at $[\Pi_*]$. By the strict positivity of **GFF-ST5** (K_{ab} is positive-definite at the global minimum), each fluctuation mode contributes $\frac{1}{2} T_{\text{GFF}}$ to the free energy, the standard equipartition result. An empirical violation of equipartition for any boundary moduli mode at scales below M_{KK} would falsify **GFF-ST7**.

Falsifier 2 (F-GRIP-2): fluctuation-dissipation near equilibrium

Standard fluctuation-dissipation theory gives

$$\langle \delta F_{\text{GFF}}^2 \rangle = T_{\text{GFF}} \chi_{\text{GFF}}, \quad (143)$$

where $\chi_{\text{GFF}} = -\partial^2 F_{\text{GFF}} / \partial h^2$ is the boundary susceptibility to an external perturbation h . This is the GG analog of the Einstein relation between mobility and diffusion coefficient. Empirical violations would falsify **GFF-ST7**.

Falsifier 3 (F-GRIP-3): Onsager reciprocity

The cross-coupling matrix between different boundary observables (the Onsager response coefficients) must be symmetric: $L_{ij} = L_{ji}$. This follows from microscopic reversibility plus the gradient-flow form of the GDOS map. Empirical violations would falsify the underlying CPTP structure of GDOS.

Remark 25.9 (Physics intuition: laboratory tests at extreme scales). The three falsifiers above are familiar to any condensed-matter or statistical physicist: they are simply the universal predictions

of linear-response theory near equilibrium. But in GG-Theory they are not phenomenological — they are forced by the dual-layer structure of GFF and the boundary GDOS flow.

The most stringent laboratory tests come from extreme regimes:

- *Ultra-cold atomic systems*: equipartition tests at nano-Kelvin scales probe the boundary fluctuation spectrum;
- *Quantum optomechanical systems*: fluctuation-dissipation tests at the standard quantum limit probe the GFF-Helmholtz identity;
- *Mesoscopic biophysical systems*: Onsager-reciprocity tests in molecular motors and active matter probe the gradient-flow form of GDOS (Friston, 2010).

A coherent failure of any of these tests across multiple platforms would falsify the dual-layer GFF structure. Conversely, the universal success of these laboratory predictions across decades of experiment is one of the strongest empirical pieces of evidence in favor of the GFF dual-layer structure.

26 Theorem T3 (T-GRIP-3): GFF Morse-positivity from CCI and thermodynamic monotonicity from GDOS

We now state and prove Theorem **T3**, the principal theorem of Part IV. The theorem is presented in two parts: **T3.1** (algebraic Morse-positivity, the original statement) and **T3.2** (thermodynamic monotonicity, new in to capture the physical content of GFF).

Dual-layer architecture and the compatibility relation $\mathfrak{C}_{\text{GFF}}$. Per the L16 architecture, T3.1 is the *algebraic layer* (Morse-positivity of $f_{\text{GFF}}|_{\Sigma}$ on every stratum on which the Hopkins–Singer descent of the CCI class is non-trivial), and T3.2 is the *operational/physical layer* (monotone decrease of the Helmholtz-type free energy F_{GFF} along the GDOS Lindbladian flow). The two layers are tied together by the named compatibility relation

$$\boxed{\mathfrak{C}_{\text{GFF}} \equiv \text{T3.1} \implies \text{T3.2}} \quad (144)$$

(via the GDOS CPTP boundary projector and the strict Morse alternating-index sum).

$\mathfrak{C}_{\text{GFF}}$ above is a one-line proposition with the proof: the algebraic Morse-positivity (T3.1) on the appropriate stratum guarantees that the critical-point cascade of f_{GFF} has the correct alternating signature; the CPTP boundary flow with detailed balance with respect to F_{GFF} (condition (2) of T3.2) then descends each critical-point neighborhood monotonically, which is the second law in its GFF-ST6 incarnation (T3.2).

26.1 Statement of T3.1 (algebraic Morse-positivity)

Theorem 26.1 (T3.1: GFF Morse-positivity from CCI). *Let $\Sigma \subset \mathcal{M}_{\text{bdy}}$ be a stratum of the boundary moduli space on which the Hopkins–Singer descent of the CCI class $[X_4] \in H^4(M_6; \mathbb{Z})$ is*

non-trivial, and let Σ be admissible in the sense of Definition 25.6: compact, connected, Kähler, with the boundary gauge connections flat along it, carrying a Hamiltonian action of a torus by Kähler isometries with a direction whose closure-torus has isolated fixed points, and with $\mathcal{F}_{\text{GFF}}|_{\Sigma}$ the moment-map component of that direction. The first four clauses are geometry; the fifth is an identification of the functional the boundary delivers with the moment map, graded as such (Remark 25.8), and on the family stratum it is derived from the paper’s own transport proposition in Theorem 24.1. Then, by Frankel’s theorem (Frankel, 1959) every Morse index of $\mathcal{F}_{\text{GFF}}|_{\Sigma}$ is even, and the strict alternating Morse-index sum

$$a^+(\Sigma) - a^-(\Sigma) \tag{145}$$

equals the Hirzebruch signature $\sigma(\Sigma)$. Consequently, if $\dim_{\mathbb{R}} \Sigma \equiv 2 \pmod{4}$ then (145) vanishes, so strict positivity requires $\dim_{\mathbb{R}} \Sigma \equiv 0 \pmod{4}$; and strict positivity holds precisely when $\sigma(\Sigma) > 0$. In real dimension four the descent gives $\sigma = p_1[\Sigma]/3 \neq 0$, but not its sign; in higher dimensions the sign and the non-vanishing must both be checked on the stratum.

For the family-projector stratum $\Sigma_{\text{fam}} = \mathbb{CP}^{N-1}$ of dimension $2(N_g - 1)$ the check is exact and the statement sharpens to an equivalence: $\sigma(\mathbb{CP}^{N-1}) = 1$ for N_g odd and 0 for N_g even, so strict positivity holds if and only if N_g is odd. This is the form the family-count argument of Part V consumes, and the only form in which the biconditional is claimed.

Remark 26.2 (Physics intuition: Morse-positivity is what makes GFF a free energy in the modern sense). T3.1 says that the restriction of f_{GFF} to any stratum on which the Hopkins–Singer descent of $[X_4]$ is non-trivial has a *strict* alternating Morse-index sum — the signature of a non-degenerate Morse function whose critical points form a finite, well-ordered cascade. The physics reading is that GFF, restricted to each boundary moduli stratum, behaves like a classical free-energy landscape with isolated minima and well-separated saddle-point energy barriers. This is what allows the cascade picture of GRIP (Part V) to be written as a sequence of geometric attractor transitions: without strict Morse-positivity, the cascade would be ill-defined. The compatibility relation $\mathfrak{C}_{\text{GFF}}$ then guarantees that the algebraic Morse landscape is operationally realized as a thermodynamic free-energy under the GDOS CPTP flow (T3.2).

26.2 Outline of the proof

The proof proceeds in four steps, mirroring the chain established in Arisaka (2026, P2-N3, Appendix E).

Scope clarification. Theorem T3.1 is stated for strata Σ that are Kähler and on which the boundary gauge connections F_5, F_B, F_{Q_i} are flat-along- Σ (so that the product structure $\mathcal{M}_{\text{bdy}} = \Sigma \times \Sigma_{\text{gauge}}$ used in the proof of $\mathcal{H}_1$ Part 2 holds). This restriction covers the family-projector stratum $\Sigma_{\text{fam}} = \mathbb{CP}^{N-1}$ on which T4.1 acts, but is stronger than the original formulation. Strata with non-trivial Wilson lines along the compact θ -direction (e.g., the Wilson–Hosotani gauge-orbit moduli of Part II) require a separate argument that is deferred to a companion paper.

Morse function specification. By *Morse function* in what follows we mean the equivariant moment-map Morse function of a direction of the maximal torus of $\text{SU}(N_g + 1)$ acting on $\mathbb{CP}^{N-1}$

— on the family stratum, the dilatation direction of Theorem 24.1 — (or, more generally, of a torus direction in the relevant symmetry group on the Kähler stratum), in the sense of Frankel (1959) and Atiyah (1982). With this Morse function, the alternating Morse-index sum equals the Hirzebruch signature $\sigma(\Sigma)$ via the localization formula and Hodge filtration matching, not the Euler characteristic $\chi(\Sigma)$. For $\mathbb{CP}^2$, $\sigma(\mathbb{CP}^2) = 1$; for $\mathbb{CP}^1$, $\sigma(\mathbb{CP}^1) = 0$.

Step 1. The alternating Morse-index sum of $f_{\text{GFF}}|_{\Sigma}$ equals the topological signature $\sigma(\Sigma)$ when Σ is a Kähler manifold with the equivariant moment-map Morse function specified above.

Step 2. By Hirzebruch’s signature theorem, $\sigma(\Sigma) = \int_{\Sigma} L_m(p_1, \dots, p_m)$ for $\dim_{\mathbb{R}} \Sigma = 4m$, and is zero for $\dim_{\mathbb{R}} \Sigma \equiv 2 \pmod{4}$.

Step 3. By the bulk-to-boundary descent of X_4 , the integrated L-genus on Σ is non-zero iff $[X_4] \in H^4(M_6; \mathbb{Z})$ is non-trivial. This is the $\mathcal{H}_2$ **theorem** of Section 27.

Step 4. Combining Steps 1–3: strict positivity holds iff $\dim_{\mathbb{R}} \Sigma \in 4\mathbb{Z}$, the parity condition of **T3**.

26.3 Detailed proof

Proof of Theorem 26.1. Step 1: alternating sum equals signature. Let Σ be a Kähler manifold of complex dimension n (real dimension $2n$), with the height function $f([\psi]) = \sum_a \lambda_a |\psi_a|^2 / \sum |\psi_a|^2$ as Morse function. By a direct count (Arisaka, 2026, P2-N3, Lemma E.2.1), the critical points of f have even Morse indices $0, 2, \dots, 2n$, and the alternating sum $\sum_a (-1)^{\text{ind}_a/2}$ equals the topological signature $\sigma(\Sigma)$. This is a textbook content of Morse theory on Kähler manifolds.

Step 2: Hirzebruch’s signature theorem. By Hirzebruch’s signature theorem (Hirzebruch, 1966), for a closed oriented manifold Σ of real dimension $4m$,

$$\sigma(\Sigma) = \int_{\Sigma} L_m(p_1, p_2, \dots, p_m), \quad (146)$$

where L_m is the Hirzebruch L-polynomial. For $\dim_{\mathbb{R}} \Sigma \equiv 2 \pmod{4}$, the L-polynomial of the relevant degree does not exist (Pontryagin classes live only in degrees that are multiples of 4), and $\sigma(\Sigma)$ is identically zero.

Step 3: bulk-to-boundary descent gives non-zero L-genus. By the $\mathcal{H}_2$ theorem (Theorem 27.2 of Section 27 below), if the CCl class $[X_4] \in H^4(M_6; \mathbb{Z})$ is non-trivial, then its Hopkins–Singer descent onto the stratum Σ is non-trivial; and by the $\mathcal{H}_1$ theorem (Theorem 27.1 of Section 27), the descended class equals $c_R \cdot p_1(\Sigma)$ for some non-zero rational coefficient c_R . Therefore $p_1(\Sigma)$ is non-trivial. In real dimension four this already settles the matter, since $\sigma(\Sigma) = p_1[\Sigma]/3$ on a closed oriented four-manifold, so the L-genus integral (146) is non-zero; its *sign* is a further fact about Σ and is not supplied by the descent. In real dimension $4m$ with $m \geq 2$ non-vanishing of p_1 does not by itself force $\int_{\Sigma} L_m \neq 0$ — the L-polynomial mixes p_1 and p_2 — so there the signature must be evaluated on the stratum. For $\Sigma = \mathbb{CP}^{N-1}$ both are immediate: $p_1 = N_g h^2$ and $\sigma = 1$ for N_g odd.

Step 4: synthesis. Combining Steps 1–3:

- If $\dim_{\mathbb{R}} \Sigma \in 4\mathbb{Z}$: the alternating sum equals the non-zero signature, hence $a^+ - a^- > 0$.
- If $\dim_{\mathbb{R}} \Sigma \equiv 2 \pmod{4}$: the signature is identically zero, hence $a^+ - a^- = 0$ (no strict positivity).

This proves **T3**. In particular, for $\Sigma = \mathbb{CP}^{N-1}$ of dimension $2(N_g - 1)$, strict positivity holds iff $2(N_g - 1) \in 4\mathbb{Z}$, i.e., N_g is odd. $\blacksquare$

26.4 The N=3 application of T3

For $N_g = 3$, the family-projector stratum is $\mathbb{CP}^2$ of real dimension $4 = 4 \cdot 1$, which is divisible by 4. By **T3**, strict Morse-positivity holds:

$$a^+(\mathbb{CP}^2) - a^-(\mathbb{CP}^2) = \sigma(\mathbb{CP}^2) = 1 > 0, \quad (147)$$

which is exactly Postulate 3.5 of the N=3 proof of Arisaka (2026, P2-N3). In of GG-2026-Generations this was a postulate; in it was reduced to two named hypotheses $\mathcal{H}_1, \mathcal{H}_2$; in the present paper, **T3** makes it a full theorem of the Gi-4 framework.

For $N_g = 2$ (excluded by the lower bound) and $N_g = 4$ (excluded by **T3**), the dimension $2(N - 1)$ is $\equiv 2 \pmod{4}$, and strict Morse-positivity fails. This matches the parity result of the N=3 proof.

26.5 Statement and proof of T3.2 (thermodynamic monotonicity)

Indeed, the second half of Theorem **T3** packages the *thermodynamic monotonicity theorem*, which packages **GFF-ST6–GFF-ST7** into a precise conditional statement linking the algebraic Morse cascade to the boundary second law, with the condition stated in the theorem.

Theorem 26.3 (T3.2: Thermodynamic monotonicity from GDOS). *Let $\mathcal{F}_{\text{GFF}}: \mathcal{M}_{\text{bdy}} \rightarrow \mathbb{R}$ satisfy **GFF-ST1–GFF-ST5** (algebraic block) on the stratified boundary moduli space of the TGG structure of Part II. Let $t \mapsto [\Pi(t)] \in \mathcal{M}_{\text{bdy}}$ be a trajectory generated by a Lindbladian $\mathcal{L}$ acting on the boundary state space (the GDOS flow), and suppose that the following two conditions hold:*

- (1) **Geometric First Law.** *The identifications $U_{\text{GFF}} = \mathcal{F}_{\text{GEP}}$ and $-T_{\text{GFF}} S_{\text{GFF}} = \mathcal{F}_{\text{CEP}}$ of Theorem 25.1 hold on the stratum, so that the stationary state of $\mathcal{L}$ there is the Gibbs state $\rho_* = e^{-U_{\text{GFF}}/T_{\text{GFF}}} / Z$;*
- (2) **Boundary CPTP structure with detailed balance.** *The Lindbladian $\mathcal{L}$ generating the boundary flow is completely positive and trace-preserving (the boundary dynamics is genuinely open-system, as OGGEF requires) and satisfies the quantum detailed-balance condition of Kossakowski, Frigerio, Gorini and Verri (Kossakowski et al., 1977) with respect to ρ_* — the condition a weakly coupled bath in a KMS state supplies by their theorem, and which the Davies limit of GG-A5-GDOS (Arisaka, 2026, A5-GDOS) provides for the boundary generator.*

Then the boundary Helmholtz-type free energy

$$F_{\text{GFF}}([\Pi(t)]) = U_{\text{GFF}}([\Pi(t)]) - T_{\text{GFF}} S_{\text{GFF}}([\Pi(t)]) \quad (148)$$

of **GFF-ST6** satisfies the monotonic decrease

$$\frac{dF_{\text{GFF}}}{dt} \leq 0, \quad (149)$$

with equality exactly at the critical points of $\mathcal{F}_{\text{GFF}}$ on the stratum in the gradient-flow representation of Theorem 25.3, and asymptotically $F_{\text{GFF}} \rightarrow F_{\text{GFF}}([\Pi_*])$, where $[\Pi_*] \in \mathcal{M}_{\text{bdy}}$ is the unique global minimum of **GFF-ST7**, whenever the semigroup relaxes (Frigerio & Verri, 1982). Condition (2) is what makes F_{GFF} the potential of a gradient flow; the descent of the relative entropy to ρ_* needs only a stationary state (Lindblad, 1975, Spohn, 1978) and does not imply detailed balance, so the theorem is a conditional and not an equivalence.

Proof. Assume conditions (1) and (2). By Theorem 25.3, the detailed-balance condition of (2) gives the gradient-flow representation of the boundary flow, under which $dF_{\text{GFF}}/dt = -|\nabla \mathcal{F}_{\text{GFF}}|_{g_\Sigma}^2 \leq 0$ with equality exactly at the critical points. The identification (1) is what makes the descending function the Helmholtz free energy rather than a bare relative entropy: when ρ_* is the Gibbs state of U_{GFF} at T_{GFF} , $T_{\text{GFF}} D(\rho \| \rho_*) = F_{\text{GFF}}(\rho) - F_{\text{GFF}}(\rho_*)$ is an identity, not an approximation. Equation (149) follows. By Theorem 25.5, the asymptotic limit is the unique global minimum.

Two things the theorem does not say, and an earlier statement of it did. Monotonicity does not imply detailed balance: the relative entropy $D(\rho_t \| \rho_*)$ is non-increasing under every CPTP semigroup with a stationary state ρ_* (Lindblad, 1975, Spohn, 1978), detailed balance or not, so neither the decomposition of F_{GFF} nor condition (2) can be read off from monotonicity, and the converse of the theorem is empty. And detailed balance does not by itself make the entropy production strictly positive off the stationary state: a qubit with Hamiltonian σ_x and the single jump operator σ_z satisfies the condition of Kossakowski *et al.* (1977) with respect to the maximally mixed state and has vanishing entropy production on every state diagonal in σ_z , none of which is stationary. The strictness the theorem has is that of the gradient-flow representation on the stratum, and its convergence clause is relaxation (Frigerio & Verri, 1982), not a property of monotonicity. ■

Remark 26.4 (Logical structure of T3: T3.1 + T3.2). Theorem **T3** as stated in is the *conjunction* of **T3.1** and **T3.2**. T3.1 fixes the algebraic side: the strict Morse-positivity $a^+(\Sigma) - a^-(\Sigma) > 0$ is forced by the integrality of the CCI class $[X_4]$ via Hopkins–Singer descent. T3.2 fixes the physical side: the monotonic decrease $dF_{\text{GFF}}/dt \leq 0$ is forced by the CPTP structure of the GDOS flow under detailed balance, via the gradient-flow representation of Theorem 25.3, with Lindblad’s monotonicity theorem supplying the descent of the relative entropy that the representation turns into a free energy. The two halves are linked by the Geometric First Law of Theorem 25.1: the same functional $\mathcal{F}_{\text{GFF}}$ that supports the algebraic Morse cascade *is* the thermodynamic free energy that satisfies the second law.

Remark 26.5 (Physics intuition: why T3.2 is the deeper half). A graduate student in physics may at first think that T3.1 is the “deeper” theorem because it produces specific Morse numbers and underlies the N=3 family proof. But T3.2 is in fact the deeper statement: it is the assertion that those Morse numbers *must* correspond to physical equilibrium states under the boundary second law. T3.1 alone only says that the algebraic Morse cascade exists — one could imagine objections that the cascade is “just topology” with no physical realization. T3.2 says what the descent *is* when the boundary dynamics has the stated regularity: under detailed balance the flow is the gradient flow of the free energy, so the Morse numbers of T3.1 are the critical points of a thermodynamic potential and not of an arbitrary function. And Lindblad’s theorem says that even without that

regularity something descends — the relative entropy to whatever state is stationary — so that irreversibility on the boundary is not optional; only its identification with F_{GFF} is conditional.

In this sense, T3.2 is the *empirical* arm of T3: the Morse numbers of T3.1 are exposed to data as the critical points of a thermodynamic potential and not as topology alone. The convergence of T3.1 (algebraic) and T3.2 (physical) on the same free-energy functional is a consistency check of GG-Theory, under the hypothesis T3.2 now states.

Corollary 26.6 (Laboratory falsifier from T3.2). *Combining Theorem 26.3 with the laboratory-testable equipartition and fluctuation-dissipation relations of **GFF-ST7**, the universal validity of equipartition $\langle \delta F_{\text{GFF}}^2 \rangle = T_{\text{GFF}} \chi_{\text{GFF}}$ across all boundary observables is a kinematic prediction of GG-Theory. A coherent empirical violation would falsify the dual-layer GFF structure and hence the entire GG-6 framework.*

Proof. This is the linear-response specialization of T3.2 with the identifications of Theorem 25.1. ■

27 The descent structure on the family strata: theorem, obstruction, and the residue M6

The N=3 proof of Arisaka (2026, P2-N3, Appendix E) relied on two named hypotheses about anomaly inflow:

- $\mathcal{H}_1$: the $\text{Spin}(1,5)$ curvature descent on the family-projector moduli is $c_R \cdot p_1$;
- $\mathcal{H}_2$: non-trivial $[X_4]$ forces non-zero L-genus on the descended stratum.

The capacity architecture of Arisaka (2026, P2-N3) settles the status of both, and this section records the settlement in Gi-4 vocabulary. $\mathcal{H}_1$ splits into an obstruction and a theorem: no functorial descent of the bulk data can supply it (the rigidity lemma of that paper — of which the naive Chern–Weil restriction is the special case), while the descended geometry *is* supplied, at the Born theorem’s own grade, by the Berry/Fubini–Study structure that the boundary runtime’s adiabatic theorems force on the family strata. $\mathcal{H}_2$ ’s salvageable content is integrality — the integrality theorem L-INT of that paper, whose topological residue R-INT is precisely the η -invariant computation recorded here as open problem M6, in sharpened form.

27.1 $\mathcal{H}_1$ as a theorem

Theorem 27.1 ($\mathcal{H}_1$: $\text{Spin}(1,5)$ curvature descent). *Let $\Sigma \subset \mathcal{M}_{\text{bdy}}$ be a smooth stratum of the boundary moduli on which the bulk-to-boundary projector $\Pi : M_6 \rightarrow M_4$ extends smoothly, and let $\check{\rho}_{\text{bdy}} : \check{H}^4(M_6; \mathbb{Z}) \rightarrow \check{H}^4(\Sigma; \mathbb{Z})$ be the Hopkins–Singer descent map. Then the descent of the TGG curvature 4-form is*

$$\check{\rho}_{\text{bdy}}([\text{tr } R_{\text{Spin}(1,5)} \wedge R_{\text{Spin}(1,5)}]) = c_R \cdot p_1(\Sigma) \in H^4(\Sigma; \mathbb{R}), \quad (150)$$

where $c_R \in \mathbb{Q}_{>0}$ is the rational coefficient set by the $\text{Spin}(1,5)$ Killing-form sign. The gauge-curvature contributions $\text{Tr } F_5^2, F_B^2, F_{Q_i}^2$ pull back to zero on Σ for any stratum transverse to the gauge-orbit moduli.

Sketch, re-sourced. The proof has two parts, and the first must be stated with care, because the obvious route is provably closed.

Part 1: the descended geometry, from the boundary runtime. What can *not* carry this step is any restriction of the bulk spin connection: on the canonical bulk $\mathrm{tr} R_{\mathrm{Spin}(1,5)} \wedge R_{\mathrm{Spin}(1,5)}$ vanishes pointwise (conformal flatness), and the rigidity lemma of Arisaka (2026, P2-N3) (its Appendix G) extends the vanishing to every connection induced functorially from the bulk data — the Chern–Weil restriction that a first attempt would write down is the special case that exposed it. The descended class is instead supplied by the stratum’s own quantum geometry: the boundary runtime’s adiabatic theorems (T-AD-1–T-AD-3 there, at the Born theorem’s own premises — the GDOS forward semigroup and the collision-level contraction) force Berry transport on the family-projector bundle, whose quantum geometric tensor equips Σ with the Fubini–Study package; the descended 4-class is the stratum’s own p_1 , with the coefficient fixed by the Fubini–Study normalization conventions (its Appendix F), not by the bulk Killing form.

Part 2: Vanishing of gauge contributions. The $\mathrm{SU}(5)$, $\mathrm{U}(1)_B$, and $\mathrm{U}(1)_{Q_i}$ bundles are pulled back from the gauge-orbit moduli Σ_{gauge} , which is by hypothesis transverse to Σ . The product structure $\mathcal{M}_{\mathrm{bdy}} = \Sigma \times \Sigma_{\mathrm{gauge}} \times \cdots$ implies that the gauge field strengths are constant along Σ at the level of cohomology, so their pull-back as 4-forms vanishes. Specifically,

$$\mathrm{Tr} F_5^2|_{\Sigma} = F_B^2|_{\Sigma} = F_{Q_i}^2|_{\Sigma} = 0 \quad (151)$$

in $H^4(\Sigma; \mathbb{R})$.

Combining the two parts gives (150). ■

This is the honest Gi-4 restatement of the $\mathcal{H}_1$ content from Arisaka (2026, P2-N3, Appendix E): an obstruction on the bulk side (rigidity), a theorem on the boundary side (Berry/Fubini–Study at Born-grade), and the transversality of the gauge-orbit moduli carrying Part 2 unchanged. Nothing here rests on restricting the bulk curvature, which would yield zero.

27.2 $\mathcal{H}_2$ as a theorem

Theorem 27.2 ($\mathcal{H}_2$: non-degeneracy of Hopkins–Singer descent). *Let $\check{x}_4 \in \check{H}^4(M_6; \mathbb{Z})$ be the differential cohomology class of X_4 (Definition 19.1). If $[X_4]$ is non-trivial in $H^4(M_6; \mathbb{R})$, then the descended class $\check{\rho}_{\mathrm{bdy}}(\check{x}_4)|_{\Sigma}$ is non-trivial in $H^4(\Sigma; \mathbb{R})$ for every stratum $\Sigma \subset \mathcal{M}_{\mathrm{bdy}}$ of real dimension divisible by four on which the descent is well-defined.*

Sketch. The Hopkins–Singer descent map $\check{\rho}_{\mathrm{bdy}}$ is, by construction (Hopkins & Singer, 2005), a morphism of differential cohomology theories. Its non-degeneracy follows from the η -invariant pairing on differential cohomology. Specifically:

Step 1. The differential cohomology class $\check{x}_4$ has ordinary cohomology image $[X_4] \in H^4(M_6; \mathbb{Z})$ which is non-trivial by **CCI-ST5**.

Step 2. The descent map respects the long exact sequence of differential cohomology (107). Therefore the ordinary cohomology image of $\check{\rho}_{\mathrm{bdy}}(\check{x}_4)$ is the descent of $[X_4]$.

Step 3. By Theorem 27.1 ($\mathcal{H}_1$), the descent of $[X_4]$ on a stratum Σ of real dimension $4m$ is

$c_R \cdot p_1(\Sigma)$. The Pontryagin class $p_1(\Sigma)$ is non-trivial on any compact closed Kähler manifold (it equals $N_g \cdot h^2$ for $\Sigma = \mathbb{CP}^{N-1}$, etc.).

Step 4. Therefore the descended cohomology class is non-trivial, and by the long exact sequence of differential cohomology, the descended differential cohomology class $\check{\rho}_{\text{bdy}}(\check{x}_4)|_\Sigma$ is non-trivial in $\check{H}^4(\Sigma; \mathbb{Z})$.

For the rigorous step 1–4 implementation in the 6D bulk, the key technical fact is that the η -invariant pairing on $\check{H}^4(M_6; \mathbb{Z})$ is non-degenerate (Freed & Hopkins, 2000, Hopkins & Singer, 2005). This non-degeneracy is inherited by the descent. ■

This is the Gi-4 restatement of the $\mathcal{H}_2$ content from Arisaka (2026, P2-N3, Appendix E), with its honest residue displayed: what the argument consumes from Hopkins–Singer theory is *integrality* of the descended periods — the integrality theorem L-INT of that paper, whose residue R-INT equals this paper’s M6 — and the degree counting; “non-degeneracy” as an unconditional property is not claimed, being underivable from either side alone (the rigidity lemma and the boundary-relativity proposition of that paper close the two routes).

27.3 Consequences for the N=3 proof

The promotion of $\mathcal{H}_1$ and $\mathcal{H}_2$ from hypotheses to theorems has direct consequences for the rigor of the N=3 proof of Arisaka (2026, P2-N3):

- In GG-2026-Generations, the strict GDOS Morse-positivity postulate was reduced (Appendix E) to two technical hypotheses $\mathcal{H}_1$ and $\mathcal{H}_2$, with the explicit acknowledgment that proving them was an open problem.
- In the present paper, both statements are settled in Gi-4 vocabulary (Theorems 27.1 and 27.2, as re-sourced above): the descended geometry at Born-grade, the bulk-only route closed by rigidity, the residue the integrality computation M6 (= R-INT of the integrality theorem L-INT).
- Therefore strict GDOS Morse-positivity stands as a consistency-grade statement — holding as a theorem at the admissible value ($\sigma(\mathbb{CP}^2) = 1$) — with the exclusion work carried by the capacity theorem.
- The N=3 derivation of Arisaka (2026, P2-N3) is, in light of the present Gi-4 formalization, a derivation from the GG-Constitutional axioms that is *reduced* to a single named residue — the integrality computation M6 (= R-INT) — together with the coupling clause and premises shared with the Constitutional Born Theorem, beyond the standard machinery of differential cohomology and the Hirzebruch signature theorem.

This is the principal payoff of formalizing GFF: the last residual softness in the N=3 derivation is the single named residue M6 (= R-INT), with everything else theorem-grade or closed by obstruction. On the benchmark background of the intended class, M6 is discharged: Appendix H of Arisaka (2026, P2-N3) proves the κ -bridge theorem — the seam’s Hopkins–Singer pushforward equals the boundary’s Berry bundle — closes R-INT there, and closes the warped and relative refinements at class level (its rigidity theorems T-INF-7 and T-INF-8). The detailed re-presentation of the N=3 derivation in Gi-4 vocabulary is given in Part VII.

28 Worked examples

We now work through four explicit examples of GFF on different strata of the boundary moduli.

28.1 Example 1: GFF on family-projector moduli $\mathbb{CP}^{N_g-1}$

The family-projector moduli stratum is $\Sigma_{\text{fam}} = \mathbb{CP}^{N-1}$ with the Fubini–Study metric.

Energetic part $\mathcal{F}_{\text{GEP}}$. By (125) restricted to $\mathbb{CP}^{N-1}$, $\mathcal{F}_{\text{GEP}}$ is the integrated X_4 on the family-projector slice. By Theorem 27.1, this integrates to $c_R \int_{\mathbb{CP}^{N-1}} p_1 = c_R \cdot N_g$ (where $N_g = N_g \cdot \int h^{N_g-1}$ under the standard normalization).

Morse function and critical points. The Morse function on $\mathbb{CP}^{N-1}$ is the height function (134). N_g critical points at $[\delta_{ab}]$, indices $0, 2, \dots, 2(N-1)$.

Strict Morse-positivity. By Theorem 26.1 (T3), strict positivity holds iff N is odd.

For $N = 3$: $\dim_{\mathbb{R}} \mathbb{CP}^2 = 4 \in 4\mathbb{Z}$, $\sigma(\mathbb{CP}^2) = 1$, strict positivity holds. This is the $N=3$ case of the proof of Arisaka (2026, P2-N3).

For $N = 2$: $\dim_{\mathbb{R}} \mathbb{CP}^1 = 2 \notin 4\mathbb{Z}$, $\sigma = 0$, strict positivity fails; $N_g = 2$ is excluded by both the lower bound (CP) and the parity constraint.

For $N = 4$: $\dim_{\mathbb{R}} \mathbb{CP}^3 = 6 \notin 4\mathbb{Z}$, $\sigma = 0$, strict positivity fails; $N_g = 4$ is excluded.

28.2 Example 2: GFF on Higgs moduli

The Higgs moduli stratum Σ_H is the moduli of admissible Wilson–Hosotani Higgs configurations.

Morse function. The Coleman–Weinberg / Wilson–Hosotani effective potential. Two critical points: the symmetric phase ($V = 0$, unstable) and the broken phase ($V = v \simeq 246 \text{ GeV}$, stable GA).

Cascade chain. Symmetric phase (saddle, index 1 in this discrete count) decays to broken phase (GA, index 0) via electroweak phase transition. The instanton-action cost is set by the depth of the broken-phase minimum.

28.3 Example 3: GFF on gauge-orbit moduli

The gauge-orbit moduli stratum Σ_{gauge} is the discrete moduli of admissible orbifold breakings of $\text{SU}(5)$.

Morse function. Discrete: each conjugacy class has a fixed energy, with the lowest-energy class being the SM-compatible $\text{SU}(3) \times \text{SU}(2) \times \text{U}(1)_Y$ broken phase.

GA catalog. The unique stable broken phase is the SM gauge structure. No cascade is needed since the moduli is discrete.

28.4 Example 4: GFF on Wilson-line moduli

The Wilson-line moduli stratum Σ_{θ} is the moduli of the $\text{U}(1)_B$ and $\text{U}(1)_{Q_i}$ phases.

Morse function. Periodic potential on the Wilson-line phase. Critical points at the discrete OGN-5 phases: $2\pi M/L = 2\pi/5$ for $\text{U}(1)_B$.

GA catalog. Five admissible phases for $\text{U}(1)_B$ (corresponding to $L = 5$ via the OGN-5 ledger), the ledger’s being fixed as background data on the anomaly paper’s class; a continuous slow-roll

on the stratum, earlier assigned to $U(1)_{Q_i}$ as the Kaxion field, now of an occupant left open (Remark 11.1). These are the boundary GA's Arisaka (2026, C1-Kaxion) uses.

29 Why there is such a thing as downhill

This Part supplies the paper with a landscape, and a landscape is a more consequential thing to supply than it first appears.

Everyone who has taken a thermodynamics course knows the picture: a quantity that decreases, a system that runs down until it can run down no further, and equilibrium as the bottom of a valley. What is rarely asked is why that picture is available at all. Why should the possible states of a system arrange themselves into something with a *shape* — with heights and slopes and bottoms — rather than into a bare list with no notion of one state being downhill from another? In ordinary thermodynamics this is put in by hand: one writes down a free energy because experience says there is one. Here it is not put in by hand. The landscape is constructed from the geometric data of the preceding Parts, and its most important property — that it curves upward in every direction away from its critical points, so that valleys are genuinely valleys and not flat shelves — is forced by the consistency condition of Part III rather than assumed.

That link deserves to be said slowly, because it is unusual. The requirement that a certain anomaly cancel is a statement about the internal consistency of a quantum field theory. The requirement that a free-energy landscape have real minima is a statement about whether matter can settle down. These are not obviously the same kind of statement, and in most of physics they belong to different subjects taught by different people. Here the first implies the second. Whether the world can come to rest is, on this account, a question with an answer in the geometry.

The functional splits into two pieces with recognizable meanings: one counts what a configuration costs in energy, the other what it costs in information — how much has to be specified, how much order is being maintained against the tendency to spread out. Anything that persists is paying both bills at once. A protein holding its fold, a star holding itself against its own weight, a cell holding a concentration gradient across a membrane: in each case something is being spent, and the accounting in this Part is meant to apply to all of them in the same units. *That the same functional prices a folded protein and a gauge-orbit stratum is the framework's most aggressive claim, and the laboratory tests named in this Part are aggressive for the same reason.*

A word about what a landscape is not. It is not a place. Nothing is literally sitting on a hillside; the heights are values of a functional and the positions are configurations, and the picture is a way of thinking rather than a map of anywhere. The reason it is nonetheless trustworthy is that the mathematics behind it has the properties the picture suggests — there really are minima, they really are separated by barriers, and the descent really is monotonic. When an intuition and a theorem agree this closely it is usually because the intuition was extracted from the theorem in the first place.

What is still missing is time. A landscape says where the low places are; it does not say that anything goes there, or by what route, or whether it stops. A ball on a hillside is a statement about a hillside, not yet a history. The next Part supplies the motion — and with it the reason the world

has three families of matter rather than any other number.

Part V

GRIP

The Gauge-Runtime Invariance Pipeline: Pillar 4 of Gi-4

This is the Part the paper is named for, and the one a reader with limited time should read.

The preceding Part gave the universe a landscape. This one describes how the universe walks across it, and the answer is the mechanism the whole cohort turns on. A symmetric starting configuration does not slide smoothly to the world we observe. It descends in a finite, ordered sequence of stages: it begins with every possibility still open; it is handed a landscape; it reaches a fork and takes one branch, at which point a symmetry that was real a moment before is gone; and it is captured in a basin, where it stays. Four moves. The remark opening the Part's first section gives the picture — a ball going down stairs rather than down a ramp — and the picture is exact rather than decorative, because each step costs a definite amount and the total that can be spent is finite.

That last clause is where the Part earns its keep. Because the budget is finite and the cost per step is fixed by the geometry of the preceding Parts, the number of steps the cascade can take is bounded — and the number of steps is the number of particle families. Three generations of matter, the fact that has resisted explanation since the discovery of the muon in 1936, comes out here as an arithmetic consequence of a budget rather than as an input. The bound is stated with its conditions attached: it leans on a capacity theorem proved in the three-generations companion ([Arisaka, 2026, P2-N3](#)), and on one topological question this paper leaves open and names rather than assumes.

The Part also draws a distinction that matters far beyond particle physics, and it is the one a reader from biology or neuroscience should look for. The engine runs in two regimes, separated by a single question: does the landscape itself change? In one regime the landscape holds still and the system descends on it until it reaches a basin with no exit — a settling, and it finishes. In the other, a transient event reorganizes the landscape and hands the next descent a different starting point. One stroke opens; the other closes. A universe with only the first would have settled once and stopped; one with only the second would never settle at all. The alternation is what lets structure keep appearing at new scales, and it is the reason this paper claims to be about more than the Standard Model. The claim that the same two strokes run in chemistry, in biochemistry, and in cognition is the cohort's wager, carried by the capstone ([Arisaka, 2026, A11-Origin](#)); what is proved here is the mechanism, sharply enough that the wager could be shown to fail.

Two further results close the Part. The descent leaves a signature — a ladder of weights falling by a fixed factor from one family to the next, which orders the generations — and the Part says what that ladder is and is not: an equilibrium weight on the skeleton, not a decay rate, the earlier reading of it as a rate being withdrawn. And the internal consistency of the whole picture requires that the compensating structure identified in Part III survive every step of the descent, which the Part establishes and which is a real constraint rather than a formality.

30 GRIP: motivation, dual-layer structure, and Structure Theorems

Remark 30.1 (Reader’s bridge: a ball rolling down stairs, not a ramp). If GFF gives the universe its preferred landscape, GRIP describes how the universe walks across it. Roughly speaking, the boundary universe does not slide smoothly from a high-energy configuration to the Standard-Model vacuum; it descends in a finite, ordered *cascade* of stages, very much like a ball rolling down a flight of stairs. Each step from one stage to the next costs a definite quantum of action (an integer multiple of π in natural units), and the total number of steps is bounded by a finite budget inherited from the BI-6 closure of Part III. Three named landings span the cascade: DGI at the top (the unstable configuration), GSSB in the middle (the broken-symmetry intermediate), and GA at the bottom (the Standard-Model vacuum). The mathematical name for this is a Morse cascade; the empirical signature is the generation weight ladder — the three charged leptons ordered by height, the electron at the floor — with the gaps fixed by the cascade-step tolls. The lifetimes are a different matter: the τ lepton lives longer than the W boson but much shorter than the muon, and that is phase space, not the cascade. The rest of this Part makes this rolling-down picture mathematically precise.

This Part develops GRIP (the *Gauge-Runtime Invariance Pipeline*), the fourth and final Gi-4 Pillar. Like TGG (Part II), CCI (Part III), and GFF (Part IV), GRIP is intrinsically a *dual-layer* object: it is simultaneously

- (i) an **algebraic/cascade structure** — a finite, ordered chain of Morse critical points $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$ on each stratum, generated by gradient flow of $\mathcal{F}_{\text{GFF}}$ in the Riemannian metric inherited from TGG, with each transition saturating the Bogomolnyi bound $S_{[2k] \rightarrow [2(k-1)]} = \pi k$ and the total cascade action bounded by the BI-6 budget $S_{\text{BI6}}^{\max} = 3\pi$, taken as input from the capacity theorem T-N3-6 of Arisaka (2026, P2-N3) (open problem M10, demand met); and
- (ii) a **physical/gauge-runtime invariance pipeline** — a boundary observability protocol that asserts that every cascade transition *preserves all four boundary Ward currents* (electric, baryonic, informational, metric) at runtime, and that the unique global GA is the only configuration consistent with the empirically observed gauge structure of the Standard-Model vacuum. The runtime checks produce the generation-ordered weight ladder along the family chain ($\tau \rightarrow \mu \rightarrow e$), whose controlled-kinematics ordering is the phenomenological falsifier of GRIP.

The two layers are not independent: they are linked by the *Gauge-Cascade Compatibility Theorem* (introduced below in Section 33), which guarantees that the Morse-cascade flow of **GRIP-ST1**–**GRIP-ST5** does not generate gauge anomalies along the trajectory — equivalently, that the PBI-4 compensator field C_μ of CCI extends smoothly across every saddle. In short:

$$\boxed{\underbrace{S_{\text{tot}} \leq S_{\text{BI6}}^{\max} = 3\pi}_{\text{algebraic cascade closure}} \xrightarrow{\text{Gauge-Cascade Compatibility}} \underbrace{\nabla_\mu J_a^\mu = \text{anomaly}_a}_{\text{Ward at runtime}}} \quad (152)$$

The arrow is precisely the Gauge-Cascade Compatibility Theorem: it carries the algebraic Bogomolnyi cascade to its physical avatar.

Section structure of Part V. Section 30 (the present section) states the seven GRIP axioms (**GRIP-ST1** through **GRIP-ST7**), organized into the algebraic block (**GRIP-ST1–GRIP-ST5**) and the physical block (**GRIP-ST6–GRIP-ST7**); gives the dual-layer formal definition; and lists the first immediate consequences. Section 31 develops the three stages of the cascade in detail. Section 32 formalizes the gradient-flow dynamics on the GFF Morse landscape. Section 33 (new in) develops the physical/gauge-runtime layer: gauge-invariance preservation along the cascade, runtime Ward-identity checks, the generation weight ladder of unstable particles, and laboratory-accessible falsifiers. Section 34 states and proves Theorem **T4**, now in two parts: **T4.1** (algebraic cascade closure forcing $N \leq 3$) and **T4.2** (runtime gauge-invariance preserved iff PBI-4 compensator extends). Section 35 connects GRIP to the generation weight ladder of unstable particles. Section 36 contains worked examples.

Remark 30.2 (Physics intuition: GRIP as the GG analog of the renormalization group). Readers familiar with quantum field theory will find an immediate analogy. In the Wilsonian renormalization group (RG), one integrates out high-energy degrees of freedom in a sequence of discrete steps, each step generating a new effective theory at a lower energy scale. The RG flow is irreversible (information about high-energy modes is lost), the flow is monotonic in some “c-function” (the Zamolodchikov c -theorem), and the IR fixed point is the universal long-distance theory.

In GG-Theory, GRIP plays the same logical role but for the boundary moduli space. DGI is the deepest unbroken (UV) phase; each GSSB saddle is an intermediate effective theory at a lower “Morse scale”; GA is the IR attractor. The c -function role is played by $\mathcal{F}_{\text{GFF}}$ (which decreases monotonically by Theorem **T3.2**); the cascade is the GG analog of the RG flow.

The conceptual upgrade in GRIP is twofold. First, the cascade has *finite* total action $S_{\text{tot}} = \pi N_g(N_g - 1)/2$, which is bounded by $S_{\text{BI6}}^{\text{max}} = 3\pi$ and forces $N_g \leq 3$. This is sharper than the RG, where the flow length is unbounded. Second, the runtime gauge-invariance preservation provides an empirical falsifier — the generation-ordered weight ladder — that has no analog in the RG framework.

Remark 30.3 (Why broader scientists should care about GRIP). The cascade structure $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$ is the universal template for *any* discrete, ordered, irreversible relaxation process in nature. In chemistry, it describes multi-step reaction networks with definite intermediate states. In condensed-matter physics, it describes successive symmetry-breaking phase transitions (e.g., paramagnetic $\rightarrow$ nematic $\rightarrow$ ferromagnetic). In biology, it describes signaling cascades (e.g., MAPK pathway with a fixed sequence of phosphorylation steps). In machine learning and neuroscience, it describes layered hierarchical computation (e.g., deep neural network layers with monotonically decreasing information content per layer, by the data-processing inequality).

What GRIP claims is that all of these distinct cascade structures are *boundary projections* of the same upstream object: the gradient flow of $\mathcal{F}_{\text{GFF}}$ on the boundary moduli, with finite cascade toll bounded by the BI-6 budget. Here too the theorem and the wager should be told apart. That every such process has a Morse decomposition and a complete Lyapunov function is Conley’s theorem, and the decomposition continues under small perturbation of the flow (Franzosa, 1988); that is why cascades are seen everywhere, and it is not this paper’s result. What this paper adds on an admissible stratum is stronger than continuation: the cascade is robust under *every* deformation

within the stratum, not only small ones, because its count of critical points is a cohomology class and its floor’s uniqueness is Frankel’s theorem. That the decompositions of chemistry, biology and cognition are one landscape is the cohort’s wager (§31.6), carried by the capstone. The runtime gauge-invariance preservation provides a *computability* of the cascade: at every stage, one can check whether the boundary configuration respects the underlying symmetry content.

30.1 Motivation: from static GFF to dual-layer cascade

The GFF structure of Part IV provides a Morse landscape on each stratum of the boundary moduli, with critical points labeled by Morse index. By itself, this is a static picture: it tells us the catalog of critical points and the global asymptotic equilibrium, but does not say how the system moves between them, which transitions are admissible, what the timescale of each transition is, or how the gauge structure of each intermediate configuration is preserved.

GRIP adds two complementary layers of dynamical content.

Question 1 (algebraic). GFF catalogs the critical points but does not specify the cascade structure. Multiple non-equivalent ordered chains $c_{\text{top}} \rightarrow c_1 \rightarrow \cdots \rightarrow c_{\text{bottom}}$ could in principle connect the same endpoints; we need a principle that selects the *unique* cascade chain compatible with gradient flow on the TGG-induced Riemannian metric, with each transition saturating the Bogomolnyi bound and the total action bounded by the BI-6 budget. This is the **algebraic content** of GRIP.

Question 2 (physical). GFF guarantees existence of a unique global GA but does not specify how the boundary observers *verify* that the GA is the empirically realized vacuum. Verification requires a runtime mechanism: at every stage of the cascade, the boundary Ward currents (J_e, J_B, J_{Qi}, J_g) must remain closed (modulo the calculable PBI-4 anomaly inflow), and the unbroken gauge subgroup at each saddle must be passed through to the next saddle by an admissible instanton transition. This runtime gauge-invariance preservation is testable in the laboratory through the generation weight ladder along the family chain $\tau \rightarrow \mu \rightarrow e$: each generation gate carries the geometric weight $w_{[2k] \rightarrow [2(k-1)]} = e^{-\pi k}$ along $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$, which orders the generations; the lifetimes are fixed by phase space in Arisaka (2026, A9-SM). This is the **physical content** of GRIP.

Why both layers must be in the same Pillar. The two questions are inseparable. The Bogomolnyi-saturating algebraic cascade (algebraic side) is precisely the unique chain that preserves all four Ward currents at every saddle (physical side) — because every Bogomolnyi-violating instanton would generate a boundary anomaly inflow of magnitude $\Delta S_{\text{tot}} \neq \pi k$, which would break gauge invariance at runtime. Conversely, the empirical observation of the generation weight ladder (physical side) directly constrains the cascade structure (algebraic side), because the ordering of the generations by height is the ordering of the weights $e^{-S_{i \rightarrow i+1}}$ — though not, as an earlier statement had it, because lifetime ratios encode action ratios; they do not (**GRIP-ST7**). This bidirectional coupling *mandates* that both layers be axiomatized together.

GRIP is the formalization of this dual content. It asserts:

- (**Algebraic, GRIP-ST1–GRIP-ST5.**) The cascade dynamics on each stratum is the gradient flow of $\mathcal{F}_{\text{GFF}}$ in the Riemannian metric inherited from TGG via bulk-to-boundary descent; each transition is an instanton with action $S_{[2k] \rightarrow [2(k-1)]} = \pi k$; the total cascade action is bounded

by $S_{\text{BI6}}^{\text{max}} = 3\pi$; the cascade terminates at a stable GA.

- (**Physical, GRIP-ST6–GRIP-ST7.**) The cascade preserves all four boundary Ward currents at every saddle (gauge-runtime invariance); the unique global GA is the only configuration empirically observed via the generation weight ladder of unstable particles ($\tau \rightarrow \mu \rightarrow e$).

Once GRIP is given, the static GFF Morse landscape is upgraded to a dual-layer dynamical cascade with finite total action, discrete attractor endpoints, and runtime-checkable gauge-invariance preservation.

Remark 30.4 (Physics intuition: why gauge-invariance preservation is automatic). A graduate student in physics may wonder why the cascade should *automatically* preserve gauge invariance at every saddle — after all, gauge-invariance violation is a generic feature of intermediate configurations in many dynamical systems. The answer is structural: each instanton transition $c_i \rightarrow c_{i+1}$ is a gradient-flow trajectory of $\mathcal{F}_{\text{GFF}}$, and $\mathcal{F}_{\text{GFF}}$ inherits its gauge invariance from the bulk TGG action via bulk-to-boundary descent. Therefore the gradient flow is itself gauge-equivariant, and gauge invariance is preserved along every trajectory.

The non-trivial content of **GRIP-ST6** is that this preservation is not just a mathematical curiosity — it is *operationally testable* via the boundary Ward identities. The runtime checks of **GRIP-ST6** make the algebraic gauge-invariance into a laboratory-verifiable empirical statement.

30.1.1 Historical background (with citations)

The gauge-runtime invariance pipeline that GRIP formalizes rests on three independent traditions in modern physics. The first is the renormalization-group flow. Stueckelberg and Petermann (1953) first identified the freedom of regularization choice as a continuous symmetry; Gell-Mann & Low (1954) produced the β -function formalism in QED; Kadanoff (1966) introduced the block-spin coarse-graining picture for statistical-mechanical systems; and Wilson (1971) unified these into the modern Wilsonian renormalization-group framework (Wilson, 1974) that underlies every contemporary treatment of scale-dependent gauge dynamics.

The second tradition is gauge symmetry breaking on compact manifolds. Hosotani (1983) showed that on a manifold with a non-contractible S^1 factor, a non-trivial Wilson line $\langle e^{i \oint A} \rangle$ spontaneously breaks the gauge symmetry without an explicit Higgs field — the Wilson–Hosotani mechanism that the GRIP GSSB step identifies as the mathematical content of the GG-Theory electroweak transition. Anderson (1972) framed the broader principle in his “More is different” essay: genuinely new collective behavior emerges at each scale and is not reducible to the underlying microphysics.

The third tradition is the instanton description of tunneling. ’t Hooft (1993) produced the original instanton calculation for the θ -vacuum structure of Yang–Mills theory, and Callan & Coleman (1977), Coleman (1977) developed the semi-classical bounce calculus, with action $S_{\text{inst}} = \pi k$ for adjacent saddle decay on a compact moduli space. GRIP closes the loop by combining these three traditions into the named cascade $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$: dynamical gauge invariance preserves gauge structure under runtime evolution, geometric spontaneous symmetry breaking selects the family-projector moduli, and the geometric attractor catalog terminates the cascade at a finite-budget closure that forces $N \leq 3$ (Theorem 34.1; the upper bound conditional on open problem M10).

30.2 The seven **GRIP** Structure Theorems (algebraic block + physical block)

The seven GRIP Structure Theorems partition naturally into two blocks:

Algebraic block: GRIP-ST1 – GRIP-ST5 These Structure Theorems fix the cascade structure, gradient-flow generation, instanton transitions, Bl-6 action budget, and cascade closure. They live on the stratified moduli space $\mathcal{M}_{\text{bdy}}$ and are Morse-homological in character.

Physical block: GRIP-ST6 – GRIP-ST7 These Structure Theorems state the runtime gauge-invariance preservation along the cascade and the generation-ordered weight ladder. They live in real (laboratory) time and are operationally testable.

The two blocks are linked by the Gauge-Cascade Compatibility Theorem of Section 33.

Algebraic block (GRIP-ST1 through GRIP-ST5)

Structure Theorem 22 (GRIP-ST1: Cascade structure). *On each stratum $\Sigma \subset \mathcal{M}_{\text{bdy}}$, GRIP is a finite, ordered sequence of transitions*

$$c_{\text{top}}(\Sigma) \rightarrow c_1(\Sigma) \rightarrow \cdots \rightarrow c_{\text{bottom}}(\Sigma), \quad (153)$$

where each c_i is a non-degenerate critical point of $\mathcal{F}_{\text{GFF}}|_{\Sigma}$ and the indices satisfy $\text{ind}(c_{\text{top}}) > \text{ind}(c_1) > \cdots > \text{ind}(c_{\text{bottom}})$. The top critical point c_{top} has index equal to $\dim_{\mathbb{R}} \Sigma$ (the deepest unstable phase, DGI); the bottom critical point c_{bottom} has index 0 (the stable GA). The intermediate points $c_1, \dots, c_{n-1}$ are the symmetry-breaking saddles (GSSB stages).

Structure Theorem 23 (GRIP-ST2: Gradient flow generation). *The cascade is generated by the gradient flow of $\mathcal{F}_{\text{GFF}}$ in the Riemannian metric g_{Σ} inherited from TGG via bulk-to-boundary descent. The flow satisfies*

$$\dot{x} = -g_{\Sigma}^{-1}(\nabla \mathcal{F}_{\text{GFF}})(x), \quad x \in \Sigma, \quad (154)$$

and is irreversible: the flow descends in $\mathcal{F}_{\text{GFF}}$ value and never ascends.

The one-way character of this descent is not an accident of detailed balance: it is the trajectory-level face of the law-level irreversibility of the AX-3 projection — the OGGEp channel is CPTP and non-invertible, records accumulate, and no boundary operation runs it backward. Custody is stated once for the whole cohort and is consumed here, not re-derived: AX-3 owns the arrow constitutionally (the AX-3 synthesis of Arisaka, 2026, A1-Constitution); Part IX of the CPT member paper (Arisaka, 2026, P5-CPT) is the canonical law-level derivation of the arrow from that axiom, over a bulk that is itself exactly T-symmetric; and the GDOS runtime (Arisaka, 2026, A5-GDOS) is its dynamical implementation — of which the gradient descent above is the Gi-4 geometric face.

On the family stratum the flow (154) has a name in another field. Write $p_{\ell} = |\psi_{\ell}|^2 / \|\psi\|^2$ for the share of the winding- ℓ mode, the moment coordinates of the torus; then with the height of Theorem 24.1, in the doubled Fubini–Study form in which the dilation runs at unit rate, the descent is

$$\dot{p}_{\ell} = -2p_{\ell}(\Delta_{+}(\ell) - \langle \Delta_{+} \rangle), \quad \langle \Delta_{+} \rangle = \sum_m \Delta_{+}(m) p_m, \quad (155)$$

the replicator equation (Hofbauer & Sigmund, 1998) with fitness $-2\Delta_+(\ell)$: the mean scaling dimension descends at twice its own variance, which is Fisher’s fundamental theorem of natural selection (Fisher, 1930). Across the throat’s $K = 35$ e-folds the shares of the losing windings fall by $e^{-2K[\Delta_+(\ell)-\Delta_+(1)]}$: 9.8×10^{-19} for the second winding and 2.3×10^{-42} for the third, the infrared share over the ultraviolet share. That is a statement about the relaxation of a population of modes toward the floor, and not the statement that every electron is identical to every other, which §31.5 makes on other grounds.

Structure Theorem 24 (GRIP-ST3: Instanton transitions between adjacent saddles). *Each transition $c_i \rightarrow c_{i+1}$ in the cascade is an instanton trajectory with definite real action $S_{i \rightarrow i+1} \in \mathbb{R}_{>0}$, and each gate it crosses carries a toll: the inflow term of §23.4 evaluated on the saddle. On the family-projector stratum the gate between the critical points of indices $2k$ and $2(k-1)$ has inflow charge $\Delta c_2 = k$ (Remark 23.5) and toll*

$$S_{[2k] \rightarrow [2(k-1)]} = \pi k. \quad (156)$$

The paper keeps the name Bogomolnyi toll for this quantity, because the Bogomolnyi-saturated action of the connecting orbit equals it when the height is the pair invariant (Theorem 32.1); under the transport proposition the height is the dilatation charge and the two differ (Theorem 24.1, Remark 32.2), and wherever this paper writes “Bogomolnyi toll πk ” the quantity meant is the inflow toll of this theorem, which is what **GRIP-ST4** consumes.

Structure Theorem 25 (GRIP-ST4: Bl-6 instanton-action budget). *The total cascade action accumulated along (153) is bounded above by a budget $S_{\text{BI6}}^{\max}$ inherited from CCI:*

$$S_{\text{tot}}(\Sigma) := \sum_i S_{c_i \rightarrow c_{i+1}} \leq S_{\text{BI6}}^{\max}, \quad S_{\text{BI6}}^{\max} = 3\pi. \quad (157)$$

The budget value $S_{\text{BI6}}^{\max} = 3\pi$ is the capacity theorem T-N3-6 of Arisaka (2026, P2-N3) — in its sharper per-family form $S_{\max}(N_g) = \pi N_g$ — derived there from the family stratum’s quantum geometry, conditional on the coupling clause and the integrality theorem L-INT (residue R-INT; open problem M10 below, whose N -independence demand that theorem meets). What is derived is the budget’s unit and not its coefficient: the CCI class is integral (**CCI-ST4**), which with the descent normalization $\Phi_\omega = \pi$ of §23.4 quantizes admissible action on a lattice of step $\pi/2$; nothing upstream selects the multiple. The reader should therefore treat 3π as an input of stated conjectural grade throughout this paper.

Structure Theorem 26 (GRIP-ST5: Cascade closure to a stable GA). *The cascade terminates at a stable Geometric Attractor $c_{\text{bottom}}(\Sigma)$, which is a local minimum of $\mathcal{F}_{\text{GFF}}|_\Sigma$ (Morse index 0). After cascade closure, no further dynamics is possible on Σ at the cascade level: any further transitions require leaving the stratum or coupling to other strata.*

Remark 30.5 (Local by **GRIP-ST5**, global by Corollary 25.7). **GRIP-ST5** asserts only what the cascade structure alone can support: the terminus is a *local* minimum, Morse index 0. That is the honest algebraic statement, and on a general stratum it is all that holds — a landscape may carry

several index-0 basins, and which one a trajectory finds is then a question about initial data. On an *admissible* stratum (Definition 25.6) the two statements collapse into one: by Corollary 25.7 there is exactly one index-0 critical point, so the local minimum of **GRIP-ST5** is the global minimum, and the terminus is the GGA. Wherever this paper or its companions need the global statement, the corollary is the warrant and **GRIP-ST5** alone is not.

The five Structure Theorems **GRIP-ST1–GRIP-ST5** constitute the *algebraic* block of GRIP. Axioms **GRIP-ST1–GRIP-ST3** are essentially the standard Morse homological structure on a Kähler manifold; the new mathematical content is **GRIP-ST4** (the BI-6 budget) and **GRIP-ST5** (the closure condition), which together force $N_g \leq 3$ on the family-projector stratum via Theorem T4.1 (conditional on open problem M10).

Physical block (GRIP-ST6 and GRIP-ST7)

The physical block translates the algebraic cascade structure of GRIP into operational, observer-accessible runtime form. It consists of two Structure Theorems.

Structure Theorem 27 (GRIP-ST6: Runtime gauge-invariance preservation). *Along every instanton transition $c_i \rightarrow c_{i+1}$ of the algebraic cascade (153), the four boundary Ward currents (J_e, J_B, J_{Qi}, J_g) remain closed (modulo the calculable PBI-4 anomaly inflow):*

$$\boxed{\nabla_\mu J_a^\mu(t) = -\alpha_a C^\mu(t) \cdot \rho_a(t), \quad a \in \{e, B, Qi, g\}, \quad t \in [t_i, t_{i+1}],} \quad (158)$$

where C_μ is the PBI-4 compensator field of CCI-ST6, the α_a are the **CCI-ST2** coefficients, and $\rho_a(t)$ are the time-dependent local densities along the instanton trajectory. Equivalently: the unbroken gauge subgroup $H_i \subset G_{\text{tot}}$ at saddle c_i is passed to the unbroken subgroup $H_{i+1} \subset H_i$ at saddle c_{i+1} by the instanton transition — there is no spurious gauge-symmetry generation or destruction along the trajectory.

The runtime statement (158) is a kinematic identity following from the gauge-equivariance of the gradient flow **GRIP-ST2** on the TGG-induced metric, together with the boundary covariant conservation law **CCI-ST6**.

Structure Theorem 28 (GRIP-ST7: Generation-ordered weight ladder). *The boundary signature of the GRIP cascade is the family ladder itself. Each generation gate $[2k] \rightarrow [2(k-1)]$ on the family-projector stratum $\mathbb{CP}^{N-1}$ carries the toll $S_{[2k] \rightarrow [2(k-1)]} = \pi k$ of **GRIP-ST3**, and the exponential of the toll is the gate's geometric weight:*

$$\boxed{w_{[2k] \rightarrow [2(k-1)]} = e^{-S_{[2k] \rightarrow [2(k-1)]}} = e^{-\pi k}.} \quad (159)$$

The weight is an equilibrium weight and not a transition rate. On an admissible stratum the small-noise descent has one asymptotically stable floor and no Markov chain of metastable states between the generations: the Bogomolnyi-saturated orbit of Theorem 32.1 is the downhill connecting orbit, whose Freidlin–Wentzell quasipotential is zero, and the exponential attached to it is the Gibbs weight of the reverse excursion — the stationary density at the higher critical point relative to the floor, $\exp[-(\mathcal{F}_{\text{GFF}}(p) - \mathcal{F}_{\text{GFF}}(\text{floor}))/T_{\text{GFF}}]$ at $T_{\text{GFF}} = 1$ in height units, with no prefactor; with the height

of Theorem 24.1 the exponents are $c[\Delta_+(k+1) - \Delta_+(k)]$ and $e^{-\pi k}$ is the weight of the toll, a different exponential. No Eyring–Kramers formula applies, because a Kähler stratum has no index-1 saddle and the floor’s basin has a complement of codimension two; a rate ladder between generations would need either a stratum with several attractors, which is not admissible, or a mechanism other than small-noise escape, and this paper supplies neither. Along the charged-lepton ladder $\tau \rightarrow \mu \rightarrow e$ (and the quark ladder $t \rightarrow b \rightarrow c \rightarrow s \rightarrow u/d$) the weights order the generations by height; the absolute lifetimes are Standard-Model quantities dominated by the species-specific phase space (Sargent’s rule $\Gamma \propto G_F^2 Q^5$), and where Arisaka (2026, A9-SM) and Arisaka (2026, P2-N3) carried the factor $e^{-\pi k}$ into a width, that reading is now held at the grade of this theorem: a weight, whose entry into any width would need the mechanism named above.

The robust, **GRIP**-level content is therefore the existence and generation-ordering of the weight ladder, not a lifetime law. The falsifier that survives is the ordering: a family ladder whose generations are not ordered as the weights order them would falsify **GRIP**; the controlled-kinematics decay test of earlier statements tests a rate reading this paper no longer derives, and is kept below at that grade.

The seven Structure Theorems **GRIP-ST1–GRIP-ST7** together constitute the full *Gauge-Runtime Invariance Pipeline*. The algebraic block forces the cascade structure and bounds $N_g \leq 3$ (conditionally on open problem M10). The physical block makes the cascade operationally testable through the boundary Ward identities and the generation-ordered weight ladder. The two halves are unified by Theorem **T4** in Section 34.

30.3 Dual-layer formal definition of **GRIP**

We now consolidate the algebraic and physical blocks into a single dual-layer definition.

Definition 30.6 (GRIP: the Gauge-Runtime Invariance Pipeline, dual-layer form). The *Gauge-Runtime Invariance Pipeline* on the boundary moduli space $\mathcal{M}_{\text{bdy}}$, equipped with the GFF Morse landscape and the PBI-4 compensator field C_μ of CCI, is the **septuple**

$$\text{GRIP} = \left(\underbrace{\{c_i\}_{i=0}^n, \dot{x} = -g^{-1} \nabla \mathcal{F}_{\text{GFF}}, S_{[2k] \rightarrow [2(k-1)]} = \pi k, S_{\text{BI6}}^{\text{max}} = 3\pi, \text{GA} = c_{\text{bottom}};}_{\text{algebraic block (GRIP-ST1–A5)}}, \underbrace{\nabla_\mu J_a^\mu = -\alpha_a C^\mu \rho_a, w_{[2k] \rightarrow [2(k-1)]} = e^{-\pi k}}_{\text{physical block (GRIP-ST6–A7)}} \right), \quad (160)$$

where the seven components correspond to the Structure Theorems **GRIP-ST1–GRIP-ST7** respectively.

The algebraic part is the collection of cascades

$$\text{GRIP}_{\text{alg}} = \bigsqcup_{\Sigma \subset \mathcal{M}_{\text{bdy}}} (c_{\text{top}}(\Sigma) \rightarrow c_1(\Sigma) \rightarrow \cdots \rightarrow c_{\text{bottom}}(\Sigma)), \quad (161)$$

one for each stratum Σ of $\mathcal{M}_{\text{bdy}}$, satisfying the Structure Theorems **GRIP-ST1–GRIP-ST5**. The cascade on each stratum is abbreviated as $\text{DGI}(\Sigma) \rightarrow \text{GSSB}(\Sigma) \rightarrow \text{GA}(\Sigma)$, where $\text{DGI}(\Sigma) = c_{\text{top}}$,

$\text{GSSB}(\Sigma)$ is the chain of intermediate saddles, and $\text{GA}(\Sigma) = c_{\text{bottom}}$. Definition 31.4 below refines the terminus: on an admissible stratum the index-0 minimum is unique (Corollary 25.7), and the cascade output is that *unique deepest* basin, the Ground Geometric Attractor. The chain is written $\text{DGI} \rightarrow (\text{GFF}) \rightarrow \text{GSSB} \rightarrow \text{GA}$ throughout — GA names the stage — with the GGA named as its terminus.

The physical part (158)–(159) asserts that the cascade preserves all four boundary Ward currents at runtime and that the generation weight ladder is the direct boundary projection of the cumulative cascade toll. The two parts are linked by the Gauge-Cascade Compatibility Theorem (Theorem 33.1 of Section 33).

One classification, introduced program-wide in the capstone **GG-A11-Origin** and defined at runtime level in §31.6 below, should be recorded already here: the septuple (160) is the **GRIP-In** regime — the runtime on a fixed stratified landscape. The complementary regime, **GRIP-Out**, reorganizes the landscape itself and supplies each stratum’s initial datum c_{top} ; it is deliberately *not* an element of the septuple, and Definition 31.6 states precisely what is delegated to **GG-A11-Origin**.

Remark 30.7 (Logical content of GRIP). The intuitive content of Definition 30.6 is twofold:

- **Algebraic side.** GRIP fixes the cascade structure on each stratum (a unique chain of Morse critical points connected by Bogomolnyi-saturating instantons), bounds the total cascade action by $S_{\text{BI6}}^{\text{max}} = 3\pi$, and forces termination at a stable GA. This is the upstream structure that produces the upper bound $N_g \leq 3$ in Theorem T4.1 (conditional on open problem M10).
- **Physical side.** GRIP packages the runtime gauge-invariance preservation into an operationally testable sequence of Ward-identity checks and produces the generation-ordered weight ladder via the toll formula (159), which orders the generations. This is the downstream content that makes GRIP a phenomenologically falsifiable theory.

The two sides are linked by the Gauge-Cascade Compatibility Theorem: the same cascade structure that produces $N = 3$ algebraically also produces the runtime-checkable Ward identities physically. This is the duality of GRIP.

Remark 30.8 (Open-system character of GRIP). The cascade dynamics of GRIP is intrinsically open-system. In a closed-system Hamiltonian theory, dynamics is unitary and reversible: all states evolve forever, and there is no privileged “cascade” direction. GRIP requires a finite instanton-action budget, an irreversible gradient flow, a discrete catalog of attractor endpoints, and a runtime-checkable gauge-invariance pipeline — all four of which are open-system features that have no closed-system analog. This is the structural reason that the upper bound $N \leq 3$ cannot be derived from any closed-system framework, no matter how elaborate, as we discussed in Arisaka (2026, P2-N3, Section 12).

Remark 30.9 (Physics intuition: why both layers are needed). A graduate student in physics may at first think that the algebraic side (cascade closure) is sufficient — after all, it gives the upper bound $N_g \leq 3$ that closes the $N=3$ family proof. But the algebraic side alone does not provide an empirical pathway to test the cascade. The physical side (**GRIP-ST6–GRIP-ST7**) is what makes GRIP testable: the generation weight ladder orders the generations by cumulative toll, and the ordering is read in the laboratory from the mass hierarchy.

Conversely, the physical side alone (a phenomenological ordering of the generations) does not specify the upstream cascade structure or explain the $e^{-\pi k}$ ordering of the family gate weights. Without the algebraic side, one would have only an empirical fit-formula with no fundamental origin.

The dual-layer structure of GRIP is therefore essential — both layers are needed, and the Gauge-Cascade Compatibility Theorem is what makes them consistent.

30.4 First consequences of the GRIP Structure Theorems

We list four immediate consequences of **GRIP-ST1–GRIP-ST5**.

Proposition 30.10 (Cascade length equals stratum index range). *On a stratum Σ of real dimension $2n$ supporting only even Morse indices (Kähler-Hamiltonian case), the cascade (153) has length n (i.e., contains $n + 1$ critical points and n transitions).*

Proof. By **GFF-ST3** of Part IV, the Morse function on Σ is perfect. The Morse polynomial of a Kähler manifold of complex dimension n is $1 + t^2 + t^4 + \cdots + t^{2n}$, indicating exactly $n + 1$ critical points with indices $0, 2, \dots, 2n$. The cascade $c_{\text{top}} \rightarrow \cdots \rightarrow c_{\text{bottom}}$ visits all of them in descending index order, giving n transitions. ■

Proposition 30.11 (Total cascade action: triangular sum). *On the family-projector stratum $\Sigma_{\text{fam}} = \mathbb{CP}^{N-1}$ of complex dimension $N_g - 1$, the total cascade action is*

$$S_{\text{tot}}(\mathbb{CP}^{N-1}) = \sum_{k=1}^{N_g-1} \pi k = \pi \frac{N_g(N_g - 1)}{2}. \quad (162)$$

Proof. By Proposition 30.10, the cascade has $N - 1$ transitions. By **GRIP-ST3**, the action of the k -th transition (between indices $2k$ and $2(k - 1)$) is πk . Summing: $\sum_{k=1}^{N_g-1} \pi k = \pi N_g(N_g - 1)/2$, the triangular sum. ■

Proposition 30.12 (Cascade asymptotic time). *The gradient flow (154) reaches the stable GA only asymptotically as $t \rightarrow \infty$, but the cumulative toll is finite (equal to $S_{\text{tot}}(\Sigma)$). The exponential $e^{-S_{i \rightarrow i+1}}$ attached to each gate is the equilibrium weight of the upper critical point relative to the lower under boundary noise, not a rate of passage: the passage itself is the deterministic downhill orbit, reached in infinite time and at zero quasipotential.*

Proof. Standard Morse-flow theory gives the first clause: the gradient flow has critical points as asymptotic limits, approached exponentially at the rate of the Hessian's eigenvalues. The second clause is **GRIP-ST7**: with one asymptotically stable floor and no index-1 saddle on a Kähler stratum, small-noise theory has no escape rate to assign, and the Boltzmann weight is what the exponential is. Witten (1982) supplies the supersymmetric Morse formalism and Schwarz (1993) the homological version; Witten's Morse-complex weights are entries of the connection matrix, which on an admissible stratum vanishes because no two critical points have adjacent index. ■

Proposition 30.13 (GRIP as the formal realization of OGGE cascade). *The existence of GRIP satisfying **GRIP-ST1–GRIP-ST5** is equivalent to the cascade content of OGGE: the irreversible bulk-to-boundary projection, with finite-budget attractor selection.*

Sketch. OGGEF states that observable boundary content is related to bulk geometric content by an irreversible CPTP bulk-to-boundary projection. GRIP formalizes the dynamical content of this projection: the cascade is the iterated application of the boundary projector, the gradient flow is the dissipative dynamics, and the GA's are the fixed points of the iterated projection. Conversely, given GRIP satisfying its five Structure Theorems, one recovers the CPTP-projection content of OGGEF via the gradient flow. ■




Figure 7: The GRIP cascade $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$ (Pillar 4 of Gi-4) with the BI-6 budget closure. Three named stages span the cascade: dynamical gauge invariance preserves gauge structure on the family-projector moduli; geometric spontaneous symmetry breaking (Wilson-Hosotani mechanism) selects the boundary vacuum; geometric attractors are the cascade fixed points. Each cascade step costs πk units of instanton action, accumulating to $S_{\text{tot}}(\mathbb{CP}^{N-1}) = \pi N_g(N_g - 1)/2$ on the family-projector stratum $\Sigma_{\text{fam}} = \mathbb{CP}^{N-1}$. The cumulative action must not exceed the BI-6 budget $S_{\text{BI6}}^{\text{max}} = 3\pi$ (the capacity theorem T-N3-6 of companion paper Arisaka (2026, P2-N3); open problem M10, demand met), which forces $N_g \leq 3$ and is saturated at $N_g = 3$. Runtime gauge-invariance preservation along the cascade (T4.2) holds iff the PBI-4 compensator $C_\mu(t)$ extends smoothly across each instanton stratum (open problem M11 tracks the topological non-obstruction).

31 The DGI – GSSB – GA cascade in detail

We now develop the three stages of the GRIP cascade — DGI, GSSB, GA — in detail.

31.1 DGI: Dynamical Gauge Invariance

Definition 31.1 (DGI: Dynamical Gauge Invariance). The *Dynamical Gauge Invariance* configuration on a stratum Σ is the critical point $c_{\text{top}}(\Sigma)$ of $\mathcal{F}_{\text{GFF}}|_{\Sigma}$ with maximum Morse index, equal to $\dim_{\mathbb{R}} \Sigma$: the top of the cascade, at which every gauge equivalence is intact and nothing has yet committed. It corresponds to the unbroken phase of the boundary projector on Σ , with the full TGG composite-connection symmetry dynamically preserved.

Physical interpretation. In the family-projector stratum $\Sigma_{\text{fam}} = \mathbb{CP}^{N-1}$, the DGI point is the family configuration that has the most symmetry: equal weight on all N_g family directions. This

corresponds to a maximally symmetric flavor basis with no preferred family. In the Higgs stratum Σ_H , the DGI point is the symmetric (unbroken) electroweak phase $V = 0$.

Stability properties. The DGI point is a Morse maximum, so it is unstable in all directions: small perturbations drive the system into the cascade. Its weight relative to the next point down is set by the toll of the first cascade transition, $S_{[\dim \Sigma] \rightarrow [\dim \Sigma - 2]} = \pi(N_g - 1)$ on $\mathbb{CP}^{N-1}$; the descent from it is downhill and has no timescale of its own beyond the Hessian's.

31.2 GSSB: Geometric Spontaneous Symmetry Breaking

Definition 31.2 (GSSB: Geometric Spontaneous Symmetry Breaking). The *Geometric Spontaneous Symmetry Breaking* chain on a stratum Σ is the sequence of intermediate saddle points $c_1, c_2, \dots, c_{n-1}$ in the cascade (153), with Morse indices $\dim \Sigma - 2, \dim \Sigma - 4, \dots, 2$. Each saddle corresponds to a partial breaking of the boundary projector symmetry.

Physical interpretation. Each saddle in the GSSB chain represents a phase with intermediate symmetry breaking. In the family-projector stratum $\mathbb{CP}^{N-1}$, the saddles correspond to flavor configurations with progressively less family symmetry: the index- $2k$ saddle has k “frozen” family directions and $N_g - k$ “free” ones.

Cascade dynamics. Each saddle decays into the next by an instanton of action πk (where k is the Morse-index drop). The GSSB chain therefore traces out a sequence of partial symmetry breakings, each with its own toll and weight.

Connection to fermion masses. In the application to the Standard Model (Part VII), each saddle in the GSSB chain on $\mathbb{CP}^2$ (the $N_g = 3$ family stratum) corresponds to a generation: the index-4 DGI point is associated with the third generation (maximally unstable), the index-2 saddle with the second generation (intermediate), and the index-0 GA with the first generation (stable). This is the structural origin of the $e^{-n\pi}$ cascade in the GG-A9-SM Yukawa codebook.

31.3 GA: Geometric Attractor

Definition 31.3 (GA: Geometric Attractor). The *Geometric Attractor* on a stratum Σ is the critical point $c_{\text{bottom}}(\Sigma)$ of $\mathcal{F}_{\text{GFF}}|_{\Sigma}$ with Morse index 0, i.e., a local minimum of the GFF Morse function. It corresponds to a stable boundary observable.

Physical interpretation. The GA is the asymptotic endpoint of the cascade: the dynamically selected, stable boundary observable. In the family-projector stratum, the GA corresponds to the first generation (electron, up, down) — the lightest, most stable fermions. In the Higgs stratum, the GA is the broken electroweak vacuum $V = v \simeq 246$ GeV. In the Wilson-line stratum, the GA is one of the discrete OGN-5 phases.

Stability properties. The GA is a Morse minimum, so it is locally stable: small perturbations relax back to the minimum. No further cascade transitions are possible from the GA; the cascade on Σ has terminated.

Catalog of physical GA's. The catalog of GA's across all strata of $\mathcal{M}_{\text{bdy}}$ constitutes what we observe as the “physical world” — the catalog of stable boundary observables. This catalog includes: the three Standard-Model fermion generations (family-projector stratum), the broken

electroweak vacuum (Higgs stratum), the gauge couplings $\alpha_3, \alpha_W, \alpha_Y$ (gauge stratum), the IR scale $M_{\text{lock}} = E_{\text{Pl}}/2e^{35} \simeq 3.85 \text{ TeV}$ (radial-scale stratum), and the Kaxion field value (Wilson-line stratum).

31.4 GGA: Ground Geometric Attractor

A Morse landscape generally admits *many* index-0 minima — many GA’s — and Definition 31.3 alone therefore does not yet say *which* stable boundary observable the cascade realizes. The missing clause is the ground-state refinement, and the present paper owns it: Corollary 25.7 establishes, on any admissible stratum, that the descent halts at a unique deepest basin. The definition below names that object. Its generalization across the layers of the program-wide cascade is the capstone’s contribution (Arisaka, 2026, A11-Origin); the runtime statement is here.

Definition 31.4 (GGA: Ground Geometric Attractor). Let $\{\text{GA}^{(1)}, \text{GA}^{(2)}, \dots\}$ be the set of Geometric Attractors (Definition 31.3) of $\mathcal{F}_{\text{GFF}}$ on a stratum Σ . The *Ground Geometric Attractor* of Σ is the global minimum:

$$\text{GGA}(\Sigma) := \text{GA}^* = \arg \min_{x \in \Sigma} \mathcal{F}_{\text{GFF}}(x). \quad (163)$$

The cascade of Definition 30.6 runs $\text{GRIP} := \text{DGI} \rightarrow (\text{GFF}) \rightarrow \text{GSSB} \rightarrow \text{GA}$, and on an admissible stratum it terminates not merely on *a* GA but on *the* GGA: the stage is the Geometric Attractor, the terminus is the Ground Geometric Attractor. Existence and uniqueness are Corollary 25.7, and no filter beyond the landscape’s own geometry is used in establishing them. Dynamically, the GGA is *where GRIP-In halts* (Definition 31.5 below): the unique attractor from which no admissible exit channel remains.

Physical interpretation. Chemistry made the distinction a century ago. A molecule’s electronic landscape has many stationary configurations, but chemistry is built on the *ground state*: every water molecule is identical, and indefinitely stable, because its electrons occupy not merely *a* stable configuration but *the* lowest one. The GGA is that idea promoted to a constitutional principle of the runtime: the cascade does not settle into *some* attractor; it settles into the ground state of the attractor manifold.

31.5 GGA and the stability of matter

The refinement from GA to GGA is what converts the catalog of stable boundary observables from a description into an *explanation*. A GA is stable against small perturbations (its Hessian is positive definite); only a GGA is stable against *tunneling*, because on its stratum there is no deeper basin to decay to. The stability of matter is, clause by clause, a GGA statement:

- **The electron is stable** because the charged-lepton stratum’s GGA is the configuration $(m_e, -e, \frac{1}{2}\hbar)$: there is no lower-lying charged-lepton basin, and every cascade realization, anywhere in spacetime, selects this same unique minimum — which is simultaneously why every electron is rigorously identical to every other. That last clause is worth stating as the mechanism it is: identity is not an extra postulate about particles but a consequence of uniqueness, because a basin has no interior degrees of freedom for two readouts of it to

differ by. The charged-lepton stratum is $\mathbb{CP}^2$ and is admissible, so the electron’s case rests on Corollary 25.7 directly (§25.5).

- **The proton is stable** because, with $U(1)_B$ gauged (GG-A9-SM), the baryonic stratum’s ground attractor is the lightest $B = 1$ configuration: gauge invariance removes every operator that would connect it to a deeper basin, so “no deeper basin” holds not approximately but exactly. The baryonic stratum’s Morse structure is not developed in the present paper, so this reading is carried at the grade of its own stratum rather than derived from Corollary 25.7; the dedicated treatment is the companion Arisaka-GG-P6-Proton.
- **The Kaxion is stable** (and with it the dark matter of the predictions branch) because it sits in a locked phase of a one-dimensional periodic stratum, whose only loss channel is the portal-suppressed leak off the stratum, with a lifetime tens of orders of magnitude beyond the age of the universe (quantified in the companion paper Arisaka-GG-C1-Kaxion). This case is stated more carefully than the other two, and deliberately. The Wilson-line stratum is the class Theorem 26.1 defers, its potential is periodic over the discrete OGN-5 phases, and a periodic potential may carry several minima of equal depth; so Corollary 25.7 is not available here and no uniqueness claim is made. Which field’s phase the stratum carries is also not settled: the consequence paper now derives the Kaxion as the tensor remnant rather than as the $U(1)_{Q_i}$ Wilson line this paper earlier placed on Σ_θ , and the anomaly paper fixes both Abelian circle holonomies as background data (Remark 11.1). What is claimed — stability of the locked phase against every channel within a one-dimensional periodic stratum — needs neither uniqueness nor the identity of the phase, and is what the dark-matter argument uses.

The same clause explains the *reproducibility* of physics: the gear ratios and coupling locks of the OGN-5 ledger are properties of *the* GGA of each stratum — as the Rydberg is a property of *the* hydrogen ground state — rather than averages over a population of attractors that could drift or disperse. Wherever this paper’s earlier sections say “the GA of the stratum” in a context where uniqueness is doing work (the first-generation fermions, the electroweak vacuum, the Kaxion phase), the reader should henceforth read the sharper statement: *the* GGA.

31.6 The two regimes of the runtime: GRIP-In and GRIP-Out

The runtime defined in this paper operates in exactly two regimes, distinguished by one question: *does the stratified landscape change?* Since the present paper is the parent of the GRIP concept, the *definitions* belong here, at runtime level, together with the one structural fact that distinguishes them (Remark 31.7). The capstone GG-A11-Origin (Arisaka, 2026, A11-Origin) carries the Lyapunov descent theorem T-GRIP-In-Lyapunov-Descent and the full seventeen-event theory of landscape reorganization.

Definition 31.5 (GRIP-In: the fixed-landscape regime). GRIP-In is the dynamics of the septuple of Definition 30.6 on a *fixed* stratified landscape $(\mathcal{M}_{\text{bdy}}, \mathcal{F}_{\text{GFF}})$: the descent that carries any initial datum down the chain $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$ of its stratum, terminating at that stratum’s ground attractor. Its anatomy is piecewise. *Within* a basin the dynamics is the Riemannian gradient flow (154); *between* basins the descent proceeds by the instanton trajectories of **GRIP-ST3**, paying the

tolls $S = \pi k$ at the GSSB gates. On an admissible stratum GRIP-In terminates, and terminates uniquely: by Corollary 25.7 the descent halts exactly at the stratum's ground state — equivalently, at the unique attractor from which no admissible exit channel remains. The GGA is, dynamically, *where GRIP-In halts*.

Definition 31.6 (GRIP-Out: the landscape-reorganizing regime). GRIP-Out is the complementary regime: a transient, out-of-equilibrium event that *reorganizes the stratified landscape itself* — opening strata and basins that did not exist on the prior landscape — and thereby supplies the initial datum $c_{\text{top}} = \text{DGI}$ from which a new GRIP-In descent begins. GRIP-Out is not generated by the gradient flow and is deliberately not an element of the septuple (160). Its necessity for cascade ascent is Remark 31.7 below, and its standing with respect to Conley's theorem — outside it as a single flow on the strata, inside its hybrid extension — §31.7; its full theory — the sixteen landscape-reorganizing events from the GUT exit to the conscious episode — resides in GG-A11-Origin (Arisaka, 2026, A11-Origin). Within the present paper, GRIP-Out appears in exactly one role: it is where each stratum's c_{top} comes from.

Remark 31.7 (Why a second regime is necessary at all). The necessity of GRIP-Out is a one-line topological fact, and since the present paper owns the runtime it should be carried here rather than deferred. The GRIP-In generator of **GRIP-ST2** is a vector field on a stratum, so its integral curves remain on that stratum for as long as they exist; and by **GFF-ST1** the strata are labeled by the *discrete* data of the boundary projector class, data that no continuous curve can change. A gradient trajectory therefore never leaves the stratum it starts on — which is the same statement **GRIP-ST5** makes at the far end, that after cascade closure any further transition requires leaving the stratum. Consequently a cascade cannot ascend onto a stratum that did not previously exist by gradient flow alone: some non-gradient, out-of-equilibrium event must supply the new landscape and the new c_{top} . That event is GRIP-Out, and this is the whole of what the present paper needs from it. Of the two reasons just given, the first — a vector field on a stratum keeps its integral curves on that stratum — holds in every topology and is enough; the second is a statement about the disjoint-union topology. What neither supplies is the space on which a single flow could carry the event: it would have to be larger than the strata, and this paper has not defined it (§31.7, problem M14). Which events they were, and in what order, is the capstone's subject.

The alternation. Every cascade history is an alternation of the two regimes: a GRIP-Out event opens a landscape; GRIP-In finishes the journey, halting at the GGA; the settled GGA becomes the substrate on which the next GRIP-Out can act (Cascade-Self-Sustenance, GG-A11-Origin). In one sentence: *GRIP-Out opens the landscape; GRIP-In finishes the journey*. The electroweak transition and the family ladder make the division concrete. The selection of the electroweak vacuum was a GRIP-Out event — it created the Yukawa landscape on which the three fermion generations are the critical points — while every weak decay since ($\tau \rightarrow \mu \rightarrow e$; the valley of β stability) is GRIP-In on that fixed landscape, as the next subsection develops and as GG-A9-SM deploys across the entire Standard Model.




Figure 8: The GRIP cascade and its two regimes, in one picture. The central column is the runtime of Definition 30.6: from the fully symmetric top (DGI, every gauge equivalence intact), the system descends the fixed landscape $\mathcal{F}_{\text{GFF}}$, selects branches at the GSSB saddles, and is captured into a stable basin — a Geometric Attractor (GA). The attractor tier has two levels, and the distinction carries the two-regime story of §31.6: from a *non-ground* GA (a way station: an excited atom; a muon on its own charge-superselected stratum, which on the family stratum itself is the index-2 critical point, §25.5) the descent *continues* — GRIP-In (Definition 31.5) proceeds through the basin’s exit channels, the weak interaction being Nature’s implementer on the matter strata (§31.10) — whereas at the GGA, the unique deepest basin of **GFF-ST7** (Theorem 21), the exit-channel set is empty and GRIP-In *halts*. After the halt, only GRIP-Out (Definition 31.6) can act: a transient out-of-equilibrium event reorganizes the landscape itself and supplies the initial datum of the next descent. In one sentence: GRIP-Out opens the landscape; GRIP-In finishes the journey; the GGA is where the journey ends. (Adapted from the four-move pipeline figure of GG-A11-Origin (Arisaka, 2026, A11-Origin), which deploys this alternation through eighteen cascade layers.)

Two boundary clarifications. Each is recorded here because each invites a misreading. *(i) GSSB is a gate, not a regime.* GSSB names the symmetry-breaking saddle; such gates are crossed both within GRIP-In chains (the intermediate saddles of $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$) and at GRIP-Out shocks (the electroweak transition). The regime label attaches to the dynamics — is the landscape fixed or reorganizing? — not to the gate. *(ii) Port-driven excitation is not GRIP-Out.* Pumping a state uphill through the boundary ports (producing a τ lepton or a top quark in a collider) is an excitation *within* the fixed landscape, promptly followed by GRIP-In relaxation; it reorganizes nothing. Only events that change the landscape itself are GRIP-Out. *(iii) A bifurcation is not GRIP-Out.* A drift of the landscape’s parameters within a stratum that merges two basins, or splits one, is a bifurcation of the descent on that stratum; Conley’s theory covers it, since Morse decompositions continue under perturbation (Franzosa, 1988), and it changes no discrete label. On an admissible stratum the floor cannot be lost this way at all: its uniqueness is a cohomology class, not a feature of the parameters, so there is no fold at which it disappears. Only a change of the discrete data is a stratum jump.

31.7 The descent and the fundamental theorem of dynamical systems

Conley’s theorem (Conley, 1978, Norton, 1995) states that every continuous flow on a compact metric space admits a continuous complete Lyapunov function: non-increasing along every orbit, strictly decreasing away from the chain-recurrent set, and separating the chain-recurrent components. Its semiflow form (Hurley, 1995, Lewis, 2025, Patrão, 2011) is the one a forward-only boundary dynamics needs. The theorem asserts existence and is silent about identity. The descent of this paper sits inside it at three levels, and it is worth saying at each level what the geometry adds and what it does not.

On the boundary state space the Lyapunov function is identified by Lindblad’s and Spohn’s theorems for open quantum systems (Lindblad, 1975, Spohn, 1978) as the relative entropy to the stationary state, and — when that state is the Gibbs state of the boundary Hamiltonian at the bath temperature, which detailed balance, supplied by a thermal weakly coupled bath (Kossakowski *et al.*, 1977), guarantees — as the free energy $U_{\text{GFF}} - T_{\text{GFF}} S_{\text{GFF}}$. That is Theorem **T3.2** with its hypothesis.

On an admissible stratum the free energy is the boundary dilatation charge, the Hamiltonian of the dilatation direction of the cell’s torus (Theorem 24.1). Its chain-recurrent set is its finite critical set, so it is itself a complete Lyapunov function; its Morse decomposition is the decomposition into the stable sets of the critical points — the cells of Białynicki-Birula (Białynicki-Birula, 1973); its connection matrix vanishes, because no two critical points have adjacent index (Frankel, 1959); and its floor is unique (Atiyah, 1982). The descent within the stratum is, in moment coordinates, the replicator flow of the scaling dimension on the moment simplex (the remark beside **GRIP-ST2**), and the thermal weights on its skeleton are equilibrium weights, with the partition function localized to the critical points (**GRIP-ST7**).

Where a stratum carries two floors of equal depth, the free energy remains a Lyapunov function but is no longer complete — it does not separate the two floors — and Conley’s complete function is a different, non-thermodynamic object. This is the situation of the Wilson-line stratum, on which no uniqueness claim is made (§31.5).

GRIP-Out lies outside the theorem’s hypotheses as a single flow that is the descent on the strata, and the reason is the one Remark 31.7 gives: each stratum is invariant under its own descent, in every topology, and on the disjoint union of **GFF-ST1** — or the product with a discrete factor of §24.1, which agrees with it — the discrete label is locally constant besides. A single flow that carried a GRIP-Out event would therefore have to fail to be the descent during the event, and to pass through points that lie in no stratum: it would live on a space larger than the strata, which this paper has not defined, and whose global topology it registers as open (problems M5 and M14). What *is* a theorem is the hybrid extension of Conley’s result (Kvalheim *et al.*, 2021): a flow on the strata with a closed guard set, reached in finite time, and a continuous reset map, glued into one compact space on which a complete Lyapunov function exists. Read that way GRIP-Out is the reset, at a price that must be stated: with a closed guard the reset fires before the descent halts, since no orbit reaches the ground attractor in finite time, so the halting statement of Corollary 25.7 holds on the hybrid system to within the guard’s radius, and the jump must not be defined as firing at the ground attractor. Whether the runtime is that hybrid system, a skew product over a slow clock, or an autonomous slow–fast system on a completion of the strata is the choice the capstone (Arisaka, 2026, A11-Origin) must make; its T-GRIP-In-Lyapunov-Descent is the Lindblad–Spohn level of the present hierarchy, and inherits the hypothesis of **T3.2**.

What the geometry adds to the theorem is not the descent but its identity: the stratum, the torus, the charge, the integers that pin the count of critical points, and the stratification that forces the second regime.

31.8 The two regimes and the two fibers: the engine’s strokes

The two-regime structure just defined — distinguished by one question, *does the landscape change?* — has a geometric root that the cohort’s recent campaign has made exact (Arisaka, 2026, W1-Symmetry; Arisaka, 2026, W2-Holonomy; Arisaka, 2026, A5-GDOS; Arisaka, 2026, P5-CPT), and stating it here upgrades Remark 31.7 from a dynamical observation to a topological one. The instrument is one statement: *holonomy requires closed loops, and an interval has none*. The bulk’s compact circle S^1_θ is nothing but closed loops — everything on it is holonomy, quantized and protected, and its transport is a group; the radial interval I_r can hold no holonomy at all — its transport composes only forward, a semigroup. The runtime’s two regimes are these two fibers, operating: *GRIP-In is the circle’s stroke* — dynamics on a *fixed* landscape, cycling the four moves within a stratum to the unique ground attractor, with the outcome deposited in exactly the quantities the circle protects (charges, windings, parities: the GGA’s labels); *GRIP-Out is the interval’s stroke* — the transient, one-way, landscape-reorganizing event that no gradient flow generates and no operation reverses, the semigroup acting at runtime level. The engine is the bulk in miniature: one stroke that closes and remembers, one stroke that advances and forgets.

Read this way, the necessity of the second regime (Remark 31.7) acquires its shortest proof sketch: a runtime with only GRIP-In would be all circle — it would relax, once, to one stratum’s ground state and cycle there forever, deciding nothing further; a runtime with only GRIP-Out would be all interval — perpetual reorganization with no convergence, nothing definite ever deposited. The cascade of nature requires exactly one stroke of each kind for the same topological reason the

bulk carries exactly one fiber of each kind: *becoming* needs a semigroup and *being* needs holonomies, and the alternation $\text{GRIP-In} \rightarrow \text{GRIP-Out} \rightarrow \text{GRIP-In}$ is the minimal dynamics in which a world can do both. The campaign’s determinism result then reads as a statement about this paper’s own objects: the vacuum manifold is a quantized dial (the circle’s alphabet), the GFF landscape has a unique floor per stratum (Corollary 25.7 — the Morse mechanism), and the between-stratum motion cannot re-choose (the interval’s semigroup) — three mechanisms, all defined in this paper, jointly closing every door through which randomness enters the standard accounts of symmetry breaking. The downstream payoffs are carried at their homes: the eighteen-layer cascade and its chronology — uniform in the geometry, exponential in the clock, $d \ln t / dr = 2$ — in the capstone (Arisaka, 2026, A11-Origin) and the CPT companion’s treatment of time (Arisaka, 2026, P5-CPT); the measurement arrow, where GRIP-In’s selection and the interval’s write meet in the laboratory, in the quantum companion (Arisaka, 2026, A8-QM). Within the present paper the sentence to keep is the definitional one: *the runtime has two regimes because the bulk has two fibers, and the question that separates the regimes — does the landscape change? — is the question of which fiber is doing the work.*

31.9 The two stages and the two fibers: custody of transport, custody of data

The mapping just stated classifies the two *regimes* by where their outcomes live: GRIP-In banks its result in the classes the circle protects; GRIP-Out changes the landscape and banks nothing. A second, finer mapping classifies the *stages* of the cascade itself, and it runs along a different axis — not where the outcome is kept, but what algebra the dynamics composes under. The two mappings are consistent, but only if a distinction is kept that this subsection makes explicit: *custody of transport* and *custody of data* are different things, and the cascade’s middle stage holds them on different fibers.

The descent is interval transport. The first half of the stage mapping is a composition of two results already established in the boundary-law companion, and is stated at the grade that composition earns.

Proposition 31.8 (The $\text{GSSB} \rightarrow \text{GA}$ leg realized as the scale fiber’s semigroup). *On the boundary, the descent that executes the $\text{GSSB} \rightarrow \text{GA}$ leg of the cascade — the GRIP-In flow of Definition 31.5 — is realized as a completely positive, trace-preserving forward semigroup whose generator is the dilatation generator along I_r projected to the boundary. In particular the selection performed at each GSSB gate composes under the interval’s semigroup law: forward only, with no inverse.*

Proof. Composition of two results of Arisaka (2026, A5-GDOS). First, the GDOS-to-Gi-4 bridge identifies the cascade with the long-time GDOS flow on the boundary moduli, stage for stage, with the four boundary Ward identities preserved; the identification itself is unconditional, while the $N_g \leq 3$ closure inherits the conditionality of the capacity theorem it rides on (Arisaka, 2026, P2-N3). Second, the same companion identifies the dissipative part of the GDOS (GKLS) generator as the boundary image of motion along the scale fiber — the dilatation generator along I_r , projected — which is exactly what makes the family of boundary channels a semigroup rather than a group.

Composing the two identifications gives the statement. The irreversibility clause of **GRIP-ST2** — the flow descends in $\mathcal{F}_{\text{GFF}}$ and never ascends — is thereby traced to its geometric root: the fiber the descent runs on has no closed loops, hence no inverse transport. ■

The invariance is the circle's. At the other end of the cascade, DGI holds a symmetry intact, and the quantum companion fixes exactly which geometry that symmetry belongs to: the phase direction S_θ^1 is a genuine principal $U(1)$ -bundle over the boundary — connection, holonomy, and all — while the scale fiber I_r is *not* a gauge bundle, its “structure group” being the one-way dilatation semigroup (Arisaka, 2026, A8-QM). Gauge structure, in this cohort, lives on the circle and only on the circle. The equivalences that DGI keeps intact therefore compose as the circle's transport — a group — and the constitutional ladder of POS-1 (LGI/DGI/GGI) acquires a geometric reading: LGI is the circle's invariance stated in a closed theory; DGI is the same invariance carried through the open dynamics, the equivalence class of boundary generators under induced gauge transformations in the operator form of the information companion (Arisaka, 2026, A7-Qi4). This identification is carried at derived-conditional grade: it rests on the bundle asymmetry of Arisaka (2026, A8-QM) and on reading the Morse-top stabilizer of Definition 31.1 as the boundary gauge group stratum by stratum — immediate on the gauge and Wilson-line strata, open in general (problem M12 of §57.1).

The middle stage crosses, and the crossing is the content. GSSB reads circle data and runs on the interval. Its order parameters are Wilson-line phases on S_θ^1 (Section 7.5); its execution — the one-way act of selection — is the interval's semigroup (Proposition 31.8); and its outcome is deposited back into the classes the circle protects, the GGA's labels. This is not a defect of the correspondence; it *is* the correspondence. Symmetry breaking, in the cohort's reading, is symmetry *reading* (Arisaka, 2026, W1-Symmetry): the alternatives and the record belong to the fiber that remembers; the choice belongs to the fiber that forgets. The shape is already in print elsewhere — the measurement write of the quantum companion is interval-side while the record is circle-kept (Arisaka, 2026, A8-QM) — so GSSB stands to symmetry breaking exactly as the write stands to measurement: one seam-event shape, two instances. The electron paper exhibits both custodies on a single object — charge and identity circle-side, mass and decay rate interval-side, the descent realized as an explicit GKLS semigroup whose terminus carries no jump operator (Arisaka, 2026, P3-Electron) — and compresses the correspondence to a sentence: an electron is a holonomy class wearing a mass. This paragraph is carried at the reading grade: an organization of proved structures, with the general theorem deferred to M12.

Why the two mappings do not collide. The stroke mapping of §31.8 and the stage mapping of this subsection classify the same engine along orthogonal axes. The strokes are named for their *products*: GRIP-In is the circle's stroke because its outcome is banked in circle classes; GRIP-Out is the interval's because it changes the landscape and banks nothing. The stages are named for their *transport*: the DGI equivalences compose as the circle's group; the descent composes as the interval's semigroup. Both classifications are pullbacks of the one fork — holonomy requires closed loops, and an interval has none — one along the outcome axis, one along the dynamics axis. In particular, “GRIP-In is the circle's stroke” and “the GRIP-In descent is interval transport” are both true, and

Table 4: The four stages of the GRIP engine, each with the transport algebra it runs on, the fiber that keeps its data, and the reading that follows. The third column is the one to read down: the circle keeps what is protected and the interval carries what is irreversible, and every stage divides its data between the two the same way. The table is the engine’s custody ledger — where a quantity lives decides whether it can be undone — and it is what makes symmetry breaking and measurement two instances of one seam-event rather than two separate phenomena.

Stage	Transport algebra	Data custody	Reading
DGI	the circle’s group: gauge equivalences intact	bulk-wide symmetry; the phases made dynamical are S_θ^1 ’s	the invariance is the circle’s; the dynamism is the interval’s
GFF	none — the landscape, not a transport	fed by both fibers: GEP through the Wilson line; $\Phi_\omega = \pi$ from warp $\times$ circumference	the seam’s price list
GSSB	the interval’s semigroup (Prop. 31.8)	order parameter and alternatives on the circle; outcome deposited to circle classes	symmetry breaking as symmetry reading
GA/GGA	the halt of the semigroup	labels circle-protected (charges, windings, parities)	where the interval stops and the circle keeps

neither sentence can replace the other: the first says where the result is kept, the second says what algebra the getting-there composes under. The sentence to keep is the custody sentence: *the circle states the alternatives; the interval performs the selection; the attractor is the circle’s receipt.*

31.10 GRIP ending at the GGA, realized in Nature: the weak interaction

The cascade of this paper is not an abstraction awaiting application; Nature runs it daily, and the runner has a name. Among the four boundary forces, the strong and electromagnetic interactions are flavor-diagonal: they *build* the basins of the matter strata (hadrons, nuclei, atoms) and relax excitations *within* a basin ($\rho \rightarrow \pi\pi$; an excited atom radiating to its ground state), but they cannot carry a state from one family basin to another. The *weak charged current* is the unique flavor-connecting channel of the boundary group — and therefore the unique physical implementer of the fixed-landscape regime defined above: the weak charged current is how Nature runs GRIP-In (Definition 31.5) on the matter strata. Two clean clauses summarize the categorization: every inter-basin descent observed in Nature is weak, and every weak flavor transition is a step of the $\mathcal{F}_{\text{GFF}}$ descent toward the stratum’s GGA. The division of labor is exactly that of Definition 31.4: the GGA fixes the *direction* of relaxation, the open-system boundary law supplies the *arrow*, and the weak coupling sets the *rate*.

The canonical example deserves to be named here because it is the most photographed Morse landscape in science: the valley of β stability. At fixed mass number A , the nuclear binding energy is a parabola in proton number Z , and β^- , β^+ , and electron-capture transitions step a nuclide down that parabola, one weak transition at a time, to the most-bound isobar — the isobaric GGA. The chart of nuclides *is* a stratified GFF landscape; β decay *is* its gradient flow. The same channel

tracks the *conditional* ground state when the stratum changes: in a neutron star, where the free energy includes density and degeneracy pressure, the very same interaction runs “uphill” in vacuum masses ($p + e^- \rightarrow n + \nu$) because neutron matter is the GGA of *that* stratum — the strongest evidence that the weak interaction descends $\mathcal{F}_{\text{GFF}}$ rather than a fixed mass table. The family ladders ($\tau \rightarrow \mu \rightarrow e$; $t \rightarrow b \rightarrow c \rightarrow s \rightarrow u/d$) are the elementary instances, developed with their lifetime and ordering consequences in GG-A9-SM, which devotes a dedicated subsection to this reading and its examples; the dark sector completes the pattern, the Kaxion condensate relaxing toward its GGA by Hubble friction itself (the companion paper Arisaka-GG-C1-Kaxion). In one sentence: *the strong and electromagnetic forces build the landscape; the weak force is how matter walks downhill in it; and the walk ends, always, at the GGA*. And in regime language (§31.6): every weak process in Nature is a step of GRIP-In; no weak process is a GRIP-Out. The one GRIP-Out of the electroweak story happened once, when the vacuum was selected and the landscape built; the weak interaction has been walking that landscape ever since.

32 Cascade dynamics: gradient flow and instanton trajectories

We now formalize the dynamical content of GRIP in detail. Everything in this section concerns the GRIP-In regime (Definition 31.5): the two halves that name the section — gradient flow within basins and instanton trajectories between them — are precisely the piecewise anatomy of GRIP-In. The landscape itself is fixed throughout; GRIP-Out events sit outside the septuple (Definition 31.6) and enter only as the suppliers of the initial data c_{top} .

32.1 The gradient flow on stratified GFF

The cascade (153) is generated by the gradient flow (154) of $\mathcal{F}_{\text{GFF}}$ on each stratum. Concretely, for a point $x \in \Sigma$ that is not a critical point of $\mathcal{F}_{\text{GFF}}|_{\Sigma}$, the flow

$$\dot{x} = -g_{\Sigma}^{ab}(x) \partial_b \mathcal{F}_{\text{GFF}}(x) \partial_a, \quad (164)$$

descends along the steepest-descent direction of the Riemannian metric g_{Σ} . The flow has the following properties:

- **Monotonicity.** $d\mathcal{F}_{\text{GFF}}(x(t))/dt \leq 0$, with equality iff x is a critical point. This is the irreversibility of the cascade.
- **Asymptotics.** As $t \rightarrow \infty$, $x(t)$ approaches a critical point of $\mathcal{F}_{\text{GFF}}|_{\Sigma}$. Generically, the asymptotic limit is the unique GA of the basin containing the initial condition.
- **Trajectory connecting saddles.** Between two adjacent critical points c_+, c_- with $\text{ind}(c_+) - \text{ind}(c_-) = 2$, there exists (generically) a one-parameter family of trajectories from c_+ to c_- along the unstable manifold of c_+ . These are the “connecting trajectories” of Morse homology.

32.2 Instanton action between adjacent saddles

Theorem 32.1 (Instanton action on Kähler strata). *Let Σ be a Kähler stratum equipped with the canonical height-function Morse form $\mathcal{F}_{\text{GFF}}(x) = \sum_a \lambda_a |\psi_a|^2 / \sum |\psi_a|^2$. The instanton action of the gradient-flow trajectory from the index- $2k$ critical point to the index- $2(k-1)$ critical point is*

$$S_{[2k] \rightarrow [2(k-1)]} = \pi k. \quad (165)$$

This action saturates the Bogomolnyi bound for gradient flows on Kähler manifolds.

Proof. We derive $S_{[2k] \rightarrow [2(k-1)]} = \pi k$ by direct computation on $\mathbb{CP}^{N-1}$, working out the Fubini–Study symplectic area enclosed by the Bogomolnyi-saturated gradient trajectory between adjacent fixed points.

Step 1: Bogomolnyi-type inequality. The gradient flow on a Kähler manifold with Hamiltonian-Morse form $\mathcal{F}_{\text{GFF}}$ satisfies

$$S \geq \int_{c_+}^{c_-} d\mathcal{F}_{\text{GFF}} = \mathcal{F}_{\text{GFF}}(c_+) - \mathcal{F}_{\text{GFF}}(c_-), \quad (166)$$

with equality saturated by the gradient trajectory itself. This is the standard Bogomolnyi bound for Euclidean gradient flow: $S = \frac{1}{2} \int (|\dot{x}|^2 + |\nabla \mathcal{F}|^2) dt \geq \int |\nabla \mathcal{F}| |\dot{x}| dt = \int d\mathcal{F}$, with equality iff $\dot{x} = -\nabla \mathcal{F}$.

Step 2: Height function on $\mathbb{CP}^{N-1}$ in homogeneous coordinates. Let $[\psi_0 : \psi_1 : \dots : \psi_N]$ be homogeneous coordinates on $\mathbb{CP}^{N-1}$ and let $f([\psi]) = \sum_{a=0}^N \lambda_a |\psi_a|^2 / \sum_{a=0}^N |\psi_a|^2$ be the moment-map Morse function induced by the maximal-torus action of $U(N+1)$ with weights $\lambda_0 < \lambda_1 < \dots < \lambda_N$. The critical points of f are the $N+1$ coordinate axes $[0 : \dots : 1_k : \dots : 0]$, with λ_k the critical value and Morse index $2k$ (counting the unstable complex directions). The index- $2k$ critical point hosts the winding- $\ell = k+1$ generation (three generations = three windings). The derivation below takes the Morse height of that point to be the pair invariant of its winding, $\lambda_k = \binom{k+1}{2} = \frac{1}{2}k(k+1)$, so that $\lambda_k - \lambda_{k-1} = k$ and the spacings grow linearly down the chain; the pair invariant is the inflow charge of the saddle, $\Delta c_2 = k$, derived in Remark 23.5 from the Whitney pair count of the winding bundle, and the unit-spaced alternative is excluded there on the toll arithmetic. That the inflow charge is also the Morse height is the identification on which Steps 4–5 rest, and it is the identification Theorem 24.1 excludes under the paper’s transport proposition: the status remark after this proof says what survives.

Step 3: Bogomolnyi-saturated trajectory between adjacent critical points. The gradient-flow trajectory from the index- $2k$ fixed point $p_k = [0 : \dots : 1_k : \dots : 0]$ to the index- $2(k-1)$ fixed point p_{k-1} traverses the projective line $\mathbb{CP}^1 \subset \mathbb{CP}^{N-1}$ spanned by ψ_{k-1} and ψ_k . In affine coordinates $z = \psi_{k-1}/\psi_k$ on this projective line, the moment map restricts to

$$f|_{\mathbb{CP}^1}(z) = \frac{\lambda_k + \lambda_{k-1}|z|^2}{1 + |z|^2} = \lambda_k - (\lambda_k - \lambda_{k-1}) \frac{|z|^2}{1 + |z|^2} = \lambda_k - k \Omega(z), \quad (167)$$

where $\Omega(z) = |z|^2/(1 + |z|^2) \in [0, 1]$ is the symplectic potential of the Fubini–Study form on $\mathbb{CP}^1$.

Step 4: Symplectic area of the trajectory. The Bogomolnyi bound (166) equals

$$S_{[2k] \rightarrow [2(k-1)]} = f(p_k) - f(p_{k-1}) = (\lambda_k) - (\lambda_{k-1}) = k \quad (168)$$

in the Fubini–Study units of Step 2; we now convert to action units. By the Duistermaat–Heckman formula, the symplectic volume of the trajectory’s $\mathbb{CP}^1$ projective line is

$$\int_{\mathbb{CP}^1} \omega_{\text{FS}} = \pi, \quad (169)$$

with ω_{FS} the standard Fubini–Study Kähler form. Each Bogomolnyi-saturated step traverses one such $\mathbb{CP}^1$, so each step carries action π *per unit eigenvalue spacing*.

Step 5: Conversion to action units. This step is a choice of units for the height, and is named as one. Duistermaat–Heckman (Duistermaat & Heckman, 1982) pins the honest moment map of a circle of weight jump k on a line of area π at a drop of $k/2$; an action of πk therefore requires the height to be π times the pair invariant, that is, the Fubini–Study form scaled by 2π , and Step 2’s “not a normalization convention” applies to the weights, not to this unit. With that unit, combining Step 4’s drop k with the area quantum π per unit eigenvalue spacing, the action of the single transition $[2k] \rightarrow [2(k-1)]$ saturates the Bogomolnyi bound at

$$S_{[2k] \rightarrow [2(k-1)]} = \pi k, \quad (170)$$

directly: the factor k is the weight jump $\lambda_k - \lambda_{k-1} = \binom{k+1}{2} - \binom{k}{2}$ of Step 2 — equivalently, the inflow charge $\Delta c_2 = k$ of Remark 23.5 — and no stacking or depth-counting convention is invoked. Descending the full chain, the accumulated actions are 0, 2π , 3π , saturating the BI-6 budget 3π exactly; a fourth generation ($\ell = 4$, $\lambda = 6$) would cost a further 3π , overrunning the budget — the $N_g \leq 3$ forcing in its sharpest form. This completes the derivation.

The derivation reproduces Witten (1982, §4) on $\mathbb{CP}^{N-1}$ and is a special case of the Atiyah–Bott / Frankel localization principle for equivariant moment-map Morse functions on Kähler manifolds (Atiyah, 1982, Frankel, 1959). ■

Remark 32.2 (Status of the theorem after Theorem 24.1). Steps 1, 3 and 4 are theorems about any moment-map height on $\mathbb{CP}^{N-1}$: the action of the connecting orbit is the height difference, and the orbit runs along a projective line of area π . Step 2’s assignment of the pair invariant to the height, and Step 5’s unit, are the two identifications, and Theorem 24.1(c) excludes the first under Proposition 31.8: the height on the family stratum is $c \Delta_+(\ell)$, with gaps $0.592c$ and $0.777c$, and no unit makes those π and 2π , because their ratio is 1.31 and not 2. Under that proposition, therefore, the Bogomolnyi action of the connecting orbit from the index- $2k$ point is $c [\Delta_+(k+1) - \Delta_+(k)]$, and the number πk is the inflow toll of the gate, $\pi \Delta c_2$ — a Chern number that the crossing pays whatever the height is. Everything this paper derives from the value πk — **GRIP-ST3**, the budget closure of **GRIP-ST4**, the triangular sum and $N_g \leq 3$ of Theorem T4.1 — consumes it as the toll and survives; what does not survive is the sentence that the toll is the height difference, and the “three π ’s” of Remark 23.5 become two. The alternative — keeping the pair invariant as the height and giving up hypothesis (i) — is recorded as the other horn of open problem M13. The unit of the

height itself, the c above and the temperature T_{GFF} in which it is measured, is not derived in this paper (Remark 25.2), and the weights on the skeleton are numbers only at $T_{\text{GFF}} = 1$ in height units.

32.3 Saturation of the Bogomolnyi bound

The instanton action (165) is the *minimum* action of any trajectory from index $2k$ to index $2(k-1)$. Higher-action trajectories exist (involving detours through other saddles or oscillations around critical points), and under boundary noise they carry the smaller weight, e^{-S} in height units at $T_{\text{GFF}} = 1$ (**GRIP-ST7**); the deterministic descent follows the saturating trajectory itself.

This is what makes the cascade “minimal”: the system follows the shortest path from DGI to GA, accumulating the smallest possible total action, and the toll budget of **GRIP-ST4** is stated for that path.

33 The physical layer of GRIP: gauge-runtime invariance and the generation weight ladder

This section develops the *physical* content of GRIP (Structure Theorems **GRIP-ST6** and **GRIP-ST7**) in operational detail: the *Gauge-Cascade Compatibility Theorem* that guarantees runtime gauge-invariance preservation along the cascade, the derivation of the gate weight (159) and of what it is not, the numerical confrontation with the empirical Standard-Model lifetimes, and laboratory-accessible falsifiers. The mathematics here combines instanton calculus with boundary Ward-identity analysis; the conceptual novelty is that the algebraic cascade and the physical generation weight ladder are the same object — the skeleton with its tolls — viewed once as a descent and once as an equilibrium distribution.

33.1 The Gauge-Cascade Compatibility Theorem

Theorem 33.1 (Gauge-Cascade Compatibility Theorem). *Let $\Sigma \subset \mathcal{M}_{\text{bdy}}$ be a stratum supporting the algebraic cascade $c_{\text{top}} \rightarrow c_1 \rightarrow \cdots \rightarrow c_{\text{bottom}}$ of **GRIP-ST1–GRIP-ST5**. Let $C_\mu(t)$ be the PBI-4 compensator field of **CCI-ST6** evaluated along the cascade trajectory $t \mapsto [\Pi(t)] \in \Sigma$. Then the four boundary Ward currents (J_e, J_B, J_{Qi}, J_g) satisfy the runtime relation*

$$\nabla_\mu J_a^\mu(t) = -\alpha_a C^\mu(t) \cdot \rho_a(t), \quad a \in \{e, B, Qi, g\}, \quad (171)$$

*along the entire cascade $t \in [t_0, t_\infty]$, with the α_a the **CCI-ST2** coefficients. Moreover, the unbroken gauge subgroup $H_i \subseteq G_{\text{tot}}$ at saddle c_i is contained in the unbroken subgroup at the previous saddle: $H_i \subseteq H_{i-1}$.*

Proof. The gradient flow **GRIP-ST2** is generated by the gauge-equivariant functional $\mathcal{F}_{\text{GFF}}$, which inherits its gauge invariance from the bulk TGG action via bulk-to-boundary descent (Part II). Therefore the flow vector field $-g_\Sigma^{-1} \nabla \mathcal{F}_{\text{GFF}}$ is itself gauge-equivariant, and any gauge symmetry of the initial configuration c_{top} that is preserved by $\mathcal{F}_{\text{GFF}}$ is also preserved by the flow. This gives the descending chain of unbroken subgroups $H_0 \supseteq H_1 \supseteq \cdots \supseteq H_n$.

For the runtime Ward divergence (171), we apply the Master Conservation Theorem **T2.2** of Part III at each instant t along the cascade. By **T2.2**, the four sector divergences are linked by $C_\mu(t)$ via Proposition 17.4. This gives the desired runtime relation (171). The continuity of $C_\mu(t)$ across each instanton transition has two parts. The *global* (topological) part — the integrality of the boundary class $[\tilde{X}_2]$ inherited from the bulk $[X_4]$ via KK reduction — is established in Arisaka (2026, A3-BI6). The *local* smoothness (monodromy-freeness) of $C_\mu(t)$ across each instanton transition is, however, not established in the present paper or elsewhere in the cohort; we take it here as a standing hypothesis, recorded as open problem M11 in §57. The runtime relation (171) — and with it **T4.2** and the present theorem — is therefore conditional on M11. ■

Remark 33.2 (Physics intuition: gauge invariance is a flow-equivariance statement). The Gauge-Cascade Compatibility Theorem is essentially the statement that gauge invariance is preserved under any flow that preserves the free-energy functional. This is a familiar theorem in classical mechanics (Noether’s theorem applied to gradient systems): if $\mathcal{F}_{\text{GFF}}$ has a continuous symmetry group G , then the gradient flow $\dot{x} = -\nabla \mathcal{F}_{\text{GFF}}$ commutes with the G -action, so trajectories stay in G -equivariant orbits.

The new content in GG-Theory is that the gauge group is the *boundary* part of the bulk $G_{\text{tot}} = \text{Spin}(1, 5) \times \text{SU}(5) \times \text{U}(1)_B \times \text{U}(1)_{Q_i}$, and the runtime gauge invariance is testable via the boundary Ward identities of Part III. This is what makes **GRIP-ST6** an empirical statement rather than a purely mathematical equivariance principle.

33.2 Cascade-action integration: from $S_{c_i \rightarrow c_{i+1}}$ to the gate weight

We now connect the cascade gate tolls of **GRIP-ST3** to the weights on the family skeleton, recovering the geometric factor $e^{-\pi k}$ of (159) and saying what it is and is not; the generation reading follows in §35.2 and the absolute lifetimes, which are Standard-Model quantities, in Arisaka (2026, A9-SM).

Setup: the gate weight is an equilibrium weight

Perturb the descent on an admissible stratum by boundary noise of strength $\sqrt{2T_{\text{GFF}}}$, isotropic in the Kähler metric of the flow. Over any fixed time window a noisy trajectory shadows the deterministic orbit to within a deviation of order $\sqrt{T_{\text{GFF}}}$; and because the deterministic flow has one asymptotically stable critical point — the floor — the noisy dynamics has no Markov chain of metastable states to hop between. It equilibrates to the Gibbs weight $e^{-\mathcal{F}_{\text{GFF}}/T_{\text{GFF}}}$, and the stationary density at a critical point p relative to the floor is

$$w(p) = \exp[-(\mathcal{F}_{\text{GFF}}(p) - \mathcal{F}_{\text{GFF}}(\text{floor}))/T_{\text{GFF}}], \quad (172)$$

with no prefactor. Because $\mathcal{F}_{\text{GFF}}$ is a torus moment-map component, Duistermaat–Heckman (Duistermaat & Heckman, 1982) makes the partition function exact: $\int_{\Sigma} e^{-\mathcal{F}_{\text{GFF}}/T_{\text{GFF}}} \omega^n/n!$ is $(2\pi T_{\text{GFF}})^n$ times the sum over the critical points of $e^{-\mathcal{F}_{\text{GFF}}(p)/T_{\text{GFF}}}$ divided by the product of the isotropy weights at p — on the family stratum $\prod_{j \neq k} (\Delta_j - \Delta_k)$ at c_k , negative at the index-2 point. The terms

at the saddle and the top are signed coefficients of the partition function, not occupations; the equilibrium share of generation k is $\langle p_k \rangle = -T_{\text{GFF}} \partial \ln Z / \partial \Delta_k \simeq T_{\text{GFF}} / (\Delta_k - \Delta_1)$ to leading order, linear in the temperature and not exponential. The factor of (159), $e^{-\pi k}$, is this weight with the toll as the exponent and $T_{\text{GFF}} = 1$ in height units; with the height of Theorem 24.1 the exponent is $c[\Delta_+(k+1) - \Delta_+(k)]$ instead, and which of the two exponentials the mass hierarchy tracks is GG-A9-SM's question. A lifetime

$$\tau_{i \rightarrow i+1} = \tau_0 \cdot e^{\pi(i+1)}, \quad \tau_0 \sim 1/M_{\text{KK}} \sim 10^{-28} \text{ s}, \quad (173)$$

which earlier statements of this paper read off the weight as a WKB rate $A_i e^{-S}$ (Callan & Coleman, 1977, Coleman, 1977), has no mechanism behind it on an admissible stratum: Coleman's false vacuum is a local minimum behind a barrier, and the index- $2k$ point is a saddle, unstable at linear order, from which the descent is downhill and deterministic. The paper therefore does not derive (173), and carries it only so that its downstream uses can be read at this grade.

Cumulative gate toll to depth i

The tolls add along the chain: from the symmetric top ($i = 0$) to the depth- i point the cumulative toll is

$$S_{\text{tot}}^{(i)} = \sum_{k=1}^i \pi k = \frac{\pi i(i+1)}{2}, \quad w^{(i)} = e^{-S_{\text{tot}}^{(i)}}, \quad (174)$$

and $w^{(i)}$ is the weight of the depth- i point relative to the top in toll units, the product of the gate weights. It is a height arithmetic, not a timescale: the earlier reading of $\tau_0 e^{S_{\text{tot}}^{(i)}}$ as a coherent multi-stage tunneling time is withdrawn with the rate reading above.

From gate tolls to the family ladder

The physical content of the gate tolls is the ordering of the weights $e^{-\pi k}$ along the *generation* cascade on $\mathbb{CP}^2$ — the charged-lepton chain $\tau \rightarrow \mu \rightarrow e$ (and the quark chain $t \rightarrow b \rightarrow c \rightarrow s \rightarrow u/d$) of §35.2. There is a *single* cascade in GRIP: the family cascade on the family-projector stratum. The absolute widths and lifetimes of the generations are Standard-Model quantities, dominated by the species-specific phase space (Sargent's rule $\Gamma \propto G_F^2 Q^5$) and computed in Arisaka (2026, A9-SM); whether the gate weight enters any width at all is the open mechanism question of **GRIP-ST7**. Ordinary unstable states that are *not* family-projector critical points — the W boson, the free neutron — are not cascade-index particles; they decay as standard weak (GRIP-In) relaxations, and the cohort quotes their lifetimes as empirical facts rather than as cascade outputs. The worked generation numbers, including the phase-space resolution of the τ/μ ratio, are given in §35.2 and Arisaka (2026, A9-SM).

The structural content of **GRIP-ST7** is therefore the *ordering* of the generation weights: each gate $[2k] \rightarrow [2(k-1)]$ contributes $e^{-\pi k}$, so the deeper (heavier-generation) point carries the smaller weight. Because the species-specific phase-space factors vary by many orders of magnitude — and reverse the naive geometric ordering for the charged-lepton lifetimes (§35.2) — nothing about the weight can be read off from absolute lifetimes. The controlled-kinematics test of earlier statements —

a class of decays in which the Sargent factor is a uniform multiplier across gates, so that a residual $e^{-\pi k}$ ordering would survive — is a test of the rate reading, and is carried below at the grade of that reading.

Remark 33.3 (Physics intuition: the gate weight is species-independent). The geometric content of the cascade is that each generation gate $[2k] \rightarrow [2(k-1)]$ carries the same toll πk , independently of which family member sits behind it. The phase-space factors $G_F^2 Q^5$ set the absolute decay rate per species and are the Standard Model's, not the cascade's; what is universal in the cascade is the height ordering of the generations, which every family ladder shares. The collapse of weak-decay rates onto a single Q^5 curve is Sargent's rule, and this paper no longer claims a geometric factor to be divided out of it.

33.3 Laboratory-testable predictions and falsifiers

The physical block of GRIP produces three families of laboratory falsifiers, all derivable from **GRIP-ST6–GRIP-ST7**.

Falsifier 4 (F-GRIP-4): the generation ordering, and the ratio of the gaps

Equation (159) assigns each generation gate the weight $e^{-\pi k}$, and Theorem 24.1 assigns the critical points their heights. The testable content has two halves, and GG-A9-SM executes both (Arisaka, 2026, A9-SM). *The ordering*: in each charged sector the generations must be ordered by mass as the heights order them along $\tau \rightarrow \mu \rightarrow e$ and $t \rightarrow c \rightarrow u$, $b \rightarrow s \rightarrow d$; a sector whose generations were not so ordered would falsify **GRIP-ST7**. No sector is. *The ratio*: if the Yukawa of a generation is the equilibrium weight e^{-S} with S the height difference from the top, the two log-mass gaps of a sector stand in the ratio of the height gaps of that sector's action, and for this paper's own action — the dilatation charge — that ratio is 1.31. Read at one scale ($\overline{M_S}$ at M_Z) the sectors give 0.53, 0.90 and 1.33 for the charged leptons, the up quarks and the down quarks: the down quarks meet the ratio inside the spread of the published compilations, 1.30–1.33, which is interesting, and is not proof; the up quarks do not; and the charged leptons, whose skeleton GG-A9-SM reads as dyadic rather than exponential, are not on this ladder. What the ratio half says is therefore not that the weights are wrong but that one action does not serve all sectors — the family action must differ sector by sector, with the isospin factor entering the weights themselves — which is the content of GG-A9-SM's cascade-depth lemma as now restated. A sector whose gaps stood in a ratio no torus action on its stratum could produce would falsify the reading outright; none does. The controlled-kinematics decay test of earlier statements — a decay class in which the Sargent phase space is a uniform multiplier, so that a residual $e^{-\pi k}$ ordering of the widths would survive — tests the rate reading of the weight, which this paper no longer derives (§33.2); it stands at the grade of that reading, as a falsifier of a mechanism the paper does not supply.

Falsifier 5 (F-GRIP-5): runtime Ward-identity violations

Equation (171) predicts that all four boundary Ward currents remain closed (modulo the calculable PBI-4 anomaly inflow) along every instanton transition. Empirically this corresponds to:

conservation of electric charge in each radioactive decay; near-conservation of baryon number (broken only by the calculable instanton-mediated $B + L$ anomaly of the Standard Model); conservation of total information (no information loss at runtime); and conservation of stress-energy. An unambiguous empirical observation of any *spontaneous* (i.e., not explained by known anomaly inflow) Ward-identity violation in any decay process would falsify **GRIP-ST6**.

Falsifier 6 (F-GRIP-6): cascade-action sum rule

A more subtle falsifier comes from the sum rule

$$\sum_{i=1}^{N_g-1} S_{\text{tot}}^{(i)} = \frac{\pi N_g (N_g - 1)(N_g + 1)}{6}, \quad (175)$$

obtained by summing $\sum_i i(i+1)/2 = N(N-1)(N+1)/6$. For $N_g = 3$ this gives $\sum = 4\pi$, a toll budget summed over depths and not a timescale; the earlier remark that it is comparable to the age of the universe in dimensionless units compared two numbers with no common unit, and is withdrawn.

Remark 33.4 (Physics intuition: what the falsifiers test). The three falsifiers test the three pillars of the physical block: Falsifier 4 (F-GRIP-4) tests the *cascade structure* (Bogomolnyi action quantization); Falsifier 5 (F-GRIP-5) tests the *runtime gauge invariance* (Ward identity preservation); Falsifier 6 (F-GRIP-6) tests the *cascade closure* (finite total action $S_{\text{Bl6}}^{\text{max}} = 3\pi$). Together, these three falsifiers cover the entire physical content of **GRIP-ST6–GRIP-ST7**.

The most direct empirical handle is Falsifier 4 (F-GRIP-4): the generations must be ordered as the weights order them, and the log-mass gaps of a sector must stand in the ratio of its action's height gaps; the mass hierarchy is where both are read, and GG-A9-SM reads them (Arisaka, 2026, A9-SM) — the ordering holds in every sector, the ratio of this paper's own action in one. The residual-width form of the test — a controlled-kinematics decay class in which the gate factors would have to order as $e^{-\pi k}$ once the phase space is divided out — is a test of a rate reading of the weight, and a measured residual ordering contradicting it would challenge that reading, which this paper does not derive, before it challenged GRIP.

34 Theorem T4 (T-GRIP-4): GRIP cascade closure forces $N \leq 3$ and gauge-runtime invariance

We now state and prove Theorem **T4**, the principal theorem of Part V. The theorem is presented in two parts: **T4.1** (algebraic cascade closure forcing $N \leq 3$, the original statement) and **T4.2** (runtime gauge-invariance preservation, introduced to capture the physical content of GRIP).

Dual-layer architecture and the compatibility relation $\mathfrak{C}_{\text{GRIP}}$. Per the L16 architecture, T4.1 is the *algebraic layer* (cascade closure on the family-projector moduli $\Sigma_{\text{fam}} = \mathbb{CP}^{N-1}$ with total instanton action $S_{\text{tot}} \leq S_{\text{Bl6}}^{\text{max}} = 3\pi$, forcing $N_g \leq 3$), and T4.2 is the *operational/physical layer* (runtime gauge-invariance preservation along the cascade, expressed as an iff statement involving

the PBI-4 compensator extension across each instanton crossing). The two layers are tied together by the named compatibility relation

$$\boxed{\begin{array}{c} \mathfrak{C}_{\text{GRIP}} \equiv \text{T4.1} \implies \text{T4.2} \\ \text{(via the PBI-4 compensator extension across the instanton stratum } \Sigma_{\text{inst}} \subset \bar{\Sigma}_{\text{fam}} \text{).} \end{array}} \quad (176)$$

$\mathfrak{C}_{\text{GRIP}}$ above is a one-line proposition with the proof: the algebraic closure of the cascade at $N = 3$ (T4.1) forces the cumulative cascade action to saturate the BI-6 budget, leaving no room for an instanton-crossing discontinuity in the PBI-4 compensator C_μ ; runtime gauge-invariance preservation (T4.2) then holds iff C_μ extends smoothly across each instanton stratum. The smooth-extension is a topological non-obstruction hypothesis: the monodromy of C_μ around the codimension-2 instanton stratum must vanish, which is established case-by-case for the cascade saddles on $\mathbb{CP}^{N-1}$ (open problem M11 below tracks the general topological argument).

34.1 Statement of T4.1 (algebraic cascade closure)

Theorem 34.1 (T4.1: GRIP cascade closure on family-projector moduli). *Let $\Sigma_{\text{fam}} = \mathbb{CP}^{N-1}$ be the family-projector stratum of the boundary moduli, equipped with the GFF Morse landscape and the GRIP cascade structure. Cascade closure under **GRIP-ST4–GRIP-ST5** requires*

$$S_{\text{tot}}(\mathbb{CP}^{N-1}) = \pi \frac{N_g(N_g - 1)}{2} \leq S_{\text{BI6}}^{\text{max}} = 3\pi, \quad (177)$$

which forces

$$N(N - 1) \leq 6 \iff N \leq 3. \quad (178)$$

Combined with the lower bound $N \geq 3$ from CCI (Theorem T2 of Part III) and the strict Morse-positivity from GFF (Theorem T3 of Part IV), the family number is forced to be uniquely $N_g = 3$ (the upper bound $N_g \leq 3$ conditional on the budget value $S_{\text{BI6}}^{\text{max}} = 3\pi$, open problem M10; see Remark 34.2).

Remark 34.2 (Status: T4.1’s budget input and its provenance). The bound $N \leq 3$ derived above rests on the input value $S_{\text{BI6}}^{\text{max}} = 3\pi$ for the BI-6 instanton-action budget — now the capacity theorem T-N3-6 of Arisaka (2026, P2-N3), conditional on its coupling clause and the integrality theorem L-INT (= M6 of this paper; residue R-INT). The natural axiomatization would posit this cap as a named conjecture; the capacity architecture of Arisaka (2026, P2-N3, Appendix C) records the two viable formulations and their completion (the sharpened per-family form is $S(N_g) = \pi N$), and the value is taken as input in the present paper. An in-paper derivation of $S_{\text{BI6}}^{\text{max}} = 3\pi$ that is independent of $N = 3$ at every intermediate step — so as to definitively rule out any circularity in the cascade-closure $N \leq 3$ argument — is recorded as open problem M10 in Section 57, whose demand is met by the capacity theorem T-N3-6 of Arisaka (2026, P2-N3). T4.1 is accordingly a conditional theorem that imports its decisive numerical input from a companion paper. This honest status note is added in to ensure that no downstream reader mistakes the $N \leq 3$ bound for an absolute, in-paper result.

Remark 34.3 (Physics intuition: why the finite BI-6 budget forces $N_g \leq 3$). T4.1 reads, in physics terms, as a hard upper bound on the number of fermion families. The BI-6 closure allows the bulk to absorb a finite amount of topological action — $S_{\text{BI6}}^{\text{max}} = 3\pi$, set by the OGN-5 ledger $K = 35$. The k -th instanton step in the family- projector cascade on $\mathbb{CP}^{N-1}$ costs πk units of action (the steps are not equal in cost), so the total cascade action on $\mathbb{CP}^{N-1}$ is $S_{\text{tot}}(\mathbb{CP}^{N-1}) = \pi \sum_{k=1}^{N_g-1} k = \pi N_g(N_g - 1)/2$, which fits inside the budget $S_{\text{BI6}}^{\text{max}} = 3\pi$ only for $N_g \leq 3$. $N_g = 4$ would need $S_{\text{tot}}(\mathbb{CP}^3) = 6\pi$, double the budget. Indeed, this is the deepest answer to the question “why are there three generations?” The number is not a free parameter of the Standard Model; it is forced by the finite topological budget of the BI-6-closed bulk. The compatibility relation $\mathfrak{C}_{\text{GRIP}}$ then carries this algebraic bound to its operational counterpart in T4.2: gauge-runtime invariance along the cascade is preserved only if the PBI-4 compensator extends smoothly across each instanton crossing.

34.2 Proof of T4

Proof of Theorem 34.1. Step 1: Total cascade action on $\mathbb{CP}^{N-1}$. By Proposition 30.11, the total instanton action along the cascade $c_{\text{top}} \rightarrow \cdots \rightarrow c_{\text{bottom}}$ on $\mathbb{CP}^{N-1}$ is

$$S_{\text{tot}}(\mathbb{CP}^{N-1}) = \sum_{k=1}^{N_g-1} \pi k = \pi \frac{N_g(N_g - 1)}{2}. \quad (179)$$

Step 2: BI-6 budget. By GRIP-ST4, the cascade is constrained by the budget

$$S_{\text{tot}}(\mathbb{CP}^{N-1}) \leq S_{\text{BI6}}^{\text{max}} = 3\pi. \quad (180)$$

The value $S_{\text{BI6}}^{\text{max}} = 3\pi$ is the capacity theorem T-N3-6 of the companion paper Arisaka (2026, P2-N3) (its Appendix C records the two viable formulations; the extensive form is $S(N_g) = \pi N$), conditional on the coupling clause and the integrality theorem L-INT (= M6; residue R-INT). The present paper takes $S_{\text{BI6}}^{\text{max}} = 3\pi$ as input from that companion paper rather than re-deriving it here; an in-paper re-derivation is deferred to a future revision and is recorded as an open task (see Section 57, M10).

Step 3: Diophantine bound. Combining (179) and (180):

$$\pi \frac{N(N-1)}{2} \leq 3\pi \iff N(N-1) \leq 6.$$

The integer solutions of $N(N-1) \leq 6$ in $\mathbb{Z}_{>0}$ are:

- $N = 1$: $0 \leq 6$. ✓
- $N = 2$: $2 \leq 6$. ✓
- $N_g = 3$: $6 \leq 6$. ✓ (*saturates the bound*)
- $N_g = 4$: $12 > 6$. ✗
- $N_g \geq 4$: $N_g(N_g - 1) > 6$. ✗

Therefore $N_g \leq 3$.

Step 4: Combination with lower bound. By Theorem **T2** of Part III, the OGN-5 ledger forces $N \geq 3$. Combined with the upper bound $N \leq 3$ from Step 3, we get $N = 3$.

Step 5: Saturation observation. At $N_g = 3$, the cascade exactly saturates the budget set by the companion paper: $S_{\text{tot}}(\mathbb{CP}^2) = 3\pi = S_{\text{BI6}}^{\text{max}}$. The saturation is structural: the companion-paper derivation of $S_{\text{BI6}}^{\text{max}}$ — the capacity theorem T-N3-6 of Arisaka (2026, P2-N3) — is independent of $N_g = 3$ at every step, consuming neither $K = 35$ nor $J = 7$, which is exactly the demand recorded as open problem M10 in Section 57; subject to that theorem’s premises (the coupling clause and M6 = R-INT), the $N = 3$ saturation is the natural fixed point at the edge of admissibility. ■

34.3 Comments on the saturation at $N_g = 3$

The saturation $S_{\text{tot}}(\mathbb{CP}^2) = 3\pi = S_{\text{BI6}}^{\text{max}}$ at $N_g = 3$ is structurally tight: $N_g = 2$ leaves BI-6 action unspent; $N_g = 4$ requires $S_{\text{tot}}(\mathbb{CP}^3) = 6\pi$, double the budget. The Bogomolnyi triangular sum $\pi N_g(N_g - 1)/2$, the budget $S_{\text{BI6}}^{\text{max}} = 3\pi$, and the lower bound $N_g \geq 3$ from CCI all derive from independently fixed structural inputs (Kähler geometry, OGN-5 ledger $K = 35, J = 7, \mathbb{H}^2$ normalization), each documented in its own section above. The reader is referred to the post-theorem Physics-Intuition remark and to Step 5 of the T4.1 proof for the full saturation discussion.

34.4 Combination with T2 and T3: the three-leg N=3 derivation

Corollary 34.4 (Three-leg derivation of $N_g = 3$). *Combining Theorems **T2**, **T3**, and **T4** of the present paper, the family number is forced to equal exactly 3:*

$$N_g = 3. \tag{181}$$

Proof. **T2** gives the lower bound $N \geq 3$ as the MSMF-minimum solution of the OGN-5 closure constraints with $L = 5, M \geq 1, N \geq 3$ (lower-bound side from CP violation). **T3** gives the parity constraint that strict Morse-positivity on Σ_{fam} requires $\dim_{\mathbb{R}} \Sigma_{\text{fam}} \in 4\mathbb{Z}$, i.e., N odd. **T4** gives the upper bound $N \leq 3$ from cascade closure. The intersection of these three constraints is uniquely $N_g = 3$. ■

This is the Gi-4 version of the three-leg N=3 proof of Arisaka (2026, P2-N3). In of GG-2026-Generations, the proof rested on three legs (L1, L2, L3) with residual hypotheses $\mathcal{H}_1, \mathcal{H}_2$. In the present Gi-4 formalization, all hypotheses are theorems and the three legs become Theorems **T2**, **T3**, **T4** of the present paper.

34.5 Statement and proof of T4.2 (runtime gauge-invariance preservation)

Indeed, the second half of Theorem **T4** packages the *runtime gauge-invariance preservation theorem*, which packages **GRIP-ST6–GRIP-ST7** into a precise iff statement linking the algebraic cascade closure to the boundary weight ladder.

Theorem 34.5 (T4.2: Runtime gauge-invariance under PBI-4 compensator smoothness). *Let $c_{\text{top}} \rightarrow c_1 \rightarrow \cdots \rightarrow c_{\text{bottom}}$ be the algebraic GRIP cascade of **GRIP-ST1–GRIP-ST5** on the family-projector stratum $\Sigma_{\text{fam}} = \mathbb{CP}^{N-1}$, with the BI-6 budget $S_{\text{BI6}}^{\text{max}} = 3\pi$ saturated by $N_g = 3$.*

Standing hypothesis (M11). The PBI-4 compensator field $C_\mu(t)$ of Definition 17.3 extends smoothly across each instanton transition $c_i \rightarrow c_{i+1}$, i.e., C_μ and its first derivative are continuous as the cascade trajectory traverses the codimension-2 instanton stratum $\Sigma_{\text{inst}} \subset \bar{\Sigma}_{\text{fam}}$. Pointwise smoothness is established case-by-case for the cascade saddles on $\mathbb{CP}^{N-1}$; the global topological non-obstruction (vanishing of the monodromy of C_μ around the instanton stratum) is recorded as open problem M11 in Section 57.

Then, under the standing hypothesis, the runtime gauge-invariance preservation of **GRIP-ST6**, expressed as

$$\nabla_\mu J_a^\mu(t) = -\alpha_a C^\mu(t) \cdot \rho_a(t), \quad t \in [t_0, t_\infty], \quad a \in \{e, B, Qi, g\}, \quad (182)$$

holds along the entire cascade **if and only if** the cascade gradient flow is gauge-equivariant: the flow $\dot{x} = -g_\Sigma^{-1} \nabla \mathcal{F}_{\text{GFF}}$ commutes with the boundary G_{tot} -action, i.e., the unbroken gauge subgroups satisfy the descending chain $H_0 \supseteq H_1 \supseteq \dots \supseteq H_n$.

Furthermore, when this iff is satisfied, the generation-ordered weight ladder of **GRIP-ST7**, $w_{[2k] \rightarrow [2(k-1)]} = e^{-\pi k}$, is the boundary observable consistent with the cascade.

Remark 34.6 (Status: T4.2 is conditional on open problem M11). T4.2 as stated promotes the smoothness of C_μ across instantons from one of the iff conditions (the form used in earlier drafts) to a standing hypothesis of the theorem. The reason is honest conditionalization: the global topological non-obstruction of the smooth extension — vanishing of the monodromy of C_μ around the codimension-2 instanton stratum — is open problem M11, now sharpened to a spectral-flow form with a verified benchmark instance (M11 entry, §57). T4.2 thus holds in two cases: pointwise (where the smooth extension is verified case-by-case on $\mathbb{CP}^{N-1}$ for $N_g \leq 3$, sufficient for the SM application) and globally (subject to M11 being resolved). Absent M11 a future revision will be needed for higher- N strata; the SM case is unaffected.

Proof. Assume the standing hypothesis (M11): $C_\mu(t)$ extends smoothly across each instanton transition $c_i \rightarrow c_{i+1}$.

($\Leftarrow$) Assume gauge-equivariance of the cascade flow. The gradient flow then preserves the gauge content of each configuration along the trajectory. By the Gauge-Cascade Compatibility Theorem 33.1, the runtime relation (182) follows. By the standing hypothesis, $C_\mu(t)$ is smooth across instantons, so the relation holds without discontinuities along the cascade. The generation gate-ordered weight ladder $w_{[2k] \rightarrow [2(k-1)]} = e^{-\pi k}$ then follows from the gate tolls, as the equilibrium weights of Section 33.2.

($\Rightarrow$) Conversely, assume the runtime relation (182) holds along the entire cascade. We must show that the cascade gradient flow is gauge-equivariant. By Master Conservation (Theorem T2.2 of Part III), the runtime relation requires the existence of a PBI-4 compensator $C_\mu(t)$ at every t ; the standing hypothesis ensures this compensator is smooth across instantons, so no delta-function source arises in (182). Gauge-equivariance of the flow then follows because any non-equivariant flow would violate the conservation of one of the Ward currents at runtime, which contradicts (182). ■

Remark 34.7 (Logical structure of T4: T4.1 + T4.2). Theorem T4 as stated in is the *conjunction* of T4.1 and T4.2. T4.1 fixes the algebraic side: the unique cascade structure on $\mathbb{CP}^{N-1}$ is forced by

the gradient-flow generation of **GRIP-ST2** and the BI-6 budget of **GRIP-ST4**, with closure forcing $N_g \leq 3$. T4.2 fixes the physical side: the runtime gauge-invariance preservation holds *conditional* on the smooth extension of the PBI-4 compensator across instantons (the smooth-extension hypothesis, open problem M11), and the generation-ordered weight ladder is the observable signature. The two halves are linked by the Gauge-Cascade Compatibility Theorem 33.1.

This is the precise sense in which the algebraic side of GRIP *constrains* the physical side: the BI-6-bounded cascade structure is the data on which the PBI-4 compensator’s smooth extension across instantons (open problem M11) would act, yielding the generation-ordered weight ladder.

Remark 34.8 (Physics intuition: why T4.2 is the deeper half). A graduate student in physics may at first think that T4.1 is the “deeper” theorem because it produces the three-leg $N=3$ family-count proof. But T4.2 is in fact the deeper statement: it is the assertion that the algebraic cascade structure carries an *observable* signature — the generation-ordered weight ladder $w_{[2k] \rightarrow [2(k-1)]} = e^{-\pi k}$. T4.1 alone only says that the cascade closes in finite toll — philosophically attractive but not, by itself, observable. T4.2 ties the cascade to the family structure: the generations must be ordered as the weights order them, on pain of either a different cascade structure (excluded by BI-6 budget closure) or a runtime gauge anomaly (conditional on M11); the earlier form of this clause, an ordering of residual decay widths in a controlled-kinematics class, tests a rate reading the paper no longer derives (**GRIP-ST7**).

In this sense, T4.2 is the *empirical* arm of T4: the $N_g = 3$ family count of T4.1 is exposed to data not through aesthetics (Bogomolnyi saturation) but through the family structure — the three-rung charged-lepton ladder $\tau \rightarrow \mu \rightarrow e$ (and its quark counterpart) whose existence and ordering the cascade predicts, with the lifetimes ($\tau_\tau \approx 2.9 \times 10^{-13}$ s, $\tau_\mu \approx 2.2 \times 10^{-6}$ s, e stable) being phase space, fixed in Arisaka (2026, A9-SM). The convergence of T4.1 (algebraic) and T4.2 (physical) on the same cascade structure is one of the strongest internal consistency checks of GG-Theory.

Corollary 34.9 (Generation gate-ordering falsifier from T4.2). *Combining Theorem 34.5 with the gate weight (159), the generation ordering — the family ladder $\tau \rightarrow \mu \rightarrow e$ with per-gate weight $w_{[2k] \rightarrow [2(k-1)]} = e^{-\pi k}$ — is a structural prediction of GG-Theory, the lifetimes being phase-space quantities fixed in Arisaka (2026, A9-SM). A family ladder whose generations were not ordered as the weights order them would constitute a falsification of GRIP. The residual-width form of the test — a controlled-kinematics decay class in which a residual ordering departing from $e^{-\pi k}$ by more than one toll unit (factor $e^\pi \sim 23$, after the phase-space factors are divided out) would count against the cascade — is carried at the grade of the rate reading it presupposes, which this paper does not derive.*

Proof. This is the empirical specialization of Theorem 34.5 to the family-projector cascade with the OGN-5-determined Bogomolnyi actions. ■

35 The generation weight ladder from the cascade

We briefly comment on the empirical signature of the GRIP cascade: the generation-ordered weight ladder.

35.1 Gate weights from gate tolls

By **GRIP-ST3** and Proposition 30.12, the weight of each gate from the critical point of index $2k$ to the next of index $2(k-1)$ on the family-projector stratum is

$$w_{[2k] \rightarrow [2(k-1)]} = e^{-S_{[2k] \rightarrow [2(k-1)]}} = e^{-\pi k}. \quad (183)$$

For $\mathbb{CP}^2$ (the $N = 3$ stratum), the cascade is [index 4] $\rightarrow$ [index 2] $\rightarrow$ [index 0] with

$$w_{[4] \rightarrow [2]} = e^{-2\pi}, \quad w_{[2] \rightarrow [0]} = e^{-\pi}. \quad (184)$$

The upper gate's weight is smaller than the lower's by a factor $e^{-\pi} \simeq 0.043$: the DGI point sits higher. These are equilibrium weights (**GRIP-ST7**); an earlier statement of this section read them as WKB decay rates, and that reading is withdrawn in §33.2.

35.2 Identification with physical generation lifetimes

Identifying the three Morse critical points of $\mathbb{CP}^2$ with the three Standard-Model fermion generations:

- **Index 4 (DGI, third generation):** τ, t, b . Behind the gate of toll $S = 2\pi$; the decays themselves are weak processes with their own phase space and weak-boson factors.
- **Index 2 (GSSB, second generation):** μ, c, s . Behind the gate of toll $S = \pi$.
- **Index 0 (GA, first generation):** e, u, d . The floor; stable.

The empirical lifetimes of the unstable charged leptons are $\tau_\tau = 2.9 \times 10^{-13}$ s (third generation) and $\tau_\mu = 2.2 \times 10^{-6}$ s (second generation), with $\tau_\mu/\tau_\tau \simeq 7.6 \times 10^6 \simeq e^{15.8}$.

The naïve cascade prediction $\Gamma_\tau/\Gamma_\mu \sim e^{-2\pi}/e^{-\pi} = e^{-\pi} \simeq 0.043$, i.e., $\tau_\tau/\tau_\mu \sim e^\pi \simeq 23$, has the measured ratio backwards: the tau is the shorter-lived of the two, by a factor 7.6×10^6 . The lifetime ratio is not a cascade observable, and it should not be read as one. It is carried by the two Standard-Model factors that dominate any weak decay — Sargent's rule $\Gamma \propto m^5$, and the fact that the electron channel is only part of the tau's width — giving $\Gamma_\tau/\Gamma_\mu \simeq (m_\tau/m_\mu)^5/\mathcal{B}(\tau \rightarrow e\bar{\nu}_e\nu_\tau) \simeq (16.8)^5/0.178 \simeq 7.5 \times 10^6$ against the measured 7.6×10^6 ([Particle Data Group, 2024](#)). That agreement is the Standard Model's, not this paper's. What the cascade action fixes is the *pattern* — which generations are unstable at all, index 4 and index 2 decaying and index 0 not — and not the numerical ratio of their lifetimes.

The detailed treatment of phase-space factors and weak-boson exchange is the subject of GG-A9-SM ([Arisaka, 2026, A9-SM](#)), and it is the whole of the lifetime calculation: the gate weight $e^{-\pi k}$ is not a factor in a width until a mechanism carries it into one (**GRIP-ST7**), and where the companion papers carried it as one, that is to be read at the present grade.

36 Worked examples

We work through three explicit GRIP cascades on internal moduli, and then a fourth on the orbital stratum — the golden case, in which the pipeline is not only exhibited but computed in closed form and measured.

36.1 Example 1: GRIP on $\mathbb{CP}^2$ (the $N = 3$ case)

The cascade on $\Sigma_{\text{fam}} = \mathbb{CP}^2$ (real dimension 4, three critical points of indices 0, 2, 4):

$$\boxed{[\text{index } 4, \tau] \xrightarrow{e^{-2\pi}} [\text{index } 2, \mu] \xrightarrow{e^{-\pi}} [\text{index } 0, e]} \quad (185)$$

Total action. $S_{\text{tot}}(\mathbb{CP}^2) = \pi + 2\pi = 3\pi$, exactly saturating the BI-6 budget.

Cascade depth. $N_g - 1 = 2$ transitions; the deepest unstable phase is at index 4 (the τ, t, b generation).

Stability. The index-0 point (e, u, d) is the unique GA, stable in all directions. The index-4 point is unstable in 4 directions (all of $\mathbb{CP}^2$); the index-2 point is unstable in 2 directions and stable in 2 directions.

36.2 Example 2: GRIP on the Higgs moduli

The cascade on Σ_H (the Wilson–Hosotani Higgs moduli):

- DGI: symmetric Higgs phase $V = 0$ (Morse index 1 in this discrete count).
- GA: broken Higgs phase $V = v \simeq 246 \text{ GeV}$ (Morse index 0).

The cascade is a single transition: a Coleman–Weinberg / Wilson–Hosotani descent from the symmetric to the broken phase. This is the electroweak phase transition in cosmological history.

Action. The instanton action is set by the Higgs effective potential and the scale $M_{\text{lock}} \simeq 3.85 \text{ TeV}$ (Arisaka, 2026, P1-KK). In units of the BI-6 budget, it is much smaller than 3π because the Higgs cascade is one-step rather than a chain.

36.3 Example 3: GRIP on the gauge-orbit moduli

The cascade on Σ_{gauge} (the orbifold-conjugacy moduli):

- DGI: unbroken $\text{SU}(5)$ phase.
- GA: $\text{SU}(3) \times \text{SU}(2) \times \text{U}(1)_Y$ broken phase (Standard Model gauge structure).

The cascade is again single-step — the orbifold projection breaks $\text{SU}(5)$ in one move via the Wilson-line holonomy (33).

36.4 Example 4 (the golden case): GRIP-In on the orbital stratum — synchrotron damping and binary inspiral, solved to the coefficient

Examples 1–3 run the cascade on internal moduli, where the gate actions are exact ($S = \pi k$) but the mobilities — the rate prefactors of the within-basin flow — remain program-level. The fourth example is the converse, and it is the reason we call it golden: a GRIP-In instance in which every element of Definition 31.5 is computed in closed form from the derived field equations and checked against precision experiment. The full derivation is the radiative completion of GG-A5-GDOS’s Part VI together with its quantitative-closure subsection (Arisaka, 2026, A5-GDOS); here we read that result in GRIP vocabulary. One fence first, so that strata are never conflated: this GRIP-In runs on the *orbital* stratum — its landscape is orbital energy, its descent is the loss of motion, and it transmutes nothing — whereas Example 1 (and the lepton cascade it encodes) runs on the family stratum and transmutes identity. The two-sector partition of the companion electron paper (Arisaka, 2026, P3-Electron) — *the phase sector moves the electron, the scale sector transmutes it, and only what sheds energy descends* — is preserved exactly. What this example adds is the demonstration that GRIP-In is *stratum-agnostic*: the same septuple runs on both, and on this stratum it runs analytically.

- **Fixed landscape (the GRIP-In premise).** The storage ring’s magnetic lattice, or the binary’s orbital geometry: a stratified $(\mathcal{M}_{\text{bdy}}, \mathcal{F}_{\text{GFF}})$ that the descent itself does not reorganize. Definition 31.5’s premise is satisfied by construction, not by idealization.
- **Within-basin gradient flow (154) — with the mobility *derived*.** Radiation damping is the Riemannian gradient flow of the orbital-energy landscape, and GG-A5-GDOS’s field-trace closure delivers the mobility in closed form: $g^{EE}(E) = P(E)$, carrying Larmor’s coefficient $q^2/6\pi\epsilon_0 c^3$, for the synchrotron; and $g^{aa}(a) = \frac{128}{5} G_N^2(m_1+m_2)/(c^5 a)$, reproducing Peters’ coefficient $\frac{64}{5} G_N^3/c^5$ exactly, for the inspiral. The within-basin flow of GRIP-ST2, which Examples 1–3 invoke structurally, here has its rate constants computed.
- **Gates (GRIP-ST3): none.** The orbital stratum is gate-free: the descent is purely within-basin, and the relaxation is power-law rather than exponentially tolled. The contrast is itself diagnostic — *gate-free strata relax; gated strata decay* — and the bubble-chamber asymmetry of the electron paper (Arisaka, 2026, P3-Electron) (the muon’s track ends in a kink, the electron’s never ends) is this line drawn by nature.
- **The GFF-ST7 halt.** The synchrotron beam halts at its equilibrium emittance — the orbital stratum’s GA — and the fluctuation–dissipation clause of GFF-ST7(2) is realized at coefficient level: the universal quantum constant $C_q = \frac{55}{32\sqrt{3}} \hbar/mc = 3.832 \times 10^{-13}$ m (electron) governs the measured equilibrium energy spread of every electron storage ring (Sands, 1970). GRIP-In “terminates, and terminates uniquely” is, on this stratum, a number on a console.
- **A stratum with no interior GA.** The binary’s landscape is unbounded below all the way to contact, so GRIP-In cannot halt: the descent exits at the merger. The merger *functions* as a landscape-reorganizing event — it closes the binary’s stratum and opens the remnant’s — and the ringdown that follows is the new stratum’s GRIP-In, relaxing to *its* GA, the Kerr attractor.

A single observed waveform — inspiral, merger, ringdown (Abbott *et al.*, 2016) — thus draws the GRIP loop end to end: descent, reorganization, descent. Five decades of binary-pulsar timing anchor the inspiral leg to a few parts in 10^3 (Hulse & Taylor, 1975).

Status of Example 4. The quantitative spine — the canonical GDOS current decomposition of the reduced dynamics, the drift as retarded flux, the closed-form mobilities, the fluctuation–dissipation pairing, and now the covariant non-Markovian kernel itself (closed-form spectral sum over the lock tower; causal; runaway-free by passivity) — is theorem-grade in GG-A5-GDOS, imported here by citation; GG-A5-GDOS’s Open Problem O7 is Resolved. Two readings are this subsection’s own, and we price them as interpretive. First, calling the merger *GRIP-Out-functioning*: by Definition 31.6, GRIP-Out is not generated by the gradient flow, whereas the merger is the descent reaching the edge of its stratum — functionally a reorganization, generatively not a GRIP-Out; the distinction is kept explicit. Second, the diagnostic slogan *gate-free strata relax, gated strata decay*. Within those two fences, this is the one GRIP instance in the cohort where the pipeline is analytic from landscape to halt and measured from mobility to noise — which is what a golden case is for.

37 From a landscape to a history

With this Part the four pillars are complete, and the paper has something it did not have before: a story with an order of events in it.

The mechanism is worth stating without machinery. A system begins with its options open — nothing has been chosen, and every direction is still available. The landscape of the previous Part tells it which way is down. It descends. At some point it arrives at a fork, and takes one branch, and a symmetry that was exactly true a moment earlier is simply gone: not violated, not approximate, gone, because the system is now somewhere that does not respect it. It continues down until there is nowhere lower to go, and stops. Then the resting place becomes the starting condition for the next landscape, and the whole thing runs again at a different scale.

Two things follow that are worth carrying away. The first is that the cascade terminates, and where it terminates is countable. Run it and you get three families of matter — not because three was chosen, and not because a parameter was set to three, but because the descent runs out of room. A fourth family would need somewhere to go and there is nowhere. This is the same kind of statement as the integers of Part III, arrived at by a different route, and the two agreeing is either a coincidence or the framework working.

The second is about broken symmetry, which is one of the genuinely strange facts about the world. The laws appear to be more symmetric than the things they describe. A magnet has a direction; the equations governing it do not. The usual response is to say the symmetry is hidden, and that is right but incomplete, because it does not say why the hiding happened when it did. Here the answer is that symmetry breaking is not an event that befalls a system but a step in a descent: it happens at the point where continuing downhill requires choosing, and the choice is what removes the symmetry. *Structure is what a falling system leaves behind.*

There is a further consequence that this Part states carefully and does not overclaim. Descent

on a fixed landscape has an end. Everything would settle once and stay settled — which describes a universe of stable matter perfectly well, and describes nothing that continues. What is needed for a history rather than a conclusion is a way for the landscape itself to be reorganized from inside. This paper establishes the engine; it does not establish that. The companion that does is named, the claim is left where it belongs, and the reader should notice that the boundary has been drawn rather than blurred.

One further remark, because readers coming from thermodynamics will want it. The descent described here is irreversible, and the irreversibility is not put in by hand or borrowed from a coarse-graining argument. It is a property of the flow: the functional decreases along it and cannot increase, so the sequence of events has an order that cannot be run backwards. Whether that is *the* arrow of time is a question this paper does not settle. That it is *an* arrow, pointing the same way everywhere in the construction, is established here.

The next Part asks whether the four pillars just built are four things or one.

Part VI Synthesis

The unified dual-layer Gi-4 system, Theorem T5, and bridges to OGGEF/MSMF and Qi-4

By this point the paper has built four things and it is fair to ask whether they are four things at all.

The suspicion is reasonable, and the remark opening this Part meets it head on: if the companion charter already states four axioms, why spend five Parts constructing a second, parallel system of four pillars? Are they a rival theory, or a redundancy? The answer this Part proves is neither. The four pillars are the same content read geometrically — the charter says what must be true, and the pillars say what the geometry is that makes it true. The two descriptions are equivalent, and the equivalence is the Part’s principal result rather than an assertion of good faith.

Establishing that has a practical consequence worth stating plainly. If the four pillars were an addition to the charter, the paper would be introducing new physical assumptions, and every result downstream would inherit them. Because the two are equivalent, the paper adds no axiom: it makes explicit a structure that was already implied. That is a weaker and much more defensible claim than the one a reader might expect from a paper of this length, and it is the claim actually made.

The Part then builds two bridges. One runs back to the constitutional charter, showing pillar by pillar which axiom each rests on. The other runs sideways, to the companion paper that treats the same four structures in the language of quantum information rather than geometry ([Arisaka, 2026, A7-Qi4](#)). The second bridge is the one a reader from outside physics may find most useful, because the informational reading of these objects is often the more intuitive of the two: it speaks of what can be known, distinguished and recorded, where the geometric reading speaks of curvature and connection. They are two vocabularies for one thing, and this Part is where the dictionary is written. The proofs are short; a reader who accepts the equivalence can take the statement and move on to Part VII, where it is used.

38 The unified Gi-4 system

Remark 38.1 (Reader’s bridge: why two presentations of the same theory matter). At this point a careful reader may wonder: if the GG-Constitution of A1 already states four axioms (LC, GGEP, OGGEF, MSMF) and a closure (BI-6), why have we spent five Parts constructing a second, parallel system of four pillars? Are TGG, CCI, GFF, and GRIP a competing theory, or are they redundant? Neither. Gi-4 is the *geometric reading* of the constitution: it makes the bundle structure, the anomaly ledger, the free-energy landscape, and the cascade dynamics explicit at the level of differential geometry and Morse theory, where mathematical work can actually be done. The constitutional presentation, by contrast, is the *foundational reading*: it states what the framework must achieve at the level of physical principles — locality, gauge-gravity equivalence, irreversible openness, manifold

minimality — without committing to the particular geometric realization. Theorem T5 below proves that the two readings are equivalent. Part VI synthesizes the four-pillar system, states and proves T5, and articulates the bridges back to the constitution and across to the Qi-4 information-theoretic pillars of A7.

We have now developed all four Gi-4 Pillars in Parts II–V in their full *dual-layer* form, with seven Structure Theorems per pillar (5 algebraic + 2 physical) and the paired theorems {**T1.1**, **T1.2**, **T1.3**}, {**T2.1**, **T2.2**}, {**T3.1**, **T3.2**}, {**T4.1**, **T4.2**}. Part VI synthesizes these into a unified framework, states and proves the constitutional-equivalence theorem **T5**, and articulates the bridges to the GG-Constitutional axioms (OGGEP, MSMF) and to the Qi-4 Pillars of GG-A7-Qi4.

38.1 The dual-layer architecture: a unified mathematical object

The most important structural feature of Gi-4 is that *every* Gi-4 Pillar carries the same dual-layer architecture:

- An **algebraic block** (axioms A1–A5) consisting of representations-theoretic, cohomological, or Morse-theoretic content that lives in the bulk M_6 or on the stratified moduli $\mathcal{M}_{\text{bdy}}$;
- A **physical block** (axioms A6–A7) consisting of operationally testable, observer-accessible content that lives on the boundary $\partial M_6 \cong M_4$ in real laboratory time.

Within each pillar, the two blocks are linked by a named “compatibility” theorem: PBI-4 (CCI), the Geometric First Law (GFF), or the Gauge-Cascade Compatibility Theorem (GRIP). Across the four pillars, the dual-layer structure forms a coherent universal architecture summarized in Table 5.

Table 5: The dual-layer construction set out Pillar by Pillar. The second column is the algebraic block, which a mathematical physicist can audit without reference to experiment; the third is the physical block, which an experimenter can falsify without reference to the algebra; the fourth names the theorem that links them. Splitting each Pillar this way is deliberate. It lets the two audiences check the framework independently, and it localizes any failure to one layer of one Pillar rather than to the framework as a whole.

Pillar	Algebraic block (A1–A5)	Physical block (A6–A7)	Compatibility link
TGG	Composite connection $\mathcal{A}_{\text{tot}}$ on M_6 , curvature $\mathcal{R}_{\text{tot}}$, Killing form, $\mathbb{H}^2$ representation	Boundary stress tensor $T_{\text{TGG}}^{\mu\nu}$, Master Conservation Law	KK descent + PBI-4
CCI	Integral $[X_4] \in H^4(M_6; \mathbb{Z})$, OGN-5 ledger (1, 3, 5; 7, 35)	Boundary 4D conservation $\nabla_\mu T^{\mu\nu} = 0$, modified TOV	PBI-4 compensator field
GFF	Morse functional on $\mathcal{M}_{\text{bdy}}$, strict Morse-positivity	Helmholtz free energy $F = U - TS$, $dF/dt \leq 0$	Geometric First Law
GRIP	Cascade $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$, BI-6 budget $S_{\text{BI6}}^{\text{max}} = 3\pi$	Runtime gauge invariance, generation-ordered weight ladder $w_{[2k] \rightarrow [2(k-1)]} = e^{-\pi k}$	Gauge-Cascade Compatibility Theorem

Remark 38.2 (Physics intuition: the dual-layer architecture as the empirical bridge). The dual-layer architecture of provides a unified empirical bridge from the bulk geometric structure of GG-6 to

the laboratory. *Every* measurable boundary observable — conservation laws, modified TOV in compact stars, free-energy minimization in chemistry/biology, generation weight ladders — is the boundary projection of an algebraic bulk identity through the corresponding compatibility link. This is the precise sense in which **GG-Theory** is a complete predictive framework: it does not introduce phenomenological parameters at the boundary, only geometric data in the bulk that descends to operationally testable predictions.

38.2 The four pillars as a single mathematical object

The four pillars (TGG, CCI, GFF, GRIP) are not four independent constructions: they are four facets of a single mathematical object whose existence is forced by the **GG-Constitution** on the 6D bulk. We summarize their relationships in the following table.

Table 6: Each of the four Pillars in one row: the mathematical object it is built from, the constitutional content it carries, and the sub-theorems that establish it in both layers. The fourth column is what makes this a hierarchy rather than a list — every Pillar is discharged by a named pair of theorems, one algebraic and one physical, and that pairing is what the dual-layer architecture means in practice. A reader auditing a single Pillar can start from its row and follow the theorem numbers into the Parts that prove them.

Pillar	Mathematical object	Logical content	Paired theorem
TGG	Composite ($P, \mathcal{A}_{\text{tot}}, \mathcal{R}_{\text{tot}}; T_{\text{TGG}}^{\mu\nu}, \nabla$)	GGEP on M_6 + boundary master tensor	T1.1+T1.2+T1.3
CCI	Septuple of (76)	Algebraic closure + boundary conservation	T2.1+T2.2
GFF	Septuple of (122)	Morse cascade + thermodynamic monotonicity	T3.1+T3.2
GRIP	Septuple of (160)	BI-6 cascade + runtime gauge-invariance	T4.1+T4.2
Unified	The Gi-4 dual-layer system	GGEP+OGGEP+MSMF+BI-6	T5 (constitutional equivalence)

The hierarchy proceeds in a definite logical direction:

- TGG is the bulk-level upstream object: it provides the bundle, the connection, and the curvature.
- CCI is the algebraic closure condition on TGG: it specifies which 4-form is the source of BI-6 and asserts its integrality.
- GFF is the boundary-level dynamical object derived from CCI via Hopkins–Singer descent: it is the free-energy functional that drives the cascade.
- GRIP is the cascade dynamics on the GFF Morse landscape: it is the dynamical content of the open-system projection.

Together, the four pillars encode the geometric backbone of GDOS.

38.3 Logical hierarchy of the four pillars

The implication chain among the four pillars is:

$$\begin{array}{l}
 \text{GGEP} \implies \text{TGG} \\
 \text{TGG} \ \& \ \text{integrality} \implies \text{CCI} \\
 \text{CCI} \implies \text{BI-6} \\
 \text{OGGEP} \ \& \ \text{CCI} \implies \text{GFF} \\
 \text{GFF} \ \& \ \text{cascade closure} \implies \text{GRIP}
 \end{array} \tag{186}$$

This logical chain is closed in the sense that no input external to the GG-Constitution is invoked. All five implications follow from the constitutional axioms (LC, GGEP, OGGEP, MSMF) plus the BI-6 closure of Arisaka (2026, A3-BI6) plus standard mathematical machinery (differential geometry, Hopkins–Singer differential cohomology, Morse theory).

38.4 The nine dual-layer sub-theorems summarized (T1.1–T4.2)

The algebraic combination **T2.1+T3.1+T4.1** closes the derivation of $N_g = 3$ at the upstream Gi-4 level, recovering the three-leg structure of Arisaka (2026, P2-N3) but with all hypotheses replaced by theorems. The physical combination **T2.2+T3.2+T4.2** provides the matched laboratory-testable predictions: modified TOV, equipartition + fluctuation–dissipation, and the generation weight ladder of unstable particles. Together, the nine dual-layer sub-theorems (with the bridge **T5**) constitute the complete derivation: every algebraic prediction $N = 3$ is paired with an empirical falsifier.

39 Theorem T5 (T-GRIP-5): Constitutional equivalence

39.1 Statement of T5

Theorem 39.1 (T5: Constitutional equivalence of Gi-4 and GG-Constitution). *The Gi-4 Pillars system (TGG, CCI, GFF, GRIP), equipped with the four principal theorems **T1–T4** of Parts II–V, is equivalent in derivational content to the GG-Constitution (LC, GGEP, OGGEP, MSMF) together with the BI-6 closure of Arisaka (2026, A3-BI6). Specifically:*

- (1) *Every theorem provable from (LC, GGEP, OGGEP, MSMF) + BI-6 is provable from the Gi-4 axiom system.*
- (2) *Every theorem provable from the Gi-4 axiom system is provable from (LC, GGEP, OGGEP, MSMF) + BI-6.*

The two systems are different presentations of the same mathematical content.

Status. Direction (1) is established below by exhibiting an explicit Structure-Theorem pointer for each clause of GG-Constitution + BI-6. Direction (2) is established up to two auxiliary propositions (22.15 and 30.13), whose internal proofs are sketched at the level of detail of the present paper; a fully formal treatment is part of the broader GG-Theory program and is recorded as a future task

Table 7: The nine sub-theorems that discharge the four Pillars, together with the bridge theorem that closes them. Each row names its Pillar and its layer — algebraic or physical — so the table doubles as an index: a reader who wants existence and uniqueness of the composite connection goes to **T1.1**, one who wants the boundary master stress tensor goes to **T1.2**. **T5** is the bridge, and it is what makes the collection a single result rather than nine separate ones. Full statements and proofs are in the Parts the rows point to.

Theorem	Pillar / Layer	Content
T1.1	TGG (alg)	Existence and uniqueness of the composite-connection $(P, \mathcal{A}_{\text{tot}}, \mathcal{R}_{\text{tot}})$ on M_6 , up to gauge equivalence.
T1.2	TGG (phy)	Existence of the boundary Master Stress Tensor $T_{\text{TGG}}^{\mu\nu} = T_{\text{matter}} + T_{\text{gauge}} + X_{\text{GS}} + T_{\text{KK}}$ via KK descent.
T1.3	TGG (link)	Master Conservation Law $\nabla_\mu T_{\text{TGG}}^{\mu\nu} = 0$ holds iff Ward package is closed.
T2.1	CCI (alg)	The OGN-5 integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ is the unique MSMF-minimum solution of the integrality constraints with $L = 5, M \geq 1, N \geq 3$.
T2.2	CCI (phy)	Boundary covariant conservation $\nabla_\mu T_{\text{TGG}}^{\mu\nu} = 0$ iff Ward + PBI-4; modified TOV equation as astrophysical falsifier.
T3.1	GFF (alg)	Strict Morse-positivity holds on a stratum Σ iff $\dim_{\mathbb{R}} \Sigma \in 4\mathbb{Z}$, derived from CCI via Hopkins–Singer descent.
T3.2	GFF (phy)	Thermodynamic monotonicity $dF_{\text{GFF}}/dt \leq 0$ from boundary CPTP structure with detailed balance (Lindblad, Spohn; Kossakowski–Frigerio–Gorini–Verri); equilibrium uniqueness with equipartition + fluctuation–dissipation.
T4.1	GRIP (alg)	Cascade closure on $\mathbb{CP}^{N-1}$ with BI-6 budget $S_{\text{BI6}}^{\text{max}} = 3\pi$ forces $N_g(N_g - 1) \leq 6$, hence $N_g \leq 3$.
T4.2	GRIP (phy)	Runtime gauge-invariance preserved iff PBI-4 extends across instantons; the generation-ordered weight ladder $w_{[2k] \rightarrow [2(k-1)]} = e^{-\pi k}$ is the laboratory signature, read from the mass hierarchy (lifetimes are phase space, Arisaka (2026, A9-SM)).
T5	Unified	The dual-layer Gi-4 system is equivalent in derivational content to the GG-Constitution + BI-6.

(Section 57). In the remainder of this section we therefore present direction (1) as a full proof and direction (2) as a structured proof sketch inheriting the level of detail of its two supporting propositions; a fully formal treatment of direction (2) is a future task.

39.2 Proof of T5: Gi-4 implies GG-Constitution

Proof of direction (1). Given the Gi-4 axiom system (TGG, CCI, GFF, GRIP) with Structure Theorems established in Parts II–V, we recover each constitutional axiom by exhibiting the explicit Structure-Theorem datum that supplies the axiom’s content.

LC (Local Causality). TGG-ST4 endows the 6D bulk M_6 with a Lorentzian, time-orientable metric, while TGG-ST1 constructs the principal G_{tot} -bundle on which the composite connection lives. The retarded Green’s kernels of any sectoral boundary current J_a^μ ($a \in \{e, B, Qi, g\}$) are then the boundary projection of the retarded bulk propagator with support in the forward Lorentzian cone of M_6 — this is the operational content of LC. The pair (TGG-ST4 metric, TGG-ST1 bundle) supplies both ingredients required for the GG-Constitution statement of LC.

GGEP (Geometry–Gauge Equivalence). TGG-ST5 asserts that the four sub-connections ω , $A_{\text{SU}(5)}$, A_B , A_{Qi} are components of a single primitive composite connection $\mathcal{A}_{\text{tot}} = \omega \oplus A_{\text{SU}(5)} \oplus A_B \oplus A_{Qi}$ on the principal G_{tot} -bundle. This is precisely the algebraic content of GGEP: gravity and the three gauge sectors are not independent structures coincidentally co-located on M_4 , but a single geometric object on M_6 that decomposes upon descent.

OGGEP (Operational/Open Geometry–Gauge Equivalence). The CPTP projection from bulk to boundary is realized in two stages. The boundary moduli space $\mathcal{M}_{\text{bdy}}$ with its free-energy functional $\mathcal{F}_{\text{GFF}}$ is GFF-ST1 of Part IV: this is the operational Stinespring-dilation form of the open-system map. The irreversible cascade $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$ that lives on this moduli space is GRIP-ST1–ST3 of Part V: this is the dynamical content of the CPTP projection. The two together — moduli landscape plus irreversible flow — exhaust the content of OGGEP in the GG-Constitution; the equivalence is formalized in Propositions 22.15 and 30.13.

MSMF (Minimal Sufficient Manifold Form). CCI-ST5 asserts MSMF minimality as the integer-cohomology version of the anomaly closure; Theorem T2.1 of Part III then shows that this minimality is uniquely realized by the OGN-5 ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$. No additional axiom is needed: the ℓ^1 minimization is structurally already inside CCI-ST5, and the Diophantine uniqueness is theorem-level (T2.1 of Part III).

BI-6 (the bulk anomaly closure). Proposition 14.3 shows that $dH_3 = X_4$ is a consequence of CCI together with the Chern–Simons construction of H_3 from TGG data (Proposition 11.3); in algebraic terms, the integrality of $[X_4]$ from CCI-ST4 together with the descent construction of H_3 from X_4 on a simply connected M_6 supplies the closure relation.

Therefore the Gi-4 axiom system, together with the Structure Theorems established in Parts II–V, recovers each clause of the GG-Constitution as a consequence. ■

39.3 Proof of T5: GG-Constitution implies Gi-4

Proof sketch of direction (2). Given the GG-Constitution + BI-6, we recover each Gi-4 pillar by pointing at the principal theorem of the corresponding Part of A6. Two of the four steps invoke

auxiliary propositions whose own proofs are sketches in the present paper; we flag those explicitly.

TGG: by Corollary 10.5 (Part II), TGG is the formal realization of GGEP on the 6D bulk. Theorem T1 of Part II establishes existence and uniqueness up to gauge. Full proof.

CCI: the algebraic content of BI-6 together with the MSMF minimality from the GG-Constitution gives CCI via Theorem T2.1 of Part III (Diophantine uniqueness of the OGN-5 ledger) and Theorem T2.2 of Part III (boundary Master Conservation iff Ward + PBI-4). Full proof.

GFF: by Proposition 22.15 (Part IV), GFF is the formal realization of OGGEp on the boundary moduli; Theorem T3 of Part IV establishes Morse-positivity from CCI. *Sketch-level:* the proof of Proposition 22.15 is given at the level of detail of the present paper; a fully formal Stinespring-dilation argument identifying $\mathcal{M}_{\text{bdy}} + \mathcal{F}_{\text{GFF}}$ with the CPTP image of OGGEp is part of the broader GG-Theory program.

GRIP: by Proposition 30.13 (Part V), GRIP is the formal realization of the cascade content of OGGEp; Theorem T4.1 establishes cascade closure (conditional on M10, per Remark 34.2). *Sketch-level:* the proof of Proposition 30.13 is again given at the level of detail of the present paper; the irreversibility content of the cascade matches the CPTP-image second-law condition of OGGEp, but a fully formal equivalence with the Lindblad monotonicity inequality is deferred.

Therefore GG-Constitution + BI-6 implies the Gi-4 axiom system at the programmatic level recorded in the theorem statement. ■

39.4 What “equivalent in derivational content” means

Theorem T5 asserts that the Gi-4 and constitutional axiom systems are equivalent. The practical meaning is:

- **Anything derivable in one is derivable in the other.** The N=3 derivation, the 28-parameter codebook, the hierarchy derivation, the dark-matter prediction, and the experimental falsifiers can all be obtained from either axiom system.
- **The Gi-4 presentation is the geometric reading.** It makes the bundle structure, the curvature, the Morse landscape, and the cascade dynamics explicit and citable. This is the presentation appropriate for technical mathematical work.
- **The constitutional presentation is the foundational reading.** It makes the four axioms (LC, GGEP, OGGEp, MSMF) explicit at the level of physical principles. This is the presentation appropriate for foundational discussion.
- **Neither system replaces the other.** They are complementary presentations of the same mathematical content.

This is the methodological payoff of Part VI: the Gi-4 reorganization of the present paper is not a new theory; it is a *geometric reading* of the GG-Constitution that makes the mathematical structure explicit.

40 How the four pillars discharge the constitutional axioms

The OGGE content of the GG-Constitution is distributed across all four Gi-4 pillars: the CPTP projection lives in GFF-ST1; its irreversibility lives in the GRIP gradient flow; the discrete GA-catalog is the joint output of GFF+GRIP; the bulk invariance is TGG; the algebraic compatibility between bulk and boundary is CCI. MSMF enters at multiple points as the minimum-complexity criterion: in CCI-ST5 (the OGN-5 ledger), in GRIP-ST4 (the BI-6 budget $S_{BI6}^{\max} = 3\pi$), and in GFF-ST4 (boundedness-below selects a finite GA catalog). The remaining axioms GGEP (the single composite-connection origin of the four sub-connections) and LC (Lorentzian causality of M_6) are upstream of the Gi-4 construction and live in TGG-ST1, ST4, ST5. The four-clause-by-four-pillar recovery is established in the proof of Theorem T5 direction (1) (§39.2); the reader is referred there for the full constitutional-equivalence argument.

41 The geometric backbone and its informational twin

41.1 The four-by-four pairing

The Gi-4 Pillars (geometric backbone) are paired with the Qi-4 Pillars (information-theoretic backbone) in GG-A7-Qi4 (Arisaka, 2026, A7-Qi4). The pairing is:

Table 8: How each geometric Pillar pairs with its quantum-information counterpart in the companion paper. The third column names the structure the pair shares and the fourth says what each side contributes to it, so every row reads as one statement in two registers: the same object seen geometrically and seen informationally. The pairing is not an analogy — each row is discharged by a bridge theorem — and it is the reason the two Pillar papers are siblings rather than a paper and its commentary.

Gi-4 Pillar	Qi-4 Pillar	Pairing structure	Content of the pairing
TGG	Qi-QUA	Master generator	Geometric (composite connection) + Information (entanglement of bulk states)
CCI	Qi-EP	Equivalence/closure principle	Geometric (anomaly polynomial integrality) + Information (information-conservation under unitary bulk evolution)
GFF	Qi-CON	Boundary functional/conservation	Geometric (free-energy Morse landscape) + Information (mutual-information / relative-entropy landscape)
GRIP	Qi-HAL	Cascade/bulk-to-boundary descent	Geometric (Morse cascade DGI->GSSB->GA) + Information (Stinespring dilation / decoherence cascade)

The two sets of four pillars are not redundant: each captures content that the other does not. Gi-4 captures the bulk-to-boundary geometric structure (gauge connections, anomaly polynomials,

Morse landscapes, attractor cascades). Qi-4 captures the bulk-to-boundary information-theoretic structure (entanglement, mutual information, decoherence, measurement outcomes). Together they constitute a unified foundational framework for GDOS, in which both the geometric and the information-theoretic content of physical observation are made foundational at the same time.

A note on the ground attractor and the pillars. One entry of the pairing above deserves a sentence of its own, because the literature of this cohort has carried it loosely. The Ground Geometric Attractor of Definition 31.4 is fixed by the landscape alone: Corollary 25.7 produces it from the Morse structure of the stratum with no admissibility filter of any kind. The four Qi-4 pillars of GG-A7-Qi4 are therefore not *selecting* the GGA from among the Geometric Attractors — there is nothing left to select once the corollary has run. What they do is *characterize* it: the object the descent halts at is also the object that carries the locked four-current ledger, extremizes the unified variational principle, transports coherently across scales, and is where an architecture can halt. That the two descriptions pick out the same object is a substantive claim and a theorem of GG-A7-Qi4 (Arisaka, 2026, A7-Qi4), which is its proper home; the present paper neither needs it nor assumes it.

41.2 Geometric vs information-theoretic backbones

The complementarity between Gi-4 and Qi-4 may be summarized as follows.

Gi-4 answers the question: *What is the geometric structure of the bulk-to-boundary projection?* The answer involves principal bundles, composite connections, anomaly polynomials, free-energy functionals, and Morse cascades. These are objects that can be written down in differential geometry.

Qi-4 answers the question: *What is the information-theoretic structure of the bulk-to-boundary projection?* The answer involves density operators, CPTP maps, entropy functionals, and decoherence channels. These are objects that can be written down in quantum-information theory.

Both questions must be answered simultaneously to give a foundational framework. Geometric structure without information structure is unobservable; information structure without geometric structure is disembodied. The two are intrinsically tied together by the bulk-to-boundary projection of OGGEF.

41.3 The unified Gi-4 + Qi-4 framework

The combined Gi-4 + Qi-4 framework constitutes the rigorous foundational layer of GDOS. Together with GG-A5-GDOS (GDOS Foundation), they form the GDOS trio: three foundational papers whose combined content axiomatizes the open-system / boundary-projection character of physical observation.

This unified framework is the formal language of GDOS. When the N=3 proof of Arisaka (2026, P2-N3) cites “GFF Morse-positivity”, when GG-A9-SM cites “GRIP cascade”, when the hierarchy derivation cites “ $M_{\text{lock}} = E_{\text{Pl}}/2e^{35}$ ”, they are all citing constructions that live in this unified Gi-4 + Qi-4 language.

The completion of GG-A7-Qi4 (Qi-4 Pillars), in preparation, will complete this language.

42 Four faces, one object

The theorem in this Part says that the four pillars and the constitutional axioms are the same content in two vocabularies: each implies the other, with nothing left over on either side.

It is worth being careful about what a claim of equivalence is worth, because it is easy to hear as either more or less than it is. It is not a discovery of new physics — no prediction follows from it that did not follow from the pillars themselves. Nor is it a triviality. An equivalence is a statement that a body of work has *closed*: that the four objects built in the preceding Parts are not four convenient constructions that happen to sit together, but four descriptions of one thing, and that the axioms one began from can be recovered from them. A framework that cannot prove such a statement about itself has not finished, whatever else it has achieved, because it cannot rule out the possibility that its pieces are independent and merely compatible.

There is an older instance of the same move that undergraduates meet early. The wave mechanics of Schrödinger and the matrix mechanics of Heisenberg were, for a while, two theories — different objects, different mathematics, mutually suspicious partisans. They were then shown to be one theory in two notations, and the demonstration produced no new prediction whatever. It was nonetheless among the most important results of the period, because it converted two rival accounts into a single subject with two calculational styles. *The value of an equivalence is that it removes a question, and removed questions are what a mature theory is made of.*

The same Part records a second pairing, and its status is different. The geometric backbone built here has an informational twin developed in a companion paper, and the two are matched pillar by pillar. That pairing is not proved here; it is stated, with the bridge theorems named and their home identified. The distinction between what this paper proves and what it points at is maintained deliberately throughout, and it is maintained here too.

It is worth saying what an equivalence theorem does not buy, since the temptation to overread one is strong. It does not show that the axioms are true, or that the arena is six-dimensional, or that any of this describes the world. It shows that if the axioms hold then the pillars hold and conversely — a statement entirely internal to the construction. The external question, whether nature satisfies any of it, is untouched by anything in this Part and is settled only by the falsifiers. A framework can be perfectly self-consistent and wrong, and the history of physics is not short of examples.

There is one more thing an equivalence buys, and it is the reason this Part exists rather than being left as a remark. It tells a reader where to attack. If the four pillars and the axioms are the same content, then anybody who wants to break the framework may do so from whichever end is more convenient — refute a pillar and the axioms fall with it, or refute an axiom and the pillars do. A construction whose parts were merely compatible would offer no such handle: one could remove a piece and the rest would stand, slightly diminished, indefinitely. *The value of proving that your work hangs together is that it can then be taken apart in one blow.*

What remains is to see whether any of this survives contact with the world. The next Part takes the machine down the ladder of the sciences — particles, chemistry, life, mind — and asks at each floor what the framework actually says and what would show it to be wrong.

Part VII

Applications

The Gi-4 framework applied to the spectrum of nature ($N = 3$, the 28 SM parameters, the mass scales) and to the dynamics of descent (radiative GRIP-In)

A framework earns its place by what it computes. This Part is the accounting, and it takes the only form the claim allows: a single walk up a single ladder.

The preceding Parts built an engine and argued that it is the engine behind every symmetry breaking in nature. An engine of that kind cannot be audited result by result, because the claim is not that it explains several things — it is that the several things are one thing seen at different depths. So this Part climbs. It begins where the cascade begins, at the first instant, and it goes up rung by rung: the particles, the atoms, the molecules, the cell, the mind, and finally the act of writing the constitution down. At every rung it asks the same three questions, in the same order.

What opened it? — the GRIP-Out shock that reorganized the landscape.

What settled? — the GGA, the ground state that the descent halted in.

Who owns it, and how strongly? — which companion paper does the derivation, and at what grade it currently stands.

That third question is why this Part is longer than a list of successes would be. The grades are not uniform. At the bottom of the ladder they are as strong as this program has: closed-form numbers with pulls against measurement and nothing fitted. At the top they are frankly weaker: named theorems with empirical anchors and derivations still being written. A survey that blurred that gradient would be worse than no survey, because it would make the strongest results harder to find and the weakest ones harder to attack. So the gradient is printed, rung by rung, in the same five words throughout.

One more thing before the climb. Almost nothing in this Part is proved here. This is the paper that builds the mechanism; the companions are the papers that run it. **What this Part is, therefore, is a preview of three of them** — the Standard-Model paper, the cosmology paper and the capstone — read through the vocabulary of Parts II–VI, so that a reader who has just watched the engine be assembled can see immediately what it has been used for.

43 The ladder, and how to read it

43.1 Two pictures of one claim

The capstone paper of this cohort draws the ladder twice, once at full length and once compressed to its skeleton. Both are reproduced here unchanged, on the two pages that follow, because they are

the shape of the claim this paper exists to support, and a reader of GG-A6-GRIP should meet them at the point where the machinery has just been built rather than in a paper three volumes away.

43.2 The reading rule

Both figures encode one alternation, and it is the alternation of Part V. A floor settles: the descent runs on a landscape that is not moving, finds the deepest basin available, and halts there. That is GRIP-In, and its terminus is a GGA. Because the landscape is fixed by the geometry rather than by local history, every copy of that object anywhere in the universe lands in the same basin — which is why every electron has the same mass. Then the settled floor *is itself* the disturbance that reorganizes the landscape above it: a universe full of stable electrons and nuclei is a universe in which atoms become possible. That is GRIP-Out, and it opens the next floor, which settles in its turn.

The engine performing that alternation is the one drawn in Figure 8 of Part V, and a reader who wants the mechanism rather than its consequences should keep that page open alongside this one. The capstone states the recursion in a single sentence worth quoting, because it is a claim about this paper’s machinery and not about biology: the transition from one floor to the next *is not a fifth operational move; it is the same four moves operating on a perturbed landscape, with the perturbation supplied by the layers below*.

Two consequences of that follow immediately, and they are what make the ladder a mechanism rather than a metaphor.

The cascade carries its own clock. Nothing external schedules the transitions. Each GRIP-Out is triggered by the completion of the GRIP-In below it, so the ordering of the eighteen layers is not a historical accident that could have gone otherwise but a consequence of which floors must exist before which others can. There is no designer in the picture and no timer beside it.

Every floor is a ground state, not a resting place. A merely stable configuration is stable against a nudge; a ground attractor is stable against eternity, because there is nothing below it to fall to. Only a floor made of ground states can be built on, because only a ground state is still there when the next floor needs it. A cascade of way stations would decay out from under itself.

43.3 The five grades, and why they are printed

The capstone assigns every rung of Figure 9 one of five confidence tags. This Part adopts them, and prints the tag on every rung below.

The capstone is blunt about what this adds up to, and so is this Part: four of the eighteen rungs are closed strictly, seven of the eighteen once an empirical anchor is allowed to stand in for a derivation, and nothing above the atomic rungs is claimed at codebook standard. What is *not* graded down is the empirical exposure: a rung tagged ASSERTED still carries falsifiers with numbers and instruments attached, and failing one of them ends the framework whatever the surrounding tag says.

Inevitable 18-Layer Cascade of GRIP-Out




Figure 9: The eighteen-layer cascade of GRIP-Out. The seventeen GRIP-Out shocks and their approximate ages, floor by floor, from the Standard-Model gauge group to the act of articulation. Each box is a layer that settled; each arrow is the event that reorganized the landscape and opened the layer above. The right-hand column names the mechanism of each shock and the left-hand column its approximate date. This Part is a walk up this figure. Carried unchanged from GG-A11-Origin (Arisaka, 2026, A11-Origin), where the row-by-row version — mechanism, timescale and empirical anchor for every transition — is given in full. One string is changed in transit and no other: Layer 18’s annotation, which in the capstone reads “the present paper”, is resolved here to the paper it names, since the top rung of the cascade is that paper and not this one.

The Cascade of GRIP-Out with GGA

one mechanism, five floors — from the constitution to the mind that reads it




Figure 10: The cascade of GRIP-Out with GGA — the same claim in five floors. Each floor is a Ground Geometric Attractor: the deepest basin of its stratum, which is why the objects on that floor are identical wherever one looks. GRIP-In (left, in cyan) is the patient locking — each layer relaxes to its unique ground state, the same everywhere in the sky, by geometry rather than by shared history. The settled ground state then *becomes* the GRIP-Out shock (orange) that opens the floor above, so the cascade carries its own clock and needs neither an external designer nor an external timer. The first lock, the electron across the whole sky, is the identity-horizon problem of GG-A10-Cosmos; the last, at the top, closes the loop back onto the constitution the whole structure started from. Carried unchanged from GG-A11-Origin (Arisaka, 2026, A11-Origin).

Table 9: The five grades, as GG-A11-Origin (Arisaka, 2026, A11-Origin) defines them, adopted here as this Part’s key. They are ordered by strength, and the order matters more than any single tag: the ladder’s grade falls as it rises, and a reader deciding where to attack this framework should attack where the grade is highest, because that is where a failure would be unrecoverable.

Grade	What it promises
DERIVED	The chain from the constitution to the result is closed in the published record. The result may be used.
ANCHOR-DERIVED	The empirical anchor is done; the derivation is deferred to a companion in preparation. The number may be used, the derivation of it may not.
CONDITIONAL	The chain closes once one named, stated lemma is proved. The result is forced only if that lemma survives.
ASSERTED	The framework expects the result and can say what would refute it, but the derivation to codebook standard is not in hand. A stated expectation, not a result.
EXISTENCE- DEMONSTRATED	Exhibited rather than derived.

43.4 The spine of this Part

Table 10 is the whole Part in one page. Everything after it expands one of its rows. A reader with ten minutes should read the table and stop.

44 Floor one — Cosmos: the particles and the sky

The first three rungs are where this program’s grade is highest, and they are the three the rest of this Part leans on. They are also the three that most of the cohort’s arithmetic lives in. This section is correspondingly the longest, and it is the one to read if only one is read.

44.1 The landscape itself: five integers

Before any rung can settle there has to be a landscape for it to settle on, and the landscape is not free. Part III’s closure condition admits exactly one integer ledger, $(M, N, L; J, K) = (1, 3, 5; 7, 35)$, with $K = LJ = M^2 + N^2 + L^2$. Everything in this Part that carries a number is, in the end, a function of those five together with π and e .

It is worth being concrete about how far that reaches, because the reach is the argument. The same five integers fix the inverse fine-structure constant at the Z pole, the number of fermion families, the electron mass, the epoch of the first atoms, the fraction of the universe that is ordinary matter — and, as §45.2 will show, the number of amino acids that biology uses. A ledger that short cannot be adjusted to fit any one of them without breaking the others. That is what “no free parameters” means in practice, and it is also what makes the framework easy to kill.

Table 10: The spine of the cascade. Every rung of Figure 9, with the GRIP-Out event that opens it, the GGA that settles into place, the paper of this cohort in which the derivation is done, and the grade that derivation currently carries. The rows owned by GG-A11-Origin point to the closing capstone; the grade column is what carries their present standing, and it is stated here in full rather than deferred. Bands follow the brackets of the figure. The sections of this Part are a walk down this table’s rows.

Layer	Opened by	Settles as	Owner	Grade
<i>Cosmos (§44)</i>				
1 Gauge group	The exit from the unified edge	$SU(3) \times SU(2) \times U(1)$ + gravity	GG-A9-SM	DERIVED
2 Leptons & quarks	Cosmic cooling to the electroweak scale	Three families; $v \approx 246$ GeV	GG-A9-SM, GG-P2-N3, GG-P3-Electron, GG-P6-Proton	DERIVED
3 Periodic table	Recombination at $z \approx 1090$	Atomic ground states; $E_{\text{Ry}} = 13.6$ eV	GG-A10-Cosmos	DERIVED
<i>Chemistry (§45)</i>				
4 Small molecules	Stellar nucleosynthesis and supernovae	Water, carbon dioxide, ammonia, methane	GG-A11-Origin	CONDITIONAL
5 Amino acids, sugars	Cometary delivery; interstellar photochemistry	Twenty L-amino acids; the D-sugars	GG-A11-Origin	CONDITIONAL
6 Polymers	Wet-dry cycling; ultraviolet pulses on clay	Oligopeptides, oligonucleotides, fatty acids	GG-A11-Origin	CONDITIONAL
<i>Life (§46)</i>				
7 Bioenergetic anchor	Hydrothermal vents; serpentinization	ATP at $2N_A Q_B \approx 30$ kJ/mol	GG-A11-Origin	ANCHOR-DER.
8 The RNA world	Vent-coupled ignition	A four-letter alphabet, $ \Sigma = N + 1$	GG-A11-Origin	CONDITIONAL
9 Information closure	The RNA-to-DNA takeover	The archival duplex; error floor 10^{-9} /bp	GG-A11-Origin	ANCHOR-DER.
10 The first cell	Membrane closure; the autocatalytic loop	A bounded, metabolizing cell	GG-A11-Origin	ANCHOR-DER.
<i>The multicellular and the mind (§47)</i>				
11– Domains, eukaryotes, multicellularity	Planetary transition; oxygenation; Snowball Earth	Two primary lipidomes; the chimeric cell; the eight-rung ladder	GG-A11-Origin	CONDITIONAL
14– Nervous system to conscious episode	Ediacaran radiation; the Cambrian; encephalization	A ring attractor; a four-dimensional map; a 400 ms episode	GG-A11-Origin	ASSERTED
<i>Articulation (§48)</i>				
18 Articulation	Writing, mathematics, the accumulated corpus	The theory that names its own origin	GG-A11-Origin	EXIST.-DEM.

44.2 Why exactly three families [DERIVED]

Of everything the Standard Model does not explain, the family count is the plainest. Ordinary matter is built from one family of quarks and leptons. Nature made two heavier copies of it and then stopped. Nothing in the Standard Model’s equations prefers three to two or to seventeen; the number has been written in by hand since the muon turned up uninvited in 1936.

The three-generations companion (Arisaka, 2026, P2-N3) derives it, and the derivation is the sharpest single application of the cascade this paper builds — sharp enough that Part V’s own bound is stated as conditional on it. The argument is a pincer, closing on one value from three independent directions.

44.2.1 The lower bound: matter and antimatter

The first leg is the oldest. Matter and antimatter behave differently, and the difference is measured: the Jarlskog invariant of the quark sector is about 3.1×10^{-5} and is definitely not zero. A CP-violating phase can only exist in a mixing matrix of at least three generations — the count of physical phases is $(N_g - 1)(N_g - 2)/2$, which vanishes for one family and for two. Under the minimality axiom this gives $N_g \geq 3$ outright, resting on one measured number and no geometry beyond it.

44.2.2 The upper bound: a descent that cannot afford a fourth step

The second leg is the one that consumes Part V, and it is where the family count becomes a statement about the cascade.

The boundary’s family configuration is a rank-one projector, and its moduli space is the complex projective space $\mathbb{CP}^{N-1}$. The free-energy functional of Part IV, restricted to that space, is a Morse function with exactly one critical point at each even index — the landscape has exactly N levels. The descent of Part V runs down them, paying at each gate: the toll for the step from index $2k$ to index $2(k - 1)$ is $S = \pi k$, so a cascade through N_g levels costs $S_{\text{tot}}(N_g) = \pi N_g(N_g - 1)/2$ in total.

Against that cost stands a capacity. The companion proves that the absorptive capacity of the family stratum is $\text{Cap}(N_g) = \pi N_g$ for $N_g \geq 3$, obtained by pairing a characteristic class of the stratum’s own geometry against the generator of its four-dimensional homology — a generator that exists only when $N \geq 3$. The arithmetic then closes with no room in it:

Table 11: The two sides of the closure condition, evaluated at each candidate N . Cost is the total action the stratum’s own geometry has to pay; capacity is what the four-dimensional homology generator can supply, and it is zero below $N = 3$ because that generator does not exist there. So below three there is nothing to balance against: at $N = 2$ the descent already costs π and the capacity to absorb it is zero. At $N_g = 3$ the two are equal and non-zero. From $N_g = 4$ the cost has already outrun the capacity, and the gap widens at every further step.

N_g	1	2	3	4	5	6
Cost $S_{\text{tot}}(N_g)/\pi$	0	1	3	6	10	15
Capacity $\text{Cap}(N_g)/\pi$	0	0	3	4	5	6

One family passes vacuously, having nothing to pay for. Two families fail: the homology generator does not exist, so the capacity is zero and the single step has no payer. Three families saturate the budget exactly. Every $N_g \geq 4$ fails, because the cost grows quadratically and the capacity only linearly. The admissible set is $\{1, 3\}$, and the first leg has already removed the 1.

The physical reading is the cascade’s own. For four or more families the toll exceeds the capacity, which means the ground attractor becomes *dynamically unreachable*: matter would strand permanently in non-ground basins, and the constitutional guarantee that every stratum has a reachable ground state would be violated. A fourth family is not forbidden by a selection rule. It is forbidden because the road down to it could not be paid for.

44.2.3 The third leg, and the fence around all three

The third leg is arithmetic consistency. Fed back into the mass codebook, the same integer reproduces the charged-lepton ratios — $m_\tau/m_\mu = 16.817$ and $m_\mu/m_e = 206.768$ — to within eight tenths of a percent, while $N_g = 2$ and $N_g = 4$ miss by factors greater than five.

The companion’s own fence is carried here unaltered. Its closing theorem keeps the word *conditional* in its title. It leans on two empirical inputs, the measured CP violation and the two mass ratios; on the anomaly closure of Part III’s companion; on premises it shares with the constitutional Born theorem, so that this leg and that theorem stand or fall together by design; and on a coupling clause whose form is asserted and whose magnitude, class and non-vanishing are then derived. The integrality statement underwriting the capacity has been discharged on a benchmark background by an inflow bridge carried as that paper’s own Appendix H, with two refinements open: the warping of the physical seam, and the relative cohomology of the bounded half-seam. Neither moves the value of N ; what they would cost is the enforcement of the coupling clause, not the prediction.

44.2.4 What would end it

A fourth charged lepton or quark below about ten TeV. A fourth light neutrino at multi-sigma significance, against the invisible-width measurement $N_\nu = 2.984 \pm 0.008$ that currently excludes both two and four. The tau-to-muon mass ratio off 16.817 by more than half a percent. Zero CP violation in the B or kaon systems. Or a proof that the Morse-positivity of Part IV is incompatible with the closure of Part III.

44.3 Application: the $N=3$ derivation in Gi-4 vocabulary

Cohort map. The dedicated derivation is the prediction-branch paper [Arisaka \(2026, P2-N3\)](#); the 3π budget is the capacity theorem T-N3-6 of [Arisaka, 2026, P2-N3](#) (conditional on its coupling clause and L-INT, residue R-INT; open problem M10, demand met), the OGN-5 ledger by [Arisaka \(2026, A2-GG6\)](#) under the MSMF parsimony axiom of [Arisaka \(2026, A1-Constitution\)](#), and the empirical CP input ($\mathcal{J}_{\text{CKM}} \neq 0$) by the flavor data carried in [Arisaka \(2026, A9-SM\)](#).

The headline downstream consequence of the four Gi-4 pillars is that the family number $N_g = 3$ is forced — not chosen. Within the present paper the three-leg derivation is Corollary [34.4](#) (§34.4), which combines the three principal theorems of Parts III–V:

- **Leg L1** (lower bound, $N \geq 3$). Theorem **T2** of Part III plus the empirical input of CP violation (Jarlskog invariant) fixes the constitutional lower bound.
- **Leg L2** (parity constraint, N_g odd). Theorem **T3** of Part IV (GFF Morse-positivity) restricts the family stratum to real-dimension $4\mathbb{Z}$, hence N_g odd.
- **Leg L3** (upper bound, $N_g \leq 3$). Theorem **T4.1** of Part V (GRIP cascade closure on $\mathbb{CP}^{N-1}$ saturating the Bl-6 budget $S_{\text{Bl6}}^{\text{max}} = 3\pi$, conditional on M10) delivers the upper bound.

The intersection of the three legs is uniquely $N_g = 3$. The full expository chain, the hypothesis-to-theorem conversion of the two “three-leg” hypotheses $\mathcal{H}_1, \mathcal{H}_2$ (now Theorems 27.1 and 27.2), and the empirical cross-checks against the observed CKM (Kobayashi & Maskawa, 1973) and PMNS (NOvA, 2022, T2K, 2020) structure live in the dedicated downstream paper Arisaka (2026, P2-N3); the present paper records only the constitutional output of the three-leg proof.

44.3.1 Physics intuition: families as basins, the weak force as the gate

Before the formulae, the picture. Imagine the space of possible flavor configurations as a landscape of hills and valleys. The most symmetric configuration — equal weight on all N_g family directions, no generation singled out — sits on the very top of a hill: it is the DGI point, and like any hilltop it is unstable, ready to roll. The system rolls downhill, and each valley it can settle into is a *generation*. The mountain passes between valleys are the GSSB saddles, and crossing a pass is not free: it costs a definite amount of “action,” the geometric analog of the energy a ball needs to get over a ridge. The decisive physical fact is that only *one* of the four forces can carry a state across such a pass — the weak interaction, the chiral, short-range, generation-changing force. Gravity, electromagnetism, and the strong force conserve the labels that distinguish the valleys and can only move a state *within* one. So the weak force is the “gate” of the family landscape, and the number of generations is simply the number of valleys the landscape can afford to carve before its gate-crossing budget is spent. “Why three families?” becomes “how many times can this landscape descend before it runs out of budget?” — a counting question, with a counting answer.

44.3.2 The GRIP mechanism: a budgeted gradient flow

Now the landscape made precise. The family-projector stratum is the complex projective space $\mathbb{CP}^{N-1}$, and the GFF supplies a Morse height function $\mathcal{F}_{\text{GFF}}|_{\mathbb{CP}^{N-1}}$ whose critical points are the N_g coordinate axes, carrying Morse indices $0, 2, 4, \dots, 2(N-1)$ (Theorem 32.1 and its proof). The index- $2(N-1)$ maximum is the DGI top; the index-0 minimum is the GA; the intermediate even-index saddles are the GSSB chain. The cascade is literally the gradient flow of this function,

$$\dot{x} = -g_{\mathbb{CP}^{N-1}}^{ab}(x) \partial_b \mathcal{F}_{\text{GFF}}(x) \partial_a, \quad \frac{d\mathcal{F}_{\text{GFF}}(x(t))}{dt} \leq 0, \quad (187)$$

the monotonicity being the irreversibility of the descent: the landscape is fixed, and the configuration only ever rolls *down*. Between two adjacent critical points the descent pays a Bogomolnyi-saturated

instanton toll, computed in closed form on the Fubini–Study geometry,

$$S_{[2k] \rightarrow [2(k-1)]} = \pi k \quad \implies \quad S_{\text{tot}}(N_g) = \sum_{k=1}^{N_g-1} \pi k = \frac{\pi}{2} N_g(N_g - 1). \quad (188)$$

This is where the budget bites. The BI-6 anomaly-inflow ledger caps the *total* action a closing cascade may spend at $S_{\text{BI6}}^{\text{max}} = 3\pi$ (the OGN-5 closure $K = 35$, $J = 7$), so a consistent cascade requires $S_{\text{tot}}(N_g) \leq 3\pi$, i.e. $N_g(N_g - 1) \leq 6$. The integer solutions are stark:

$$S_{\text{tot}}(2) = \pi < 3\pi, \quad S_{\text{tot}}(3) = 3\pi = 3\pi, \quad S_{\text{tot}}(4) = 6\pi > 3\pi. \quad (189)$$

Two families leave the budget unspent and the cascade fails to close; three spend it exactly; four would need twice what the geometry can pay, and are inconsistent. A second constraint, requiring *no* conjecture, removes the even values independently: the alternating Morse sum on $\mathbb{CP}^{N-1}$ equals the Hirzebruch signature, $a^+(N_g) - a^-(N_g) = \sigma(\mathbb{CP}^{N-1})$, which is 1 for odd N_g and 0 for even N_g , so an even- N_g cascade has no consistent index balance at all. Combined with the lower bound $N_g \geq 3$ from the existence of a physical CP phase (a non-zero Jarlskog invariant needs at least three generations), the three legs intersect on a single integer, $N_g = 3$ — the upper bound resting on the capacity theorem that fixes $S_{\text{BI6}}^{\text{max}} = 3\pi$ (M10’s demand met; residue M6 = R-INT), the lower bound on the measured reality of CP violation. The full conditioning is carried in [Arisaka \(2026, P2-N3\)](#).

44.3.3 Why the GRIP reading helps

The value of casting $N = 3$ this way is sharpest by contrast. *Against Froggatt–Nielsen*: there the flavor hierarchy is built from a suppression ϵ^q with both the small parameter ϵ and the integer charges q chosen to fit the data; GRIP fixes the base of the suppression at the Bogomolnyi value $e^{-\pi}$ — not adjustable — and replaces the free charges by Morse indices that are *counted*, not assigned. *Against the string landscape*: a continuum of $\sim 10^{500}$ vacua predicts no family number at all, because it carries no selection principle; the BI-6 budget *is* a selection principle, and it selects $N_g = 3$. *Against the anthropic principle*: that explains the family number by observer selection after the fact, whereas the cascade derives it before any observer is invoked. The reading converts an empirical accident into a counting statement — three is the number of downhill steps a 3π budget can pay for — and, crucially, it is falsifiable at the root: a fourth chiral family would demand $S_{\text{tot}}(4) = 6\pi$, exactly twice the available budget, so a single confirmed fourth generation would not dent the framework but break its central mechanism outright.

44.4 The twenty-eight parameters [6 DERIVED, 14 CONDITIONAL, 8 ASSERTED]

The Standard Model is the most accurate theory ever written and the least explanatory. It predicts the anomalous magnetic moment of the electron to twelve figures, and it cannot say why the electron has the mass it has. Twenty-eight quantities — the coupling strengths, the fermion masses, the mixing angles, the CP phases — must be measured and written in by hand.

The Standard-Model companion ([Arisaka, 2026, A9-SM](#)) argues that those twenty-eight are not

free choices of nature but *observations*: values that a fixed six-dimensional geometry, projected onto the boundary we inhabit, must return.

44.4.1 What the companion is allowed to use

Two short lists and nothing else: five dimensionful constants — the speed of light, Planck’s constant, Newton’s constant, Boltzmann’s constant and the Arisaka constant — and the five integers of §44.1, together with π and e . Explicitly forbidden as inputs are the Fermi constant, the W , Z and Higgs masses, the strong coupling, the weak mixing angle, and every fermion mass and mixing parameter. The discipline matters more than any individual result: one short ledger has to serve all twenty-eight quantities at once, so a value that comes out right cannot have been arranged, because arranging it would break something else.

44.4.2 Every parameter is one of four things

The structural claim is the one Part V makes possible. Each of the twenty-eight is exactly one of four objects in the cascade: a Geometric Attractor; the Ground Geometric Attractor of its stratum; a GSSB toll, that is, a Bogomolnyi action $S = \pi k$ appearing as a suppression exponent; or an algebraic descendant of a selected branch. The companion states the reading in one line — every parameter Part of that paper is the same construction, a GRIP cascade, read on a different boundary stratum — and it assigns the regimes explicitly: the electroweak vacuum selection is the single GRIP-Out event, and every fermion mass, mixing angle and lifetime is a GRIP-In readout on the landscape that event built.

Table 12: The assignment of the twenty-eight Standard-Model parameters to boundary strata, as the Standard-Model companion (Arisaka, 2026, A9-SM) makes it. The table is the reason that paper can describe itself as one construction repeated: every Part of it runs the cascade of Part V on a different stratum, and the parameter class is fixed by which of the four cascade objects the readout is. Only the first row is a GRIP-Out event; everything below it is GRIP-In on the landscape that event built.

Parameter class	Stratum	Cascade object
Gauge couplings, the Higgs sector	The electroweak lock seam	The selected branch of the single GRIP-Out event
Charged leptons; quarks	The family-projector moduli $\mathbb{CP}^2$	Critical points at depths 0, 1, 2; the depth-2 point is the GGA, the only basin
The mass ratios within a family	The same stratum, between gates	GSSB tolls: suppression exponents $S = \pi k$
Neutrino masses	The rank-two Majorana block	Attractors on a stratum of different rank
Mixing angles and CP phases	The misalignment between projectors	Algebraic descendants of the selected branches
The strong-CP angle	The topological ledger	A winding, not a basin

Table 13: Representative entries from the parameter codebook of GG-A9-SM (Arisaka, 2026, A9-SM), quoted with that paper’s own first-order values, measurements and pulls. Every closed form is built from the five ledger integers together with π and e , and from nothing else. The point of the table is not any single row but the constraint the rows are under: one ledger has to serve all twenty-eight quantities simultaneously, so no row can be adjusted without disturbing the others.

Quantity	Closed form	First order	Measured	Pull
$\alpha^{-1}(M_Z)$	2^J	128.0	127.930(8)	+0.99
$\alpha^{-1}(0)$	$2^J + N^2$	137.0	137.036	−1.00
$\sin^2 \theta_W(M_Z)$	$L^2/4N^3$	0.23148	0.23155(4)	−1.70
$\alpha_s(M_Z)$	$1/\pi e$	0.11710	0.1180(9)	−0.70
v [GeV]	$M_{\text{lock}} (2/L)^N$	246.33	246.220	+1.00
λ_H	2^{-N}	0.125	—	—
m_e [MeV]	$(v/\sqrt{2}) N/(L^3 2^{J+L+1})$	0.51030	0.51100	−0.88
$\sin \theta_{12}^{\text{CKM}}$	$ e^{-\pi/2} - e^{i3\pi/5} e^{-\pi} $	0.22502	0.22501	+0.01

44.4.3 Some of what comes out

Two features of Table 13 are worth pointing at. The closed forms are made of the ledger integers and nothing else: 2^J , $L^2/4N^3$, 2^{-N} . And the electron mass sits at the bottom of a chain of suppressions, $L^3 2^{J+L+1}$, which is exactly what Part V’s cascade produces — a descent that pays a toll at each gate arrives exponentially far below where it started.

44.4.4 What the companion does not claim

Three statements are carried here because a survey that omitted them would misrepresent that paper.

Its own triage: of the twenty-eight, six are graded Derived, fourteen are conditional on a single named lemma each, and eight — the charged-fermion Yukawa couplings other than the top — remain Asserted. The companion is explicit that its headline claim, that not one of the twenty-eight is a free input, is an ontological statement about what the parameters *are*, not a claim that all of them have been computed.

A tension it records rather than hides: the weak mixing angle agrees with the effective leptonic value at the level shown in the table; against the $\overline{\text{MS}}$ value the same prediction sits at about six and a half standard deviations, and the companion prints that in the open.

And the status of the family count, which §44.2 has already given at its own grade.

44.4.5 What would end it

Six channels are pre-registered: a resonance cluster near $M_{\text{lock}} \approx 3.85$ TeV; a dark-matter line at 8.1 μeV ; a neutrino package of which an established inverted ordering would falsify the assignment outright; absolute proton stability; a fifth force of stated range; and a cosmological constant with $w = -1$ exactly, two-sided.

44.5 Application: the 28 SM parameters in Gi-4 vocabulary

Cohort map. The full twenty-eight-parameter codebook with its four-significant-figure fits is [Arisaka \(2026, A9-SM\)](#); the anomaly polynomial X_4 and its Hopkins–Singer descent come from [Arisaka \(2026, A3-BI6\)](#), the OGN-5 ledger (with $2^J L = 640$) from [Arisaka \(2026, A2-GG6\)](#) and [Arisaka \(2026, A1-Constitution\)](#), and the electron ground-state reading from [Arisaka \(2026, P3-Electron\)](#).

The 28-parameter codebook of GG-A9-SM ([Arisaka, 2026, A9-SM](#)) maps onto the Gi-4 vocabulary as cascade outputs. Each Standard-Model parameter is a depth-labeled Geometric Attractor on a specific stratum of $\mathcal{M}_{\text{bdy}}$, and the codebook’s structural ingredients have a one-to-one Gi-4 reading:

- *Yukawa hierarchies* as GA cascade depths on the family-projector stratum $\Sigma_{\text{fam}} = \mathbb{CP}^2$ (third-, second-, first-generation = the index-4, 2, 0 critical points; the two adjacent Bogomolnyi tolls are 2π then π descending from the heavy end (Theorem 32.1), so the accumulated action is 0, 2π , 3π , saturating the BI-6 budget $S_{\text{BI6}}^{\text{max}} = 3\pi$);
- *Angular-winding suppression* $e^{-n\pi}$ in the GG-A9-SM quark codebook as the cascade exponential $e^{-S_{\text{tot}}(n)}$;
- *QD-completion rationals* (e.g. $\mathcal{R}_\tau = 637/640$, $\mathcal{R}_\mu = 159/160$) as integrality coefficients of the Hopkins–Singer descent of X_4 , with denominators $2^J L = 640$ directly visible in the OGN-5 ledger.

The full 28-parameter codebook, including all numerical fits to four-significant-figure precision, the QD-completion derivation as a Hopkins–Singer descent theorem, and the empirical confrontation with Particle Data Group values, is the subject of GG-A9-SM ([Arisaka, 2026, A9-SM](#)). The present paper records only the Gi-4-side translation of the codebook’s structural ingredients.

44.5.1 Physics intuition: a parameter is the depth of a basin floor

The orthodox Standard Model hands the experimenter twenty-eight numbers and asks that they be measured and inserted by hand. The Gi-4 reading says something different and, once seen, hard to unsee: each of those numbers is not a dial but an *altitude* — the depth at which a particular descent comes to rest on its own stratum. A Yukawa coupling is how far down the family cascade a generation’s valley floor lies. A mixing angle is how far two cascades, running on neighboring strata, are rotated out of register with each other. A gauge coupling is the value the running settles to once it reaches its attractor. The codebook of twenty-eight parameters is therefore a single catalog of basin floors, and the altitude of a floor is fixed by integer Morse data — which axis, how many gates down — not chosen on a whim. The parameters stop being a list of accidents and become a topographic map of one landscape.

44.5.2 The GRIP mechanism: cascade depth becomes exponential hierarchy

The map has a precise scale. A generation reached at the index- $2k$ critical point is arrived at only after the gradient flow has paid the accumulated Bogomolnyi tolls of (188), and the resulting

Yukawa is the *exponential* of that accumulated action,

$$y_{\text{gen}} \sim e^{-\Delta S}, \quad \Delta S \in \{\pi, 2\pi, 3\pi, \dots\}, \quad (190)$$

so the toll quantum π sets a fixed geometric spacing $e^{-\pi} \approx 0.043$ between rungs of the spectrum. On the charged-lepton stratum $\mathbb{CP}^2$ the three generations occupy the index-0, 2, and 4 critical points; descending from the index-4 top the two adjacent tolls are 2π then π (Theorem 32.1), so the action accumulated from the heavy end runs 0, 2π , 3π and saturates the budget 3π of (188) — the lighter the generation, the deeper it sits. The spectrum is thus laid on a *fixed* geometric ladder whose base $e^{-\pi}$ is not a fitted parameter but the Bogomolnyi action of one gate-crossing. The finer four-significant-figure structure rides on top of this skeleton: the quark angular-winding suppressions $e^{-n\pi}$ are the same cascade exponentials read on the Σ_θ stratum, and the QD-completion rationals (for instance $\mathcal{R}_\tau = 637/640$, $\mathcal{R}_\mu = 159/160$) are the integrality coefficients of the Hopkins–Singer descent of the anomaly polynomial X_4 , with denominators $2^J L = 640$ read directly off the OGN-5 ledger. The complete numerical confrontation with Particle Data Group values is the business of GG-A9-SM (Arisaka, 2026, A9-SM); what the Gi-4 reading contributes is the structural statement that *every* entry of the codebook has the form $e^{-(\text{integer} \times \pi)}$ times a ledger rational.

44.5.3 Why the GRIP reading helps

A free parameter hides its degrees of freedom; a cascade-depth parameter publishes them, as an integer Morse index that a measurement can contradict. The reading therefore does two things at once. It *disciplines*: turning “twenty-eight unexplained numbers” into “twenty-eight basin floors on a fixed $e^{-\pi}$ ladder,” each carrying a falsifiable integer rather than an adjustable real. And it *unifies*: a mass, a mixing angle, and a gauge coupling are no longer three different kinds of input but the same kind of object — the floor of an attractor — distinguished only by which stratum of the boundary they live on. A fit buys agreement by spending continuous freedom; here the only freedom is which integer, and the base of the hierarchy is forced. That is the difference between describing the spectrum and deriving it.

44.6 The electron: a descent that halts [DERIVED]

Every electron in the universe has the same mass, the same charge and the same spin. None has ever been seen to decay, or to show an inner structure at any energy yet reached. Exactly two heavier copies of it exist, and only two. These are ordinary facts, taught early and rarely questioned, and the Standard Model accommodates all of them without explaining any.

The electron companion (Arisaka, 2026, P3-Electron) derives them from one statement, and the statement is the terminus of this paper’s cascade: *the electron is the unique Ground Geometric Attractor of the charged-lepton stratum of the boundary*.

44.6.1 Identity, and why it is not a coincidence

The landscape of Part IV admits exactly one deepest basin on the charged-lepton stratum. Every electron is a readout of that single attractor: its mass is the value of the free-energy functional at

the minimum, its charge the integer-normalized eigenvalue on the unique stable leg of the cascade. Two electrons cannot differ in any intrinsic property, because there is only one basin for them to be in. Two electrons on opposite sides of the sky have never been in causal contact, so their agreement cannot have been negotiated — and it was not. It was never in question. That is the identity-horizon problem of §44.9.4.

44.6.2 Stability: the basin with no exit gate

This is the cleanest application in the cohort of the distinction Part V draws between an attractor and a *ground* attractor.

On the family stratum $\mathbb{CP}^2$ the landscape has three critical points, of Morse index four, two and zero, and they are the three charged leptons: the tau at depth zero, the muon at depth one, the electron at depth two. The observed decay chain $\tau \rightarrow \mu \rightarrow e$ is therefore not three unrelated weak processes but a single GRIP-In descent that pays the tolls 2π and π in order — the total 3π that §44.2.2 showed exactly saturates the stratum’s capacity.

At the bottom the descent stops, for a reason that needs no machinery: the electron basin is the global minimum of its stratum and possesses no exit gate. There is no admissible GSSB gate, of any index and any finite toll, leading out of it. The electron cannot decay — not because some conservation law forbids each channel of an otherwise unexplained spectrum, but because the landscape contains no downhill direction and no gate.

The runtime dual is the weight ladder of Part V. A particle with an exit gate has a jump operator and therefore a width; what the gate toll fixes is the particle’s weight on the skeleton, not its lifetime, which is phase space (**GRIP-ST7**). The muon has one gate and the familiar width follows. The electron has no gate, so its reduced generator is purely unitary, its width is exactly zero, and its lifetime is infinite. The companion adds that the numerical muon lifetime is derived modulo the muon depth, which its own triage grades as Asserted, so the $2.2 \mu\text{s}$ is not advanced as a parameter-free postdiction — only the *reading* is unconditional.

44.6.3 The mass, and the grade it carries

The electron Yukawa coupling comes out of the ledger as $y_e = N/(L^3 2^{J+L+1}) = 3/1,024,000$. With the second-order completion of the codebook the mass is 0.511097845 MeV against a measured 0.51099895069(16) MeV — agreement at the 1.9×10^{-4} level, with the residual in the fourth figure declared rather than absorbed. The companion grades this as codebook-level.

Two honesty clauses belong with it. The electron’s anomalous magnetic moment is engaged only qualitatively, as the empirical anchor of structurelessness; a quantitative statement is an open problem. And the strict Morse-positivity on which the uniqueness of the deepest basin rests is listed as unsettled in GG-A5-GDOS’s own open problems.

44.6.4 What would end it

A single observed electron decay, in any channel: there is no conditioning available to appeal to. A confirmed structural form factor at any energy. A fourth sequential charged lepton lighter

than the electron, which would remove the ground state from under everything above. A confirmed asymmetry between the electron and the positron.

44.7 The proton: a winding that cannot be undone [DERIVED]

Almost everything anyone has ever weighed is protons and neutrons, and almost none of that weight is the quarks inside them. The proton presents a single exact mass with zero width, a charge equal and opposite to the electron's to at least twenty-one decimal places, and a spin of exactly one half that its own constituents famously fail to add up to. It has never been observed to decay, and the experimental floor under its lifetime now stands above 2.4×10^{34} years.

The proton companion (Arisaka, 2026, P6-Proton) runs the electron paper's method one stratum over: *the proton is the Ground Geometric Attractor of the baryonic stratum, and its baryon number is not a bookkeeping label but a winding of the compact geometry itself.*

44.7.1 Why baryon number could not have been a label

The companion's first result is negative and it forces the rest. On the split unified spectrum, baryon number cannot be realized as a linear combination of the available charges: three conditions on two unknowns, with no solution. The combination that *is* realizable, $B - L$, does not protect the proton at all, because every dimension-six proton-decay operator conserves it. So baryon number is either accidental — in which case the proton's stability is an accountant's coincidence that the next operator could spoil — or it is something other than a charge.

44.7.2 The twist, and the four locks

It is something other than a charge. The descent that breaks the unified symmetry does so by a holonomy — a twist of the compact direction — and that same twist shifts the winding lattice of the boundary fields into thirds. The chiral zero modes then carry winding exactly equal to their baryon number: $+\frac{1}{3}$ for the quark doublet, $-\frac{1}{3}$ for the singlets, zero for the leptons. Integrating over the compact angle annihilates every baryon-number-violating vertex, while the one vertex that survives conserves baryon number. One twist, four locks.

The resulting stability theorem states that every local boundary operator whose baryon-number violation is not a multiple of three has identically vanishing coefficient to all orders, so proton decay and neutron–antineutron oscillation are both forbidden outright. It carries three named pins, and the companion states its own grade plainly: this is not “the proton is stable, full stop”, but “the proton's stability has gone from an accountant's accident to a structural theorem with named, bounded residues”.

44.7.3 Two grounds, one grammar

This is the part of the floor that most repays a slow reading, because it shows the mechanism of Part V doing two different jobs with one vocabulary.

Both particles are absolutely stable, and both because they are the ground attractor of their stratum. But the *reason* the basin has no exit differs. The electron gets its ground-state status

from the shape of its landscape: there is no deeper charged-lepton valley, and the descent halts for want of anywhere to go. The proton earns its from a gauged topological law: its baryon number is an integer winding, and an integer is not protected by the absence of a lower state but by the impossibility of continuous motion. Nothing continuous acts on the integers. A symmetry can be broken; a winding has nothing to break.

The companion’s one-sentence summary is worth quoting: the electron is immortal because its basin has no exit gate; the proton is immortal because its winding has no unwinding operator. One framework, two pillars of matter, two different kinds of forever.

Table 14: The electron and the proton, read through the same vocabulary. Every row but one is a parallel; the exception is the fourth, and it is the whole physics of the proton companion. Compiled from [Arisaka \(2026, P3-Electron\)](#) and [Arisaka \(2026, P6-Proton\)](#). Neither column is derived here; both are quoted at the strength their own papers give them.

	The electron	The proton
Stratum	The charged-lepton stratum, $\mathbb{CP}^2$	The color-singlet baryonic stratum
The one fact	It is the unique deepest basin of its stratum	It is the unique deepest basin of its stratum
Why every copy is identical	One basin, one readout; identity is not negotiated	The same, transposed directly from the electron paper
Why it cannot decay	No deeper charged-lepton valley exists; GRIP-In halts for want of anywhere to go	A gauged winding removes every operator that could reach a deeper state
The way stations above it	The muon and the tau: attractors with exit gates, hence widths	The $\Delta(1232)$ and the free neutron: one open gate each
The twin	The positron, the same fact read in the mirror	The antiproton, the reversed wrap
The mass, and its grade	0.5111 MeV at 1.9×10^{-4} ; codebook grade	≈ 900 MeV to about five percent; no strong-sector input

Between them, the two establish the ground attractors from which all ordinary matter is assembled. Bind one of each and the atomic stratum opens: hydrogen, which is Layer 3 of Figure 9 and the floor on which everything above this section is built.

44.7.4 The mass, the way stations and the twin

The proton’s mass is not a Yukawa coupling; it is almost entirely the energy of the gluon field, and the route to it runs through the strong coupling. The ledger fixes $\alpha_s(M_Z) = 1/\pi e$, the renormalization-group flow carries it down to the confinement scale, and the proton mass follows at roughly 900 MeV — correct to about five percent, with no experimental input from the strong sector at any point in the chain. One pure number in that chain is imported from the lattice rather than derived, and five percent is the accuracy of a first-order registration.

The cascade vocabulary applies to the unstable states too. The $\Delta(1232)$ sits about 294 MeV above the nucleon with a width of some 117 MeV: a way station, not a ground state, and its width is the rate of the descent out of it. The free neutron sits 1.293 MeV above the proton with a lifetime

near 878 seconds; it has one open gate, and its beta decay is a GRIP-In descent through that gate — the same apparatus the electron companion used for the muon.

44.7.5 What would end it

A single observed proton decay, in any channel — the rare falsifier that grows stronger with every year of silence. A single confirmed free fractional charge. An observed neutron–antineutron oscillation. A detected long-range force coupled to baryon number.

44.8 Application: hierarchy, KK resonance, and dark matter

The Gi-4 framework is also the formal foundation of three other key results of GG-Theory: the hierarchy derivation, the KK resonance prediction, and the Kaxion dark-matter candidate.

44.8.1 Hierarchy from the radial-scale stratum

Cohort map. The scale profile e^{-K} is established in Arisaka (2026, A1-Constitution); the scale chain $E_{\text{Pl}} \rightarrow E_{\text{GUT}} \rightarrow M_{\text{lock}}$ and the reduced bulk scale M_6 in Arisaka (2026, A2-GG6) and Arisaka (2026, A4-Euclid); the electroweak vev over-constraint $v = M_{\text{lock}}(2/L)^N$ in Arisaka (2026, A9-SM).

The hierarchy $E_{\text{EW}} \ll E_{\text{Pl}}$ is derived in Arisaka (2026, A1-Constitution) via the scale profile e^{-K} with $K = 35$. In Gi-4 vocabulary, this is the GA of the radial-scale stratum Σ_r : the unique stable boundary configuration of the scale profile at $r = K$, with $M_{\text{lock}} = E_{\text{Pl}}/2e^{35} \simeq 3.85 \text{ TeV}$.

The OGN-5 integer $K = 35$ enters directly as the scale exponent.

Physics intuition. The radial coordinate is a long, gentle downhill slope — the renormalization-group axis dressed as geometry. The electroweak scale is not a tiny number that has been delicately tuned against the Planck scale; it is simply the altitude at the *bottom* of a slope whose length is an integer, $K = 35$. Nothing is being canceled to thirty-two decimal places. The smallness is a descent depth, the way the floor of a deep mine is far below the surface not by coincidence but by the number of levels dug.

The mechanism, quantitatively. The radial-scale stratum Σ_r carries its own GA at $r = K$, and the scale that survives there is

$$M_{\text{lock}} = \frac{1}{2} E_{\text{Pl}} e^{-K}, \quad K = M^2 + N^2 + L^2 = 1 + 9 + 25 = 35, \quad (191)$$

so the exponent is not a free input but the squared norm of the OGN-5 ledger $(M, N, L) = (1, 3, 5)$. Each unit of K is one e-fold of radial descent; thirty-five of them carry $E_{\text{Pl}} \simeq 1.22 \times 10^{19} \text{ GeV}$ down to $M_{\text{lock}} \simeq 3.85 \text{ TeV}$, and the same descent fixes the electroweak vacuum value $v = M_{\text{lock}}(2/L)^N \simeq 246 \text{ GeV}$ as an over-constraint.

Why this reading helps. The gauge hierarchy problem — “why is the weak scale sixteen orders of magnitude below the Planck scale without a fine-tuning of the Higgs mass to one part in 10^{32} ?” — simply dissolves. There is no tuning because there is no cancellation: a small number written as e^{-35} is the exponential of an integer, and the integer is the ledger’s own squared norm. The problem was an artifact of treating the weak scale as an additive remainder; in the cascade it is a multiplicative descent depth, and depths are counted, not balanced.

44.8.2 KK resonance from the Wilson-line moduli

Cohort map. The collider spectroscopy of the lock cluster is the prediction-branch paper [Arisaka \(2026, P1-KK\)](#); its scale is the M_{lock} derived in [Arisaka \(2026, A2-GG6\)](#), and the same scale anchors the electroweak vacuum of [Arisaka \(2026, A9-SM\)](#).

The ~ 3.85 TeV KK resonance cluster of [Arisaka \(2026, P1-KK\)](#) is the first KK mode of the bulk, determined by the lock scale M_{lock} . In Gi-4 vocabulary, this is the energy scale of the first excitation above the GA on the radial-scale stratum.

The cluster structure (multiple closely-spaced poles rather than a single pole) reflects the boundary conditions its members carry on the interval: the spin-2 and the Dirichlet-terminated vectors sit at one Bessel zero, the Neumann octet at another, each a distinct KK tower, and the visible cluster is their superposition.

Physics intuition. Once the descent has reached the radial GA, the first rung *above* it is the lock cluster — the geometry’s lowest ringing mode, the note the radial throat sounds when struck. It is not a second free scale bolted on; it is the first excited basin, pinned to the same integer $K = 35$ that fixes the floor.

The mechanism, quantitatively. The lock cluster sits at the GA scale $M_{\text{lock}} \simeq 3.85$ TeV, and it is a *cluster* rather than a single pole because several bulk fields share one radial scale — one curvature unit, $ke^{-K} = M_{\text{lock}}/3.8317 = 1.005$ TeV. On it the spin-2 graviton and the Dirichlet-terminated leptophobic $U(1)_B$ vector are degenerate at the lock exactly, both at the first zero of J_1 ; the color-octet KK gluon, a Neumann mode, has its free-field first level a factor 1.56 below, at 2.46 TeV, and sits at the lock on the collider companion’s composite-hardened reading of the same geometry. The split within the cluster is a boundary condition, not the multi-Wilson-line structure of the Σ_θ stratum. The detailed collider spectroscopy is the subject of [Arisaka \(2026, P1-KK\)](#).

Why this reading helps. A first excited mode whose scale is a derived integer, not a scanned range, is the sharpest kind of falsifier: the framework stakes a single number, $M_{\text{lock}} \simeq 3.85$ TeV, on the next decade of LHC running, rather than the vague “somewhere in the few-TeV region” that a model with an adjustable scale would offer. The Gi-4 reading is what licenses the sharpness — the cluster is the first excitation above a GA whose depth is already fixed by (191).

44.8.3 Kaxion dark matter from the $U(1)_{Q_i}$ attractor

Cohort map. The Kaxion is derived in the consequence paper [Arisaka \(2026, C1-Kaxion\)](#); its relic abundance in the cosmology paper [Arisaka \(2026, A10-Cosmos\)](#); its $U(1)_{Q_i}$ stratum in [Arisaka \(2026, A7-Qi4\)](#); and the ground-state / GRIP-Out framing in the capstone [Arisaka \(2026, A11-Origin\)](#).

The ≈ 8.1 μeV Kaxion dark-matter candidate of [Arisaka \(2026, C1-Kaxion\)](#) is the slowly-rolling GA of a one-dimensional periodic stratum. Earlier statements of this subsection named that stratum the $U(1)_{Q_i}$ Wilson-line stratum; the consequence paper now derives the Kaxion as the tensor remnant, a single line, and which field’s phase the stratum carries is open (Remark 11.1). In Gi-4 vocabulary, this is the GA of Σ_θ , whatever its occupant.

The mass scale ≈ 8.1 μeV is set by the cascade-action budget on this stratum. In regime language (§31.6) the dark sector is the purest GRIP-In specimen in Nature: the condensate relaxes toward its GGA with Hubble friction itself as the within-basin flow, and its exit-channel set is *structurally empty*.

— the gauge-protected shift symmetry deletes any inter-basin channel — so the halt is permanent. That is the runtime reading of the Kaxion’s cosmological stability (Arisaka, 2026, C1-Kaxion).

Physics intuition. The dark sector is the cleanest specimen of descent in all of Nature: a ball released on the side of a frictionless valley that has *no exits*. It rolls toward the bottom and, finding no channel out, simply stays there — for the entire age of the universe.

The mechanism, quantitatively. A GRIP-Out event at the end of inflation leaves the compact-angle field displaced from its minimum by an initial misalignment; thereafter the field executes a pure GRIP-In descent with Hubble friction as the only within-basin flow, and the cold relic it leaves behind saturates the observed abundance. The mass $\approx 8.1 \mu\text{eV}$ is set by the cascade-action budget on the compact-angle stratum, and the gauge-protected shift symmetry makes the exit-channel set *structurally empty*, so the halt at the GGA is permanent — the runtime reading of the Kaxion’s cosmological stability (Arisaka, 2026, C1-Kaxion).

Why this reading helps. Cosmological stability stops being an extra assumption and becomes a structural fact: there is simply no channel out of the basin to decay through. This is the *empty-exit* terminus of the GRIP-In taxonomy that the next section makes systematic (§50.1); placed beside the radiative descents of the synchrotron and the inspiral, the Kaxion is the limiting case in which the drain is switched off entirely, and the contrast is what makes the taxonomy complete.

Dynamical tier: the radiative descent.

44.9 The sky: recombination, the inventory, and two redshifts [DERIVED-CONDITIONAL]

The third rung is the one that made the universe transparent. It is also where the same engine is read at the largest available scale, by the cosmology companion (Arisaka, 2026, A10-Cosmos).

44.9.1 The inventory

The contents of the universe come out as coupled powers of a single number.

Table 15: The present-epoch inventory as the cosmology companion (Arisaka, 2026, A10-Cosmos) derives it, with that paper’s own values, measurements and pulls. The companion grades the whole inventory as Derived-conditional under a named integration conjecture, and it is the *ratios* rather than the absolute epoch that the geometry fixes for all time — the one quantity genuinely read from the sky being *when* the reading is taken.

Quantity	Closed form	Value	Measured	Pull
Ω_m	$1/\pi$	0.3183	0.315 ± 0.007	+0.5
Ω_Λ	$1 - 1/\pi$	0.6817	0.685 ± 0.007	−0.5
Ω_b	$1/(2\pi^2)$	0.05066	0.0493 ± 0.0006	+2.3
Ω_{dm}	$1/\pi - 1/(2\pi^2)$	0.2676	0.2657 ± 0.0070	+0.3
n_s	$1 - 2/N_*$	0.9620	0.9649 ± 0.0042	−0.7
r	$12/N_*^2$	0.0043	< 0.036 (95%)	—
w	−1 (exact)	−1	-1.028 ± 0.032	+0.9
ω_b/ω_m	$1/2\pi$	0.15915	0.15709 ± 0.0016	+1.3

Dark energy is treated in two theorems of different strength. The first is unconditional: ghost-free canonical scalars cannot have $w < -1$, with equality exactly at a stationary point of the free-energy functional — which is to say at a Ground Geometric Attractor. The second, giving $w = -1$ identically to about one part in 10^{91} , is conditional on three named hypotheses, one of them a lemma of this paper on vacuum-flux freezing at the attractor. The kill is two-sided: the framework cannot absorb $w \neq -1$.

44.9.2 Recombination as the cosmic GRIP-Out

The regime assignment runs the opposite way from the naive expectation, and it is the part of the cosmology companion that most directly consumes Part V. The scale-setting events — the exit from the unified edge, the electroweak lock, the launch and end of inflation, and recombination itself — are GRIP-Out: they reorganize the landscape. Everything that then settles onto a landscape which has stopped moving — the relaxation to the dark-energy attractor, the frozen present-day constants — is GRIP-In. The companion states the division of labor with the Standard-Model paper in one line: where the cosmic landscape itself came from is the GRIP-Out question the cosmology paper answers, and the particle paper answered the complementary GRIP-In question.

44.9.3 Two redshifts that are not the same quantity

One point deserves care, because it is the kind of thing a survey flattens and the companion is at pains not to.

There are two distinct quantities near redshift eleven hundred. The corridor switch is $z_{\text{sw}} = e^J = e^7 \approx 1096.6$: a topological marker, exact, with no continuous parameter in it and no observational handle of its own. The atomic decoupling redshift z_{dec} is computed from the cohort's own Rydberg scale and baryon fraction through the standard recombination calculation, and lands at 1091.1 raw against a measured 1089.92 ± 0.25 .

The agreement of the atomic chain is 0.11 to 0.16 percent raw — the zero-fitted-parameter statement, already inside the modeling systematic of the recombination calculation itself — tightening to 0.03 to 0.07 percent when calibrated on the known modeling offset, which the companion states for completeness and expressly does not claim as independent precision. The kill is placed on the observable one only: the reading is refuted if the measured decoupling redshift sits off the atomic-chain value by more than half a percent. Why the geometric integer lands within four tenths of a percent of the atomic one is, in the companion's own words, either a coincidence the descent merely records or a theorem the descent has not yet yielded.

44.9.4 The second horizon problem, and the bridge upward

The oldest puzzle of cosmology is that regions of the sky which have never been in causal contact have the same temperature. Inflation answers it: they were in contact, before being stretched apart.

The cosmology companion names a second and distinct puzzle, which inflation does not address and which this paper's mechanism does. Those same causally disconnected regions do not merely have the same *conditions*; they contain the same *things*. The electron here and the electron at the

edge of the observable universe have the same mass to every digit measured. Temperature is a state, and states can be equalized by exchange. Identity is a specification — a complete parts list with couplings — and no exchange of photons can negotiate a particle’s mass between two regions.

The answer offered is the ground-attractor mechanism read cosmologically: each object has exactly one specification, the GGA of its stratum, so the agreements were never negotiated because they were never in question. The companion is careful about the logical status — it calls the argument a reduction rather than an independent proof, since it does not establish that the vacuum is unique but shows that if it is, cosmic identity follows for free — and says in terms that it would be dishonest to sell it as a standalone solution.

Carried at that grade, it is nonetheless the hinge of this entire Part. **Everything above this floor depends on it.** A chemistry that worked differently in different regions would have no periodic table; a biochemistry built on a locally varying electron mass would have no universal genetic code. The reason the ladder can continue upward at all is that its first rung is the same rung everywhere.

44.9.5 What would end it

The tensor-to-scalar ratio at 0.0038, two-sided. Any secure departure of the dark-energy equation of state from -1 . A null result for the dark-matter line across its stated band. An inventory incompatible with $1/\pi$, $1 - 1/\pi$ and $1/(2\pi^2)$. The companion labels its own mixture honestly, noting which are forward predictions and which are audits of existing data that the theory must not be allowed to overclaim.

45 Floor two — Chemistry [CONDITIONAL]

From here upward the grade changes, and the change is announced rather than absorbed. Nothing on this floor or above is claimed at the standard of the floor below it. The capstone ([Arisaka, 2026, A11-Origin](#)) is explicit that it is a preview: its business is to show that the floors are one structure and to state the mechanism that binds them, not to derive chemistry, life and mind to the standard at which this program derives particle physics. Companion papers on those floors are being written.

What follows is correspondingly brief. It is here because a Part that stopped at Layer 3 would be showing the reader a ladder and climbing three steps of it.

45.1 What opened it

Two shocks, in order. Stellar nucleosynthesis and the supernovae that disperse it made the heavy elements that Layer 3’s atoms could not supply, and cometary delivery together with interstellar photochemistry put them where chemistry could happen. The landscape of the previous floor — a universe of stable hydrogen — is what made both possible.

45.2 What settles: three rungs

Small molecules. A closed quartet on the warm, wet stratum: water, carbon dioxide, ammonia, methane, with their bond angles and dipoles fixed. The claim is closure — that this set is the ground state of its stratum rather than an accident of local abundance.

Twenty amino acids, and where the twenty comes from. This rung is the one to notice. The capstone gives the count as

$$|\text{GGA}_5| = 20 = K - NL = 35 - 15,$$

together with the twenty being left-handed and the sugars right-handed. **The same five integers that fix the inverse fine-structure constant and the electron mass count the amino acids that biology uses.** Whether that survives is a question for the chemistry companion; that it is even stated is the clearest single indication that the ladder is meant to be one structure and not a series of analogies. The empirical anchors are the Murchison meteorite’s roughly eighty amino acids, glycine measured in situ at comet 67P, uracil in the Ryugu returned samples, and ribose in Murchison and NWA 801.

Polymers. Three sub-classes emerging together — oligopeptides, oligonucleotides, fatty acids — with none of the three privileged over the others. The GRIP-Out that opens the rung is wet-dry cycling with ultraviolet pulses on clay.

45.3 The fence

All three rungs are graded **CONDITIONAL**: each closes once one named lemma is proved, and the lemmas are named — small-molecule closure, the chirality lock, polymer-class closure. (The codon mapping is a lemma too, and it belongs to rung eight, not to this floor.) The capstone asks to be told if any rung above the atomic ones reads as though it were at codebook standard, on the grounds that this would be a defect. This Part takes it at its word and marks the boundary here.

46 Floor three — Life

Four rungs: an energy currency, a replicator, an archive, and a boundary around all three. One of them matters to *this* paper more than any other rung above the atomic floor, and it is taken first.

46.1 The bioenergetic anchor: where Q_B becomes physical [ANCHOR-DERIVED]

Every paper in this cohort carries the acronym plate, and the last box on that plate is the Arisaka constant Q_B with the equation $J_E/Q_B = J_e = J_B = J_{Q_i}$ beneath it. Until this section, no Part of this paper has said what that constant *does*.

It is introduced and derived not here and not in the capstone but in GG-A7-Qi4 ([Arisaka, 2026, A7-Qi4](#)), the information-side companion to this paper, where it is defined as the energy cost of one

ideal locked boundary event — the exchange rate between energy per winding on one face of the compact circle and energy per qubit on the other, fixed by the ledger integers rather than by any temperature. That is why it is temperature-independent where the Landauer bound is not, and it is the subject of §49 below. Its value follows from the Rydberg energy and the ledger,

$$Q_B = \frac{2}{5} \frac{E_{\text{Ry}}}{K} = \frac{2}{5} \cdot \frac{13.6057 \text{ eV}}{35} \approx 0.1555 \text{ eV per qubit} \approx 15.0 \text{ kJ/mol},$$

and then makes the claim that turns it from bookkeeping into physics: the hydrolysis of ATP, the reaction that powers essentially every process in every cell on Earth, releases

$$\Delta G_{\text{ATP}}^\circ = 2N_A Q_B \approx 30.0 \text{ kJ/mol},$$

against a measured $30.5 \pm 0.5 \text{ kJ/mol}$ — an agreement of about 1.6 percent, with no quantity fitted anywhere in the chain from the Rydberg energy and the integer K to the bench measurement.

If that holds, the energy quantum of biology is the same quantity that appears in the acronym plate of a paper about six-dimensional gauge geometry, and the two are the same number for a reason. It is the single most surprising line in the cohort, and it is the reason this Part insists on the plate’s last box.

The fences are the capstone’s own and are carried without softening. The grade is ANCHOR-DERIVED: the anchor is done and the derivation from geometry to codebook standard is deferred to a companion in preparation. The load-bearing claim is *the single ATP anchor alone*; the further per-process assignments sometimes quoted — one quantum per replicated bit, three per methylation event — are described by that paper as illustrative rather than measured. And on the chemiosmotic proton-motive force the honest claim is that the elementary step is *of order* Q_B , with a stated tolerance, not a coincidence to three figures. The kill criterion is a deviation of the ATP figure by more than five percent.

46.2 Four letters, because three families [CONDITIONAL]

The genetic alphabet has four letters. The capstone’s Layer-8 theorem gives the size of the alphabet as

$$|\Sigma| = N_g + 1 = 4,$$

with N_g the same family number that §44.2 derived from an instanton budget on $\mathbb{CP}^2$. **On this reading, the genetic code has four letters for the same reason that matter has three generations.**

That is either the ladder’s most striking single closure or its most obvious overreach, and this paper is not in a position to decide which. The grade is CONDITIONAL on a named bifunctionality lemma. What can be said here is what the mechanism requires: the alphabet is the GGA of its stratum, so it should be the same alphabet everywhere — which is, so far, exactly what is observed, and which the capstone lists among its falsifiers in the form of a genetic alphabet of any other size found anywhere.

46.3 The archive, and the boundary [ANCHOR-DERIVED]

The remaining two rungs are given briefly. The RNA-to-DNA takeover installs an archival duplex whose error floor, 10^{-9} per base pair, is presented as a three-tier product rather than a single mechanism, and with it the one-way flow of the central dogma. Membrane closure then draws a boundary around a metabolizing interior and the first cell follows — the point at which the cascade first has, in the capstone’s vocabulary, a *self-owning* open system rather than merely an open one.

Both are ANCHOR-DERIVED: the numbers are anchored in the literature, the derivations from geometry are deferred.

47 Floor four — The multicellular and the mind

Seven rungs, taken as two bands. The grade here is the weakest in the Part and is marked as such; the reason for including the floor at all is that the mechanism claim is a claim about *these* rungs as much as about the ones below, and a paper that built the engine and then declined to say where it is supposed to run would be ducking its own thesis.

47.1 Life becomes the engine

One structural observation belongs to this paper rather than to biology, and it is the reason this floor is not simply a list.

Below Layer 10, every GRIP-Out shock is geological or astrophysical: stars explode, planets cool, vents open. From Layer 11 upward the shocks are increasingly produced by what the previous floor built. Oxygenation of the atmosphere — the event that opens the eukaryotic rung — was caused by photosynthesizing life. The capstone puts it directly: from Layer 11 forward, *life is the cascade’s engine of advance*.

In the vocabulary of Part V that is a statement about the recursion. The alternation of GRIP-In and GRIP-Out does not require an external driver at any rung; it requires only that a settled floor be capable of disturbing the landscape above it. Below Layer 10 the disturbance comes from outside the biosphere; above it, from inside. The mechanism does not change. Only the source of the perturbation does.

47.2 Domains, eukaryotes, multicellularity [CONDITIONAL]

Three rungs. The primary domain split gives *two* attractors rather than one — two lipid chemistries, bacterial and archaeal, with the capstone claiming no third primary lipidome is available. The eukaryotic cell is then a chimera of the two, opened by the Great Oxygenation Event around 2.4 Ga and settling some hundreds of millions of years later, with an energy-per-gene budget as the mechanism. Multicellularity follows the Neoproterozoic glaciations as an eight-rung ladder from clonal adhesion to organogenesis, arriving at the Cambrian.

All three are CONDITIONAL, on a dual-gauge lemma, an energetics lemma and a ladder-uniqueness lemma respectively.

47.3 From a nerve ring to a conscious episode [ASSERTED]

Four rungs, and the weakest grade in the Part. They are worth a page here for one reason: at this floor the vocabulary of Part V stops being transferred and starts being literal.

The head-direction system of an insect or a rodent is a *one-dimensional attractor on the circle*: a ring of neurons with lateral inhibition, supporting exactly one stable bump of activity that tracks heading and drifts by less than a degree per second. That is not a metaphor borrowed from dynamics into neuroscience. It is the same object Part V defines — a landscape with a single deepest basin, and a descent that halts in it — realized in tissue, on a stratum that happens to be $SO(2)$. The framework’s claim that the four moves run at every scale is, at this rung, checkable against a connectome.

Above it the capstone places a four-dimensional allocentric map — place, grid, time and causal history — and then the conscious episode itself, given as a closure lasting about 400 milliseconds, two cycles of the theta rhythm, with gamma-band binding inside it. The empirical anchor is the P300 event-related potential at roughly 300 ± 50 milliseconds, and the kill criterion is a reliably reported unitary conscious episode completing within a single theta cycle.

The fence, in the capstone’s own words. Every one of these four rungs is ASSERTED: the framework expects the result and can say what would refute it, but the derivation is not in hand. And on the last rung the capstone draws a sharper line still, which this Part reproduces rather than paraphrases: the *identification* of the construction with phenomenal experience — what it is like, from the inside — is an explicitly labeled structural-identity hypothesis, *defended, not deduced*, and is not a consequence read off the geometric apparatus. It is that paper’s own additional commitment. It is not inherited from any companion, and it is not claimed here.

48 Floor five — Articulation [EXISTENCE-DEMONSTRATED]

The top rung of Figure 9 is the act of writing the constitution down.

The capstone’s claim is that this is not a rhetorical flourish but the cascade’s own closure: a floor whose GRIP-Out shocks are cumulative cultural ones — writing, mathematics, the printing press — and whose GGA is a description of the geometry from inside the geometry. The loop of Figure 10 closes there: the constitution at the top of the picture and the articulation at the bottom are the same object seen from two ends.

Its grade is the only one of its kind in the Part. EXISTENCE-DEMONSTRATED means exhibited rather than derived, and what exhibits it is the cohort itself. The capstone also states that there is no Layer 19, and gives its reason: a description that includes its own description has nowhere further to recurse.

This Part takes no position on whether that closure is a theorem or a figure of speech. What can be said from inside GG-A6-GRIP is narrower and worth saying: the mechanism of Part V does not exclude such a rung. A descent that halts in a basin containing a model of the landscape is not a new kind of move. It is the same four moves, on a stratum whose configurations happen to be representations.

49 The other half: Qi-4, and why this paper is not finished without it

Every application in this Part has been read geometrically. A landscape was built, a descent ran down it, and a ground state was reached. That reading is this paper’s business and it is only half of the story — and the reader should be told which half, because the other half changes what several of the rungs above mean.

49.1 Two backbones, not one

The boundary law of GG-A5-GDOS has two pillar papers standing under it, and they are siblings. This paper supplies the *geometric* backbone: four pillars that say what the shapes are and how a system moves on them. GG-A7-Qi4 (Arisaka, 2026, A7-Qi4) supplies the *quantum-information* backbone: four pillars that say what is being counted, what is being inferred, and what is being remembered while all that shape-changing goes on. Its own summary is that the two together constitute the backbone of the open-system law, geometric half and information half, and that with both in hand the trio closes.

GG-A7-Qi4 is scrupulous about the direction of the debt, and this Part repeats its wording rather than softening it. That paper consumes the cascade; it does not define it. No information-side runtime may be introduced before the free-energy landscape, the pipeline and the stratification of this paper are in hand. It also inherits this paper’s conditionality: the cascade closure of Part V holds modulo two named open problems of ours, and every construction of GG-A7-Qi4 that rests on the cascade carries that same conditional status forward.

The debt runs the other way too, and that is the point of this section. The mapping is not one pillar to one pillar, and it would be tidier if it were.

49.2 The relation between a ground attractor and the four certificates

One statement has to be made carefully, because the careless version of it is a circle and GG-A7-Qi4 says so about its own earlier drafts.

The Ground Geometric Attractor is *not* selected by the quantum-information pillars. It is fixed by the landscape alone: the ground-state result of Part IV produces it from the Morse structure of an admissible stratum, using no admissibility filter of any kind. What GG-A7-Qi4 proves is the other direction — that the ground attractor, once the geometry has produced it, *satisfies* all four quantum-information conditions. It carries the locked four-current ledger. It is a genuine extremum of the unified extremal principle rather than a local minimum of some coarse-grained effective potential. Its scale signatures are those of a fixed point, so it survives being looked at from a different distance. And the runtime can *halt* there.

That last certificate is the one worth carrying back into this paper, because it is a statement about GRIP-In. Halting admissibility excludes basins that are deep but unreachable or unsustainable — and a basin nothing can reach, or nothing can stay in, is not a ground state of anything in practice however low it sits. Part V’s halting theorem says where the descent stops; GG-A7-Qi4’s fourth

Table 16: The inward map of GG-A7-Qi4 (Arisaka, 2026, A7-Qi4): which theorem of this paper each quantum-information pillar consumes. It is deliberately not a one-to-one pairing. The conservation ledger feeds two pillars; the free-energy functional is inherited whole and then read three ways; and one information pillar, the conformal one, has no geometric parent in this paper at all — the information side carries a structure the geometry side does not.

Qi-4 pillar	What it consumes here	What it does with it
Qi-QUA — current locking	TGG and the Master Conservation Law (T1); the ledger (T2)	Makes information the fourth boundary current, equal in standing to the electric, baryonic and energetic ones, and locks all four through a single conversion constant
Qi-EP — the extremal principle	The free-energy functional $\mathcal{F}_{\text{GFF}} = \mathcal{F}_{\text{GEP}} + \mathcal{F}_{\text{CEP}}$ (T3); the ledger (T2)	Shows that three principles usually taught as unrelated — least time, maximum entropy production, and the free-energy principle of neuroscience — are three gauges of that one functional
Qi-CON — scale coherence	—	Projects conformal isometry of the bulk into the boundary signatures of scale-free behavior, with a bounded band rather than an unbounded slope
Qi-HAL — memory architecture	The GRIP cascade (T4)	Formalizes what a memory-bearing system does with the cascade: sensing, prediction, motion, storage, and the ability to schedule its own landscape reorganizations

certificate says what it takes for stopping there to mean something.

The converse — that no *other* attractor of a stratum passes all four certificates — is **not** proved, in either paper, and is not implied by anything above. GG-A7-Qi4 registers it as an open problem and records that an earlier formulation of that section asserted it and rested it on the capstone, whose own argument invokes the pillars: a circle, found and cut.

49.3 The Arisaka constant, and the plate’s last box

§46.1 used the constant Q_B without saying where it comes from. It comes from here. GG-A7-Qi4 defines

$$Q_B = \frac{2}{5} \frac{E_{\text{Ry}}}{K}, \quad K = 35,$$

as the energy cost assigned to one ideal locked boundary event when the electric, baryonic and informational sectors are each counted in one natural unit: about 0.1555 eV per qubit, about 15.0 kJ/mol, about $5.8 k_B T$ at body temperature. It introduces no new continuous parameter — both the Rydberg energy and the integer K are already fixed upstream — and the paper is candid that the *form* E_{Ry}/K rather than some other function of the ledger is justified by minimality rather than proved, and carries that as a named open problem.

The equation printed under the last box of the acronym plate, $J_E/Q_B = J_e = J_B = J_{Q_i}$, is an equality of *normalized event counts*, not an identity among raw microscopic carriers, and GG-A7-Qi4 says so in those words. Its kill criterion is correspondingly sharp and correspondingly narrow: find one elementary catalytic event whose four-current bookkeeping disagrees with that equation by more than twenty percent, and the pillar is refuted. What is being tested is the claim that a *single*

conversion factor locks the four currents — not the agreement between Q_B and any one calorimetric measurement.

One boundary is worth stating here because this Part climbed above the chemistry floor. GG-A7-Qi4 derives and validates Q_B at the chemistry scale, at ATP hydrolysis, and it explicitly does *not* derive a per-bit quantum for neural information. Any reading that prices a thought at Q_B is not in that paper and is not in this one.

49.4 What each rung gains

The enhancement is not uniform, and it is worth being concrete about which rung gains what.

The particles. The electron and the proton acquire a second, independent description of what makes them what they are. Geometrically each is the deepest basin of its stratum; informationally each carries four certificates, and the certificates are stated as properties of the particles rather than as the instrument that found them.

The parameters and the sky. The selector that Part IV supplies as a free-energy functional is the same object GG-A7-Qi4 reads three ways. Where this paper says a configuration descends, that paper can say equivalently that it infers — and the equivalence is a theorem there, not an analogy.

The life floors. This is where the enhancement is largest, and where a reader from biology should go. The free-energy principle used across neuroscience and theoretical biology is obtained in GG-A7-Qi4 as the inferential gauge of its second pillar, by pullback of the same functional this paper builds. It is not a separate law introduced for living systems. It is one of three faces of the selector that also gives least time and maximum entropy production. And the persistence class that §46.3 gestured at — a system that owns itself rather than merely being open — is defined there, with a maintenance margin that has to be positive for the thing to go on existing.

The mind floors. The ring attractor, the map and the episode of §47.3 are, in GG-A7-Qi4's vocabulary, a memory architecture: sensing feeding prediction feeding motion feeding memory, closed into a loop. And that paper adds one thing this paper's Part V does not have. A memory-bearing system carries its own landscape-reorganizing dial: it can schedule its GRIP-Out events on evidence, and repeat them, where the cosmological ones of §44.9 happen once and cannot be recalled. An engine that can decide when to disturb its own landscape is a different kind of engine, and it is the last thing the ladder builds.

49.5 The plain statement

So, plainly, and in this paper's own voice rather than as a quotation: **GG-A6-GRIP is not finished without GG-A7-Qi4.** This paper gives the shapes and the motion on them. That paper gives what is conserved, what is inferred, and what is remembered. Either alone is a description of half of an open system, and the boundary law of GG-A5-GDOS needs both halves to be a law about anything that observes.

A reader who has followed this Part and wants the other half should read GG-A7-Qi4 next, and should read it before the capstone, because the capstone’s upper floors are written in its vocabulary.

50 Interlude: the descent rate, computed

Everywhere above, the descent has been *described*: a system starts symmetric, is handed a landscape, chooses at a saddle, and settles. This section computes one, in the single place where the arithmetic is clean enough to do it end to end.

It is the most technical section of the Part, and a reader following the ladder rather than the machinery should skip to §51 without loss. It is kept here, and not in the appendices, because it is the only place in the cohort where a GRIP-In flow is run as a rate rather than as a classification — and because the result is a small unification in its own right: synchrotron radiation and gravitational-wave inspiral, two processes usually taught in different courses, come out as one descent on one landscape, differing only in the spin of the mediator.

50.1 Application: radiative descent — synchrotron and gravitational-wave inspiral as one GRIP-In flow

Cohort map. The field-theory derivation (Larmor/Peters coefficients, the storage-ring constant C_q) is Arisaka (2026, A5-GDOS, Part VI); the mediator-spin content from the four-force-uniqueness theorem of Arisaka (2026, A9-SM); the absorbing black-hole terminus from Arisaka (2026, C3-BH); and the empty-exit comparison from Arisaka (2026, C1-Kaxion).

The three applications above are static: Gi-4 fixes the attractors of the spectrum. The runtime also governs a dynamical process of equal reach — the radiative descent of an accelerated source — and here the four pillars act together on a single worked example. The underlying field theory (both static inverse-square laws from one boundary connection; the Larmor/synchrotron and Einstein/Peters radiation laws as their retarded completion; and the exact damping coefficients, including the storage-ring constant C_q , derived from one influence-functional trace) is carried out in Arisaka (2026, A5-GDOS, Part VI). The present section records only its Gi-4 reading.

50.1.1 The leading multipole is the mediator spin: $l_{\text{lead}} = s$

The composite TGG connection (Part II) carries both radiating sectors of the boundary at once — the spin-1 gauge field and the spin-2 metric perturbation. CCI (Part III) supplies the conserved currents these sectors couple to: the charge current J^μ for the gauge sector and the energy–momentum tensor $T_{\mu\nu}$ for the metric sector. Current conservation is a Ward identity, and it makes every source multipole below the mediator spin non-radiating: charge conservation removes the monopole from the gauge sector, and energy–momentum conservation removes both the monopole and the dipole from the metric sector. The leading radiating multipole is therefore fixed by the spin alone,

$$\boxed{l_{\text{lead}} = s}, \quad L_s \propto \frac{g_s^2}{c^{2s+1}} \left\langle (\partial_t^{s+1} \mathcal{I}^{(s)})^2 \right\rangle, \quad (192)$$

with $\mathcal{I}^{(s)}$ the 2^s -pole moment of the source: $(s, l, \text{time-derivatives}, c\text{-power}) = (1, 1, 2, 3)$ gives Larmor/synchrotron (electric dipole $\ddot{\mathbf{p}}$), and $(2, 2, 3, 5)$ gives the quadrupole formula (mass quadrupole $\ddot{Q}_{ij}$). The spin s that indexes the ladder is not assumed: it is read from the bulk Spin(1, 5) Lorentz content on the boundary (the four-force-uniqueness theorem of Arisaka (2026, A9-SM)), which makes the photon the rank-1 vector of the unbroken $U(1)$ and the graviton the rank-2 symmetric-traceless tensor. One geometric fact — the mediator spin — and one conservation law — CCI — fix the entire difference between the two radiation laws.

50.1.2 One descent, three terminations: the GRIP-In taxonomy

In runtime language (§31.6) the orbit itself is conservative *circulation* and descends nothing; what descends is the radiation. The emitted power drains the source into the boundary, and the slow contraction of the orbit is a GRIP-In gradient flow on the GFF binding-energy landscape GEP, carrying the source inward and chirping its frequency upward toward a GGA. Synchrotron cooling and gravitational-wave inspiral are the $s = 1$ and $s = 2$ instances of this one flow.

Placed beside the Kaxion of §44.8.3, the three exhaust the GRIP-In taxonomy by exit channel:

- *Empty-exit GGA (Kaxion)*. The gauge-protected shift symmetry deletes every inter-basin channel; with no drain the condensate halts permanently at its protected GGA (§44.8.3; Arisaka, 2026, C1-Kaxion).
- *Relaxation GGA (synchrotron)*. The radiative drain is balanced by the diffusion of discrete photon emission, so the descent settles into a normalizable fluctuation–dissipation fixed point — the equilibrium emittance of GFF-ST7, with the universal electron constant $C_q = \frac{55}{32\sqrt{3}} \hbar/mc = 3.83 \times 10^{-13} \text{m}$ that every storage ring measures (Arisaka, 2026, A5-GDOS, Sands, 1970).
- *Absorbing GGA (inspiral)*. The binding landscape $\text{GEP} \simeq -G_N m_1 m_2 / 2a$ is unbounded below all the way to contact, so there is no normalizable floor: the descent does not settle, it exits the basin at merger into the next GGA down — a black-hole horizon, the maximal-entropy OGGEF port with entropy $S = A/4$ derived in Arisaka (2026, C3-BH).

The same GRIP-In gradient flow therefore terminates in three structurally distinct ways — a permanent halt, a thermal relaxation, and an absorbing exit — selected by whether the exit set is empty, drained-and-balanced, or drained-and-bottomless. The inspiral is the case whose terminus is the deepest attractor the cohort knows.

50.1.3 The dimensionless descent rate

Absolute powers are not comparable across the two columns (the electron charge and a solar mass are incommensurable couplings). The framework-level comparator is the dimensionless *descent rate per cycle*,

$$\zeta \equiv \frac{L T_{\text{orb}}}{|E_{\text{orb}}|} = \frac{\text{energy radiated per orbit}}{\text{orbital energy}}, \quad (193)$$

which measures how fast the GRIP-In flow advances in units of one cycle. Both systems run from $\zeta \ll 1$ (adiabatic descent) to $\zeta \sim \mathcal{O}(1)$ (the terminus, where the adiabatic approximation breaks):

for the synchrotron ζ reaches order unity at the radiation-reaction limit; for the binary, at merger. Representative values, recomputed from the standard formulae, are collected in Table 17.

Table 17: The radiative GRIP-In descent, side by side. Absolute powers are incommensurable across columns; the comparator is the dimensionless descent rate ζ of (193). Representative numbers recomputed from the standard formulae; the field-theory derivation is Arisaka (2026, A5-GDOS, Part VI).

Aspect	Synchrotron ($s = 1$)	Inspiral ($s = 2$)
Mediator (bulk Spin(1, 5))	photon, rank-1 vector	graviton, rank-2 sym. traceless
Coupling / charge	electric charge q	mass-energy $\sqrt{G_N} m$
Leading multipole $l = s$	dipole	quadrupole
Power law	$q^2 a^2 \gamma^4 / 6\pi\epsilon_0 c^3$	$(32/5)(G_N/c^5) \mu^2 a^4 \Omega^6$
GRIP-In landscape GEP	orbital energy of the beam	binding energy $-G_N m_1 m_2 / 2a$
Chirp	$\omega_c = qB/\gamma m$ rises as $\gamma \rightarrow 1$	Ω rises as $a \rightarrow 0$
Terminus GGA	relaxation (equilibrium emittance)	absorbing (black hole, Arisaka, 2026, C3-BH)
Diffusion / noise	photon shot noise; $C_q = 3.83 \times 10^{-13} \text{ m}$	graviton shot noise $\sim 10^{-79}$, negligible
Descent rate ζ	$\approx 3 \times 10^{-2}$ (LEP e^-)	$\approx 10^{-13}/\text{orbit}$ (Hulse-Taylor) $\rightarrow \mathcal{O}(1)$ (merger)
Empirical anchor	every storage ring (Sands, 1970)	Hulse-Taylor (Hulse & Taylor, 1975); LIGO/Virgo

50.1.4 What Gi-4 adds, and the open step

The radiation laws are unmodified; what the Gi-4 reading supplies is their *unity* — one composite TGG connection, one CCI conservation law fixing the multipole order through $l_{\text{lead}} = s$, and one GRIP-In flow with a three-way terminus taxonomy — together with the dissolution of the classical self-force pathology, which in this reading is the gradient sector of an open-system current rather than a local third-derivative force (Arisaka, 2026, A5-GDOS). The genuinely open step is sharply localized: the Larmor coefficient $1/6\pi\epsilon_0 c^3$ and the Peters coefficient $32G_N/5c^5$ are each derived exactly in Arisaka (2026, A5-GDOS, Part VI), but as separate evaluations; exhibiting both as the $s = 1$ and $s = 2$ values of *one* spin-indexed GFF descent functional — that is, deriving the normalization N_s and the $(s+1)$ -th-derivative structure of (192) from the bulk action — remains open. We record it as the natural successor problem to the radiative reading.

51 What the ladder shows

Eighteen rungs, six companions, five grades. It is worth saying plainly what they share, because the shared structure is this paper’s actual claim and it is easy to lose in the detail.

One mechanism, eighteen strata. Nothing above required a new mechanism. The Standard-Model parameters, the family count, the electron, the proton, the cosmic inventory, the amino acids, the genetic alphabet, the ring attractor — all of them are the same four moves, run on different boundary strata. What differs from rung to rung is the stratum, not the engine. The capstone states the recursion in the form this paper would state it: the step from one floor to the next is not a fifth move, it is the same four moves on a landscape the floor below has perturbed.

Stability is always ground-state-ness. Every absolutely stable object in the survey is the deepest basin of its stratum, and every unstable one has an exit gate whose action sets its width. That is one statement covering the electron’s infinite lifetime, the muon’s microsecond, the Δ resonance’s hundred and seventeen MeV, the dark-energy attractor’s $w = -1$, and — if the upper floors hold — the persistence of a genetic alphabet across four billion years. It is also, read cosmologically, the answer to why two electrons on opposite sides of the sky agree.

The two regimes divide the labor. The rare landscape-reorganizing events are GRIP-Out — the unified exit, the electroweak lock, recombination, the supernovae, the vents, the oxygenation, each of the seventeen shocks in Figure 9 — and everything that settles afterward is GRIP-In. Below Layer 10 the shocks come from outside the biosphere; above it, from inside. The engine does not change.

The grade falls as the ladder rises, and that is the honest shape of the claim. Four of the eighteen rungs are closed strictly; seven once an empirical anchor may stand in for a derivation; nothing above the atomic rungs is at codebook standard. This Part has printed the grade on every rung rather than averaging it, because the average would misrepresent both ends: it would understate how strong the bottom is and overstate how settled the top is.

51.1 Where to attack it

A framework of this shape is tested at its most exposed points, not its most impressive ones, and the capstone pre-registers ten of them with numbers, instruments and decade windows attached.

Six of the ten are physics and are the ones to bet on: they are the rungs where the grade is highest, the numbers sharpest and the instruments already funded. The remaining four are the price of the claim that the ladder is one structure — because a framework that extended to life and mind without exposing itself there would be extending to them for free.

That is the accounting. The engine of Part V, run on eighteen strata, produces a periodic table, a genetic code and a conscious episode from the same four moves that produce the electron; and it can be ended by about a dozen measurements, most of them due this decade.

52 A closing word, for any reader

This Part has been long, so here is what it said.

Table 18: The Top 10 falsifiers, as GG-A11-Origin (Arisaka, 2026, A11-Origin) registers them. Six belong to the physics floors, three to life, one to mind. Each carries a number and an instrument, and decisive failure on any one of them ends the framework — whatever grade the surrounding rung happens to carry. This is the right object on which to close a survey of applications: not what the framework explains, but where it can be killed.

	What is predicted	Where it fails	Instrument, window
1	Lock cluster at $M_{\text{lock}} = 3.85 \pm 0.10$ TeV	Exclusion of the spin-one pair	HL-LHC, 2030–35
2	Kaxion at $m_a = 8.1$ μeV	Null across the band at derived coupling	Haloscopes, 2027–30
3	$w = -1$ exactly	Any secure $w \neq -1$	DESI, Euclid, Roman, 2026–30
4	Proton lifetime beyond 10^{35} yr	One decay event, any channel	Hyper-K, DUNE, 2027–40
5	$r = 12/N_*^2 = 0.0043$ at leading order (0.0038 exact), two-sided	Detection above 0.01, or a null at the floor	CMB-S4, LiteBIRD, 2030–35
6	$\delta_{\text{PMNS}} \approx 196^\circ$	A converged value elsewhere	DUNE, Hyper-K, 2030–35
7	$\Delta G_{\text{ATP}}^\circ = 2N_A Q_B = 30.0$ kJ/mol	Deviation beyond five percent	Single-molecule biophysics, 2026–30
8	A genetic alphabet of exactly four letters	Any other alphabet size, anywhere	Comparative genomics, continuing
9	A two-knee $1/f$ spectrum in eukaryotic metabolism and in cortex	A single knee where two are predicted	Single-cell imaging; intracranial EEG
10	A conscious episode of about 400 ms	A unitary episode completing in one theta cycle	Psychophysics and EEG, 2026–28

There is one machine in this paper. It has four moves. A system holds every possibility open; a landscape tells it which way is downhill; at some point it reaches a fork and takes one branch, and a symmetry that was real a moment before is gone; and it comes to rest in the lowest place it can reach. Then — and this is the part that makes it a history rather than an event — the resting place becomes the disturbance that opens the next landscape, and the machine runs again, one story up.

Run that machine at the highest energies and it gives you the particles: three families of matter and not four, an electron that cannot decay because there is nowhere below it to go, a proton that cannot decay because you cannot unwind a knot. Run it as the universe cools and it gives you atoms, and a sky whose contents are fixed by a handful of whole numbers. Run it in the warm water of a young planet and, if the argument holds, it gives you the twenty amino acids, a four-letter alphabet, and a molecule that carries thirty kilojoules per mole because a constant fixed by the geometry of six dimensions says that is what one locked event costs. Run it long enough and it gives you an animal with a ring of neurons that knows which way it is facing — which is, in the exact technical sense of this paper, a landscape with one basin and a descent that halts in it. Run it once more and the animal writes the machine down.

That is the claim. This paper does not prove all of it, and the earlier sections have been careful to mark where the proof stops: four rungs of eighteen are closed strictly, and the ones above the

atoms are honest forecasts with their derivations still being written. What this paper contributes is the machine itself, stated sharply enough that someone can show it does not run.

It is worth saying what would follow if it does run, because the consequence is not the one people usually expect from a physical theory.

The oldest form of the question is why there is something rather than nothing, and the second oldest is why we are here to ask. Both invite two bad answers. One is that it was arranged — that somewhere there is an intention the universe was built to satisfy. The other is that it was luck — that we happen to occupy the corner of a vast lottery where the numbers came up habitable, and the question dissolves into a shrug.

The picture in this Part is neither. Nothing here is arranged; there is no designer in any equation, and no step waits for permission. But nothing here is lucky either, because at each stage the machine has no choice about where to stop. It stops at the bottom. The bottom is not a preference — it is a property of the shape, and the shape was fixed before the first second. Cosmos, life and mind are then not three separate miracles and not three lottery wins. They are three depths of one descent, and the reason they follow one another in that order is that each is the floor the next one needs.

If that is right, then the strangest thing about the universe is not that it produced observers. It is that it could not have avoided producing them, given how it started — and that the observers it produced turn out to be made of the same four moves as everything they are looking at. The theory does not stand outside the world it describes. It is the last rung of the ladder it is describing, written by an arrangement of matter that got there the same way the electron did.

None of which makes any of it true. That is what the ten measurements in Table 18 are for, and most of them will report this decade. A framework that reached this far and could not be killed would not be worth the paper it is written on. This one can be killed by a single proton decay, by one dark-matter search coming up empty across a narrow band of frequencies, or by a chemist measuring the energy of a single molecule and getting the wrong number.

That is the right way for an argument of this size to end: not with a claim about what must be so, but with a short list of things the world is about to be asked.

Part VIII

Discussion

What the four pillars establish, how the framework stands against the alternatives, what would falsify it, and what remains open

This Part steps back from the construction and asks the questions a reader is entitled to ask of it: what is actually new here, how does it stand against the other serious attempts at the same problem, what would prove it wrong, and what is still missing.

The order is deliberate. It begins with a compact statement of what the preceding Parts established, because significance cannot be judged before content is known. It then places those results against the principal alternatives — string compactification, the gauge/gravity correspondence, warped extra dimensions — and the comparison is meant to be evaluative rather than diplomatic: each framework is credited with what it does well and asked directly what it leaves postulated. Innovation is separated out and stated plainly, because a claim of novelty that is not stated cannot be checked. Then the falsifiers: every prediction the four pillars make that could kill the framework, with the observable that would do it and the facility that will decide. Then the open problems, given at length and without softening, and finally the directions the program takes next.

A reader who wants the short version can take the first section and the falsifier table and stop; between them they contain the claim and the exposure. A reader who intends to argue with the paper should go instead to the open problems, which is where the argument is most likely to be winnable.

Two warnings about how to read what follows. The comparisons are between *structures*, not between communities or research programs, and no claim is made about which body of work is the more valuable. And the falsifier list is deliberately unhedged: where a prediction is at risk it is stated as at risk, and where a measurement already constrains it the constraint is given rather than described.

53 What the four pillars establish

Before any claim is made about what these results are worth, here is what they are.

This paper took as given a six-dimensional arena, the group acting on it, an integer ledger, a boundary law, and four constitutional axioms — all of them established elsewhere and all of them listed in Part I — and from those inputs built four objects and proved five theorems about them. Stated without machinery, the four objects are a *single geometric structure* carrying all four forces, a *consistency condition* on its bookkeeping, a *landscape* on which the boundary world can sit, and a *descent* that moves it across that landscape and stops.

What is new here, and was not available before this paper, is the following.

- **The four forces assemble into one object, and in only one way.** Uniqueness is the load-bearing half: a merely possible unification explains nothing, because the question "why this one?" has no answer. Here the constitutional requirements leave exactly one arrangement standing.
- **That object obeys a single conservation law.** Energy and momentum of matter, gauge fields, residual higher-dimensional modes, and the anomaly-correction term add to something that does not leak out of the observable world. In a theory where the visible world is a projection of a larger one this is not automatic — projections are exactly the operation that loses things — and it is proved rather than assumed.
- **The integer ledger is forced, not chosen.** Imposing closure of the anomaly polynomial admits exactly one solution. Nothing was fitted; no datum entered. The consequence is that the ledger cannot be adjusted if a measurement disagrees, which is what makes the next item a real prediction.
- **A measurable astrophysical consequence follows from the ledger alone:** a correction to the neutron-star mass–radius relation, of a size fixed entirely by the integers, with no freedom anywhere in it.
- **The free-energy landscape has genuine minima, and this is forced by the anomaly condition.** That link is the paper's most unexpected internal result: a requirement about the consistency of a quantum field theory implies a statement about whether matter can settle down.
- **The descent terminates, and where it terminates is countable.** Running the cascade yields three families of matter — not by choice and not by fit, but because a fourth would need somewhere to go and there is nowhere.
- **The four pillars and the constitutional axioms are equivalent.** Each implies the other with nothing left over, which means the construction has closed rather than merely accumulated.

Two boundaries around that list should be stated in the same breath, because they are what make the rest of this Part meaningful. Every result above is *conditional on the declared inputs* — this paper does not derive six dimensions, it shows what six dimensions force. And several steps rest on named hypotheses that are not yet theorems; they are cataloged in §57 under their own names, and nothing in the sections between here and there quietly promotes one of them.

53.1 Problems solved by GG-A6-GRIP

We first record the principal problems that GG-A6-GRIP *settles*, each with its global identifier, the theorem that carries it, and a live empirical falsifier in the sense of POS-7 (Measurement and Falsifiability). These entries complement the open and future items that follow.

PS-GRIP-1. Axiomatic foundation for the **Gi-4** vocabulary

Resolved (this paper, Theorems T1–T5). The GFF/GRIP/cascade vocabulary used informally across the cohort is given a rigorous five-theorem basis, so that downstream papers can cite a fixed foundation. *Live falsifier*: a failure of any pillar’s structural falsifier below (F-GRIP-1...7) would force a re-examination of the corresponding pillar axioms.

PS-GRIP-2. The family number $N_g \leq 3$

Partially resolved (Theorem T4; the $N_g \leq 3$ upper bound via the capacity theorem T-N3-6 of Arisaka (2026, P2-N3), conditional on its coupling clause and M6 (= R-INT); the lower bound $N_g \geq 3$ from CP unconditional). The GRIP cascade closes against the budget $S_{\text{tot}} \leq 3\pi$, forcing $N_g \leq 3$ on the family-projector stratum. *Live falsifier*: a fourth chiral generation (cf. the LEP invisible-width bound $N_\nu = 2.984 \pm 0.008$, ALEPH *et al.*, 2006) would force a re-examination of BI-6 closure.

PS-GRIP-3. A geometric arrow of time on the boundary

Resolved (Theorem T3). The GFF free-energy functional obeys the second law $dF_{\text{GFF}}/dt \leq 0$ along every GDOS CPTP trajectory with detailed balance. Custody note: Theorem T3 is the runtime face of the one constitutional arrow, consumed rather than re-derived — AX-3 owns the arrow (the AX-3 synthesis of Arisaka, 2026, A1-Constitution); Part IX of the CPT member paper (Arisaka, 2026, P5-CPT) is the canonical law-level derivation over a bulk that is itself exactly T-symmetric; the GDOS companion (Arisaka, 2026, A5-GDOS) is the dynamical implementation. *Live falsifier*: a measured violation of equipartition, the fluctuation–dissipation relation, or Onsager reciprocity near boundary equilibrium (F-GRIP-1, F-GRIP-2, F-GRIP-3).

PS-GRIP-4. The generation gate-ordering ladder

Resolved (Theorem T4, GRIP-ST7). The gate tolls $S_{[2k] \rightarrow [2(k-1)]} = \pi k$ yield the generation-ordered geometric gate weight $w = e^{-\pi k}$, an equilibrium weight and not a rate (lifetimes are phase space, fixed in Arisaka (2026, A9-SM)). *Live falsifier*: a family ladder whose generations are not ordered as the weights order them, or whose log-mass gaps stand in a ratio no torus action on its stratum can produce (F-GRIP-4, executed in Arisaka (2026, A9-SM): the ordering holds in every sector; the ratio of this paper’s own action, 1.31, holds for the down quarks and not for the up quarks); the residual-width form of the test presupposes a rate reading the paper does not derive.

PS-GRIP-5. Runtime gauge invariance under the cascade

Partially resolved (Theorem T4.2; conditional on the compensator-monodromy conjecture M11). The boundary Ward identities are preserved along the cascade iff the PBI-4 compensator extends smoothly across each instanton stratum. *Live falsifier*: a runtime Ward-identity violation in the boundary currents (F-GRIP-5).

PS-GRIP-6. The OGN-5 integer ledger as the parameter-free backbone

Partially resolved (Theorem T2; Diophantine uniqueness imported from Arisaka (2026, A3-BI6, Theorem 10.2), conditional on the MSMF norm choice M9). MSMF closure of X_4 forces $(M, N, L; J, K) = (1, 3, 5; 7, 35)$, from which every quantitative prediction follows in closed form with zero free parameters. *Live falsifier*: a measured deviation from any ledger-derived prediction (e.g. $\alpha^{-1}(0) \neq 137$).

PS-GRIP-7. Geometric–gauge unification with a Master Conservation Law

Resolved (Theorem T1, CCI-ST7). TGG supplies a unique composite connection $\mathcal{A}_{\text{tot}} = \omega \oplus A$ and the boundary Master Conservation Law $\nabla_\mu \mathcal{T}_{\text{TGG}}^{\mu\nu} = 0$. *Live falsifier*: a precision compact-star mass–radius measurement contradicting the modified Tolman–Oppenheimer–Volkoff relation of CCI-ST7.

PS-GRIP-8. Constitutional closure of Gi-4

Partially resolved (Theorem T5). Gi-4 is shown equivalent (the converse direction a structured proof sketch conditional on the cascade conditioning — today M11’s general statement, M10 resolved externally) to the GG-Constitution of Arisaka (2026, A1-Constitution) (four axioms, seven postulates), so the framework adds no new physical axiom. *Live falsifier*: this is a structural (internal-consistency) result; its kill criterion is any falsifier of the upstream constitutional clauses.

54 Significance, uniqueness, and the alternatives

The results of §53 are worth something only in comparison. Every serious framework for fundamental physics can point at things it explains; what separates them is what each is willing to leave postulated, and that is a question with checkable answers rather than a matter of taste.

This section makes the comparison in one place. It first sets out what is structurally distinctive about the present construction, then treats the three principal alternatives — string compactification, the gauge/gravity correspondence, and warped extra dimensions — one at a time, and closes with a full side-by-side audit. The subsections are credit where credit is due: each of these programs solves problems this one does not, and the comparison is drawn between *structures*, never between communities.

54.1 Significance and uniqueness of the GDOS / Gi-4 framework

This Part is a discussion of the significance and uniqueness of the GDOS / Gi-4 framework relative to other leading approaches to foundational physics. We address three principal questions. First: why is the open-system / boundary-projection content of GDOS a “missing concept,” and how does its absence in orthodox quantum mechanics (Arisaka, 2026, A8-QM), quantum field theory, and the Standard Model account for the structural limitations of those frameworks? Second: how is GDOS fundamentally different from string theory and its phenomenological approaches (notably AdS/CFT

and the Randall–Sundrum model)? Third: how does GDOS compare against other contemporary unification programs in a side-by-side audit? We close with a comprehensive comparison table.

54.2 The five pillars of significance

Before proceeding to the comparisons, we summarize the five principal respects in which the GDOS/Gi-4 framework is structurally distinctive. These are neither stylistic preferences nor pedagogical choices: they are technical features that distinguish the framework from the alternatives.

First: the boundary projector is foundational. In GDOS, the irreversible CPTP boundary projector $\Pi : M_6 \rightarrow M_4$ is the foundational dynamical object. Unitary evolution is a special case (the bulk dynamics in isolation, before projection); “measurement” is another special case (the cascade endpoint, the GA catalog). No closed-system framework places the projector at the foundational level; without this elevation, the measurement problem of orthodox quantum mechanics cannot be structurally resolved.

Second: integers and hierarchies become derived. Because GDOS has access to cascade closure, finite instanton-action budgets, and Morse-attractor selection, it produces *sufficient* conditions on integer multiplicities and hierarchies, not merely *necessary* ones. The family number $N_g = 3$, the hierarchy $M_{\text{lock}}/E_{\text{Pl}} = 1/(2e^{35})$, the 28 SM parameters, the dark-matter scale — all become derived rather than postulated.

Third: the algebraic backbone is uniquely $d = 6$. The magic-dimension ladder $\mathbb{R} \rightarrow \mathbb{C} \rightarrow \mathbb{H} \rightarrow \mathbb{O}$ identifies $d = 6$ with the unique quaternionic Lorentz double cover $\text{Spin}(1, 5) \cong \text{SL}(2, \mathbb{H})$. This forces the bulk dimension; lower dimensions (3, 4) under-constrain, and higher dimensions (10) overshoot into octonionic non-associativity. No prior unification framework identifies a uniquely admissible bulk dimension by purely algebraic means.

Fourth: the framework is constitutionally predictive. The combination of BI-6 closure, the OGN-5 ledger, and the GRIP cascade produces sharp experimental falsifiers — ~ 3.85 TeV KK resonance cluster, ≈ 8.1 μeV Kaxion dark matter, exactly three generations (cf. the LEP lineshape bound $N_\nu = 2.984 \pm 0.008$, [ALEPH *et al.*, 2006](#), and the CMB N_{eff} concordance, [Planck Collaboration, 2020](#)), $m_\tau/m_\mu = 16.817$ to four-significant-figure precision. The framework can be tested at the LHC, at low-mass DM detectors, and at flavor facilities, with clear falsification criteria that do not require new theory to formulate.

Fifth: the framework is open at the right level. GDOS does not propose new fundamental fields, new gauge groups beyond what is forced by BI-6 closure, new dimensions beyond the 6D bulk, or new free parameters. The conceptual economy is substantial: one principle (OGGEP), one manifold (M_6), one connection ($\mathcal{A}_{\text{tot}}$), one cascade (GRIP). All structural integers and hierarchies emerge as theorems of this minimal input.

54.3 What is genuinely new

The genuinely new content of GDOS relative to its predecessors and contemporaries can be summarized as follows.

- **Open-system foundational level.** The CPTP boundary projection is at the foundational level rather than derived. This is new relative to QM, QFT, GR, and string theory.
- **Magic-dimension forced bulk.** The bulk dimension is fixed at $d = 6$ by the magic-dimension ladder. This is new relative to string theory ($d = 10$ or 11 chosen by other constraints).
- **Derived integer multiplicities.** $N = 3$ is forced rather than chosen. This is new relative to all closed-system attempts (heterotic CY, magnetized tori, $\mathbb{Z}_3$ orbifolds, modular flavor symmetry, etc.).
- **Derived hierarchies.** The hierarchy $E_{\text{GUT}}/E_{\text{Pl}} \sim e^{-35}$ is forced by the OGN-5 ledger $K = 35$. This is new relative to the Randall–Sundrum model (which has free warp slope) and to anthropic/landscape approaches (which leave hierarchies stochastic).
- **Quantitative codebook.** All 28 SM parameters are derivable from the Gi-4 cascade structure, with four-significant-figure agreement with experiment. This is new relative to all approaches that fit SM parameters phenomenologically.
- **Resolution of the measurement problem.** The boundary-projector dissolution of the unitary/projective dichotomy is new relative to all interpretations of orthodox QM (Copenhagen, Many-Worlds, Bohmian, GRW, decoherence-only); the dissolving projection is itself always on and needs no observer (GG-A5-GDOS §23).

These six new ingredients together constitute the structural distinctiveness of GDOS.

54.4 The measurement-problem carry-over: from QM to QFT to the Standard Model

The measurement problem — how a closed-system unitary dynamics projects onto a classical record wherever a Geometric Attractor forms (the projection itself being always on, independent of any observer) — is the canonical foundational discontinuity of twentieth-century physics. GG-A5-GDOS (Arisaka, 2026, A5-GDOS, Part III, Theorems T7.1–T7.7, and the always-on-projection synthesis of §23) carries the constitutional resolution: the carry-over from QM to QFT to the Standard Model is the OGGEF bulk-to-boundary projection $\Pi : M_6 \rightarrow M_4$ acting as a CPTP channel on the bulk state space, with the four Gi-4 pillars supplying the geometric backbone and the four Qi-4 pillars supplying the information backbone. Within the present Gi-4 paper, the role of the measurement-problem carry-over is captured by Theorem T2.2 (boundary covariant conservation as an iff statement on the Ward currents) and by the GRIP cascade structure of Theorem T4.2 (runtime gauge invariance along the cascade is the operational measurement protocol). Readers seeking the full GDOS measurement-problem narrative, the historical chain from Heisenberg through Zurek (Zurek, 1981), and the seven-theorem chain T7.1–T7.7, are referred to GG-A5-GDOS Part III.

54.5 Comparison with string theory, AdS/CFT, and Randall–Sundrum

Beyond orthodox QM, QFT, and the Standard Model, the principal contemporary foundational frameworks include string theory and its phenomenological realizations. The two most heavily developed are the AdS/CFT correspondence and the Randall–Sundrum (RS) model. We compare GDOS against each in turn.

54.6 String theory in a nutshell

String theory replaces point particles with one-dimensional strings, defined on a D -dimensional target spacetime with the worldsheet action governing string dynamics. Its principal features are:

- **Bulk dimension.** $D = 10$ for superstring theory, $D = 11$ for M-theory. The dimension is selected by worldsheet conformal invariance and supersymmetry.
- **Compactification.** Six (or seven) of the dimensions are compactified on a small manifold, typically a Calabi–Yau threefold (Candelas *et al.*, 1985). The choice of CY manifold is not unique; estimates are that $\sim 10^{500}$ admissible compactifications exist (the “string landscape”).
- **Family content.** For heterotic compactifications on a Calabi–Yau threefold X , the net chirality is $|\chi(X)/2|$, so X must be chosen to have $\chi = \pm 6$ for three generations; the intersecting-brane, orbifold, and modular-symmetry alternatives face the analogous choice (Blumenhagen *et al.*, 2007, Ibáñez *et al.*, 1987, Ibáñez & Uranga, 2012, Kobayashi *et al.*, 2020).
- **Foundational framework.** Closed-system unitary quantum field theory on the worldsheet, lifted to the target spacetime. Standard QFT closed-system structure throughout.
- **Predictions.** Specific predictions depend on compactification choice; the landscape produces a vast space of admissible vacua, making sharp predictions difficult.

String theory is the most heavily developed contemporary unification program. It has produced a deep mathematical superstructure (supersymmetry, dualities, holography, modular forms, etc.) and has re-shaped much of theoretical physics. However, its predictive power is limited by the landscape problem, and its measurement treatment is inherited unchanged from orthodox QFT.

54.7 The AdS/CFT correspondence

The AdS/CFT correspondence (Maldacena, 1997) is the most precise example of holographic duality known. In its most studied form, it asserts:

$$\text{Type IIB string theory on } \text{AdS}_5 \times S^5 \cong \mathcal{N} = 4 \text{ SU}(N) \text{ Super-Yang-Mills on the 4D boundary.} \quad (194)$$

The two theories are dual descriptions of the same physical system. The five-dimensional Anti-de Sitter bulk and its four-dimensional conformal boundary are related by a precise dictionary mapping bulk field content to boundary operators and vice versa.

AdS/CFT is a beautiful illustration of holography in a tightly constrained setting. But it differs from GDOS in several structural respects:

- **Both sides are closed-system.** The bulk is unitary quantum gravity (in the AdS sense); the boundary is unitary CFT. The “holography” is a duality, not a CPTP boundary projection. No irreversible measurement structure enters.
- **Specific bulk geometry.** The bulk is $\text{AdS}_5 \times S^5$, maximally symmetric. This is far from the 6D bulk of GDOS.
- **No Standard-Model content.** AdS/CFT in its prototypical form does not contain the SM gauge group, fermion content, or hierarchies. Phenomenological adaptations exist (“AdS/QCD”, etc.) but are not the core correspondence.
- **No family-number derivation.** The correspondence does not have a notion of fermion families that can be derived; it is a duality between specific theories.
- **Closed-system measurement.** Measurement on the boundary CFT is treated by orthodox QFT methods, with no resolution of the measurement problem.

AdS/CFT is therefore a structural illustration of holography in specific theories, but it is not a foundational framework for all of physics. GDOS elevates the bulk-to-boundary projection to foundational status across all physics, in the 6D bulk relevant to our universe, with the CPTP open-system structure throughout.

54.8 The Randall–Sundrum model

The Randall–Sundrum (RS) model places a five-dimensional warped bulk between two branes, with gauge fields and gravity propagating in the bulk and matter (typically) localized on one brane. Its principal features are:

- **Bulk dimension.** $D = 5$, with one warped extra dimension. The bulk metric is approximately AdS_5 between two branes.
- **Hierarchy mechanism.** The warp factor between the UV and IR branes provides a natural hierarchy between E_{Pl} and E_{EW} . This is structurally similar to GG-6’s warp factor $E_{\text{Pl}}/M_{\text{lock}} = 2e^{35}$, but with the warp slope and brane separation as adjustable parameters.
- **Family content.** Fermion families are added by hand (typically as bulk fields with localized wave-functions). Family number is not derived.
- **KK spectrum.** Heavy KK gravitons and KK gauge bosons appear at the warp scale. These are similar to the KK spectrum of GG-6 (Arisaka, 2026, P1-KK) but without the cluster structure that the multi-Wilson-line content of GG-6 produces.
- **Foundational framework.** Closed-system QFT on a fixed warped 5D background. No CPTP boundary projection at the foundational level.

The RS model is a phenomenological warped-extra-dimension construction without a foundational completion. It shares with GG-6 the warped geometry and the KK spectrum at $\sim \text{TeV}$, but differs in the following ways:

- RS has $D = 5$; GG-6 has $D = 6$ forced by the magic-dimension ladder.
- RS treats the warp slope as a free parameter; GG-6 derives it from the OGN-5 ledger ($K = 35$, giving $M_{\text{lock}} = E_{\text{Pl}}/2e^{35}$).
- RS treats family content phenomenologically; GG-6 derives $N_g = 3$.
- RS has no measurement-foundational content; GDOS elevates the boundary projector to foundational status.
- RS has no second compact direction; GG-6 has the S_θ^1 that supports the Wilson–Hosotani Higgs and the Kaxion field.

The RS model is therefore a partial precursor to GG-6, sharing the warped-bulk intuition but lacking both the magic-dimension forcing and the open-system foundational layer.

54.9 Structural differences in summary

The structural differences between GDOS/Gi-4 and the contemporary frameworks may be summarized as four moves:

1. **Bulk dimension forced by algebra.** GDOS uses the magic-dimension ladder to force $d = 6$. String theory chooses $d = 10$ for worldsheet conformal invariance; RS chooses $d = 5$ for AdS/braneworld convenience; AdS/CFT studies $d = 4 + 1$ specifically.
2. **Open-system foundational layer.** GDOS elevates the CPTP boundary projector to foundational status. String theory and RS both use closed-system QFT throughout; AdS/CFT uses two closed-system theories related by duality.
3. **Derived integers and hierarchies.** GDOS derives $N = 3$ and the warp hierarchy from the BI-6 closure plus MSMF. String theory and RS treat both as inputs.
4. **Resolution of the measurement problem.** GDOS dissolves the orthodox unitary/projective dichotomy. String theory, AdS/CFT, and RS all inherit the measurement problem unresolved.

These four moves are not stylistic. They are the technical features that make GDOS structurally distinctive — and they are the features that allow GDOS to derive what other frameworks postulate.

54.10 Comprehensive comparison table

We close Part VIII with a comprehensive comparison table audit of GDOS/Gi-4 against the principal contemporary unification frameworks. The table is intended to be evaluative rather than diplomatic: each entry records the structural answer of each framework to a specific question, with no attempt to soften the contrasts.

Table 19: Comprehensive comparison of contemporary unification frameworks — (I) foundations and structure. The comparison continues on the next page with (II) parameters, predictions, and falsifiability. The table is evaluative rather than diplomatic: each cell records what the framework in that column actually answers to the question in that row, with no attempt to soften a contrast, and a blank or an N/A is itself a finding. The present framework is the rightmost column, placed last so that each row can be read left to right as a century of answers before arriving at what is claimed here.

Aspect	Orthodox QM/QFT/SM	Heterotic String / CY	AdS/CFT	Randall–Sundrum	GDOS / Gi-4 (this paper)
<i>Foundational framework</i>	Closed-system	Closed-system	Closed-system (duality)	Closed-system	Open-system
<i>Bulk dimension</i>	4	10	4+1	5	6 (forced by magic ladder)
<i>Lorentz Lie algebra</i>	$\mathfrak{so}(1, 3) \cong \mathfrak{sl}(2, \mathbb{C})$	$\mathfrak{so}(1, 9)$	$\mathfrak{so}(2, 4)$	$\mathfrak{so}(1, 4)$	$\mathfrak{spin}(1, 5) \cong \mathfrak{sl}(2, \mathbb{H})$
<i>Magic-dim. forcing?</i>	$\mathbb{R}$ (no)	$\mathbb{O}$ (non-assoc.)	N/A	N/A	$\mathbb{H}$ (forced)
<i>Holographic projection</i>	N/A	As specific dualities	As bulk-boundary duality	Brane-bulk via metric	Foundational CPTP
<i>Measurement framework</i>	Postulated (Born (Born, 1926) rule)	Inherited from QFT	Inherited from QFT	Inherited from QFT	Derived (GA cascade)
<i>Unitarity of dynamics</i>	Yes (always)	Yes	Yes (both sides)	Yes	Bulk only; boundary CPTP
<i>Family number $N = 3$</i>	Postulated	$ \chi/2 = 3$ chosen	N/A	Postulated	Forced (T2.1+T3.1+T4.1)
<i>Dual-layer architecture</i>	N/A (one layer only)	N/A (one layer only)	Two sides via duality	N/A (one layer only)	Algebraic (A1–A5) + Physical (A6–A7) per pillar; linked by compatibility theorems
<i>Astrophysical falsifier</i>	None (SM has no compact-star prediction)	Compactification-dependent	N/A	Maybe (KK gravitons)	Modified TOV (CCI-ST7); 0.1–1% NS shift
<i>Thermodynamic falsifier</i>	None (free-energy not foundational)	None	None	None	Equipartition + fluct.–diss. (GFF-ST7) for any boundary obs.
<i>Generation weight ladder</i>	Sargent rule (phenomenological)	N/A	N/A	N/A	$w_{[2k] \rightarrow [2(k-1)]} = e^{-\pi k}$ gate weight (GRIP-ST7); lifetimes are Sargent’s, via A9
<i>Gauge group</i>	Postulated $SU(3) \times SU(2) \times U(1)$	Compactification-dependent	$\mathcal{N}=4$ Yang–Mills	Postulated	Spin(1,5) $\times$ SU(5) $\times$ U(1)_B $\times$ U(1)_{Qi}, from BI-6

Table 19 (continued) — (II) parameters, predictions, and falsifiability.

Aspect	Orthodox QM/QFT/SM	Heterotic String / CY	AdS/CFT	Randall–Sundrum	GDOS / Gi-4 (this paper)
<i>28 SM parameters</i>	Phenomenological input	Largely free moduli	N/A	Mostly free	Derived from OGN-5 ledger (4 sig. figs.)
<i>Hierarchy $E_{\text{EW}}/E_{\text{Pl}}$</i>	Fine-tuning puzzle	Free moduli	N/A	Free warp slope	$1/(2e^{35})$ from $K = 35$
<i>Dark matter</i>	Add-on candidates	Compactification-dependent	N/A	Brane-localized candidates	$\approx 8.1 \mu\text{eV}$ Kaxion (predicted)
<i>KK spectrum</i>	N/A	Compactification-dependent	N/A	Free first-mass	$\sim 3.85 \text{ TeV}$ cluster (predicted)
<i>CP violation</i>	Phenomenological	Compactification-dependent	N/A	Phenomenological	Predicted from BI-6
<i>Free parameters</i>	28 (SM) + cosmological	$\gtrsim 100$ moduli + landscape	1 (N of $\text{SU}(N)$)	few (warp slope, etc.)	0 (after BI-6 + MSMF)
<i>Foundational principle</i>	Unitary evolution + Born rule	Worldsheet conformal invariance + SUSY	Bulk-boundary duality	Brane-bulk warp	OGGEP (operational/open geometry–gauge equivalence)
<i>Conceptual economy</i>	Many (QM postulates + SM Lagrangian)	SUSY + worldsheet + GS anomaly + landscape	Two sides + dictionary	SM + warp + branes	4 axioms (LC, GGEP, OGGEP, MSMF) + BI-6
<i>Resolves meas. problem?</i>	No	No (inherited)	No (inherited)	No (inherited)	Yes (boundary projector foundational)
<i>Resolves BH info paradox?</i>	No	Partial (via AdS/CFT)	Partial	No	Yes (info preserved at bulk)
<i>LHC falsifiable?</i>	No SM-internal predictions	Hard (landscape)	No direct	Maybe (KK gravitons)	Yes ($\sim 3.85 \text{ TeV}$ cluster)
<i>Low-mass DM falsifiable?</i>	No	Hard (landscape)	No	No	Yes ($\approx 8.1 \mu\text{eV}$ Kaxion)
<i>Flavor-precision falsifiable?</i>	No	Hard (landscape)	No	Limited	Yes (m_τ/m_μ to 0.5%)
<i>Status</i>	Phenomenologically successful, foundationally incomplete	Mathematically rich, predictively limited	Conceptually deep, narrowly applicable	Phenomenological completion of SM, no foundational layer	Unified open-system foundational framework

54.11 Reading the table

The comparison table makes explicit what the structural differences between GDOS/Gi-4 and the alternatives amount to. Several patterns emerge.

First: GDOS is the only framework with the foundational “open-system” designation. Every other framework treats unitary evolution as foundational and the boundary projection as derived (or absent).

Second: GDOS is the only framework with a forced bulk dimension. String theory chooses $d = 10$ from worldsheet conformal invariance; RS chooses $d = 5$ from AdS/braneworld convenience; AdS/CFT studies $d = 4 + 1$ for specific dualities. Only GDOS forces $d = 6$ from purely algebraic considerations (the magic-dimension ladder).

Third: GDOS is the only framework that derives $N = 3$ without external input. String theory chooses CY manifolds with $|\chi/2| = 3$; RS postulates families; orthodox SM postulates families. GDOS derives $N_g = 3$ as a theorem from the GG-Constitution + BI-6 + MSMF.

Fourth: GDOS is the only framework with no free parameters after the constitutional axioms and BI-6 are accepted. Every gauge coupling, every Yukawa, every mixing angle, every CP phase, every hierarchy is derived. This is what we mean when we say the framework is “constitutionally predictive”.

Fifth: GDOS is the only framework that resolves the measurement problem at the foundational level. Every other framework either inherits the problem unchanged from QM or addresses it only by interpretation rather than structural resolution.

Sixth: GDOS is the only framework with a coherent set of sharp experimental falsifiers across multiple energy scales — LHC KK cluster, low-mass Kaxion DM, flavor-physics precision tests on m_τ/m_μ . Each of these is a Popperian falsification target; positive results would constitute confirmation, and negative results across the central ones would constitute refutation.

These six patterns together constitute the structural distinctiveness of GDOS as a contemporary unification framework. They explain why the framework is, we have argued, the missing concept of foundational physics for the past century.

54.12 Why other approaches missed this

It is worth asking, by way of closing this discussion, why the combination of features present in GDOS was not assembled earlier by the string-theory or RS communities. Three reasons stand out.

First, the prestige of closed-system frameworks. QM and QFT have been so empirically successful as closed-system frameworks that the extension to a foundational open-system framework has seemed unnecessary. Most string-theory work has remained within the closed-system paradigm. The ingredients of GDOS (Lindblad, holography, decoherence, Hopkins–Singer) have been developed in their own technical silos rather than synthesized.

Second, the prestige of $d = 10$. Worldsheet conformal invariance forces $d = 10$ for the standard superstring, and most string-theory work has accepted this without revisiting whether other algebraic forcings might select a different dimension. The magic-dimension ladder, which selects $d = 6$ via $\text{Spin}(1, 5) \cong \text{SL}(2, \mathbb{H})$, has been known mathematically since at least the 1990s but

has not been used to motivate a different bulk dimension.

Third, the prestige of phenomenological inputs. The Standard Model’s 28 parameters have been treated as phenomenological inputs to be measured and fit, rather than as derivable numbers. Most string-theory and RS work has continued this tradition, treating compactification choices and brane positions as adjustable parameters. The notion that the 28 parameters might all be derivable from a single foundational principle has been considered too ambitious.

GDOS steps outside each of these three traditions. It treats the open-system structure as foundational (challenging the closed-system prestige). It selects $d = 6$ via the magic-dimension ladder (challenging the $d = 10$ prestige). It derives all 28 SM parameters (challenging the phenomenological-input prestige). These three challenges are why the framework is structurally distinctive — and why, if its predictions are confirmed experimentally, it will constitute the foundational shift we have argued for in [Arisaka \(2026, P2-N3, § 12\)](#) and in Part IX (below).

55 The Ten Structural Innovations of Gi-4

The Gi-4 construction is, by design, conservative at the level of physical input: it adds no new axiom to the GG-Constitution. Its contribution is structural — it converts an informally used vocabulary into a rigorous five-theorem framework. We collect here the ten structural innovations that this conversion delivers, each carrying its cohort-global identifier (§G).

Innovation 1 (I-GRIP-1): Axiomatize the informal Gi-4 vocabulary. The GFF, GRIP, GA-cascade, family-projector and anomaly-inflow vocabulary deployed informally by the $N_g = 3$ derivation ([Arisaka, 2026, P2-N3](#)), the Standard-Model derivation ([Arisaka, 2026, A9-SM](#)), the boundary cosmology ([Arisaka, 2026, A10-Cosmos](#)) and the dark-matter prediction ([Arisaka, 2026, C1-Kaxion](#)) is here given a rigorous axiomatic basis as five principal theorems. This is the geometric half of the foundational writing the program had postponed.

Innovation 2 (I-GRIP-2): The dual-layer four-pillar architecture. The four pillars TGG, CCI, GFF, GRIP are formalized as a single system in which each pillar carries a bulk (algebraic) layer and a boundary (operational) layer, for nine dual-layer sub-theorems **T1.1–T4.2**. The separation of the algebraic from the operational content is what makes each pillar’s falsifier explicit.

Innovation 3 (I-GRIP-3): The pillar forcing chain $TGG \Rightarrow CCI \Rightarrow GFF \Rightarrow GRIP$. The four pillars are not four independent posits but a single derivation spine: each pillar forces the next, so that the whole framework descends from the upstream GG-Constitution, 6D bulk, and BI-6 closure without an independent assumption at any link.

Innovation 4 (I-GRIP-4): The CPTP boundary projector at the foundational level. In the GDOS reading carried by Gi-4, the irreversible boundary projector $\Pi : M_6 \rightarrow M_4$ is the foundational dynamical object; unitary evolution is the isolated-bulk special case and “measurement” is the cascade endpoint. No closed-system framework places the projector at this level, and without that elevation the measurement problem cannot be structurally dissolved.

Innovation 5 (I-GRIP-5): Geometric–gauge unification with a Master Conservation Law (Theorem T1). TGG supplies a unique composite connection $\mathcal{A}_{\text{tot}} = \omega \oplus A$ on the principal G_{tot} -bundle over the 6D bulk, together with its boundary stress tensor and the Master Conservation Law $\nabla_\mu \mathcal{T}_{\text{TGG}}^{\mu\nu} = 0$ — gravity and gauge structure as two facets of one geometric object.

Innovation 6 (I-GRIP-6): The OGN-5 integer ledger from anomaly closure (Theorem T2). MSMF closure of the four-component anomaly polynomial X_4 forces the unique integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$. This ledger is the backbone from which every quantitative prediction of the framework follows in closed form, with zero free parameters once the GG-Constitution and BI-6 are accepted.

Innovation 7 (I-GRIP-7): A boundary second law from geometry (Theorem T3). The GFF free-energy functional $\mathcal{F}_{\text{GFF}} = \mathcal{F}_{\text{GEP}} + \mathcal{F}_{\text{CEP}}$ on the stratified boundary moduli is Morse-positive on every Kähler stratum, and the second law $dF_{\text{GFF}}/dt \leq 0$ holds along every GDOS CPTP trajectory with detailed balance — a geometric arrow of time on the boundary.

Innovation 8 (I-GRIP-8): The family number $N_g \leq 3$ as a theorem (Theorem T4). The cascade $\text{DGI} \rightarrow \text{GSSB} \rightarrow \text{GA}$ closes under the BI-6 budget $S_{\text{tot}} \leq 3\pi$, forcing $N_g \leq 3$ on the family-projector stratum (the budget delivered by the capacity theorem of Arisaka (2026, P2-N3), meeting M10’s demand; residue M6 = R-INT), a two-sided pincer with the $N_g \geq 3$ lower bound from CP. Integer multiplicities and hierarchies become *derived* rather than postulated.

Innovation 9 (I-GRIP-9): The geometric gate weight $w = e^{-\pi k}$. The cascade tolls $S_{[2k] \rightarrow [2(k-1)]} = \pi k$ deliver a generation-ordered weight ladder $w_{[2k] \rightarrow [2(k-1)]} = e^{-\pi k}$ as the boundary signature of the cascade — an equilibrium weight on the skeleton, not a rate. The ladder is the geometric *gate* factor only; absolute lifetimes are phase-space-dominated and are fixed in Arisaka (2026, A9-SM).

Innovation 10 (I-GRIP-10): The constitutional bridge (Theorem T5). Gi-4 is shown equivalent — with the converse direction a structured proof sketch conditional on the cascade conditioning of §57 — to the GG-Constitution: the four pillars together carry exactly the content of the four axioms and seven postulates of Arisaka (2026, A1-Constitution), closing the loop and confirming that the framework adds no new physical axiom.

56 Predictions for falsification

A theory of foundational physics earns its credibility through empirical falsifiability. The dual-layer architecture of produces a coherent set of laboratory- and observatory-testable predictions, each tied to a specific Structure Theorem of the physical block of one of the four Pillars. We collect the principal falsifiers in a single table for the convenience of experimentalists and observers, then briefly comment on each.

Table 20: Every prediction the four Pillars make that could kill the framework, with the observable that would do it and where the measurement currently stands. The fourth column is the honest one: some rows are already confirmed to high precision and are therefore no longer at risk, and the rest name the facility that will decide them. A framework that forces its parameters has to accept that a single failed row ends it, and this is the table where that exposure is stated rather than implied.

Pillar / ST	Prediction	Empirical observable	Status / facility
TGG-ST6, ST7	Master Conservation Law $\nabla_\mu T_{\text{TGG}}^{\mu\nu} = 0$	Tests of energy/momentum conservation in any boundary process	Confirmed to 10^{-15} in low-energy experiments
CCI-ST7	Modified TOV $\Delta_{\text{GS}} \sim 9.6 \times 10^{-3}$	0.1%–1% shift of NS mass–radius curve from standard TOV	NICER (current); next-gen X-ray timing; LIGO/Virgo NS inspirals
GFF-ST7 (1)	Equipartition $\langle \delta X^2 \rangle = T_{\text{GFF}} \chi$ for any boundary observable	Universal T -scaling of fluctuations near equilibrium	Ultra-cold atoms, optomechanics; quantum-limited sensors
GFF-ST7 (2)	Fluctuation–dissipation: $\langle \delta F^2 \rangle = T_{\text{GFF}} \chi_{\text{GFF}}$	Einstein-type relation between mobility and diffusion	Mesoscopic transport experiments
GFF-ST7 (3)	Onsager reciprocity $L_{ij} = L_{ji}$	Symmetric cross-coupling matrix in any open boundary network	Molecular motors, active matter, biophysical networks
GRIP-ST7	Generation weight ladder $w_{[2k] \rightarrow [2(k-1)]} = e^{-\pi k}$	Geometric gate weight along the family chains $\tau \rightarrow \mu \rightarrow e, t \rightarrow \dots \rightarrow u/d$ (absolute lifetimes via A9)	Generation-ordering of the gate-action ladder; controlled-kinematics decay class
GRIP-ST7	Cascade-toll sum rule	$\sum_i S_{\text{tot}}^{(i)} = \pi N_g (N_g - 1) (N_g + 1) / 6$, a toll budget, not a timescale	None by itself; the ordering is read from the mass hierarchy
BI-6 (Companion)	KK resonance cluster $M_{\text{lock}} \simeq 3.85 \text{ TeV}$	Tower of KK modes near 3.85 TeV in dijet/diphoton/dilepton	LHC Run 3, HL-LHC; FCC-hh
BI-6 (Companion)	Kaxion dark matter $m_K \approx 8.1 \text{ } \mu\text{eV}$	Ultra-light bosonic DM near μeV with axion-like couplings	ABRACADABRA, DMRadio, ADMX-EFR; cosmological constraints
GG-Cosmos	Modified Friedmann equations	Coherent shift of $f\sigma_8(z)$, $w(z) + 1 \propto \langle X_4 \rangle / M_{\text{KK}}^4$	DESI, Euclid, Roman Space Telescope; CMB-S4
GG-6F (codebook)	28 SM parameters at $\sim 0.5\%$ precision	e.g. m_τ/m_μ to 0.5% from cascade depths	Global SM precision fits (Particle Data Group (Particle Data Group, 2024))

56.1 The astrophysical falsifier (modified TOV) in detail

The modified TOV equation (74) predicts a *universal* small correction to the neutron-star mass–radius relation, with calculable amplitude $\kappa_{\text{GS}} = 42/4375 \approx 9.6 \times 10^{-3}$ from the OGN-5 ledger. This is a single-parameter prediction (no fits): the only empirical input is the assumed value of the bulk integral $\langle X_4 \rangle_{\text{bulk}}$, which is bounded by the OGN-5 integers. A coherent population study across multiple compact stars (neutron stars, white dwarfs, magnetars) at the 0.1%–1% precision level provides the strongest constraint. A null result (no shift seen across a coherent population to better than 0.1%) would constitute a serious challenge to CCI-ST7 and hence to the entire GG-6

framework.

56.2 The thermodynamic falsifiers (equipartition + FDT) in detail

The thermodynamic predictions of GFF-ST7 are universal in the sense that they apply to any boundary observable in any open system. The expected signal is the standard linear-response form of equipartition and fluctuation–dissipation, with the boundary effective temperature $T_{\text{GFF}} \sim M_{\text{KK}}$. In laboratory systems below M_{KK} , the boundary world is effectively at “zero temperature” compared to the GG temperature scale, so the Boltzmann statistics are dominated by the global minimum $[\Pi_*]$. A laboratory observation of equipartition violation, or non-symmetric Onsager cross-coupling, would falsify GFF-ST7.

56.3 The generation gate-ordering falsifier (cascade structure) in detail

The generation gate weight $w_{[2k] \rightarrow [2(k-1)]} = e^{-\pi k}$ of GRIP-ST7 assigns each family gate a fixed geometric weight, an equilibrium weight and not a suppression of any rate. The robust, testable content is the *ordering* of the generations by these weights along the charged-lepton ladder $\tau \rightarrow \mu \rightarrow e$ (and its quark counterpart), read from the mass hierarchy sector by sector in Arisaka (2026, A9-SM), together with the ratio of each sector’s two log-mass gaps, which must be the ratio of the height gaps of that sector’s action (§33.3); the lifetimes themselves are phase-space quantities and not a falsifier. A family ladder whose generations were not ordered as the weights order them would falsify GRIP-ST7; the ratios, read at one scale, say that the action must differ sector by sector. The residual-width form of the test — a controlled-kinematics decay class in which a residual ordering departing from $e^{-\pi k}$ by more than $e^\pi \approx 23$ would count against the cascade — presupposes that the weight enters a width, which needs a mechanism this paper does not supply; it is carried at that grade.

56.4 The cosmological falsifiers: deferred to GG-A10-Cosmos

The same Master Conservation Law that produces the modified TOV equation also predicts coherent modifications of the Friedmann equations at high redshift, sourced by X_{GS} . The boundary-side cosmological observables — $w(z)$, $f\sigma_8(z)$, and the density-fraction predictions $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, $\Omega_b = 1/(2\pi^2)$ — are the subject of GG-A10-Cosmos (Arisaka, 2026, A10-Cosmos) — with the Hubble-tension reading developed in Arisaka (2026, C2-Hubble) — where the full Friedmann-equation modification chain, the DESI/Euclid/Roman next-decade precision targets, and the residual cosmological-constant puzzle are worked out in detail. The Gi-4 contribution to that program is the CCI-side identity that fixes the structural form of X_{GS} ; the cosmological reduction itself is the boundary law’s application to the Friedmann-Robertson-Walker stratum, which is A10’s territory.

56.5 Falsification logic: a multi-channel commitment

A critical feature of the GG-Theory falsification program is its *multi-channel commitment*. The same OGN-5 ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ produces predictions across five independent

empirical domains (compact-object structure, laboratory thermodynamics, particle lifetimes, cosmological observables, and SM parameter values). A theory with one tunable parameter could fit one domain by hand; GG-Theory has zero free parameters after the GG-Constitution and BI-6 are accepted, and the same integers must reproduce all five. This makes the framework simultaneously *falsifiable* (any single domain can refute it) and *informative* (any single domain that confirms it constrains the others).

57 Open problems

This section lists what is not done. It is long, and its length is deliberate: an open-problem list is the only part of a framework that tells a reader where to go next, and a short one is usually evidence of inattention rather than of progress.

The entries are given under their own names, in the mathematical, physical and experimental categories, and each is marked with its status — genuinely open, partially resolved, or resolved in a named companion. Where a result earlier in this paper rests on one of these, the dependence is stated here and not softened.

57.1 Open mathematical questions

- (M1) **Higher-derivative corrections to the bulk action.** *Status: Genuinely open.* The present paper assumes the leading-order two-derivative bulk action. What is the systematic structure of higher-derivative corrections (four-derivative, six-derivative, etc.) to TGG, and how do they modify the boundary descent of GFF and GRIP? Is there a “UV completion” of GG-Theory that fixes these corrections?
- (M2) **Quantum corrections to the cascade actions.** *Status: Genuinely open.* The Bogomolnyi-saturating cascade actions $S_{[2k] \rightarrow [2(k-1)]} = \pi k$ are classical results. What are the leading one-loop corrections, and do they modify the generation gate-ordering? This is a calculation in instanton perturbation theory that has not been performed.
- (M3) **Non-perturbative corrections to PBI-4.** *Status: Genuinely open.* The PBI-4 compensator C_μ is constructed perturbatively by KK reduction. Are there non-perturbative contributions (D-branes, other solitonic objects in the bulk) that modify C_μ ? If so, do they affect the modified TOV correction?
- (M4) **Topological refinement of the OGN-5 ledger.** *Status: Partially resolved.* The Diophantine derivation is settled in Arisaka (2026, A3-BI6, Theorem 10.2); the further question of a topological-geometric construction generating the integers is open. The OGN-5 integers $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ satisfy two algebraic closure constraints. Is there a deeper topological construction (e.g., a specific Calabi–Yau or G-structure) that *generates* these integers as topological invariants?
- (M5) **Stratification of $\mathcal{M}_{\text{bdy}}$ at the global level.** *Status: Genuinely open.* The boundary moduli space $\mathcal{M}_{\text{bdy}}$ has been described locally as a disjoint union of strata. What is the precise global topology of $\mathcal{M}_{\text{bdy}}$? Does it have non-trivial π_0 or π_1 that influences the cascade dynamics?

- (M6) **Hopkins–Singer descent / η -invariant parity non-degeneracy.** *Status: Partially resolved* (in Arisaka (2026, P2-N3)). The proof of $\mathcal{H}_2$ (Theorem 27.2) sketches the η -invariant non-degeneracy argument; the detailed bulk computation is carried by Appendix H of Arisaka (2026, P2-N3), whose κ -bridge theorem executes the Hopkins–Singer pushforward on the benchmark background. The parity leg of the N=3 derivation is therefore rigorous *modulo* M6, as the Common Objections (O-GRIP-14, O-GRIP-52) state. In the sharpened form of the capacity architecture of Arisaka (2026, P2-N3), M6 is identical to the topological residue R-INT of the integrality theorem L-INT there — and that appendix discharges it on the benchmark background. Its two declared refinements — the warped and relative computations — are closed at class level by the rigidity theorems T-INF-7 and T-INF-8 there (warp rigidity; relative closure), so what remains of M6 is the off-benchmark Wu-basepoint canonicity together with two engineering items (the constructive warped representative and the Dai–Freed packaging): scope and exposition, not truth on the physical background.
- (M7) **Detailed Riemannian metric on stratified $\mathcal{M}_{\text{bdy}}$.** *Status: Genuinely open.* The stratified boundary moduli $\mathcal{M}_{\text{bdy}}$ has been described qualitatively in Part IV, but its detailed Riemannian structure (and the explicit CEP computation on each stratum) deserves a separate technical treatment.
- (M8) **Higher-order cascade corrections (beyond Bogomolnyi-saturating).** *Status: Genuinely open.* The Bogomolnyi-saturating gradient flow of GRIP gives the leading action of the connecting orbit between adjacent critical points — the height difference, $c[\Delta_+(k+1) - \Delta_+(k)]$ under Theorem 24.1 — while the gate carries the inflow toll πk (GRIP-ST3). Higher-order corrections to either (non-Bogomolnyi trajectories, fluctuations, multi-instanton processes) are sub-leading but in principle computable; no explicit calculation has yet been carried out.
- (M9) **Norm-canonicalization of MSMF.** *Status: Genuinely open — norm-independence established; the canonicity question remains.* The Diophantine uniqueness of $(1, 3, 5; 7, 35)$ in T2.1 was originally established under the ℓ^1 norm $\text{MSMF}(M, N, L; J, K) = M + N + L + J + K$, with independence from that choice left open. The independence half is now settled by an elementary dominance argument. For every admissible competitor $(M, N, 5; J, K) \neq (1, 3, 5; 7, 35)$ of the T2.1 constraints one has $M^2 + N^2 \geq 20$ (the next multiple of 5 above 10 attainable with $M \geq 1$, $N \geq 3$), hence $K = 25 + M^2 + N^2 \geq 45$ and $J = K/5 \geq 9$; moreover $K > J > \max(M, N, 5)$ (since $J - N = 5 + (M^2 + N^2)/5 - N > 0$ and likewise for M), the median of $\{M, N, 5\}$ is at least 3, and the minimum is at least 1. The sorted competitor therefore dominates the sorted ledger $(1, 3, 5, 7, 35)$ coordinatewise, strictly in the top entry. It follows that $(1, 3, 5; 7, 35)$ is the unique minimizer of *every* symmetric functional strictly increasing in each argument — every ℓ^p with $0 < p \leq \infty$, every weighted symmetric gauge, every entropy- or product-type complexity — not merely ℓ^1 (verified exhaustively for $M, N \leq 2000$ alongside a thirteen-norm battery; support archive). What remains of this problem, and keeps it open, is the canonicity question: the argument covers the class of symmetric monotone complexity functionals, and a canonical physical motivation for that class — or for the ℓ^1 representative specifically (e.g., as the integer cohomology rank required to realize the ledger) — would close the question entirely.

For the selection itself, the norm choice no longer matters.

- (M10) **In-paper derivation of the BI-6 budget** $S_{\text{BI6}}^{\text{max}} = 3\pi$. *Status: Resolved* (externally, in [Arisaka \(2026, P2-N3\)](#); residue identified with M6). The demand this problem records — a derivation independent of $N = 3$ at every intermediate step, so as to rule out any circularity in the cascade-closure $N_g \leq 3$ argument of T4.1 — is met by the capacity theorem T-N3-6 of [Arisaka \(2026, P2-N3\)](#): $\text{Cap}(N_g) = \pi N_g$, derived from the family stratum’s quantum geometry (Berry transport and the Fubini–Study package, forced at the Born theorem’s own grade by its adiabatic theorems), consuming neither $K = 35$ nor $J = 7$ and inserting $N = 3$ at no step. The residual conditionality is that theorem’s: the coupling clause with the integrality theorem L-INT, whose topological residue R-INT is precisely this paper’s η -invariant computation M6 in its sharpened, well-posed form — discharged on the benchmark background by the κ -bridge theorem of [Arisaka \(2026, P2-N3, Appendix H\)](#). Cross-references “conditional on M10” throughout this paper should accordingly be read as conditional on the coupling clause and M6 (= R-INT). An in-paper re-derivation remains available as future work but is no longer load-bearing. Two observations narrow it. First, the seam-count motivation in the form $\lfloor J/2 \rfloor \pi$ cannot be the derivation, because $J = 2N + 1$ makes $\lfloor J/2 \rfloor = N$ identically, so that expression is πN_g evaluated at the answer; the per-family form $S_{\text{max}}(N_g) = \pi N_g$ is the honest statement of the same conjecture and is free of that circularity. Second, the family count is far less sensitive to the budget than the exact value suggests: the cascade costs $\pi N_g(N_g - 1)/2$, so $N_g = 3$ costs 3π while the next candidate admitted by the parity condition, $N = 5$, costs 10π . *Any* budget in the range $[3\pi, 10\pi)$ therefore yields exactly $N_g = 3$. What requires the exact value 3π is not the family count but the saturation reading — the claim that the budget is exhausted precisely at $N = 3$ — and that reading should be quoted with the distinction in view.
- (M11) **Monodromy of PBI-4 compensator across instanton stratum (spectral-flow form).** *Status: Genuinely open — sharpened, with a verified benchmark instance.* The reverse direction of T4.2 requires the PBI-4 compensator C_μ to extend smoothly across each instanton crossing on the cascade. Pointwise this is shown; the global topological non-obstruction (vanishing of the monodromy of C_μ around the codimension-2 instanton stratum) is established case-by-case for $\mathbb{CP}^{N-1}$ saddles in this paper but lacks a general theorem. The sharpened form of the missing statement is index-theoretic: *the monodromy of C_μ around the stratum equals $\exp(\pi i \Delta \eta)$, the mod-period reduction of the η -invariant jump across it* — the same object the integrality chain L-INT/R-INT of [Arisaka \(2026, P2-N3\)](#) controls — so that vanishing monodromy is equivalent to integrality of the trapped inflow charge, and the two conditionings merge. On the benchmark background that charge is the boundary Berry unit (-1 , an integer, factor and sign) by the κ -bridge theorem of [Arisaka \(2026, P2-N3, Appendix H\)](#), which gives the monodromy statement its first verified instance there — consistent with, and now explaining, the direct case-checks of this paper. The residual hypotheses are fewer than they first appear. The apparent need for adiabaticity at the cascade saddle dissolves on inspection: the monodromy, the trapped charge, and the η -jump are each invariants of the *link* of the instanton stratum, and the link lies in the gapped complement — the stratum is the degeneracy locus, and a codimension-2 degeneracy leaves its complement gapped by definition — so the companion’s adiabatic hypotheses are

consumed only where they hold with margin, and by homotopy invariance the invariants are independent of the link radius. What the dissolution consumes is transverse nondegeneracy of the gap opening at the stratum — exactly the strict Morse-positivity of Theorem T3 — together with one named identification, GAP=HESS: that the boundary family’s spectral gap transverse to the stratum is the T3 Hessian, the natural reading under the runtime premises of the companion’s adiabatic theorems. (Measured instantiations, support archive: the link holonomy is radius-independent over five decades, unchanged at Morse anisotropy e^{-7} , and dies exactly when transverse nondegeneracy fails.) Off the benchmark, the warped and relative scopes are closed at class level by the rigidity theorems T-INF-7 and T-INF-8 of Arisaka (2026, P2-N3, Appendix H), leaving the Wu-basepoint canonicity as the one scope item. The general theorem therefore rests on GAP=HESS and the off-benchmark basepoint convention — one identification and one convention — rather than on any analytic hypothesis at the saddle, and the cascade closure hangs on one named index-theoretic statement rather than on two unrelated conjectures. Reading rule: for the Standard-Model application the finite $\mathbb{CP}^{N-1}$ case-checks stand on their own; the spectral-flow form governs the general statement, and cross-references “conditional on M11” concern that general statement only.

- (M12) **The stage–fiber correspondence at full generality.** *Status: Genuinely open (statement assembled; one stabilizer step unproven).* Subsection 31.9 states the custody correspondence: the DGI stage’s equivalences compose as the circle’s group, the GSSB→GA descent as the interval’s semigroup (Proposition 31.8), with data custody crossing at GSSB. The DGI half rests on reading the Morse-top stabilizer of Definition 31.1 as the boundary gauge group stratum by stratum. On Σ_{gauge} and Σ_{θ} this is immediate; on the family stratum $\Sigma_{\text{fam}} = \mathbb{CP}^{N-1}$ the stabilizer involves the family quantum geometry, and the Berry transport that prices the GSSB gates ($\text{Cap}(N_g) = \pi N_g$; Arisaka, 2026, P2-N3) is holonomy mathematics doing interval-side bookkeeping — the deepest crossing of the two custodies. A general stabilizer theorem, uniform across strata, would complete the correspondence at theorem grade.
- (M13) **The family weights: which premise gives way.** *Status: Genuinely open (one horn adopted).* Theorem 24.1 shows that the transport proposition (Proposition 31.8) and the reading of the gate toll πk as the height difference of the family landscape cannot both hold: the first makes the critical values the scaling dimensions $c \Delta_+(\ell)$, gap ratio 1.31, and no scaling-dimension spectrum of any bulk field has the pair invariant’s gap ratio 2. This paper adopts the transport proposition and reads πk as the inflow toll (**GRIP-ST3**, Remark 32.2). The other horn — a descent on the family stratum that is not the projected dilatation, so that the pair invariant may be the height — would require a mechanism for the descent other than the composition of GG-A5-GDOS’s two results, and none is on the books. Two things would close the problem: a derivation of the winding sign from the six-dimensional chirality of the family spinors, which would make the isolated-fixed-point clause of admissibility a theorem rather than an input; and a computation, in GG-A9-SM, of which of the candidate weights the sector-by-sector Yukawa gap ratios track, since no single ladder fits the three sectors. The unit of the height (c , and the temperature T_{GFF} in which it is measured) is part of the same problem.

- (M14) **The space of the second regime.** *Status: Genuinely open.* The strata are invariant under their descents, so a single flow that carries a GRIP-Out event must fail to be the descent during the event and must pass through points that lie in no stratum — on a space larger than the strata, which this paper has not defined (§31.7). The one extension of Conley’s theorem under which the second regime is a theorem is the hybrid one, with a closed guard set reached in finite time (Kvalheim *et al.*, 2021), at the price that the reset fires before the descent halts; whether the runtime is that hybrid system, a skew product over a slow clock, or an autonomous slow–fast system on a completion of the strata is the choice the capstone (Arisaka, 2026, A11-Origin) must make, and M5 (the global stratification) is the same problem seen from the moduli side.

57.2 Open physical questions

- (PS-GRIP-9) **The information sector $U(1)_{Q_i}$ and quantum entropy.** *Status: Resolved* (in Arisaka (2026, A7-Qi4), in preparation). The fourth gauge factor $U(1)_{Q_i}$ has been treated as a formal ingredient; its physical role in carrying information-theoretic content (Qi-4 framework) has been only sketched. What are the explicit Qi-4 Structure Theorems, and how do they pair with the present Gi-4 Structure Theorems? This is the content of the companion paper GG-A7-Qi4 (in preparation).
- (PS-GRIP-10) **Cosmological initial conditions.** *Status: Partially resolved* (in Arisaka (2026, A10-Cosmos)). The cascade $DGI \rightarrow GSSB \rightarrow GA$ describes the relaxation of the boundary world to its current equilibrium. What was the cosmological initial condition — the configuration of $\mathcal{M}_{\text{bdy}}$ at $t = 0$? Is it the maximally unstable DGI stratum, or something else? This ties GG-Theory to inflationary cosmology.
- (PS-GRIP-11) **Black-hole information and bulk-boundary unitarity.** *Status: Genuinely open.* The GG-6 framework preserves information at the bulk level while the boundary is CPTP. How does this resolve the Hawking information paradox in detail? A worked example for a specific black-hole geometry would clarify; the GG-Theory black-hole-entropy treatment — the horizon as an OGGEF port, with exact Bekenstein–Hawking $A/4$ from induced gravity — is Arisaka (2026, C3-BH).
- (PS-GRIP-12) **Connection to the cosmological constant problem.** *Status: Partially resolved* (in Arisaka (2026, A10-Cosmos), Arisaka (2026, A4-Euclid)). The geometric mass scale $m_\Lambda \sim M_{\text{lock}}^2/M_{\text{Pl}} \approx 6 \text{ meV}$ is derived; the full 122-order hierarchy from M_{Pl}^4 is not yet fully closed. The observed cosmological constant $\Lambda_{\text{obs}} \sim 10^{-122} M_{\text{Pl}}^4$ is unexplained by the SM. Does the PBI-4 compensator and the bulk scale running of GG-6 account for this hierarchy? A first calculation suggests $\Lambda_{\text{obs}}/M_{\text{Pl}}^4 \sim e^{-2K}$ with $K = 35$, giving $\sim 10^{-30}$, which is closer than SM but still off by 90 orders of magnitude. More work needed.
- (PS-GRIP-13) **Neutrino masses and the seesaw mechanism.** *Status: Partially resolved* (in Arisaka (2026, A9-SM)). The framework predicts $N_g = 3$ generations (the pincer of GG-P2-N3, not a ledger entry). Do the neutrino mass differences $\Delta m_{\text{sol}}^2, \Delta m_{\text{atm}}^2$ arise from a specific GRIP cascade in the lepton sector? This is in GG-A9-SM but deserves a focused paper.

57.3 Open experimental directions

- (PS-GRIP-14) **Multi-messenger compact-object surveys.** *Status: Genuinely open* (experimental, awaiting data). A coherent population study of neutron stars, white dwarfs, and magnetars at the 0.1% mass–radius precision level is the most direct test of the modified TOV CCI-ST7. Combining NICER X-ray data with LIGO/Virgo gravitational-wave inspirals would provide the strongest constraint.
- (PS-GRIP-15) **Long-lived heavy-particle searches at LHC and beyond.** *Status: Genuinely open* (experimental, awaiting data). A discovery of a fourth long-lived heavy particle (e.g., a SUSY-like state with $\tau \gtrsim 10^{-9}$ s) would test the cascade-structure prediction of GRIP-ST7 directly. HL-LHC and FCC-hh have the relevant reach.
- (PS-GRIP-16) **Sub-eV dark-matter direct detection.** *Status: Genuinely open* (predicted in [Arisaka \(2026, C1-Kaxion\)](#), awaiting experiment). The Kaxion prediction $m_K \approx 8.1 \mu\text{eV}$ (≈ 1.95 GHz) requires next-generation detectors (ABRACADABRA, DMRadio, ADMX-EFR). A discovery in this mass range would be a smoking gun for GG-6 at low energies.
- (PS-GRIP-17) **Cosmological precision tests.** *Status: Genuinely open* (experimental, awaiting next-decade data). DESI, Euclid, Roman Space Telescope, and CMB-S4 are expected to reach 10^{-3} precision on $f\sigma_8(z)$ and $w(z)$, the relevant level for the GG-6 cosmological predictions. A coherent shift correlated across these observables would be a strong signal.
- (PS-GRIP-18) **Precision SM parameter fits.** *Status: Genuinely open* (predicted in [Arisaka \(2026, A9-SM\)](#), awaiting precision improvement). The 28 SM parameters predicted by GG-A9-SM at $\sim 0.5\%$ precision can be tested against the global fits maintained by the Particle Data Group. As the experimental precision on each parameter improves, the test becomes increasingly stringent.

57.4 Triage summary

Of the 30 open and future-direction entries in the four lists above, 2 are *Resolved* in a named companion paper (PS-GRIP-9, in [Arisaka \(2026, A7-Qi4\)](#); M10, in [Arisaka \(2026, P2-N3\)](#)); 7 are *Partially resolved* with the central claim settled in a named companion paper and specific sub-tasks still pending (M4 in [Arisaka \(2026, A3-BI6\)](#); M6 in [Arisaka \(2026, P2-N3\)](#); PS-GRIP-10 in [Arisaka \(2026, A10-Cosmos\)](#); PS-GRIP-12 in [Arisaka \(2026, A10-Cosmos\)](#) and [Arisaka \(2026, A4-Euclid\)](#); PS-GRIP-13 in [Arisaka \(2026, A9-SM\)](#); PS-GRIP-19 in [Arisaka \(2026, A5-GDOS\)](#) and [Arisaka \(2026, A7-Qi4\)](#); PS-GRIP-24 in [Arisaka \(2026, A4-Euclid\)](#) and [Arisaka \(2026, BookI\)](#)). The remaining 21 are *Genuinely open* and constitute the live agenda for the next phase of the GG-Theory program: eleven mathematical questions (M1, M2, M3, M5, M7, M8, together with M9 norm-canonicalization of MSMF, M11 monodromy of PBI-4 compensator in spectral-flow form, M12 the stage–fiber correspondence, M13 the family weights, and M14 the space of the second regime), one physical question (PS-GRIP-11, black-hole information), the five experimental directions (PS-GRIP-14–PS-GRIP-18, all awaiting data), and four program-level extensions (PS-GRIP-20, PS-GRIP-21, PS-GRIP-22, PS-GRIP-23). Total $11 + 1 + 5 + 4 = 21$.

Indeed, the proportion of Partially-resolved entries is a useful snapshot of the program’s state:

most of the constitutionally load-bearing questions are settled in named companion papers, and the live open work concentrates either on experimental verification or on extensions to adjacent scientific domains. The methodological gain of the GG-Theory program is that “open” now means “open in the sense of awaiting data” or “open in the sense of extending an axiomatic framework to new domains,” rather than “open in the sense of having no constitutional content at all.”

58 Future prospects

Two things follow from the open-problem list: a mathematical program and an experimental one, and they run on very different clocks.

58.1 Future directions for the GG-6 program

- (PS-GRIP-19) **Complete the GDOS quartet.** *Status: Partially resolved* (GG-A5-GDOS stable in [Arisaka \(2026, A5-GDOS\)](#); GG-A7-Qi4 in preparation). The present paper is one of four foundational papers in the GDOS quartet (along with the Constitution ([Arisaka, 2026, A1-Constitution](#)), the GG-6 Foundation ([Arisaka, 2026, A2-GG6](#)), and the BI-6 closure ([Arisaka, 2026, A3-BI6](#))). The companion Qi-4 paper (GG-A7-Qi4) will complete the quartet on the information side.
- (PS-GRIP-20) **Integrate with quantum gravity.** *Status: Genuinely open.* The GG-6 bulk has a metric, but the present paper does not formalize its interaction with quantum gravity at the Planck scale. A natural extension is to embed GG-6 into a richer 11D or 12D framework where gravitational fluctuations are quantized.
- (PS-GRIP-21) **Develop the GDOS condensed-matter framework.** *Status: Genuinely open.* The GFF/GRIP cascade structure suggests that condensed-matter phase transitions (paramagnetic–ferromagnetic, normal–superconducting, etc.) are boundary projections of GG-6 cascades on appropriate strata. A systematic mapping would unify SM physics and condensed-matter universality classes.
- (PS-GRIP-22) **Develop the GDOS biology framework.** *Status: Genuinely open.* Theorem T3.2 (thermodynamic monotonicity) and the Geometric First Law (Theorem [25.1](#)) suggest a deep connection between GG-Theory and the free-energy minimization framework of theoretical biology and neuroscience. Bridging GFF to Friston’s free-energy principle, to active inference, and to non-equilibrium thermodynamics of biological systems is a major research program.
- (PS-GRIP-23) **Develop the GDOS information-theoretic framework.** *Status: Genuinely open* (partially in [Arisaka \(2026, A7-Qi4\)](#) but not yet rigorous). The boundary moduli $\mathcal{M}_{\text{bdy}}$ supports a natural entropic structure, and the cascade dynamics has the form of a discrete information flow. Mapping GRIP to information-bottleneck theory and to hierarchical Bayesian inference would clarify the relationship to machine learning and AI.
- (PS-GRIP-24) **Pedagogical and historical synthesis.** *Status: Partially resolved* (the historical synthesis lives in [Arisaka \(2026, A4-Euclid\)](#); the eventual textbook treatment is the GG-Theory Foundation Book ([Arisaka, 2026, BookI](#)), in preparation). As the framework matures, a comprehensive

textbook treatment in the spirit of Misner–Thorne–Wheeler’s *Gravitation* is needed. The present paper provides one chapter of the eventual GG-Theory textbook; the broader project is the work of years and many collaborators.

58.2 Closing thought

The GG-Theory program is in its early years. The present paper formalizes one of the four foundational documents; companion papers formalize the others. The experimental tests are years away. But the dual-layer architecture of makes the framework simultaneously *falsifiable* (each Pillar has a sharp empirical falsifier) and *universal* (the same framework subsumes physics, cosmology, biology, and information theory). If the predictions hold and the framework survives the experimental tests of the coming decade, the methodological lesson — that open-system foundations make otherwise-undetermined integers and hierarchies derivable — will have transformed the methodology of theoretical physics.

Part IX

Conclusions

Methodological synthesis, open problems, and the foundations of physics

What follows is the paper’s own account of itself: what was assumed, what was built, what was proved, and what the whole thing amounts to if it holds.

The order retraces the construction rather than summarizing the findings by importance, because the logical dependence is the point — each pillar rests on the one before, and a summary that reordered them would misrepresent how the argument works. A reader who has followed the Parts will find nothing new here. A reader who has not, and who wants the shape of the thing without the machinery, can begin here and follow the references backwards into whichever Part turns out to be the one they doubt.

Two things this Part will not do. It will not restate the falsifiers, which belong to the assessment and are given there in full; and it will not soften any of the conditional results into unconditional ones. Where a result depends on a hypothesis, the hypothesis is named here as it was named where the result was proved. A conclusion is the easiest place in a paper to lose that discipline, because it is written last and read first, and the temptation to round a conditional up to a claim is strongest exactly there.

59 Methodological synthesis

59.1 What Gi-4 has accomplished

The present paper has formalized the four pillars of the geometric backbone of GDOS – TGG, CCI, GFF, GRIP – as mathematical objects with axioms, definitions, theorems, and worked examples. The principal accomplishments are:

1. **Formal axiomatization.** Each of the four pillars is now defined by five explicit axioms with first-consequence propositions and worked examples. The vocabulary that the GG-Theory program has been using informally for more than a year now has a citable rigorous foundation.
2. **Five principal theorems.** Theorems **T1–T5** are stated and proved, modulo the two named open questions M9 and M11 (M10 resolved externally; §57). **T1** establishes existence and uniqueness of TGG. **T2** establishes the OGN-5 ledger as unique. **T3** establishes GFF Morse-positivity from CCI. **T4** establishes the cascade-closure bound $N_g \leq 3$ (the budget imported from the capacity theorem T-N3-6 of [Arisaka \(2026, P2-N3\)](#), which meets M10’s independence demand; conditional on that paper’s coupling clause and the residue M6 = R-INT). **T5** establishes the constitutional equivalence of Gi-4 with the GG-Constitution (its converse direction as a structured proof sketch, conditional on the cascade conditioning of §57).

3. **Reduction of the N=3 hypotheses.** The two technical hypotheses $\mathcal{H}_1, \mathcal{H}_2$ from [Arisaka \(2026, P2-N3\)](#) are *reduced* from free-standing assumptions of the N=3 program to consequences of the Gi-4 structure, *modulo* a stated η -invariant computation (M6) — since discharged on the benchmark by that companion’s Appendix H, with its scope refinements closed at class level (see §57). “Reduced to a named computation, discharged on the physical background” is the honest claim; “promoted to theorems in full generality” would overstate it.
4. **Foundation for downstream rigor.** The 28-parameter codebook of GG-A9-SM, the hierarchy derivation, the KK prediction, and the dark-matter candidate now all have a rigorous Gi-4 foundation. Their forthcoming revisions will cite the present paper.

59.2 The closed-to-open foundational shift

The deeper methodological achievement of the present paper, beyond the technical content, is to make explicit the closed-to-open foundational shift that is the structural diagnosis of GG-Theory. Closed-system QFT and GR can produce only necessary conditions on integer multiplicities; open-system frameworks like GDOS produce sufficient conditions through cascade closure. The Gi-4 Pillars are the formal language in which this open-system foundational shift is made explicit.

This is, we have argued in [Arisaka \(2026, P2-N3, Section 12\)](#), the missing link in the foundations of physics for the century since the establishment of quantum mechanics in 1925–1930.

59.3 Gi-4 as the language of GDOS

The four pillars constitute the language of GDOS. When we say “GDOS” from now on, we mean the framework axiomatized by:

- the four constitutional axioms (LC, GGEP, OGGEF, MSMF) of [Arisaka \(2026, A1-Constitution\)](#);
- the BI-6 closure of [Arisaka \(2026, A3-BI6\)](#);
- the Gi-4 Pillars of the present paper;
- the Qi-4 Pillars of [Arisaka \(2026, A7-Qi4\)](#) (in preparation).

This combined framework is the foundational language of GG-Theory in its mature form.

60 Conclusion: the Gi-4 Pillars and the foundations of physics

60.1 The achievement in summary

The present paper has accomplished, at the technical level, the formalization of the geometric backbone of GDOS into four named mathematical pillars (TGG, CCI, GFF, GRIP) in their *dual-layer* form, with seven Structure Theorems per pillar (5 algebraic + 2 physical) and the nine dual-layer sub-theorems **T1.1–T4.2** (with the bridge **T5**). At the methodological level, it has made explicit the closed-to-open foundational shift that distinguishes GG-Theory from the closed-system frameworks of the past century. At the structural level, it has *reduced* the residual softness in the N=3 proof to the named open problem M6, and provided the rigorous foundation for the 28-parameter codebook

of GG-A9-SM and the related results. At the empirical level, it has produced a coherent set of laboratory-testable falsifiers — modified TOV in compact stars (CCI-ST7), equipartition + fluctuation–dissipation in any boundary observable (GFF-ST7), and the generation-ordered weight ladder (GRIP-ST7) — that together provide independent experimental pathways to confirm or refute the framework.

These are technical accomplishments. We close with broader reflections on what they imply.

60.2 The dual-layer revolution: from to

The conceptual upgrade from to of this paper is the discovery that every Gi-4 Pillar admits a *dual-layer* structure — an algebraic block (axioms A1–A5) and a physical block (axioms A6–A7) — linked by a named compatibility theorem. This is not a cosmetic reformulation: it is the recognition that the four pillars do not just live in the bulk geometry but *descend operationally* to the laboratory through specific, named, derivable mechanisms.

In, each pillar was treated as a single algebraic structure with a single principal theorem. In, each pillar is a septuple, with paired theorems T*.1 and T*.2 covering the algebraic and physical sides respectively. This split has three deep consequences:

(a) Empirical accountability. Every algebraic prediction of (e.g., $N = 3$, OGN-5 ledger, Morse-positivity, Bl-6 budget) is now paired with an empirical falsifier (modified TOV, the generation ordering, equipartition, etc.). No prediction is left without an operational pathway to test.

(b) Universality across sciences. The dual-layer architecture is the cohort’s account of why the central concepts of GG-Theory appear across the natural sciences, and it is carried at the grade of the wager the introduction fences. The Master Conservation Law of CCI-ST6 is proposed as the parent of the conservation laws of chemistry, biology, and information theory. The free-energy minimization of GFF-ST6 is proposed as the parent of the relaxation processes from protein folding to neural predictive coding — each of which descends a Lyapunov function by Conley’s theorem in any case, the claim being that the functions are one. The cascade structure of GRIP-ST1 underlies every irreversible discrete relaxation in chemistry, condensed matter, and machine learning. These are not analogies; they are boundary projections of the same upstream geometric structure.

(c) Conceptual unification at the architectural level. Before, the four pillars were related by an implication chain ($TGG \Rightarrow CCI \Rightarrow GFF \Rightarrow GRIP$) but had no shared internal architecture. In, all four pillars share the *same* dual-layer architecture (5+2 axioms, paired theorems, named compatibility link), making the Gi-4 system a coherent unified mathematical object rather than a sequence of distinct constructions.

This is the precise sense in which is not just an improvement of but a structural upgrade: it identifies the dual-layer architecture as the universal organizing principle of the entire Gi-4 framework. Future papers in the GG-Theory program will inherit this organizing principle as a baseline.

60.3 Historical perspective: the four scientific revolutions

The history of theoretical physics, viewed at sufficient distance, is a history of foundational shifts. Each of the great syntheses of the past two and a half millennia identified a previously hidden foundational principle and used it to unify a previously unrelated cluster of phenomena. We may

identify four such “scientific revolutions”:

First scientific revolution (Euclid, c. 300 BCE). Euclid’s *Elements* introduced the axiomatic method to geometry and laid down the architectural template for all mathematical reasoning to follow: a small set of axioms (the five postulates and the common notions), explicit definitions, and logical derivation of theorems. This was the first instance of a *constitutional* approach to knowledge: a fixed set of foundational principles from which all consequences must be derivable by transparent steps. All subsequent revolutions inherit Euclid’s template; the present paper is consciously written in this Euclidean style.

Second scientific revolution (Newton, 1687). Newton’s *Principia* unified celestial and terrestrial mechanics under the principle of universal gravitation, with the calculus as the mathematical language. This was the first foundational identification of a single physical principle from which a vast empirical domain could be derived.

Third scientific revolution (Maxwell, Einstein, and Quantum Mechanics, c. 1865–1932). Three convergent foundational shifts in the long nineteenth and early twentieth centuries reshaped fundamental physics: **Maxwell (1865)** unified electricity, magnetism, and optics under his four equations, with vector calculus as the mathematical language. **Einstein (1905, 1915)** unified geometry and gravitation under the principle of equivalence, with tensor calculus as the mathematical language. **Heisenberg, Schrödinger, Dirac, von Neumann (1925–1932)** unified atomic spectra, chemistry, and quantum statistics under the principle of unitary evolution on Hilbert space, with operator algebra as the mathematical language. **The Standard Model (1960s–1970s)** extended these into a non-Abelian gauge theory of the strong, weak, and electromagnetic forces, with Yang–Mills theory as the mathematical language.

In each case, a foundational principle was identified, and a mathematical language was developed that allowed the principle to produce specific quantitative predictions. The principle was not self-evident before its identification; the mathematical language was not pre-existing. The two were created together, in tension, by intensive collaborative work over years or decades.

Fourth scientific revolution (GG-Theory, 2024–). **GG-Theory** is, in this lineage, the proposed fourth scientific revolution. Its foundational principle is OGGE_P, the operational/open geometry–gauge equivalence: physical observation is the irreversible bulk-to-boundary projection of a six-dimensional bulk onto a four-dimensional boundary. Its mathematical language is the combination of Gi-4 (geometric backbone) and Qi-4 (information backbone), with differential cohomology, Morse theory, and Hopkins–Singer descent as the technical tools. Its dual-layer architecture (algebraic + physical, paired theorems, named compatibility links) is the universal organizing principle of the framework, and is the principal structural innovation of the present paper.

The present paper formalizes Gi-4 as the geometric half of this language. When combined with Qi-4 (in preparation), with the Constitution, and with BI-6, the full mathematical vocabulary of GDOS will be in place.

60.4 The methodological lesson

The methodological lesson of GG-Theory, articulated already in Arisaka (2026, P2-N3, Section 12) and made formally explicit in the present paper, is that closed-system frameworks — however mathematically elaborate — cannot produce sufficient conditions on the structural integers and hierarchies of physics. Closed-system QFT can constrain N and $M_{\text{lock}}/E_{\text{Pl}}$ and the QCD scale and the electroweak vacuum, but it cannot derive any of them. Each requires an additional input, set by hand.

Open-system frameworks like GDOS are different. Because they include irreversible bulk-to-boundary projection as a foundational ingredient, they have access to cascade closure, finite instanton-action budgets, and Morse-attractor selection. These are the structural ingredients that produce sufficient conditions: $N \leq 3$ from cascade closure, $M_{\text{lock}}/E_{\text{Pl}} = 1/(2e^{35})$ from scale closure, the twelve charged-fermion masses from cascade depths.

The closed-to-open foundational shift is therefore not a stylistic choice; it is the structural prerequisite for deriving the integers and hierarchies of the Standard Model from first principles. This is the deepest lesson of the GG-Theory program, and the present paper provides the formal vocabulary in which the lesson can be stated rigorously.

60.5 The unity of physical reality, the role of mathematics, and the road ahead

Indeed, the deepest payoff of the Gi-4 reorganization is the structural unity it makes visible. The same four-pillar architecture — master-object, conservation-ledger, free-energy landscape, irreversible cascade — recurs across the magnificent edifice of physics: from particle physics ($N_g = 3$ generations from GRIP cascade closure) to gravitation (modified TOV from CCI-ST7) to chemistry and condensed matter (free-energy minimization from GFF) to neuroscience and biology (the free-energy principle as a boundary projection of OGGE; see GG-A7-Qi4 for the full Qi-4 boundary information backbone). Mathematics serves this unity not as ornamental language but as structural skeleton: the bundle structure, the Morse-theoretic landscape, the Diophantine ledger, and the Bogomolnyi-saturating cascade are all rigid algebraic facts that, once written down, leave no room for adjustable parameters. After all, the Gi-4 Pillars are not a new theory imposed on physics from outside; they are the geometric reading of what physics has been telling us, in different rooms of the same building, all along. Looking forward, the predictions cataloged in §56 and the open problems of §57, especially M9 (the canonicity of the complexity class) and M11 in its spectral-flow form (the general monodromy statement, its residues reduced to one identification and one convention), are the concrete agenda for the next phase of the program. The framework stands or falls on the modified-TOV signature, on the generation-ordering signature, and on the empirical reach of the 28 SM parameter predictions of GG-A9-SM. The next century of empirical work will decide.

Acknowledgments

This work was developed through extensive iterative collaboration with three large language models: ChatGPT (OpenAI) for exploratory drafting, narrative development, and broad-scope ideation; Perplexity for literature synthesis, citation verification, and recent experimental cross-referencing; and Claude (Anthropic) for structural organization, mathematical consistency checking, LaTeX typesetting, and final manuscript preparation. The sequential refinement of the GG-Theory framework (GG-Constitution / GG-6 / BI-6 / GDOS / Gi-4 Pillars / Qi-4 Pillars) across hundreds of conversation sessions with these three systems was essential to the coherence of the present paper.

The author thanks Aaron Blaisdell and Hirotaka Iwasaki for stimulating discussions on the program's interdisciplinary scope and for critical feedback on the accessibility of the presentation. The author also thanks hundreds of collaborators in the Elegant Mind Club over the past decade, including Javier Carmona, Elizabeth Milles, Mingda He, Syed Bashir Hydari, Brian Ta, Chris Dao, Umaima Afifa, Isabella Bestanoby, Caominh Le, and Natsuko Yamaguchi.

All scientific concepts (especially the foundation of the GG-Constitution, GGEP/GG-6/BI-6, and OGGE/ GDOS/Gi-4/Qi-4 Pillars), theoretical judgments, physical interpretations, and any remaining errors are the sole responsibility of the author.

Part X

Technical Appendices

Technical machinery, worked examples, and indices

The Appendices gather the technical machinery and the reference material that supports the main-text theorems of the Gi-4 pillars. Roughly speaking, this is where the reader who wants to verify a specific derivation, run a specific worked example, or trace a specific symbol will spend time — while the reader who wants only the axiomatic content and the principal theorems T1–T5 of the main text can stop after Part IX and treat the appendices as a self-contained companion volume. Appendices A and B provide the abbreviations and the glossary of key symbols. Appendix C unfolds the differential-cohomology preliminaries in full pedagogical detail. Appendix D supplies the Hopkins–Singer machinery used in the CCI closure. Appendix E provides Morse theory on stratified spaces. Appendix F collects the detailed worked examples for the four pillars. Appendix G is the pedagogical tour through the dual-layer architecture in plain language. Appendix H gives the step-by-step Diophantine analysis for OGN-5. Appendix I supplies the free-energy-principle bridge to life sciences and neuroscience. Appendix J gives the explicit derivation of the modified TOV equation. Appendices K and L provide the Hopkins–Singer descent and Morse-theory worked examples on $\mathbb{CP}^{N-1}$. Appendix M unfolds the generation weight ladder from cascade actions to laboratory observations. Appendix N is the life-science and neuroscience glossary. Finally, the inline bibliography with full Relevance annotations follows the appendices and closes the volume. Indeed, the nine main-text Parts plus fourteen appendices together realize the Gi-4 framework in full: Parts I–VI develop and synthesize the pillar theorems with proofs, Parts VII–IX treat applications, discussion, and conclusions, and the appendices supply the technical and bibliographic infrastructure that makes the whole document self-checking.

A Summary of Abbreviations

This appendix collects every abbreviation used in the present paper. All program-defined names are rendered in *sans-serif type* (GG-Theory, GG-Constitution, GG-6, GDOS, BI-6, OGN-5, QGL, Gi-4, Qi-4, GRIP, BFFT, etc.) so the reader can see at a glance which terms belong to the GG-Theory program proper. Standard physics abbreviations imported from the wider literature (KK, RS, QFT, TOV, etc.) are rendered in plain uppercase Roman, following the convention of the major physics journals. Mathematical objects are in roman or italic in the usual way. The third column gives the *primary link*: the paper in the Arisaka-GG series and the section where each program-defined name is first introduced in detail.

Table 21: Summary of abbreviations used in the present paper. Program-defined names are rendered in sans serif; standard physics abbreviations in plain uppercase Roman; mathematical objects in roman/italic in the usual way. The third column points at the paper or the section where each name is fixed, so an abbreviation met in the middle of an argument can be resolved without leaving the appendix.

Symbol	Meaning	Primary link
<i>Framework name and core program objects</i>		
GG-Theory	Grand Geometric Theory (also: Geometry-Gauge Equivalence Theory)	GG-A1-Constitution, Preface
GG-Constitution	The constitutional charter of the GG-Theory program: four axioms (AX-1–AX-4) and seven postulates (POS-1–POS-7)	GG-A1-Constitution, §2–§3
GG-6	Geometry-Gauge locking in six dimensions: the 6D bulk	GG-A2-GG6, §8 (T1-GG6)
GDOS	Geometric Dynamics of Open Systems: the boundary law implied by OGGEF	GG-A5-GDOS
BI-6	Green–Schwarz-modified Bianchi identity at 6D	GG-A3-BI6, §3
OGN-5	Occam-Gauge-Normalization 5-integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$	GG-A3-BI6, §7; T6-GG6
QGL	Quadruple Gauge Locking: the four-factor bulk gauge group $\text{Spin}(1, 5) \times \text{SU}(5) \times \text{U}(1)_B \times \text{U}(1)_{Qi}$	GG-A2-GG6, §10 (T3-GG6)
Gi-4	Geometric-Invariance four-pillar structure (TGG, CCI, GFF, GRIP)	this paper
Qi-4	Quantum-Information four-pillar structure (Qi-QUA, Qi-EP, Qi-CON, Qi-HAL)	GG-A7-Qi4
<i>Four axioms (AX-1–AX-4)</i>		
AX	Axiom (axiomatic primitive; non-negotiable)	GG-A1-Constitution, §1.5
LC	AX-1 Locality and Causality	GG-A1-Constitution, §2.2
GGEP	AX-2 Geometry–Gauge Equivalence Principle	GG-A1-Constitution, §2.3
OGGEF	AX-3 Operational/Open Geometry–Gauge Equivalence Principle (boundary observation dictionary)	GG-A1-Constitution, §2.4
MSMF	AX-4 Maximum Symmetry, Minimum Freedom (no hidden continuous knobs)	GG-A1-Constitution, §2.5
<i>Seven postulates (POS-1–POS-7)</i>		
POS	Postulate (admissibility filter; testable in principle)	GG-A1-Constitution, §1.5
GI	POS-1 Gauge invariance across scales (LGI / DGI / GGI)	GG-A1-Constitution, §3.3
UEP	POS-2 Unified Extremal Principle (PLA $\Leftrightarrow$ PLT / MEPP / FEP)	GG-A1-Constitution, §3.4
AF	POS-3 Anomaly Freedom and cancellability by inflow	GG-A1-Constitution, §3.5
CP	POS-4 Closure and Positivity (CPTP-admissible open dynamics)	GG-A1-Constitution, §3.6
RG	POS-5 Renormalization-Group covariance (scale covariance)	GG-A1-Constitution, §3.7
OGN	POS-6 Occam Gauge Normalization (canonical integer/rational normalization)	GG-A1-Constitution, §3.8
MF	POS-7 Measurement and Falsifiability (operational contract + kill criteria)	GG-A1-Constitution, §3.9

continued on next page

Table 21 continued from previous page

Symbol	Meaning	Primary link
LGI	Local Gauge Invariance (microscopic / closed-system layer of GI)	GG-A1-Constitution, §3.3
GGI	Global Gauge Invariance (macroscopic / ledger layer of GI)	GG-A1-Constitution, §3.3
<i>Gi-4 pillars</i>		
TGG	Total Geometric–Gauge object (Pillar I of Gi-4)	this paper
CCI	Covariant Conservation Identity (Pillar II of Gi-4)	this paper
GFF	Geometric Free-Energy Functional (Pillar III of Gi-4)	this paper
GRIP	Gauge-Runtime Invariance Pipeline (Pillar IV of Gi-4)	this paper
GEP	Geometry–Gauge Energy Principle (energetic block of GFF)	this paper
CEP	Curvature–Entropy Principle (entropic block of GFF)	this paper
DGI	Dynamical Gauge Invariance (initial stage of GRIP)	this paper
GSSB	Geometric Spontaneous Symmetry Breaking (middle stage of GRIP)	this paper
GA	Geometric Attractor (final stage of GRIP)	this paper
<i>Qi-4 pillars</i>		
Qi-QUA	Quantum information for QGL (Quadruple Gauge Locking) – Pillar I of Qi-4	GG-A7-Qi4
Qi-EP	Quantum information for the UEP (Unified Extremal Principle) – Pillar II of Qi-4	GG-A7-Qi4
Qi-CON	Quantum information for Conformal Invariance – Pillar III of Qi-4	GG-A7-Qi4
Qi-HAL	Quantum information for the Holographic Attractor Lattice – Pillar IV of Qi-4	GG-A7-Qi4
OH	Outer Hologram (boundary projection of bulk data; outward face of EIC)	GG-A7-Qi4
IH	Inner Hologram (inner self-modeling layer; inward face of EIC)	GG-A7-Qi4
EIC	Evolving Inner Cosmos (boundary completion of Qi-4 as the OH+IH composite)	GG-A7-Qi4
<i>Qi-HAL acronyms</i>		
MePMoS	Memory–Prediction–Motion Sensing	GG-A7-Qi4, Qi-HAL
NHT	Neural Hologram	GG-A7-Qi4, Qi-HAL
HAL	Neural Holographic Tomography	GG-A7-Qi4, Qi-HAL
Q-HAL	Quadratic / Cubic HAL	GG-A7-Qi4, Qi-HAL
BFFT	Binary Fourier / Fractal Fourier Transform	GG-A7-Qi4, Qi-HAL
<i>GDOS-internal acronyms (program-defined)</i>		
SIP	System Instance Packet (canonical GDOS input unit)	GG-A5-GDOS
SeOS	Selfish Open System	GG-A5-GDOS + GG-A7-Qi4

continued on next page

Table 21 continued from previous page

Symbol	Meaning	Primary link
<i>Program-defined mathematical symbols</i>		
Q_B	Quantum-information Bit – the bioenergetic Arisaka constant $Q_B \approx 0.1555$ eV per qubit; matches ATP hydrolysis at biological temperatures	GG-A2-GG6, §15 (T8-GG6)
<i>Standard physics abbreviations (plain Roman, journal convention)</i>		
<i>Open-system / quantum-channel formalism</i>		
CPTP	Completely Positive Trace-Preserving (channel admissibility for open evolution; AX-3/POS-4)	—
GKLS	Gorini–Kossakowski–Lindblad–Sudarshan (Gorini, Kossakowski & Sudarshan, 1976) (master equation for Markovian CPTP semigroups; POS-4)	—
GNS	Gelfand–Naimark–Segal (construction)	—
NZ	Nakajima–Zwanzig (memory-kernel projection for non-Markovian open dynamics)	—
POVM	Positive Operator-Valued Measure (generalized measurement model; AX-3/POS-7)	—
<i>Quantum and classical foundational theories</i>		
BRST	Becchi–Rouet–Stora–Tyutin (cohomology)	—
GR	General Relativity	—
QCD	Quantum Chromodynamics	—
QED	Quantum Electrodynamics	—
QFT	Quantum Field Theory	—
QM	Quantum Mechanics	—
<i>Standard Model and beyond</i>		
EFT	Effective Field Theory	—
EW	Electroweak	—
GUT	Grand Unified Theory	—
SM	Standard Model of particle physics	—
SSB	Spontaneous Symmetry Breaking (closed-system; generalized by GSSB)	—
SUSY	Supersymmetry	—
<i>Higher-dimensional and holographic constructions</i>		
AdS	Anti–de Sitter spacetime	—
AdS/CFT	Anti–de Sitter / Conformal Field Theory correspondence	—
CFT	Conformal Field Theory	—
KK	Kaluza–Klein (mode expansion / dimensional reduction)	—
RS	Randall–Sundrum (warped extra-dimension model)	—
<i>Quantum gravity and cosmology</i>		

continued on next page

Table 21 continued from previous page

Symbol	Meaning	Primary link
Λ CDM	Standard cosmological model with cosmological constant + cold dark matter	—
LQG	Loop Quantum Gravity	—
TOV	Tolman–Oppenheimer–Volkoff (equation)	—
<i>Foundations of quantum mechanics</i>		
CHSH	Clauser–Horne–Shimony–Holt (inequality)	—
EPR	Einstein–Podolsky–Rosen (paradox / argument)	—
<i>Open-system / extremal-principle sub-acronyms</i>		
FEP	Free Energy Principle (boundary inference avatar of UEP)	GG-A7-Qi4, Qi-EP
MEPP	Maximum Entropy Production Principle (heuristic boundary avatar of UEP)	GG-A7-Qi4, Qi-EP
PLA	Principle of Least Action (closed/bulk avatar of UEP; standard variational dynamics)	GG-A7-Qi4, Qi-EP
PLT	Principle of Least Time (operational/boundary avatar of UEP; fastest lawful path)	GG-A7-Qi4, Qi-EP

Figure 11: hierarchy of program-defined acronyms in the GG-Theory program. Top: the framework name (GG-Theory) and its constitutional charter (GG-Constitution). Tier 2: the four axioms (AX-1–AX-4) and seven postulates (POS-1–POS-7) that together define the admissibility predicate. Tier 3: three bulk objects forced by the constitution — GG-6 (6D bulk), BI-6 (modified Bianchi identity), OGN-5 (integer ledger). Tier 4: the boundary law GDOS that the bulk projects to via OGGEp, with input acronym SIP. Tier 5: the two four-pillar structures, Gi-4 (geometric flank, with sub-acronyms GEP/CEP blocks of GFF and DGI/GSSB/GA stages of GRIP) and Qi-4 (quantum-information flank, with Qi-EP sub-acronyms PLT/MEPP/FEP, Qi-CON sub-concepts Conformal/Lévy/Fractal/ $1/f$, Qi-HAL sub-acronyms MePMoS/NHT/HAL/BFFT, the boundary completion OH+IH=EIC, and the runtime carrier SeOS (Selfish Open System) underneath). Bottom: the bioenergetic Arisaka constant Q_B (derived secondary quantity), with the Arisaka equation $J_E/Q_B = J_e = J_B = J_{Q_i}$ that locks the four boundary currents through one conversion factor.




Figure 11: How the program’s own names are related to each other. The chart is a hierarchy, not a list: the charter sits at the top, the geometry it forces below it, the boundary law below that, and the derived pillars at the foot, so the position of a name on the chart tells the reader what it depends on. Names at the same level are siblings introduced by the same construction. The purpose is orientation rather than definition — the meanings are in the abbreviations table — and a reader who has lost track of which name belongs to which layer can recover it here in one glance.

B Glossary of Key Symbols

Table 22: Every mathematical symbol this paper uses repeatedly, with its meaning. The list is deliberately restricted to symbols that recur across Parts rather than every symbol that appears once: a reader meeting M_6 , I_r or K in a later Part should be able to recover it here without returning to the section that introduced it. Where a symbol carries both a bulk and a boundary reading, both are given, since one letter doing two jobs is the commonest source of confusion in the cohort’s notation.

Symbol	Meaning
N	The OGN-5 ledger’s second entry, $N = 3$. Every ledger identity is written with it
N_c	Color triplicity, the three colors of $SU(3)_C$ — what that entry counts
N_g	The generation (family) count, 3. <i>Not</i> a ledger entry; fixed by GG-P2-N3
N_f	The quark-flavor count of a beta function, $N_f = 2N_g = 6$ at high scale
$M_6, M_4, I_r, S_\theta^1$	6D bulk, 4D boundary, radial interval, compact angular direction
r, θ	Radial coordinate $\in [0, K]$, angular $\in [0, 2\pi)$
K	Dual-lock integer = 35
$g_{AB}, g_{\mu\nu}$	6D bulk metric, 4D boundary metric
R_6, R_4	6D Ricci scalar, 4D Ricci scalar
$M_{Pl}, M_6, M_{\text{lock}}, v$	Reduced Planck mass, 6D fundamental mass, KK lock scale, EW vev
Λ_4	4D cosmological constant
G_N	Newton’s gravitational constant
E_{Pl}, E_{GUT}	Planck and GUT energy scales
$\Phi(x, r, \theta)$	Generic bulk field
$\phi^{(n, \ell)}(x)$	4D KK-mode field (radial index n , angular index ℓ)
$f_n(r)$	Radial profile of n -th KK mode
$m_{n, \ell}^2$	KK mass squared = $m_n^2 + \ell^2$ in natural units
$\mathcal{H}_{\text{Bulk}}, \mathcal{H}_{\text{Bdy}}, \mathcal{H}_{\text{env}}$	Bulk, boundary, environment Hilbert spaces
$\rho_{\text{Bulk}}, \rho_{\text{Bdy}}$	Bulk, boundary density operators
$\mathcal{P}, \mathcal{Q}$	Nakajima–Zwanzig projection superoperators
$\mathcal{L}, \mathcal{L}_{\text{GDOS}}$	Liouville superoperator, GDOS generator
$H_{\text{Bulk}}, H_{\text{eff}}$	Bulk Hamiltonian, effective boundary Hamiltonian
L_α, γ_α	Jump operators and rates of GKLS
$\Phi_{t,s}$	Boundary CPTP evolution channel
$\mathcal{I}_\alpha, E_\alpha$	Instrument and POVM effect
$K(t, s), \mathcal{I}(t)$	NZ memory kernel and inhomogeneous initial-correlation term
$\mathcal{R}$	Boundary reduction map
$\mathcal{D}_{\text{OGGEP}}$	OGGEP dictionary = $\mathcal{R} \circ \text{Tr}_{\text{env}}$
$T^{\mu\nu}, T_{\text{TGG}}^{\mu\nu}$	Boundary stress tensor, unified TGG stress tensor
$\Delta T_{\mu\nu}^{\text{GS}}, \Delta T_{\mu\nu}^{\text{KK}}$	GS-tensor and KK-tower inflow
X_4, H_3	BI-6 anomaly polynomial and Green-Schwarz two-form
$\mathcal{A}_{\text{tot}}$	Composite bulk connection
G_{tot}	Bulk gauge-gravity group $\text{Spin}(1, 5) \times SU(5) \times U(1)_B \times U(1)_{Qi}$
Q^c, T^a, Y, B, Q_{Qi}	$SU(5)$ generators, hypercharge, baryon, informational charges
$F_{\mu\nu}$	Gauge field strength
s_6, s_4	6D and 4D BRST operators
$Z_4[J], W_4[J], \Gamma_4$	Generating functional, connected generating functional, effective action
$\text{Tr}_{\text{env}}, \text{Tr}_{\text{Bdy}}, \text{Tr}_{\text{Bulk}}$	Partial and full traces
$\boxed{\cdot}$	Boxed equation = key result
$\blacksquare$	End of proof

C Differential cohomology preliminaries (pedagogical)

Reading goal. This appendix is written for first-year graduate students in physics, theoretical life science, or neuroscience, who may have seen vector calculus on $\mathbb{R}^3$ but not algebraic topology. We build up from the divergence theorem of freshman calculus, through de Rham cohomology, to the differential cohomology refinement that is the backbone of the CCI closure of Part III and the Hopkins–Singer descent of Part IV. Every step is illustrated with explicit examples on small manifolds (the circle S^1 , the sphere S^2 , the torus T^2 , and the projective space $\mathbb{CP}^2$). No previous exposure to algebraic topology is assumed.

For a complete development see [Freed & Hopkins \(2000\)](#), [Hopkins & Singer \(2005\)](#); for a more pedestrian introduction to de Rham cohomology, see [Bott & Tu \(1982\)](#).

C.1 Why differential forms (a five-minute review)

Recall from vector calculus on $\mathbb{R}^3$ that:

- A scalar function $f(x, y, z)$ has a *gradient* ∇f , which is a vector field.
- A vector field $\mathbf{F}$ has a *curl* $\nabla \times \mathbf{F}$, which is also a vector field.
- A vector field $\mathbf{F}$ has a *divergence* $\nabla \cdot \mathbf{F}$, which is a scalar function.

The three operations satisfy $\nabla \times (\nabla f) = 0$ and $\nabla \cdot (\nabla \times \mathbf{F}) = 0$.

The language of *differential forms* unifies these operations. A scalar function f is called a *0-form*; a vector field written as $\mathbf{F} = (F_1, F_2, F_3)$ corresponds to a *1-form* $\omega^{(1)} = F_1 dx + F_2 dy + F_3 dz$; a vector field also corresponds to a *2-form* $\omega^{(2)} = F_1 dy \wedge dz + F_2 dz \wedge dx + F_3 dx \wedge dy$; and a scalar function corresponds to a *3-form* $\omega^{(3)} = f dx \wedge dy \wedge dz$.

In this language, ∇ , $\nabla \times$, and $\nabla \cdot$ are all the *same* operation, the *exterior derivative* d :

$$df = \omega_{\nabla f}^{(1)}, \quad d\omega^{(1)} = \omega_{\nabla \times \mathbf{F}}^{(2)}, \quad d\omega^{(2)} = \omega_{\nabla \cdot \mathbf{F}}^{(3)}. \quad (195)$$

The two identities $\nabla \times (\nabla f) = 0$ and $\nabla \cdot (\nabla \times \mathbf{F}) = 0$ become the single statement

$$\boxed{d^2 = 0}, \quad (196)$$

the *exterior calculus identity*. Equation (196) holds in arbitrary dimensions and is the central fact that motivates de Rham cohomology.

Remark C.1 (Physics intuition: why $d^2 = 0$?). $d^2 = 0$ is a consequence of the antisymmetry of mixed partial derivatives: $\partial_i \partial_j = \partial_j \partial_i$. When this symmetry is contracted against the antisymmetric wedge product $dx^i \wedge dx^j = -dx^j \wedge dx^i$, the result vanishes. Geometrically: the boundary of a boundary is empty ($\partial^2 = 0$), and d is dual to ∂ under Stokes’s theorem.

C.2 Ordinary cohomology and de Rham cohomology

For a smooth manifold X , the de Rham cohomology $H_{\text{dR}}^k(X; \mathbb{R})$ classifies closed k -forms modulo exact ones:

$$H_{\text{dR}}^k(X; \mathbb{R}) = \frac{\ker(d : \Omega^k(X) \rightarrow \Omega^{k+1}(X))}{\text{im}(d : \Omega^{k-1}(X) \rightarrow \Omega^k(X))}. \quad (197)$$

In words: a closed k -form ω ($d\omega = 0$) represents a non-trivial cohomology class iff it is *not* the exterior derivative of any $(k-1)$ -form. Two closed k -forms represent the same class iff they differ by an exact form.

By the de Rham theorem (de Rham, 1931), $H_{\text{dR}}^k(X; \mathbb{R})$ is naturally isomorphic to the singular cohomology $H^k(X; \mathbb{R})$ with real coefficients. This allows us to compute cohomology either by analysis (via differential forms) or by topology (via cell complexes).

Worked example 1: the circle S^1 . On S^1 parametrized by angle $\theta \in [0, 2\pi)$, the constant 1-form $d\theta$ is closed but not globally exact (it is locally $d\theta = d(\theta)$, but θ is not a globally defined function on S^1). Its integral $\int_{S^1} d\theta = 2\pi \neq 0$ detects the non-trivial class. Therefore $H^1(S^1; \mathbb{R}) \cong \mathbb{R}$, generated by $[d\theta/2\pi]$.

Worked example 2: the 2-sphere S^2 . On S^2 embedded in $\mathbb{R}^3$, the area 2-form $\omega = \sin \theta \, d\theta \wedge d\phi$ in spherical coordinates is closed and integrates to $\int_{S^2} \omega = 4\pi \neq 0$. Its class generates $H^2(S^2; \mathbb{R}) \cong \mathbb{R}$. By contrast, $H^1(S^2; \mathbb{R}) = 0$ (every closed 1-form on S^2 is exact, because every loop on S^2 contracts to a point).

Worked example 3: the 2-torus $T^2 = S^1 \times S^1$. The 1-forms $d\theta_1$ and $d\theta_2$ are both closed and non-exact, generating $H^1(T^2; \mathbb{R}) \cong \mathbb{R}^2$. Their wedge $d\theta_1 \wedge d\theta_2$ generates $H^2(T^2; \mathbb{R}) \cong \mathbb{R}$. This example shows how cohomology counts “independent loops/holes” in a space.

C.3 Integer cohomology and quantization

The integer cohomology $H^k(X; \mathbb{Z})$ is the singular cohomology with integer coefficients. The natural map $H^k(X; \mathbb{Z}) \rightarrow H^k(X; \mathbb{R})$ has kernel the torsion sub-group $\text{Tor}(H^k(X; \mathbb{Z}))$.

A closed differential form $\omega \in \Omega^k(X)$ is said to have *integer cohomology class* if its de Rham class lies in the image of $H^k(X; \mathbb{Z}) \rightarrow H^k(X; \mathbb{R})$. Equivalently, all periods $\int_C \omega$ over closed k -cycles $C \subset X$ are integers.

Why integrality matters in physics. Integrality is the mathematical statement of *quantization*. Examples:

- The Dirac monopole quantization $\int_{S^2} F/2\pi \in \mathbb{Z}$ for the magnetic-monopole field strength F (Dirac, 1931).
- The Aharonov–Bohm phase $\oint A/2\pi \in \mathbb{Z}$ for a Wilson loop around a flux tube (Aharonov & Bohm, 1959).

- The first Chern class $c_1 \in H^2(X; \mathbb{Z})$ of a complex line bundle, which classifies the topology of monopole-like configurations.
- The CCI integrality $[X_4] \in H^4(M_6; \mathbb{Z})$ of the present paper, which selects the OGN-5 ledger.

Worked example 4: integrality on the projective line $\mathbb{CP}^1$. The Fubini–Study form on $\mathbb{CP}^1 \cong S^2$ is

$$\omega_{\text{FS}} = \frac{i}{2\pi} \partial \bar{\partial} \log(1 + |z|^2) = \frac{1}{\pi(1 + |z|^2)^2} dx \wedge dy. \quad (198)$$

Its integral is $\int_{\mathbb{CP}^1} \omega_{\text{FS}} = 1$, an integer. This generates $H^2(\mathbb{CP}^1; \mathbb{Z}) \cong \mathbb{Z}$, with $[\omega_{\text{FS}}]$ the positive generator.

C.4 Differential cohomology as a refinement

Why a refinement is needed. Ordinary cohomology $H^k(X; \mathbb{Z})$ remembers only the integer *class* of a closed differential form, not the form itself. Two different forms representing the same class are indistinguishable. But in physics, we often need to remember *both*: the integer (topological) information and the specific connection (geometric) information. The differential cohomology $\check{H}^k(X; \mathbb{Z})$ is the precise refinement that remembers both simultaneously.

The triple structure. Differential cohomology $\check{H}^k(X; \mathbb{Z})$ is a refinement of ordinary cohomology that combines topological and differential-form data. An element $\check{x} \in \check{H}^k(X; \mathbb{Z})$ may be thought of as a triple

$$\check{x} \sim ([x], \omega, \eta), \quad (199)$$

where:

- $[x] \in H^k(X; \mathbb{Z})$ is the underlying integer cohomology class (the *topology* of $\check{x}$);
- $\omega \in \Omega_{\text{closed}}^k(X)$ is the curvature representative satisfying $[\omega] = [x] \otimes \mathbb{R}$ (the *geometry* of $\check{x}$);
- η is a transgression / connection datum that relates the two.

The precise definition is more technical; see [Hopkins & Singer \(2005, Section 2\)](#).

Physical example: a magnetic monopole. Consider a magnetic monopole at the origin of $\mathbb{R}^3$. The flux through a sphere S^2 enclosing the monopole is quantized: $\int_{S^2} F = 2\pi n$ for some integer n (the Dirac quantization). The integer n is the topology ($[F] \in H^2(S^2; \mathbb{Z}) \cong \mathbb{Z}$), and the explicit form F (e.g., $F = (n/2) \sin \theta d\theta \wedge d\phi$) is the geometry. The differential cohomology class $\check{F} \in \check{H}^2(S^2; \mathbb{Z})$ remembers both n and the specific form F — which is necessary to compute, e.g., the Aharonov–Bohm phase $\oint A$ for a Wilson loop on the boundary of S^2 .

Worked example 5: differential cohomology on S^1 . On S^1 , the differential cohomology $\check{H}^1(S^1; \mathbb{Z})$ classifies $U(1)$ connections (gauge potentials) up to gauge equivalence. An element $\check{A} \in \check{H}^1(S^1; \mathbb{Z})$ has:

- underlying integer class $\oint A/2\pi \in \mathbb{Z}$ (the winding number);

- curvature representative $dA = 0$ on S^1 (since $\dim S^1 = 1$, there is no F);
- Wilson-line phase $\eta = \exp(i \oint A) \in U(1)$.

This shows how differential cohomology naturally encodes both the quantized winding and the continuous Wilson-line phase — exactly what is needed for the boundary projector classes of GDOS.

C.5 The exact sequence

The differential cohomology fits into the four-term exact sequence

$$0 \rightarrow H^{k-1}(X; \mathbb{R}/\mathbb{Z}) \rightarrow \check{H}^k(X; \mathbb{Z}) \xrightarrow{\text{curv}} \Omega_{\text{closed}, \mathbb{Z}}^k(X) \rightarrow 0, \quad (200)$$

where:

- $\text{curv} : \check{H}^k \rightarrow \Omega_{\text{closed}, \mathbb{Z}}^k$ extracts the curvature representative;
- $\Omega_{\text{closed}, \mathbb{Z}}^k(X)$ is the space of closed k -forms with integer cohomology class;
- the kernel of curv is the “flat” sub-group, isomorphic to $H^{k-1}(X; \mathbb{R}/\mathbb{Z})$.

C.6 Pontryagin and Chern classes

The Pontryagin classes $p_i(E) \in H^{4i}(X; \mathbb{Z})$ of a real vector bundle $E \rightarrow X$ are integer cohomology classes constructed from the curvature of any connection on E via Chern–Weil theory:

$$p_1(E) = -\frac{1}{8\pi^2} \text{tr}(R \wedge R), \quad (201)$$

where R is the curvature of any connection on E . The class is independent of the connection choice up to exact forms.

The Chern classes $c_i(F) \in H^{2i}(X; \mathbb{Z})$ of a complex vector bundle $F \rightarrow X$ are similarly constructed:

$$c_1(F) = \frac{1}{2\pi i} \text{Tr}(F_F), \quad c_2(F) = \frac{1}{8\pi^2} [\text{Tr}(F_F \wedge F_F) - \text{Tr}(F_F) \wedge \text{Tr}(F_F)], \quad (202)$$

where F_F is the curvature of any connection on F .

The Pontryagin and Chern classes are integer cohomology classes that can be lifted to differential cohomology via the Chern–Weil map: $\check{p}_1 \in \check{H}^4(X; \mathbb{Z})$, etc.

C.7 Relevance to BI-6 and CCI

The four-form anomaly polynomial X_4 of BI-6 is

$$X_4 = \text{tr} R_{\text{Spin}(1,5)}^2 - \text{Tr} F_5^2 - \frac{2}{5} F_B^2 - \frac{2}{5} F_{Qi}^2, \quad (203)$$

which combines Pontryagin (gravity), Chern (gauge), and squared first Chern (Abelian) contributions. Its lift to differential cohomology $\check{x}_4 \in \check{H}^4(M_6; \mathbb{Z})$ is the foundational object for the Hopkins–Singer descent in Appendix B.

D Hopkins–Singer machinery for GDOS

This appendix sketches the Hopkins–Singer construction adapted to the 6D bulk of GG-Theory. The full mathematical content is in [Hopkins & Singer \(2005\)](#); we extract here only what is needed for the Hopkins–Singer descent of $\check{x}_4$ to the boundary moduli used in Theorem 27.2 of Part IV.

D.1 Differential function spaces

For a smooth manifold X and integer $k \geq 0$, the differential function space $\text{Func}^k(X; \mathbb{Z})$ is the space of continuous maps $X \rightarrow K(\mathbb{Z}, k)$ (an Eilenberg–MacLane space) equipped with a differential refinement. An element $f \in \text{Func}^k(X; \mathbb{Z})$ has underlying topological data (the homotopy class of the continuous map) plus differential data (a closed k -form representative of the class).

Differential cohomology is then defined as the homotopy classes of differential functions:

$$\check{H}^k(X; \mathbb{Z}) \cong \pi_0 \text{Func}^k(X; \mathbb{Z}). \quad (204)$$

D.2 The descent map

For a smooth fibration $\pi : X \rightarrow Y$ with closed fibers of dimension n , the descent map for differential cohomology is

$$\tilde{\pi}_! : \check{H}^k(X; \mathbb{Z}) \rightarrow \check{H}^{k-n}(Y; \mathbb{Z}), \quad (205)$$

which integrates a differential cohomology class along the fibers.

For the 6D bulk $\pi : M_6 \rightarrow \mathcal{M}_{\text{bdy}}$ projecting onto the boundary moduli, the descent map sends $\check{x}_4 \in \check{H}^4(M_6; \mathbb{Z})$ to a differential cohomology class on $\mathcal{M}_{\text{bdy}}$ (or, restricted to strata, on each stratum). When the fiber dimension is 0 (i.e., when restricting to a section of the projection), the descent preserves cohomological degree.

D.3 The η -invariant pairing

The fundamental non-degeneracy property of the Hopkins–Singer descent is the η -invariant pairing, which is a refinement of the intersection pairing on integer cohomology. For two differential cohomology classes $\check{x} \in \check{H}^k(X; \mathbb{Z})$ and $\check{y} \in \check{H}^{n-k+1}(X; \mathbb{Z})$ on a closed oriented manifold X of dimension n , the η -invariant pairing is

$$\langle \check{x}, \check{y} \rangle_\eta \in \mathbb{R}/\mathbb{Z}. \quad (206)$$

This pairing is non-degenerate in the sense that $\check{x} = 0$ iff $\langle \check{x}, \check{y} \rangle_\eta = 0$ for all $\check{y}$.

D.4 Application to BI-6 anomaly inflow

The Hopkins–Singer descent applied to the CCI class $\check{x}_4 \in \check{H}^4(M_6; \mathbb{Z})$ gives a class $\check{\rho}_{\text{bdy}}(\check{x}_4)$ on the boundary moduli $\mathcal{M}_{\text{bdy}}$, which restricts to each stratum $\Sigma \subset \mathcal{M}_{\text{bdy}}$ as a class in $\check{H}^4(\Sigma; \mathbb{Z})$. By Theorem 27.1 of Part IV, this restriction is $c_R \cdot \check{p}_1(\Sigma)$, where $c_R \in \mathbb{Q}_{>0}$ is the rational Killing-form coefficient.

The non-degeneracy theorem of $\mathcal{H}_2$ (Theorem 27.2) follows from the η -invariant non-degeneracy: if $\check{x}_4 \neq 0$, then $\check{\rho}_{\text{bdy}}(\check{x}_4) \neq 0$ on any stratum supporting a non-trivial pairing.

D.5 Forward to integral L-genus

For a closed oriented $4m$ -manifold Σ , the integrated L-genus

$$\sigma(\Sigma) = \int_{\Sigma} L_m(p_1, \dots, p_m) = \int_{\Sigma} L_m(c_R^{-1} \check{\rho}_{\text{bdy}}(\check{x}_4), \dots) \quad (207)$$

is non-zero by the combination of Hirzebruch's signature theorem and the Hopkins–Singer non-degeneracy. This is the technical content underlying the strict Morse-positivity of GFF via Theorem T3.

E Morse theory on stratified spaces

This appendix collects the basic facts about Morse theory used in Parts IV and V. The standard references are Milnor (1963), Schwarz (1993) for smooth Morse theory and Goresky & MacPherson (1988) for the stratified case.

E.1 Standard Morse theory

Let M be a closed smooth n -manifold and $f : M \rightarrow \mathbb{R}$ a smooth function. f is a *Morse function* if:

- all critical points $\{c \in M : df(c) = 0\}$ are non-degenerate, i.e., the Hessian $\text{Hess}(f)(c)$ is non-singular;
- the Morse index $\text{ind}(c) \in \{0, 1, \dots, n\}$ is the number of negative eigenvalues of $\text{Hess}(f)(c)$.

The *Morse polynomial* of f is

$$\mathcal{P}_f(t) = \sum_{c \in \text{Crit}(f)} t^{\text{ind}(c)}. \quad (208)$$

The Morse inequalities relate $\mathcal{P}_f(t)$ to the Poincaré polynomial $\mathcal{P}_M(t)$ of M : $\mathcal{P}_f(t) \geq \mathcal{P}_M(t)$ in the sense of polynomial coefficients, with equality for “perfect” Morse functions.

E.2 The height function on $\mathbb{CP}^{N-1}$

The height function $f([\psi]) = \sum_a \lambda_a |\psi_a|^2 / \sum |\psi_a|^2$ on $\mathbb{CP}^{N-1}$ is a perfect Morse function with critical points $c_a = [\delta_{ab}]$ and indices $\text{ind}(c_a) = 2(a-1)$. The Morse polynomial is

$$\mathcal{P}_f^{\mathbb{CP}^{N-1}}(t) = 1 + t^2 + t^4 + \dots + t^{2(N-1)}, \quad (209)$$

which agrees with the Poincaré polynomial of $\mathbb{CP}^{N-1}$.

E.3 Stratified Morse theory

A *Whitney stratified space* is a topological space Z with a decomposition $Z = \bigsqcup_a \Sigma_a$ into smooth manifolds (strata) satisfying Whitney’s regularity conditions (Goresky & MacPherson, 1988). A *stratified Morse function* $f : Z \rightarrow \mathbb{R}$ is one whose restriction $f|_{\Sigma_a}$ is a Morse function on each stratum, with appropriate compatibility at the strata boundaries.

The Morse polynomial of a stratified Morse function is the sum over all critical points across all strata. For the boundary moduli $\mathcal{M}_{\text{bdy}} = \prod_a \Sigma_a$ (a product of strata), the total Morse polynomial factorizes:

$$\mathcal{P}_f^{\mathcal{M}_{\text{bdy}}}(t) = \prod_a \mathcal{P}_{f|_{\Sigma_a}}^{\Sigma_a}(t). \quad (210)$$

E.4 Connecting trajectories and instanton actions

Between two critical points c_+, c_- with $\text{ind}(c_+) - \text{ind}(c_-) = 1$ in the smooth case (or $= 2$ in the Kähler-Hamiltonian case), the gradient flow admits a moduli of *connecting trajectories* from c_+ to c_- . The action of each connecting trajectory is

$$S = \int_{\gamma} |\nabla f| \, ds = f(c_+) - f(c_-), \quad (211)$$

saturating the Bogomolnyi bound.

For the height function on $\mathbb{CP}^{N-1}$ with the Fubini–Study metric, the action between adjacent critical points is πk where the index drops from $2k$ to $2(k-1)$. This is the result used in Theorem 32.1 of Part V.

F Cross-paper coherence map

The Gi-4 framework of the present paper does not stand alone: it imports its constitutional inputs from Arisaka (2026, A1-Constitution), its bulk geometry from Arisaka (2026, A2-GG6), its anomaly closure from Arisaka (2026, A3-BI6), and its boundary-law projector structure from Arisaka (2026, A5-GDOS); in turn it exports its four pillars to the downstream papers in the program, principally Arisaka (2026, A7-Qi4) (the parallel Qi-4 pillars), Arisaka (2026, A9-SM) (the 28-parameter codebook), and Arisaka (2026, A10-Cosmos) (boundary cosmology). This section makes the imports and exports explicit by naming, for each cross-paper invocation, the specific theorem number or section location where the referenced content lives.

Note on version and Phase numbers. Companion-paper versions are independent revision tracks. Each row of the coherence tables below records the version of the *cited companion paper* as of the date on the title page of the present paper (May 2026). When a companion paper is later revised, the rows below should be re-walked against the revised paper’s actual internal theorem numbering before the present paper is itself revised; vague pointers (e.g., “the central trace-ratio theorem of BI-6”) are not acceptable, and every entry below names a specific theorem number, section, or appendix location.

Audit trail. The internal theorem numbering of the cited companion papers, used to populate the tables below, is as follows:

- **GG-A1-Constitution.** Four Axioms AX-1–AX-4 (LC, GGEP, OGGE, MSMF) in Sections 3–6; seven Postulates POS-1–POS-7 (GI, UEP, AF, CP, RG, OGN, MF) in Sections 7–13; eight central theorems T1–T8 in the body; the Arisaka equation $J_E/Q_B = J_e = J_B = J_{Qi}$ as central boundary identity.
- **GG-A2-GG6.** Eight principal theorems including Theorem 1 dimensional minimality ($d_{\text{bulk}} = 6$); Theorem 3 canonical 6D metric $ds_6^2 = e^{-2r}\eta_{\mu\nu}dx^\mu dx^\nu + dr^2 + d\theta^2$; Theorem 4 four-component BI-6 coefficient vector $(a, b, c, d) = (1, -1, -2/5, -2/5)$; Theorem 6 OGN-5 Diophantine cascade; Theorem 8 Arisaka equation and Quadruple Gauge Locking.
- **GG-A3-BI6.** Theorem 9.1 OGN-5 trace-ratio lock $|a/c| = |a/d| = 5/2$; Theorem 4.1 global anomaly (Witten, 1985) freedom; Section 3 Chern–Simons construction of H_3 from TGG data; § 8.7 “BI-6 as the GGEP completion of GS: the geometrization reading” (the supportive treatment of the Green–Schwarz legacy invoked in the present paper).
- **GG-A4-Euclid.** Historical chain from Euclid through Newton, Riemann, Maxwell, Einstein to GG-Theory; parameter-audit framework recording every dimensional constant as a conversion factor and every dimensionless integer as a Diophantine output.
- **GG-A5-GDOS.** GDOS boundary law $\mathcal{D}_{\text{OGGE}} = \mathcal{R} \circ \text{Tr}_{\text{env}}$ (CPTP form); the SIP construction; the OH+IH=EIC framework on the boundary; the OGGE-projection from M_6 to M_4 as the foundational dynamical move.
- **GG-A7-Qi4, GG-A8-QM, GG-A9-SM, GG-A10-Cosmos, GG-P1-KK, GG-C1-Kaxion** consume the GRIP cascade downstream — the Qi-4 pillars, the quantum-measurement reading, the Standard-Model parameter block, the cosmological descent, the collider phenomenology, and the dark-matter sector respectively; cross-citations to them name registered theorem identifiers.

F.1 Inward map: imports into GG-A6-GRIP from the foundation block

The following longtable lists, for each theorem-level fact that Gi-4 imports from a foundation-block paper, the specific companion-paper location and the role the imported fact plays in the present manuscript.

Table 23: Inward coherence map: GG-A6-GRIP imports from A1, A2, A3, A5. Each row names one imported fact, the exact place in the companion paper where it is established, and what this paper does with it. The point of the table is auditability: every fact taken on trust from elsewhere appears here once, so a reader can check that nothing is borrowed and then cited back as support for its own source.

Imported fact	Companion-paper location	Role in GG-A6-GRIP
Four Axioms AX-1–AX-4 (LC, GGEP, OGGE, MSMF)	GG-A1-Constitution, Sections 3–6 (Arisaka, 2026, A1-Constitution)	Constitutional source of TGG (GGEP), CCI (AF+GI), GFF (MSMF+OGGE), and GRIP (OGGE+RG).

continued on next page

(continued from previous page)

Imported fact	Companion-paper location	Role in GG-A6-GRIP
Seven Postulates POS-1–POS-7 (GI, UEP, AF, CP, RG, OGN, MF)	GG-A1-Constitution, Sections 7–13 (Arisaka, 2026, A1-Constitution)	Constitutional source of the Ward package of TGG-ST7 , the integrality of $[X_4]$ via AF, and the cascade structure via RG.
Arisaka equation $J_E/Q_B = J_e = J_B = J_{Q_i}$	GG-A1-Constitution, central boundary identity (Arisaka, 2026, A1-Constitution)	Establishes the four-component boundary current structure that the Master Conservation Law of T1.3 enforces; underlies the L12 “empirically preferred” framing of $U(1)_{Q_i}$.
6D bulk metric $ds_6^2 = e^{-2r} \eta_{\mu\nu} dx^\mu dx^\nu + dr^2 + d\theta^2$	GG-A2-GG6, Theorem 3 (canonical metric) (Arisaka, 2026, A2-GG6)	Bulk arena on which the composite connection $\mathcal{A}_{\text{tot}}$ of T1.1 is defined.
$d_{\text{bulk}} = 6$ from the magic-dimension ladder; $\text{Spin}(1, 5) \cong \text{SL}(2, \mathbb{H})$	GG-A2-GG6, Theorem 1 (dimensional minimality) (Arisaka, 2026, A2-GG6)	Fixes the bulk dimension; gates the existence of the quaternionic Lorentz double cover used in the Section on TGG and the magic-dimension ladder of A6.
Four-component BI-6 coefficient vector $(a, b, c, d) = (1, -1, -2/5, -2/5)$	GG-A2-GG6, Theorem 4; GG-A3-BI6, Theorem 9.1 (trace-ratio lock) (Arisaka, 2026, A2-GG6; Arisaka, 2026, A3-BI6)	Input to the algebraic content of CCI; the unique solution of MSMF closure used in T2.1.
OGN-5 Diophantine cascade $(M, N, L; J, K) = (1, 3, 5; 7, 35)$	GG-A2-GG6, Theorem 6 (Diophantine cascade); GG-A3-BI6, Theorem 10.2 (Arisaka, 2026, A2-GG6; Arisaka, 2026, A3-BI6)	Direct input to T2.1; sets the lower bound $N \geq 3$ used in T4.1.
BI-6 closure $dH_3 = X_4$	GG-A3-BI6, Section 6 (Chern–Simons/Chern–Weil construction); Theorem 6.3 (Arisaka, 2026, A3-BI6)	Bulk-side closure that the algebraic content of CCI (T2.1) makes operational.
Hopkins–Singer integrality of $[X_4] \in H^4(M_6; \mathbb{Z})$	GG-A3-BI6, Section 3, with reference to Hopkins & Singer (2005)	Cohomological backbone for the Morse-positivity argument of T3.1.
GS-as-GGEP “geometrization reading”	GG-A3-BI6, § 8.7 (Arisaka, 2026, A3-BI6)	Supportive framing of the Green–Schwarz legacy invoked in the present paper, throughout the intro and the bibliography Relevance paragraph for Green–Schwarz1984.
Global anomaly freedom of BI-6	GG-A3-BI6, Theorem 4.1 (Arisaka, 2026, A3-BI6)	Justifies that the four-form X_4 on M_6 admits a globally well-defined Chern–Simons primitive, which is the algebraic content of CCI-ST3 .
GDOS boundary law $\mathcal{D}_{\text{OGGEP}} = \mathcal{R} \circ \text{Tr}_{\text{env}}$ as CPTP	GG-A5-GDOS, throughout (Arisaka, 2026, A5-GDOS)	Boundary-projection construction that underlies the operational layer T3.2 (GFF thermodynamic monotonicity) and T4.2 (GRIP runtime gauge-invariance).
SIP (System Instance Packet) construction	GG-A5-GDOS, throughout (Arisaka, 2026, A5-GDOS)	Provides the local-system bookkeeping required for the boundary stress-tensor decomposition $\mathcal{T}_{\text{TGG}}^{\mu\nu} = T_{\text{matter}} + T_{\text{gauge}} + X_{\text{GS}} + T_{\text{KK}}$ of T1.2.

(continued on next page)

(continued from previous page)

Imported fact	Companion-paper location	Role in GG-A6-GRIP
OGGEP-projection from M_6 to M_4	GG-A5-GDOS, throughout (Arisaka, 2026, A5-GDOS)	The KK zero-mode reduction that takes the bulk action of TGG-ST2 to the boundary action of T1.2.

F.2 Outward map: exports from **GG-A6-GRIP** to downstream papers

The following longtable lists, for each theorem-level fact that Gi-4 exports to a downstream paper, the present-paper theorem label or section location and the downstream paper’s anticipated use.

Table 24: Outward coherence map: GG-A6-GRIP exports to A7, A8, A9, A10, P1, C1. Each row names one result established here, the place in this paper where it is proved, and the downstream paper that consumes it and for what. Read together with the inward map, the pair states this paper’s entire dependency surface in both directions.

Exported GG-A6-GRIP fact	A6 location	Downstream paper and anticipated use
TGG existence and uniqueness on M_6 (T1.1, T1.2)	Theorem 10.1 (T1)	GG-A7-Qi4: parallel Qi-4 pillars Qi-QUA, Qi-EP defined on the same composite-connection structure; GG-A8-QM: TGG stress tensor as the geometric input to the 4D measurement framework; GG-A9-SM: gauge content of $\mathcal{A}_{\text{tot}}$ used to identify the 28-parameter codebook.
Master Conservation Law $\nabla_\mu \mathcal{T}_{\text{TGG}}^{\mu\nu} = 0$ (T1.3)	Theorem 10.1 (T1.3)	GG-A9-SM: boundary stress-tensor conservation as the underlying gauge-invariance identity; GG-A10-Cosmos: bulk-to-boundary stress flow as the geometric origin of cosmological energy-momentum balance.
Compatibility proposition $\mathfrak{C}_{\text{TGG}}$	Section 10.1, eq. (45)	GG-A7-Qi4: parallel compatibility proposition $\mathfrak{C}_{\text{Qi}}$ for the four Qi-pillars; GG-A8-QM: dual-layer template for the algebraic-vs-operational structure of measurement.
OGN-5 uniqueness from MSMF closure (T2.1)	Theorem 18.1 (T2.1)	GG-A9-SM: integer ledger (1, 3, 5; 7, 35) as the input to the closed-form derivation of the 28 Standard-Model parameters.
Master Conservation iff Ward + PBI-4 (T2.2)	Theorem 18.4 (T2.2)	GG-A9-SM: four-current Ward-identity package as the QFT-level conservation discipline; GG-A10-Cosmos: boundary stress-tensor conservation under the 6D KK reduction.
Compatibility proposition $\mathfrak{C}_{\text{CCI}}$	Section 18.1, eq. (103)	GG-A7-Qi4: parallel compatibility proposition for Qi-CON information-conservation.

continued on next page

(continued from previous page)

Exported GG-A6-GRIP fact	A6 location	Downstream paper and anticipated use
GFF Morse-positivity from CCI via Hopkins–Singer descent (T3.1)	Theorem 26.1 (T3.1)	GG-A7-Qi4: variational backbone for Qi-CON (information-theoretic CCI analog); GG-A9-SM: Morse-cascade structure on lepton- and quark-mass moduli used in the seesaw-mechanism derivation.
Thermodynamic monotonicity from GDOS (T3.2)	Theorem 26.3 (T3.2)	GG-A10-Cosmos: boundary-entropy production as the origin of the cosmological arrow of time; GG-A7-Qi4: Qi-EP entropy law for the information-pillar dynamics.
Compatibility proposition $\mathfrak{C}_{\text{GFF}}$	Section 26.1, eq. (144)	GG-A7-Qi4: parallel compatibility proposition for the algebraic-vs-thermodynamic content of Qi-EP.
GRIP cascade closure forces $N \leq 3$ (T4.1, combined with T2 lower bound)	Theorem 34.1 (T4.1)	GG-A9-SM: hard upper bound $N_g = 3$ for the family count; GG-P1-KK: cascade structure on the family-projector moduli used to predict the ~ 3.85 TeV KK resonance cluster.
Runtime gauge-invariance iff PBI-4 compensator extends (T4.2)	Theorem 34.5 (T4.2)	GG-A7-Qi4: Qi-HAL family (MePMoS, NHT, HAL, BFFT) as the information-side analog of GRIP runtime invariance; GG-P1-KK: smooth-extension hypothesis as the technical condition on the KK lock cluster prediction; GG-C1-Kaxion: cascade-runtime structure on the $U(1)_{Q_i}$ sector.
Compatibility proposition $\mathfrak{C}_{\text{GRIP}}$	Section 34.1, eq. (176)	GG-A7-Qi4: parallel compatibility proposition for Qi-HAL.
Geometric First Law of GFF ($F = U - TS$ form)	Theorem 25.1	GG-A7-Qi4: free-energy ledger for the information-pillar dynamics; GG-A10-Cosmos: thermodynamic backbone of the boundary cosmology.
GFF-stratified Morse landscape on $\mathcal{M}_{\text{bdy}}$	Section 26 and Appendix	GG-A8-QM: critical-point catalog as the structural basis of the discrete measurement outcomes; GG-A9-SM: lepton- and quark-mass strata.
GRIP instanton-action ledger $S_{[2k] \rightarrow [2(k-1)]} = \pi k$	Theorem 32.1	GG-A9-SM: cascade-step structure for the seesaw mechanism; GG-P1-KK: cascade-step structure for the KK lock cluster; GG-C1-Kaxion: cascade-step structure for the Kaxion oscillation frequency.
Gi-4 / Qi-4 pairing framework	Part VI (Synthesis) and Section 41	GG-A7-Qi4: definitional template for the four information pillars Qi-QUA, Qi-EP, Qi-CON, Qi-HAL, each pillar paired with one of TGG, CCI, GFF, GRIP.

The two maps together establish the position of Gi-4 in the GG-Theory program: the constitutional content flows downward from A1, the bulk-side technical content flows in from A2 and A3, the

boundary-projection structure flows in from A5, and the four pillars then flow outward to every downstream paper that addresses the 4D-physics, cosmology, or predictions branches of the program. The map is stated against the theorem numbering the companion papers carry, and a reader following a citation into one of them should expect the numbering, not the content, to be what shifts.

G Global naming convention for cross-cohort citation

Beginning with [Arisaka \(2026, A2-GG6\)](#) and extended across the cohort, each load-bearing item of the GG-Theory program carries a globally unique identifier of the form *Prefix-Suffix-N*, where the Prefix names the item type and the Suffix names the paper:

- **T-Suffix-N** for the *N*th Theorem (e.g., T-GRIP-4 is Theorem T4 of the present paper, the GRIP cascade closure forcing $N_g \leq 3$).
- **L-Suffix-N** for the *N*th Lemma.
- **C-Suffix-N** for the *N*th Corollary.
- **D-Suffix-N** for the *N*th Definition.
- **R-Suffix-N** for the *N*th Remark.
- **P-Suffix-N** for the *N*th Prediction.
- **O-Suffix-N** for the *N*th Objection (e.g., O-GRIP-1 is the first Common Objection of the present paper).
- **F-Suffix-N** for the *N*th Falsifier (e.g., F-GRIP-4 is the generation gate-action-ordering falsifier).
- **I-Suffix-N** for the *N*th Innovation (e.g., I-GRIP-8 is the $N_g \leq 3$ family-count theorem above).
- **PS-Suffix-N** for the *N*th Problem-Solved entry (§53).

The canonical Suffix list for the twenty-two cohort papers is: *Constitution* (A1), *GG6* (A2), *BI6* (A3), *Euclid* (A4), *GDOS* (A5), *GRIP* (A6, this paper), *Qi4* (A7), *QM* (A8), *SM* (A9), *Cosmos* (A10), *Origin* (A11), together with the phenomenology preprints *KK* (P1), *N3* (P2), *Electron* (P3), the foundations preprints *Real* (P4) and *CPT* (P5), and the consequence preprints *Kaxion* (C1), *Hubble* (C2), *BH* (C3), and the white papers *Symmetry* (W1) and *Holonomy* (W2). Inside each paper the internal numbering (“Theorem T1”, “Theorem T2”, ...) continues to be supplied locally; the global identifier is the same item viewed from outside the paper, so that a cross-cohort citation such as “GG-A6-GRIP consumes T-GDOS-3 and produces T-SM-12” is unambiguous and search-friendly. The four axioms (AX-1 ... AX-4) and seven postulates (POS-1 ... POS-7) of [Arisaka \(2026, A1-Constitution\)](#) keep their existing cohort-wide names, since those are already paper-specific and unique.

H Detailed worked examples

This appendix collects detailed explicit computations that were sketched in the main text.

H.1 Spin(1,5) generators and structure constants in detail

The 15 generators of $\mathfrak{spin}(1,5)$ in the antisymmetric basis $M^{ab} = -M^{ba}$, $a, b \in \{0, 1, 2, 3, 4, 5\}$, satisfy the commutation relations

$$[M^{ab}, M^{cd}] = i(\eta^{ac}M^{bd} - \eta^{ad}M^{bc} - \eta^{bc}M^{ad} + \eta^{bd}M^{ac}), \quad (212)$$

with $\eta = \text{diag}(+, -, -, -, -, -)$. The 15 generators decompose as 10 rotations J^{ab} ($1 \leq a < b \leq 5$) plus 5 boosts $K^a = M^{0a}$ ($a = 1, \dots, 5$).

H.2 Quaternionic doublet matrix representation

The fundamental representation of $\text{SL}(2, \mathbb{H})$ on $\mathbb{H}^2$ is realized by 2×2 quaternionic matrices acting on column vectors $(q_1, q_2)^T$ with $q_i \in \mathbb{H}$. Writing $q = \alpha + j\beta$ with $\alpha, \beta \in \mathbb{C}$, one identifies $\mathbb{H}^2 \cong \mathbb{C}^4$ as a real vector space, and the $\text{SL}(2, \mathbb{H})$ action becomes a specific real representation of dimension 8 (matching the $2^{6/2-1} = 8$ -dimensional six-dimensional Dirac spinor).

H.3 CCI integrality check on $\mathcal{M}_6$

For the 6D bulk with OGN-5 ledger (1, 3, 5; 7, 35), the integrals of the four sub-curvature 4-forms over the fundamental 4-cycle C_4 are:

$$\frac{1}{8\pi^2} \int_{C_4} \text{tr } R_{\text{Spin}(1,5)}^2 = K = 35, \quad (213)$$

$$\frac{1}{8\pi^2} \int_{C_4} \text{Tr } F_5^2 = 2L = 10, \quad (214)$$

$$\frac{1}{4\pi^2} \int_{C_4} F_B^2 = M^2 = 1, \quad (215)$$

$$\frac{1}{4\pi^2} \int_{C_4} F_{Qi}^2 = N^2 = 9. \quad (216)$$

Combining with the **CCI-ST2** coefficients:

$$\frac{1}{8\pi^2} \int_{C_4} X_4 = 35 - 10 - \frac{2}{5} \cdot 1 - \frac{2}{5} \cdot 9 = 25 - \frac{20}{5} = 25 - 4 = 21, \quad (217)$$

which is an integer. This confirms the integrality $[X_4] \in H^4(M_6; \mathbb{Z})$ for the OGN-5 ledger.

H.4 GFF Morse data on $\mathbb{CP}^2$

For the family-projector stratum $\Sigma_{\text{fam}} = \mathbb{CP}^2$ with height function $f([\psi_1 : \psi_2 : \psi_3]) = (\lambda_1|\psi_1|^2 + \lambda_2|\psi_2|^2 + \lambda_3|\psi_3|^2) / \sum |\psi_a|^2$, $\lambda_1 < \lambda_2 < \lambda_3$:

Table 25: The Morse data of the complex projective plane, critical point by critical point, with the Standard-Model generation each one carries. There are three critical points, of index 0, 2 and 4, in the order the Morse function ranks them; the generation assignment follows that order and so does the Yukawa hierarchy, smallest to largest. This table is the arithmetic behind the claim that three generations are a consequence of the geometry rather than an input: the count is the number of critical points, and nothing in the construction was free to make it three.

Crit. point	Coordinate	Index	f -value	SM gen	Yukawa y
c_1	$[1 : 0 : 0]$	0	λ_1	1 st : e, u, d	smallest
c_2	$[0 : 1 : 0]$	2	λ_2	2 nd : μ, c, s	intermediate
c_3	$[0 : 0 : 1]$	4	λ_3	3 rd : τ, t, b	largest

The cascade is $c_3 \xrightarrow{e^{-2\pi}} c_2 \xrightarrow{e^{-\pi}} c_1$ with total action 3π , exactly saturating the BI-6 budget.

H.5 GRIP cascade weights on $\mathbb{CP}^2$

The gate weights on $\mathbb{CP}^2$, in toll units, are:

$$w_{c_3 \rightarrow c_2} = e^{-S_1} = e^{-2\pi} \simeq 1/535, \quad (218)$$

$$w_{c_2 \rightarrow c_1} = e^{-S_2} = e^{-\pi} \simeq 1/23. \quad (219)$$

An earlier statement of this subsection read the inverses as cascade lifetimes and multiplied the ratio 23 by the phase-space factor $(m_3/m_2)^5 \simeq 1.34 \times 10^6$ to recover the measured $\tau_\mu/\tau_\tau \simeq 7.6 \times 10^6$ “to within order unity.” That arithmetic is withdrawn twice over: the weights are not rates (**GRIP-ST7**), and the product it formed ran the cascade factor the wrong way, since a longer bare lifetime for the third generation would have to be divided out of, not multiplied into, the measured ratio. The measured ratio is Sargent’s rule and the branching fraction, as §35.2 states.

I The free-energy principle: a bridge to life sciences and neuroscience

Reading goal. This appendix bridges the **GFF-ST6** thermodynamic content of GG-Theory to the free-energy principle of theoretical neuroscience and the non-equilibrium thermodynamics of biological systems. The intended audience is a first-year graduate student in theoretical biology or computational neuroscience who is familiar with statistical mechanics but not differential geometry.

I.1 From Helmholtz free energy to variational free energy

The Helmholtz free energy $F = U - TS$ of equilibrium statistical mechanics is the unique function whose minimization governs the equilibrium of a system in contact with a heat bath at fixed temperature T . Spontaneous processes always decrease F (the second law), and the equilibrium configuration is the unique global minimum.

In Bayesian inference, an analogous quantity is the *variational free energy*

$$F_{\text{var}}[q] = \mathbb{E}_q[\log q(s) - \log p(s, o)] = D_{\text{KL}}(q \| p(\cdot|o)) - \log p(o), \quad (220)$$

where $q(s)$ is a candidate posterior distribution over hidden states s given observed data o , and $p(s, o)$ is the joint generative model. Minimizing F_{var} over q approximates the true posterior $p(s|o)$ (the variational principle of [Hinton & van Camp, 1993](#)).

In neuroscience, Friston’s *free-energy principle* ([Friston, 2010, 2023](#)) proposes that the brain itself is an inference engine that minimizes a generalized F_{var} over its internal generative model of the sensory world. This single principle is claimed to unify perception, action, learning, and homeostasis.

I.2 How GFF subsumes Friston’s free energy

The GFF of GG-Theory is the Helmholtz-type free energy on the boundary moduli space $\mathcal{M}_{\text{bdy}}$ of GDOS. Restricting $\mathcal{M}_{\text{bdy}}$ to a specific stratum corresponding to neural representations of an external world gives Friston’s variational free energy as a special case.

More precisely:

- The GG-Theory boundary effective temperature $T_{\text{GFF}} \sim M_{\text{KK}}$ corresponds in neural systems to a thermal/noise scale set by metabolic energy availability.
- The GG-Theory internal energy U_{GFF} corresponds in neural systems to the negative log-likelihood of sensory input given the generative model.
- The GG-Theory entropy S_{GFF} corresponds in neural systems to the entropy of the posterior distribution over hidden states.
- The GG-Theory Lindblad CPTP structure corresponds in neural systems to the irreversible, dissipative dynamics of synaptic plasticity and neural firing.

The *Geometric First Law* of Theorem [25.1](#) guarantees that the algebraic Morse cascade and the thermodynamic free-energy minimization are the same dynamical process. In neuroscience terms, this means that the discrete catalog of “percept attractors” in the brain (Morse critical points of the neural GFF) corresponds to the thermodynamic equilibria of the free-energy minimization.

I.3 Empirical predictions for biological systems

The dual-layer architecture of GFF produces specific empirical predictions for biological systems:

- (1) **Equipartition for biological fluctuations.** Near a biological steady state, the variance of any fluctuating biological observable is $\langle \delta X^2 \rangle = T_{\text{bio}} \chi_X$, where T_{bio} is the metabolic temperature and χ_X is the susceptibility. This is the GG analog of equipartition.
- (2) **Onsager reciprocity in biological networks.** The cross-coupling matrix between biological fluxes (e.g., metabolic fluxes in coupled enzymatic networks) must be symmetric: $L_{ij} = L_{ji}$. Empirical violations would falsify the Lindblad/CPTP structure of the underlying GDOS flow.

- (3) **Gate-ordered weights.** Biological cascades (signaling networks, gene-regulatory networks) should exhibit gate-ordered equilibrium weights of the form $w \propto e^{-S^k}$ for some toll unit S depending on the system, as in **GRIP-ST7**; whether those weights appear as rates is, there as here, a question of mechanism.

These predictions provide concrete empirical bridges from the fundamental physics of GG-Theory to the laboratory phenomena of biology and neuroscience. Confirmation across multiple platforms would provide strong empirical support for the universal applicability of the dual-layer GG-6 framework.

I.4 Caveat: domain of applicability

It is important to caveat that the GG-Theory predictions for biological systems are *universal kinematic relations*, not *specific dynamical predictions*. The framework predicts that free-energy minimization, Onsager reciprocity, equipartition, and gate-ordered cascade weights will hold for *any* biological system that is fundamentally a boundary projection of the GG-6 bulk. It does not predict the specific time constants, susceptibilities, or rate matrices — these depend on the particular biological system.

The empirical content for biology is therefore the universal *form* of these laws (e.g., equipartition $\langle \delta X^2 \rangle \propto T$, Onsager $L_{ij} = L_{ji}$, gate weight $w \propto e^{-\pi^k}$), not the specific values. This is analogous to how thermodynamics predicts $PV = nRT$ universally without specifying the particular gas under study.

J Hopkins-Singer descent: explicit worked examples

Reading goal. This appendix illustrates the Hopkins–Singer descent of differential cohomology classes from the bulk to the boundary moduli, with three explicit worked examples. Intended audience: a first-year graduate student in mathematical physics who has read Appendix A and is comfortable with differential forms but not with abstract category-theoretic constructions.

J.1 The descent procedure: a concrete recipe

Given a differential cohomology class $\check{x} \in \check{H}^k(X; \mathbb{Z})$ on a smooth fibration $\pi: X \rightarrow Y$ with closed fibers of dimension n , the descent class $\check{\pi}_!(\check{x}) \in \check{H}^{k-n}(Y; \mathbb{Z})$ is computed in three steps:

- (1) **Choose a connection representative.** Pick a curvature representative $\omega \in \Omega_{\text{closed}}^k(X)$ in the class $\check{x}$ (i.e., $\text{curv}(\check{x}) = \omega$).
- (2) **Integrate along fibers.** For each $y \in Y$, integrate ω over the fiber $\pi^{-1}(y)$:

$$\tilde{\omega}(y) = \int_{\pi^{-1}(y)} \omega \in \Omega^{k-n}(Y).$$

The form $\tilde{\omega}$ is closed (by Stokes and the closure of ω).

- (3) **Lift back to differential cohomology.** The integer class $[\tilde{\omega}] = \pi_![\omega]$ pairs canonically with

a differential cohomology refinement $\check{\pi}_!(\check{x}) \in \check{H}^{k-n}(Y; \mathbb{Z})$.

J.2 Worked example 1: S^2 -bundle over S^2

Consider the trivial bundle $X = S^2 \times S^2$ with projection $\pi: X \rightarrow Y = S^2$ projecting onto the second factor. Each fiber is $\pi^{-1}(y) \cong S^2$ of dimension 2.

Take $\check{x} \in \check{H}^4(X; \mathbb{Z})$ represented by $\omega = \omega_1 \wedge \omega_2$, where ω_i is the area 2-form on the i -th S^2 normalized by $\int_{S^2} \omega_i = 1$. Then $\omega \in \Omega^4(X)$ has integer class $[\omega] = 1 \in H^4(X; \mathbb{Z}) \cong \mathbb{Z}$.

Integrating along the first-factor fibers:

$$\tilde{\omega}(y) = \int_{\pi^{-1}(y)} \omega_1 \wedge \omega_2 = \omega_2(y) \cdot \int_{S^2} \omega_1 = \omega_2(y).$$

Therefore the descent gives $\check{\pi}_!(\check{x}) \in \check{H}^2(Y; \mathbb{Z})$ with curvature ω_2 and integer class 1.

J.3 Worked example 2: 6D bulk to 4D boundary

Now consider the 6D bulk $M_6 = \mathbb{R}^{1,3} \times I_r \times S_\theta^1$ of GG-Theory, with projection $\pi: M_6 \rightarrow M_4 = \mathbb{R}^{1,3}$ integrating over the two non-spatial dimensions (r, θ) — the scale dimension and the gauge dimension. Each fiber has dimension 2.

Take the CCI 4-form

$$X_4 = \text{tr } R_{\text{Spin}(1,5)}^2 - \text{Tr } F_5^2 - \frac{2}{5} F_B^2 - \frac{2}{5} F_{Qi}^2,$$

which has integer class $[X_4] \in H^4(M_6; \mathbb{Z})$ (Theorem **T2.1**). Integrating along the (r, θ) fibers:

$$\tilde{X}_2(x) = \int_{(r, \theta)} X_4(x, r, \theta) \in \Omega^2(M_4),$$

which is the boundary 2-form $\tilde{X}_2$ used in Section 17.1. This is the operational definition of the PBI-4 compensator field C_μ via $dC = \tilde{X}_2$ (Definition 17.3).

J.4 Worked example 3: descent to a stratum of $\mathcal{M}_{\text{bdy}}$

The boundary moduli space $\mathcal{M}_{\text{bdy}}$ is stratified. Consider the family-projector stratum $\Sigma_{\text{fam}} = \mathbb{CP}^{N-1}$ inside $\mathcal{M}_{\text{bdy}}$. The Hopkins–Singer descent of $\check{x}_4 \in \check{H}^4(M_6; \mathbb{Z})$ to Σ_{fam} is computed by integrating X_4 over the fiber of the bundle $\pi^{-1}(\Sigma_{\text{fam}}) \rightarrow \Sigma_{\text{fam}}$.

For $N = 3$ ($\Sigma_{\text{fam}} = \mathbb{CP}^2$), the descent gives

$$\check{\pi}_!(\check{x}_4) \Big|_{\mathbb{CP}^2} = c_R \cdot \check{p}_1(\mathbb{CP}^2) \in \check{H}^4(\mathbb{CP}^2; \mathbb{Z}),$$

where $c_R \in \mathbb{Q}_{>0}$ is the Killing-form coefficient (this is Theorem 27.1 of Part IV). The class $\check{p}_1$ is the differential first Pontryagin class of $\mathbb{CP}^2$, which on $\mathbb{CP}^2$ equals $3 \cdot [\omega_{\text{FS}}]^2$ where ω_{FS} is the Fubini–Study form. This non-vanishing class is what makes the alternating Morse-index sum on $\mathbb{CP}^2$ strictly positive ($a^+ - a^- = \sigma(\mathbb{CP}^2) = 1 > 0$), confirming the strict Morse-positivity of **T3.1**.

J.5 Why the descent preserves integrality

A subtle point: the descent map $\tilde{\pi}_!$ preserves integrality because the integral of a closed form over a closed fiber is itself an integer (when the input form has integer cohomology class). This is a consequence of Stokes's theorem and the absence of boundary terms when the fiber is closed.

For the GG-Theory descent from M_6 to M_4 , the closed-fiber condition is satisfied by the compactness of the extra dimensions (the radial interval I_r has IR/UV boundaries that are compactified by the GG-6 IR boundary construction; the angular S_θ^1 is closed by definition). Therefore the boundary 2-form $\tilde{X}_2$ has integer cohomology class on M_4 , which is the precise condition needed for the PBI-4 compensator C_μ to exist globally without anomalies.

K Morse theory on $\mathbb{CP}^{N-1}$: explicit worked examples

Reading goal. This appendix carries out the Morse theory on $\mathbb{CP}^{N_g-1}$ in full pedagogical detail for $N_g = 2, 3, 4$. Intended audience: a first-year graduate student who has seen the two-sphere S^2 Morse picture (the height function on a sphere) but not the projective generalization.

K.1 The height function on $\mathbb{CP}^{N-1}$

Recall the projective space $\mathbb{CP}^{N-1} = \{[\psi_1 : \psi_2 : \cdots : \psi_N]\}$ with the equivalence $[\psi_1 : \cdots : \psi_N] \sim [\lambda\psi_1 : \cdots : \lambda\psi_N]$ for $\lambda \in \mathbb{C}^*$. The complex dimension is $N_g - 1$, real dimension $2(N_g - 1)$.

The *height function* on $\mathbb{CP}^{N-1}$ is defined by

$$f([\psi]) = \frac{\sum_{a=1}^N \lambda_a |\psi_a|^2}{\sum_{a=1}^N |\psi_a|^2}, \quad \lambda_1 < \lambda_2 < \cdots < \lambda_N. \quad (221)$$

(In coordinates this is well-defined on $\mathbb{CP}^{N-1}$ because both numerator and denominator scale by $|\lambda|^2$ under the equivalence.)

K.2 Critical points of the height function

The critical points of f on $\mathbb{CP}^{N-1}$ are exactly the N_g “corner” points

$$c_a = [0 : \cdots : 0 : 1 : 0 : \cdots : 0] \quad (1 \text{ in slot } a), \quad a = 1, \dots, N.$$

At each c_a , the value of f is $f(c_a) = \lambda_a$. Since the λ_a are all distinct, f takes N distinct values at N distinct critical points.

K.3 Morse index of each critical point

The Morse index of c_a is the number of negative eigenvalues of the Hessian $\text{Hess}(f)$ at c_a . Computing the Hessian in local holomorphic coordinates centered at c_a , one finds:

$$\text{ind}(c_a) = 2(a - 1). \quad (222)$$

So $\text{ind}(c_1) = 0$ (local minimum, the GA), $\text{ind}(c_2) = 2, \dots, \text{ind}(c_N) = 2(N-1)$ (local maximum, the DGI).

Note that the indices are all even. This is a consequence of the *Kähler* structure of $\mathbb{CP}^{N-1}$: the Hessian eigenvalues come in complex pairs $(\lambda, \bar{\lambda})$, so they always have even index.

K.4 Morse polynomial

Summing $t^{\text{ind}(c)}$ over critical points gives the Morse polynomial:

$$\mathcal{P}_f^{\mathbb{CP}^{N-1}}(t) = \sum_{a=1}^N t^{2(a-1)} = 1 + t^2 + t^4 + \dots + t^{2(N_g-1)}. \quad (223)$$

This equals the Poincaré polynomial of $\mathbb{CP}^{N-1}$ (the cohomology polynomial of $\mathbb{CP}^{N-1}$ is generated by the Kähler form $[\omega_{\text{FS}}] \in H^2$ and its powers, giving $\dim H^{2k}(\mathbb{CP}^{N-1}; \mathbb{Q}) = 1$ for $k = 0, 1, \dots, N_g - 1$). This shows that f is a *perfect Morse function*.

K.5 Worked example: $N = 2$ ($\mathbb{CP}^1 \cong S^2$)

For $N = 2$, $\mathbb{CP}^1 \cong S^2$ has two critical points c_1, c_2 with indices 0, 2. Morse polynomial $\mathcal{P}_f = 1 + t^2$. Alternating sum $a^+ - a^- = 1 - 1 = 0$ (with a^+ = number of indices $\equiv 0 \pmod{4}$ and a^- = number $\equiv 2 \pmod{4}$).

Conclusion: $\dim_{\mathbb{R}} \mathbb{CP}^1 = 2$ is *not* divisible by 4, so strict Morse-positivity *fails* on $\mathbb{CP}^1$. This rules out $N_g = 2$ for family-projector closure.

K.6 Worked example: $N_g = 3$ ($\mathbb{CP}^2$)

For $N_g = 3$, $\mathbb{CP}^2$ has three critical points c_1, c_2, c_3 with indices 0, 2, 4. Morse polynomial $\mathcal{P}_f = 1 + t^2 + t^4$. Alternating sum:

$$a^+(\mathbb{CP}^2) - a^-(\mathbb{CP}^2) = |\{0, 4\}| - |\{2\}| = 2 - 1 = 1 > 0.$$

Conclusion: $\dim_{\mathbb{R}} \mathbb{CP}^2 = 4$ is divisible by 4, and strict Morse-positivity *holds*. This is the parity argument that makes $N_g = 3$ the unique admissible family count together with the upper bound $N_g \leq 3$ from **T4.1**.

K.7 Worked example: $N_g = 4$ ($\mathbb{CP}^3$)

For $N_g = 4$, $\mathbb{CP}^3$ has four critical points with indices 0, 2, 4, 6. Morse polynomial $\mathcal{P}_f = 1 + t^2 + t^4 + t^6$. Alternating sum:

$$a^+(\mathbb{CP}^3) - a^-(\mathbb{CP}^3) = |\{0, 4\}| - |\{2, 6\}| = 2 - 2 = 0.$$

Conclusion: $\dim_{\mathbb{R}} \mathbb{CP}^3 = 6$ is *not* divisible by 4, so strict Morse-positivity fails. This rules out $N_g = 4$ from the parity side, complementing the action-budget upper bound from **T4.1**.

K.8 Pattern: parity of $\dim_{\mathbb{R}} \mathbb{CP}^{N-1}$

The pattern is $\dim_{\mathbb{R}} \mathbb{CP}^{N-1} = 2(N_g - 1)$:

- N_g odd ($N_g = 1, 3, 5, \dots$): $\dim_{\mathbb{R}} = 0, 4, 8, \dots$ (multiples of 4), strict positivity holds.
- N even ($N = 2, 4, 6, \dots$): $\dim_{\mathbb{R}} = 2, 6, 10, \dots$ (NOT multiples of 4), strict positivity fails.

So among integer $N_g \geq 1$, the family-projector dimension test selects $N_g \in \{1, 3, 5, 7, \dots\}$. Combined with the lower bound $N_g \geq 3$ (from CP violation, Theorem **T2.1**) and the upper bound $N_g \leq 3$ (from cascade closure, Theorem **T4.1**), the unique admissible value is $N_g = 3$.

K.9 Cascade actions on $\mathbb{CP}^{N_g-1}$

The cascade on $\mathbb{CP}^{N-1}$ is $c_N \rightarrow c_{N-1} \rightarrow \dots \rightarrow c_1$. The action of the k -th transition (from c_{N-k+1} at index $2(N-k)$ to c_{N-k} at index $2(N-k-1)$) is, by Theorem **T1.1**, $S_k = \pi(N_g - k)$. Total cascade action:

$$S_{\text{tot}}(\mathbb{CP}^{N-1}) = \sum_{k=1}^{N_g-1} \pi(N_g - k) = \pi \sum_{j=1}^{N_g-1} j = \pi \frac{N_g(N_g - 1)}{2}.$$

For $N_g = 3$: $S_{\text{tot}}(\mathbb{CP}^2) = 3\pi$, exactly saturating $S_{\text{BI6}}^{\text{max}} = 3\pi$. For $N_g = 4$: $S_{\text{tot}}(\mathbb{CP}^3) = 6\pi$, exceeding the budget by 3π (equivalent to one full cascade in addition). This is why $N_g \leq 3$.

L The generation weight ladder: from cascade gate tolls to laboratory observations

Reading goal. This appendix sets the cascade gate weight $e^{-\pi k}$ beside the *family* decay widths, which are Standard-Model quantities dominated by phase space, with explicit numerical comparison to laboratory data — and says, at each step, what the weight is and what it is not. There is a single cascade — the generation cascade $\tau \rightarrow \mu \rightarrow e$ of §35.2; W and the free neutron are ordinary weak (GRIP-In) decays, not cascade-index particles, and appear here only as familiar order-of-magnitude anchors for the phase-space factor. Intended audience: a first-year graduate student in particle physics or related fields. Useful background: the WKB approximation; standard weak-decay phase space (Sargent’s rule).

L.1 From gate toll to a rate: what would be needed

The decay rate of a metastable configuration through an instanton transition with Euclidean action S_E is, in the WKB approximation,

$$\Gamma = A \cdot e^{-S_E}, \quad (224)$$

where A is a one-loop prefactor of order the relevant energy scale (Callan & Coleman, 1977, Coleman, 1977). Applied with $M_{\text{KK}} \simeq 3.85 \text{ TeV}$ as the scale, $A \sim M_{\text{KK}} \sim 10^{27} \text{ Hz}$ (since $1 \text{ TeV} \approx 1.5 \times 10^{27}$

Hz), and the bare lifetime would be

$$\tau = 1/\Gamma = (1/A) \cdot e^{+S_E} \approx \tau_0 \cdot e^{S_E}, \quad \tau_0 = 1/A \sim 10^{-28} \text{ s}, \quad (225)$$

the value $\tau_0 \sim 10^{-28}$ s that Eq. (173) of Part V carries. The formula needs what Coleman's setting has and the family stratum does not: a local minimum behind a barrier. The generation points of $\mathbb{CP}^2$ are the critical points of a perfect Morse function with even indices; the index-2 and index-4 points are saddles, unstable at linear order, from which the descent is downhill and deterministic, and small-noise theory attaches to them an equilibrium weight and no escape rate (**GRIP-ST7**). So (224) is not derived here for the gates, and the lifetimes below it are carried at the grade of that statement. What *is* derived is the weight e^{-S_E} itself, with S_E the toll.

(Note: the $\tau_0 \sim 10^{-28}$ s above would be a bare timescale before species-specific phase-space corrections; the empirical $\tau_W \sim 10^{-25}$ s already includes the gauge-coupling enhancement $1/g_W^2$ and the inherited $1/m_W$ kinematic scale relative to it.)

L.2 Cumulative cascade toll

The i -th cascade stage adds the toll $\pi(i+1)$ to the running total. Cumulative toll up to stage i :

$$S_E^{(i)} = \sum_{k=1}^i \pi k = \pi \frac{i(i+1)}{2}. \quad (226)$$

The cumulative weight after descending i gates is $e^{-S_E^{(i)}}$; the timescale

$$\tau_i^{\text{geom}} = \tau_0 \cdot e^{\pi i(i+1)/2}, \quad (227)$$

which earlier statements formed from it, is the rate reading of the previous subsection and is carried at that grade: a cumulative gate weight written as a time, *not* a physical lifetime.

L.3 Numerical sequence

Setting $\tau_0 \sim 10^{-28}$ s (after accounting for phase-space factors that depend on the specific particle), we compute:

Table 26: The geometric gate weight along the generation cascade, from the tau through the muon to the electron. Each step down the cascade carries one more factor of $e^{-\pi}$ in the weight, and the third column turns that into the numbers the fourth column carries: two gates for the third generation, one for the second, none for the first — which is why the electron is the floor. The lifetimes in the last column are the measured values, which the companion paper accounts for by phase space; they are set beside the weights and are not derived from them. What the present paper supplies is the exponent of a weight, and this is where the exponent meets the measured values.

Gate k	Gate action πk	Gate factor $e^{-\pi k}$	Generation	Lifetime (via A9)
2	2π	$e^{-2\pi} \approx 1.9 \times 10^{-3}$	3rd (τ)	$\tau_\tau \approx 2.9 \times 10^{-13} \text{ s}$
1	π	$e^{-\pi} \approx 4.3 \times 10^{-2}$	2nd (μ)	$\tau_\mu \approx 2.2 \times 10^{-6} \text{ s}$
—	—	— (ground state)	1st (e)	stable

L.4 Why the empirical lifetimes are what they are

The empirical lifetimes bear no visible relation to the bare timescales of the previous subsection, and on the present reading they should not: the lifetimes are weak-decay widths, and weak-decay widths are set by the *kinematic* phase-space factors that depend on the specific particle and decay mode. An earlier statement of this appendix read the cascade factor $e^{\pi i(i+1)/2}$ as a geometric enhancement of the lifetime; that is the rate reading, withdrawn above.

For weak decays, the standard kinematic factor is Sargent’s rule:

$$\Gamma_{\text{weak}} \propto G_F^2 Q^5, \quad (228)$$

where Q is the energy release. Plugging in:

- W boson: $Q \sim m_W = 80 \text{ GeV}$, $\Gamma \sim G_F^2 m_W^5 \sim 2 \text{ GeV}$, $\tau \sim 10^{-25} \text{ s}$. (Sargent phase-space estimate matches the observed lifetime)
- Muon: $Q \sim m_\mu = 106 \text{ MeV}$, $\Gamma \sim G_F^2 m_\mu^5 \sim 0.5 \mu\text{eV}$, $\tau \sim 10^{-6} \text{ s}$. (Sargent phase-space estimate matches the observed lifetime)
- Neutron: $Q \sim \Delta m_{n-p} = 1.3 \text{ MeV}$, $\Gamma \sim G_F^2 \Delta m^5 \sim 10^{-3} \text{ eV}$, $\tau \sim 10^3 \text{ s}$. (Sargent phase-space estimate matches the observed lifetime)

The species-specific Sargent factors by themselves reproduce the empirical lifetimes to within an order of magnitude per species, and that is the whole account of the lifetimes this paper offers. What the cascade supplies is a different thing: the ordering of the generations by height, which is the ordering of the weights $e^{-\pi i(i+1)/2}$, and which the mass hierarchy, not the lifetimes, is where to read. Whether any residual of the weight survives in a width, after the phase space is divided out, would need the mechanism **GRIP-ST7** names and does not supply.

L.5 Universality across sectors

What is universal in the cascade is that the toll $\pi i(i+1)/2$ to depth i depends only on the depth, not on the sector — it is an inflow charge, a Chern number of the gate (**GRIP-ST3**), and the charged-lepton, up-quark and down-quark ladders share it. What the data do not support is a

universality of the *heights*: read at one scale, the three sectors' log-mass gaps stand in three different ratios, 0.53, 0.90 and 1.33, against the 1.31 of this paper's own action (F-GRIP-4 above; [Arisaka, 2026, A9-SM](#)), so the torus action whose height the weights read must differ sector by sector, and the isospin factor enters the family weights themselves. The skeleton is one; the ladder of heights hung on it is not. The collapse of weak-decay lifetimes onto a single Sargent curve against Q^5 is a fact about the weak interaction and phase space, and is not claimed here as a consequence of the cascade.

L.6 Falsifier 7 (F-GRIP-7): a particle outside the sequence — retired into Falsifier 4

This falsifier duplicated Falsifier 4 (F-GRIP-4) and is retired into it; the heading is kept as the record. Its content — a generation that sat outside the height ordering of the weights, a family ladder whose members were not ordered as the gate weights order them, would falsify the cascade structure of **GRIP-ST7**, and the ordering is read from the mass hierarchy, sector by sector, in [Arisaka \(2026, A9-SM\)](#) — is the ordering half of Falsifier 4 as now stated in §33.3. The residual-width form of this test — a controlled-kinematics decay class in which the Sargent phase space is a uniform multiplier and the residual ordering would have to follow $e^{-\pi k}$, with a contradiction by a factor $e^\pi \approx 23$ or more counting against it — is the form earlier statements gave, and it tests the rate reading of the weight, which this paper does not derive; the absolute lifetimes themselves are phase-space-dominated and are not a falsifier at all.

Part XI

Common Objections and Responses

The strongest objections a hostile reviewer would raise against the Gi-4 construction — on reading every conservation law off one composite connection, the cascade and its budget, and the open-problem ledger M9–M11 — answered, with pointers into the paper

A construction that claims to read every conservation law of physics off a single composite connection, and to force the family number to three by an instanton-action budget, invites the most demanding scrutiny; it would not be honest scholarship to leave the obvious objections unstated. We collect here the strongest adversarial readings of the Gi-4 thesis — the ones a particle theorist, a mathematical physicist, a string theorist, and a skeptical reviewer would each raise on a careful reading — and answer each, with explicit pointers to the theorem, table, or companion paper that bears the weight. Where the honest answer is “not yet,” we say so and name the open item (in particular the open-problem ledger M9–M11 of §57).

M For the curious reader: the big questions, plainly answered

The twenty-five questions that follow are not a specialist’s. They are the ones a thoughtful reader without a physics degree actually carries into a paper like this one: why there are three kinds of matter rather than two or four, why time runs one way when the equations do not care which, what became of a symmetry when it broke, and whether “downhill” is a fact about the world or a way of talking about it. Several of these have been settled in popular science in directions this paper does not take, and a reader is owed plain answers before the specialist objections begin. They are met first, and plainly.

The hardest of the twenty-five is the eighth — whether predicting three families of matter, when three is what everyone already knows, is a prediction at all. The response concedes the difficulty rather than dissolving it, and says which part of it the paper can carry and which part it cannot. Where a claim rests on a step that is not closed, the response says so and names the step; where the honest answer is that the framework does not reach the question, it says that too.

Question 1 (Q-GRIP-1) (on the curious reader)

“Why are there exactly three generations of matter?”

Response. Because the descent runs out of room, and the number three is what falls out of a budget rather than a preference.

The picture is worth having plainly. Structure forms here by a system settling: it starts with its options open, moves downhill on a landscape, reaches a fork, takes one branch, and continues

until there is nowhere lower. Each completed round of that process leaves behind one family of matter, and each round costs something from a fixed budget fixed by the integer ledger. Three rounds fit. A fourth does not, because there is nothing left to spend and nowhere lower to go. Two things make this a prediction rather than a story. First, the budget is not adjustable — it is a ratio of whole numbers, and there is no version of the framework with a slightly larger one. Second, the argument runs in the direction of forbidding rather than accommodating: it says a fourth family cannot exist, which is a statement that could be shown false tomorrow by a single discovery.

It is worth being exact about the status. The upper bound — no more than three — is conditional on a capacity result established in a companion paper, and that conditionality is flagged wherever the claim is made. The lower bound — at least three — is unconditional. Nobody should quote the result without the condition, and the paper does not.

Question 2 (Q-GRIP-2) (on the curious reader)

“Nobody has seen a fourth generation. Isn’t predicting three just predicting what we already know?”

Response. This is the hardest question here, and it deserves a straight answer rather than a defense.

You are right that three is not news. It has been known since the nineteen-eighties, when the decay of the Z boson was measured precisely enough to count how many light neutrino species exist, and the answer came out just under three with an uncertainty of less than one percent. Any framework producing three today is producing a number its author knew. That is a real weakness in the evidential value of the result and it should not be dressed up.

What can be said in the framework’s favor has two parts. The first is that most frameworks do not produce three at all — they accommodate it. The Standard Model does not predict the number of generations; it is an input. String compactifications can be arranged to give three, and also two, and also many other values, depending on the folding chosen. A construction that forces three and cannot be made to give anything else is doing something different from one that fits three.

The second is that the same budget argument makes claims that are not already known: an ordering of decay widths across the families, and a relationship between them fixed by the same integers. Those are not measured yet in the form the framework states them. That is where the evidential weight has to sit, and a reader is right to discount the family count accordingly.

Question 3 (Q-GRIP-3) (on the curious reader)

“Why does time run one way when the underlying equations don’t care which way it runs?”

Response. Because the equations are not the whole description, and what is missing from them is exactly what supplies the direction.

The puzzle is genuine and old. Write down the fundamental laws and they work equally well run backwards; play a film of two colliding molecules in reverse and nothing looks wrong. Play a film of a smashed cup in reverse and everything looks wrong. Something asymmetric has

entered between those two scales and it is not in the equations.

The standard answer, from Boltzmann onward, is statistical: there are vastly more disordered arrangements than ordered ones, so a system started in an ordered state will almost certainly become less ordered. That answer is correct and incomplete, because it explains the direction only by assuming the past was ordered, which is the thing one wanted explained.

What this framework adds is structural rather than statistical. Descent on the landscape is monotonic — the functional decreases along the flow and cannot increase — so the sequence of settlements has an order built into the geometry rather than into the counting of microstates. That does not by itself explain why the universe began where it did, and the paper does not claim it does. What it claims is narrower: that the irreversibility is a property of the projection rather than an accident of initial conditions, which is a different kind of answer and one that could be wrong in a checkable way.

Question 4 (Q-GRIP-4) (on the curious reader)

“What is symmetry breaking, really? Did the symmetry go somewhere?”

Response. Nowhere. It is still in the laws; what changed is the state, and keeping those two apart is the whole of the idea.

The standard example is a pencil balanced on its point. The situation is perfectly symmetric — no direction is preferred — and the pencil falls, and now there is a direction. The laws did not change; nothing chose east over west; and yet the outcome has a property the setup did not. The symmetry is not hidden or destroyed. It survives as the statement that the pencil was equally likely to fall any way, which is a fact about the ensemble and no help to the fallen pencil.

What this framework adds is a reason the falling happens *when* it does. In the usual telling, symmetry breaking is something that befalls a system as it cools, and the timing is a matter of circumstance. Here it is a step in the descent: it occurs at the point where continuing downhill requires choosing, and the choice is what removes the symmetry. The breaking is not an event that interrupts the process; it is what the process looks like at a fork.

The sentence worth carrying away is that structure is what a falling system leaves behind. Everything with a definite property — a direction, a charge, a mass — has it because something further back could have gone otherwise and did not.

Question 5 (Q-GRIP-5) (on the curious reader)

“If the structure is inevitable, why did the universe need fourteen billion years?”

Response. Because a fixed rule book says nothing about rates, and the two questions are more separable than they feel.

What the framework determines is what can exist and in what order it can form. It does not determine how long any stage takes, which depends on densities, distances and temperatures that are not fixed by the geometry in any usable way. A construction can force the existence of carbon and be entirely silent about how long stars take to make it.

What is forced is the *ordering*, and that is more than nothing. Each settling needs the previous one to have finished and needs conditions the previous one produced. Nuclei cannot precede

the quarks that make them; atoms cannot precede nuclei; chemistry cannot precede atoms; and nothing resembling biology can precede the heavy elements that only stellar deaths supply. A layered structure requires a long history whether or not the layers were inevitable.

It is worth saying what the framework does not claim, since it would be easy to read in. The particular history our universe had is not forced. Fourteen billion years is a measurement, not a prediction of this or any related construction, and anyone offering one should be asked to show the calculation.

Question 6 (Q-GRIP-6) (on the curious reader)

“Is ‘downhill’ a physical fact, or just a way of speaking?”

Response. A physical fact, and the reason it is worth asking is that in most of physics the picture is imported rather than derived.

In ordinary thermodynamics one writes down a free energy because experience says there is one, and the landscape picture that follows — valleys, barriers, a system running down — is a way of visualizing the mathematics rather than something the mathematics produced. That is perfectly respectable and it does mean the picture is doing no independent work.

Here the landscape is constructed from the geometric data, and its crucial property — that it curves upward in every direction away from its critical points, so valleys are genuinely valleys and not flat shelves — is forced by the anomaly condition rather than assumed. That link is the paper’s most surprising internal result, and it is worth stating in the strange form it deserves: whether matter can come to rest turns out to be a question with an answer in the consistency of a quantum field theory.

What the landscape is *not* is a place. Nothing sits on a hillside; the heights are values of a functional and the positions are configurations. The picture is trustworthy because the mathematics has the properties the picture suggests, which is usually because the picture was extracted from the mathematics in the first place.

Question 7 (Q-GRIP-7) (on the curious reader)

“What is a ‘cascade’, and why does it stop?”

Response. A cascade is a sequence of settlings in which each one creates the conditions for the next, and it stops when a settling produces conditions that support no further step.

The everyday version is familiar. Water vapor cools and condenses; the droplets fall and pool; the pool freezes. Each stage needed the one before and each has its own character. Nothing external decided the number of stages; the sequence ran until the conditions stopped supporting a new one.

What makes the version here more than an analogy is that the stopping condition is computable. Each round of descent consumes part of a fixed budget set by the integer ledger, and when the remaining budget cannot pay for another round the cascade halts. Three rounds fit. That is why the framework can say how many families of matter there are rather than observing how many there happen to be.

There is a limitation the paper is careful about. Descent on a landscape that does not move has an end, full stop — everything settles once and stays settled, which describes stable matter

and describes nothing that continues. For a history rather than a conclusion, something must reorganize the landscape from inside. This paper establishes the engine and does not establish that; the companion that does is named, and the boundary is drawn rather than blurred.

Question 8 (Q-GRIP-8) (on the curious reader)

“Why is matter stable? Why doesn’t the electron just decay?”

Response. Because there is nowhere for it to go, and that is a stronger statement than it sounds. Decay is not something particles choose. A particle decays when there exists a lighter combination carrying the same conserved quantities; if such a combination exists, the decay happens, and if it does not, the particle is stable forever. The electron is the lightest thing carrying its charge, so charge conservation alone makes it permanent. Nothing further is required.

What a framework of this kind can add is why the ladder of possibilities terminates where it does — why there is a lightest charged thing at all, rather than an unbounded sequence of lighter ones. That follows from the descent having an end: the cascade halts, and what it halts at is what cannot decay.

There is a second and less obvious case worth mentioning because it works differently. The proton’s stability, on this account, is not about there being nothing lighter — there is — but about the decay requiring a change in something that counts rather than something that varies. Unwinding a count is not a small change, so the process is not merely slow but structurally obstructed. Whether that obstruction is absolute is a live question and one of the framework’s sharpest exposures, since experiments have been watching large tanks of water for decades waiting for one proton to disagree.

Question 9 (Q-GRIP-9) (on the curious reader)

“Isn’t the second law about statistics, not geometry?”

Response. Both, and the two accounts are answering different questions that are easy to run together.

The statistical account says a system moves toward disorder because there are overwhelmingly more disordered arrangements than ordered ones. That is correct and it explains the overwhelming reliability of the law. What it does not explain is why the past was ordered, and without that the argument gives no direction — run the same reasoning backwards and it predicts that the past was disordered too, which is false. This has been recognized since Boltzmann and remains the honest state of the subject.

The geometric account here says something different. The functional decreases monotonically along the flow because of the structure of the flow, not because of a counting argument about arrangements. That supplies a direction without needing an assumption about how things started.

The two are compatible and they are not the same claim, and the framework does not assert it has solved the arrow of time. It asserts that one component of the asymmetry is structural rather than statistical, and it is specific about which component. A reader should be wary of any account that claims to have settled this question outright; it has resisted a great deal of talent.

Question 10 (Q-GRIP-10) (on the curious reader)**“What does ‘free energy’ mean when there is no heat bath?”**

Response. It means the same thing it means with one — a quantity that decreases and whose decrease measures what has been irreversibly spent — but the reason it exists is different, and the difference matters.

In ordinary thermodynamics free energy is defined relative to a bath at a temperature, and the bath is doing real work in the definition: it is what makes some energy available and some not. Take the bath away and the usual definition has nothing to hold on to, which is why applying thermodynamic language to a single system or to the universe as a whole has always been delicate.

The functional here is built from the geometry rather than from a bath, and it splits into two recognizable pieces: one counting what a configuration costs in energy, the other counting what it costs in specification — how much has to be pinned down, how much order is being maintained against spreading. Anything that persists pays both.

That is the framework’s most aggressive claim, and its exposure is exactly proportional. If the same functional prices a folded protein and a gauge-orbit stratum, then a set of relationships among laboratory measurements must hold that no ordinary account requires. Those relationships are stated in the assessment and are measurable on a bench, which is an unusually accessible place for a claim of this ambition to be tested.

Question 11 (Q-GRIP-11) (on the curious reader)**“Why should a landscape have a bottom at all? Why not descend forever?”**

Response. It is a fair question and the answer is not obvious — plenty of mathematically respectable landscapes have no bottom, and a system on one would fall without end.

Such a landscape describes nothing that could exist. If the functional is unbounded below, a system can always lower it further; there are no stable configurations, nothing settles, and there is no matter in any recognizable sense. So the requirement that the landscape be bounded is not an aesthetic preference. It is the condition for there being anything at all, and a framework that failed to guarantee it would be describing a world with no contents.

What is done here is to prove it rather than assume it. The boundedness follows from the same consistency condition that fixes the integer ledger — the anomaly condition again, doing a second job. That is the technically interesting part of the paper: a single requirement about the internal consistency of a field theory turns out to imply both that the ledger is unique and that the landscape has a floor.

The consequence worth carrying away is that stability is not a happy fact about our universe. On this account, any world that is consistent enough to be described has to have somewhere for things to come to rest — which converts the existence of stable matter from a coincidence into a theorem, conditional on the arena.

Question 12 (Q-GRIP-12) (on the curious reader)**“If everything descends, why is there any structure left? Shouldn’t it all have run**

down?”

Response. It has, and that is the point — the structure *is* what the running down left. The intuition that descent destroys structure comes from thinking of the landscape as smooth, so that everything ends in one featureless minimum. Real landscapes are not smooth. They have many separated valleys, and a system descending gets caught in one and stays. What we call structure — a nucleus, an atom, a molecule, a cell — is a system sitting in a valley it cannot climb out of, and it is there *because* it descended, not despite it.

So the answer is that running down and building structure are the same process seen from two sides. Every stable thing is a place where a descent stopped. The water in a lake is not evidence that water failed to flow downhill; it is evidence that it did.

There is a second half that this paper does not supply and is careful to say so. Descent on a fixed landscape happens once, and afterward nothing further occurs. For a world that keeps producing new kinds of structure, the landscape itself has to be reorganized from within, repeatedly. That requires something the four pillars here do not provide, and the companion that provides it is named.

Question 13 (Q-GRIP-13) (on the curious reader)

“Could the cascade have stopped somewhere else? What if it had run twice, or four times?”

Response. It could not, given the budget — and that is exactly what makes the count a prediction rather than an observation.

The budget is fixed by the integer ledger, which is itself the unique solution of a consistency condition. Each round of descent costs a computable fraction of it. Two rounds would leave the budget unspent, which the descent does not permit — it runs until it cannot. Four would need more than exists. Three is the only number that both fits and exhausts.

The right way to be skeptical about this is not to doubt the arithmetic but to ask what the budget depends on. It depends on the arena, which is an input. A different arena would give a different budget and could give a different count. So the claim is not that three is necessary absolutely; it is that three is necessary given a structure fixed on independent grounds. That is the honest form and it is the form the paper uses.

What makes the claim useful anyway is that the independent grounds were settled first and are written down. Nobody can adjust the arena after a measurement to recover the count, because the arena’s argument makes no reference to families of matter. If a fourth family appears, the arena was wrong, and everything built on it goes with it.

Question 14 (Q-GRIP-14) (on the curious reader)

“Does this explain why the universe isn’t uniform? Why there are galaxies rather than a smooth soup?”

Response. Partly, and the part it does not explain is worth marking, because the question mixes two things.

Structure requires two ingredients: a mechanism that amplifies small differences, and small

differences to amplify. Gravity supplies the first, and has since Newton — a region slightly denser than its surroundings pulls harder, grows denser still, and runs away. That part is not in dispute and is not this framework’s contribution.

Where the seeds came from is the harder half. The standard account attributes them to quantum fluctuations stretched to cosmic size during an early rapid expansion, and the measured pattern in the microwave background matches that account remarkably well. This framework is consistent with that story and does not replace it; the cohort paper on cosmology carries the detail.

What the present paper contributes is one step further down: given that structure formed, why it formed in *layers* — nuclei, then atoms, then molecules — rather than in one undifferentiated collapse. That is the cascade, and the layering is forced by each stage requiring the previous one’s output. A reader looking here for an account of galaxy formation will not find one, and should not; a reader asking why the world has levels at all will.

Question 15 (Q-GRIP-15) (on the curious reader)

“What is the difference between a law and a habit? How do you know the constants aren’t just drifting slowly?”

Response. This is a sharper question than it looks, and it has been taken seriously enough to be measured.

Dirac proposed in 1937 that some constants might vary over cosmic time, and the idea has never quite died. It is checkable: light from distant quasars was emitted billions of years ago, and the atomic transitions imprinted on it depend on the fine-structure constant, so comparing them with laboratory values constrains any drift. The bounds are now very tight — variation, if any, is below a part in 10^{17} per year — and a natural nuclear reactor that ran in Gabon two billion years ago provides an independent and similarly tight constraint.

What this framework adds is a reason to expect exactly zero rather than a small number. Quantities fixed by whole numbers cannot drift, because there is no continuous path from one integer to another. A framework in which constants are ratios of integers predicts no variation at all — not merely a variation too small to have been seen.

That is a strong position and an exposed one. A confirmed nonzero drift, at any level, would end this framework outright while leaving accounts that treat constants as parameters entirely comfortable. It is one of the cleanest disagreements between the two kinds of theory, and it is being measured now.

Question 16 (Q-GRIP-16) (on the curious reader)

“Isn’t our universe just one bubble among countless others?”

Response. It might be, and this paper’s contribution is to remove the question’s motivation rather than to answer it.

The multiverse exists to explain why the constants look arranged, and it works by making the arrangement a selection effect within an ensemble. That requires the constants to have been capable of taking other values.

Here they were not. The family count is what a fixed budget permits; the budget is a ratio of integers; the integers are the unique solution of a consistency condition. There is no distribution

to draw from. If that is right, the ensemble has nothing to explain and its principal job is gone. The comparison worth making is with the framework's own risk. The multiverse cannot be embarrassed by a measurement of the family count, because any count is compatible with some member of the ensemble. This framework can be ended by finding a fourth. That asymmetry is not an argument that this one is correct. It is an argument that it is the kind of thing that can be found out, which is the more useful property for a theory to have.

Question 17 (Q-GRIP-17) (on the curious reader)

“The universe looks fine-tuned for life. Doesn’t that point to a designer?”

Response. It points to something requiring explanation, and the inference to a designer needs a premise this framework denies.

The observation is real: alter certain quantities slightly and there are no stable nuclei, no chemistry, nothing that lasts. Feeling the pull of design there is not naive.

The inference requires that the quantities could have been set otherwise. A dial that was turned implies a hand. If the values are what a fixed budget permits and the budget is not adjustable, there was no dial, and asking who set them has no purchase.

This is not an argument against a designer and should not be recruited as one. Someone who believes on other grounds is untouched by anything here, and so is someone who does not. What changes is that one particular argument loses its premise if the framework holds — and whether it holds is settled by whether a fourth family of matter turns up, which is an odd place for a theological question to be decided and is where this one now sits.

Question 18 (Q-GRIP-18) (on the curious reader)

“If the cascade is forced, is my future already written?”

Response. Nothing here touches your future, and the inference does not hold however tight it feels.

What the cascade fixes is the rule book: what kinds of matter exist, what is stable, in what order structure can form. It does not fix any particular history. The distinction is between the rules of chess and the game played last Tuesday, and a rule book admitting no alternative says nothing about which of an enormous number of permitted games occurred.

There is a further point that cuts deeper. Even a fully deterministic rule book delivers no readable future, because the only system capable of computing what the rules imply for something as complicated as a person is a system at least that complicated running at least that long. There is no shortcut and no vantage point. A universe can be lawful all the way down and remain, in the only sense available to anything inside it, unwritten until it happens.

And the honest limit: this paper is about a descent and its budget. What agency is, and whether the concept survives contact with physics of any kind, is a question it does not reach.

Question 19 (Q-GRIP-19) (on the curious reader)

“Why is there something rather than nothing?”

Response. The framework does not touch it, and saying so is more useful than a gesture.

Everything here is conditional: given an arena with certain properties, the descent has this

shape and stops after three rounds. Why there is an arena at all sits outside the construction, and no rearrangement reaches it. That is not a gap further work closes; it is a different kind of question.

What can be contributed is narrower. If the framework is right, the space of possible somethings is far smaller than it appears — very nearly a single admissible structure rather than a landscape from which one was drawn. That converts “why this?” into “why anything?”, which is harder and cleaner. Narrowing a question is a real service and it is not an answer.

There is a temptation, in a paper about descent, to say that the question presumes a starting point and that a structure without one needs no beginning. That move is available, has defenders, and does not follow from anything here.

Question 20 (Q-GRIP-20) (on the curious reader)

“This explains a great deal from very little. Shouldn’t that make you suspicious?”

Response. It should and it does, and Eddington is the standing warning: a great astrophysicist who spent his final years deriving the fine-structure constant from combinatorics that convinced nobody, and whose reputation never recovered.

What separates a derivation from a coincidence is not the quality of the agreement. It is whether the machinery producing it exists for another reason and whether it commits you to anything further. A quantity reverse-engineered from a measurement will always match.

Here the descent structure is built from consistency conditions that know nothing about particle physics, and it then fixes not only the family count — which was known — but an ordering of decay widths across the families and a set of thermodynamic relations in ordinary laboratory systems, neither of which was used in building it.

That is the only defense available and it is the right one: not “too good to be chance” but “here are further consequences, please check them.” Keep the suspicion. If those consequences fail, the agreement was a coincidence, and no elegance in the construction will help.

Question 21 (Q-GRIP-21) (on the curious reader)

“Why should I take one person’s framework seriously against thousands working on other approaches?”

Response. Not because of who wrote it, and not despite it — both are the same error with opposite sign.

The practical question is fair, since attention is finite. The useful filter is exposure rather than headcount. A framework supported by thousands that forbids nothing measurable is a worse use of attention than one supported by a single person that forbids specific things and names the measurement that decides. Community size tracks funding history and fashion and does not track correctness.

There is a real cost to working outside a large community: errors are found faster where more people look, and a framework developed largely by one person carries a higher chance of a mistake a hundred readers would have caught quickly. Internal coherence does nothing to reduce that.

So do not take it on trust. The thermodynamic relations this paper predicts are measurable on

an ordinary bench, need no six-dimensional mathematics to evaluate, and either hold exactly or do not. Check one of those, and judge the rest by what you find.

Question 22 (Q-GRIP-22) (on the curious reader)

“What would convince you that you are wrong?”

Response. Several things, named rather than gestured at.

A fourth family of matter would end it: the count follows from a budget that cannot be enlarged. A measured drift in a dimensionless constant, at any level, would end it, since quantities fixed by integers cannot drift. A clean violation of the thermodynamic relations the framework says must hold exactly — the symmetry of a particular coupling matrix, a specific fluctuation-dissipation relation — would break the third pillar, and those are bench measurements rather than frontier ones. And a decay ordering across the families that disagrees with the predicted geometric factor would break the fourth.

An internal failure counts fully too: if the argument that the landscape is bounded below has a hole, the descent has no guaranteed floor and the whole structure loses its terminus.

What would not convince: a disagreement in a regime not claimed, or the failure of a result whose hypotheses are already flagged as open. A falsifier list is only honest if it excludes things, and a reader should be wary of any proponent unwilling to say which failures are survivable.

Question 23 (Q-GRIP-23) (on the curious reader)

“Why does complexity seem to increase when entropy is also increasing? Aren’t those opposites?”

Response. They are not opposites, and the belief that they are is one of the most common misreadings of thermodynamics.

The second law says that the total entropy of an isolated system does not decrease. It says nothing whatever about any part of that system. A refrigerator makes its inside colder and more ordered every day without violating anything, because it dumps more disorder into the kitchen than it removes from the cabinet. The Earth does the same on a larger scale: it receives concentrated energy from a hot sun and radiates diffuse energy to cold space, and everything that has ever grown here was paid for out of that difference.

So local complexity is not a loophole or an exception. It is what a system does when energy flows through it, and the flow is possible precisely because the overall descent is happening. Structure is not swimming upstream; it is an eddy in the current.

What this framework contributes is a way of pricing the eddy. The functional’s two parts — what a configuration costs in energy and what it costs in specification — let the trade be written as an inequality with named terms rather than described qualitatively. Whether the pricing is right is checkable in the laboratory, which is more than can be said for most statements about complexity.

Question 24 (Q-GRIP-24) (on the curious reader)

“Are the four ‘pillars’ four things or one? It sounds like a framework in search of an

architecture.”

Response. One, and the paper proves it rather than asserting it — which is the difference between an architecture and a list.

The suspicion is reasonable. Frameworks announcing a fixed number of pillars, principles or laws often turn out to be organizing a set of independently developed results after the fact, and the number is a presentational choice. The test is whether anything would break if you removed one.

Here the four are shown to be equivalent to the constitutional axioms taken together: each implies the other, with nothing left over on either side. That is a real theorem and it has a specific consequence — the four are not four convenient constructions that happen to sit together but four descriptions of one thing, and the axioms can be recovered from them.

What an equivalence buys is worth stating because it is easy to overrate. It produces no new prediction; nothing follows from it that did not follow from the pillars. What it does is tell a critic where to attack: if the four and the axioms are the same content, then refuting a pillar refutes an axiom and conversely. A construction whose parts were merely compatible would offer no such handle — one could remove a piece and the rest would stand, slightly diminished, indefinitely.

Question 25 (Q-GRIP-25) (on the curious reader)

“I’m not a physicist and I’ve got this far. What is the one thing you want me to take away?”

Response. That structure is what a falling system leaves behind, and that this is a more useful idea than it sounds.

The intuition most people carry is that order and decay are opposed — that things run down, and that the existence of anything complicated is a small miracle working against the grain. That picture is wrong in a specific and fixable way. Everything with a definite property has it because something further back could have gone otherwise and did not. A magnet has a direction because it cooled and had to choose. An atom exists because a descent stopped there. The three kinds of matter are three rounds of the same process, and there is no fourth because there was nothing left to spend.

Once you have that, several separate-looking puzzles stop being separate. Why time runs one way, why the world has levels rather than being uniform, why matter is stable, why there are three families and not four — these are one question with one answer, which is that a system running downhill on a landscape with many valleys gets caught, and what it gets caught in is what we call the world.

So take away the picture rather than the theorem. And when you next meet a claim about fundamental physics, ask what it forbids: this one forbids a fourth family and forbids the constants drifting, and both are being measured now.

N The composite connection and TGG (Pillar 1)

Objection 1 (O-GRIP-1) (the differential geometer)

“Your ‘composite connection’ $\mathcal{A}_{\text{tot}} = \omega \oplus A$ is just the direct sum of four commuting connections on a fiber-product bundle. It carries no information beyond its summands — no cross-terms, no mixing. Calling the obvious sum a single unified primitive is a relabeling, not a unification.”

Response. At the level of the connection one-form the objection is correct, and Theorem 10.1 does not claim otherwise: $\mathcal{A}_{\text{tot}}$ is a direct sum on the product structure group $G_{\text{tot}} = \text{Spin}(1, 5) \times \text{SU}(5) \times \text{U}(1)_B \times \text{U}(1)_{Q_i}$. The content of TGG is not that the sum is mysterious but that the constitution *forces exactly these four factors and no others* (the gauge-group proposition, §7, Proposition 6.6) and that the four sectors’ stress contributions assemble into one boundary tensor whose conservation is a single law (T1.3). The unification claimed is constitutional and conservational, not a claim of off-diagonal mixing in the connection. Where the objection bites, and we concede it, is that “composite connection” as a phrase oversells a direct sum; the honest statement is that the four summands are jointly forced and jointly conserved, which is what GGEP (AX-2 of (Arisaka, 2026, A1-Constitution)) asserts and what Corollary 10.5 records.

Objection 2 (O-GRIP-2) (the differential geometer)

“Your uniqueness proof kills the obstruction by asserting H^1 of the contractible bulk vanishes — yet the same construction classifies the two $\text{U}(1)$ bundles by a nontrivial H^2 generated by the compact θ -circle, with Wilson holonomy. The base cannot be contractible and simultaneously support a nontrivial H^2 with winding. Which is it?”

Response. The two statements live on different factors of the product geometry $M_6 = \mathcal{M}_4 \times I_r \times S_\theta^1$, and the proof should — and in the final-review pass does — say so explicitly. The H^1 -vanishing that fixes the connection up to gauge is on the *non-compact, contractible* $\mathcal{M}_4 \times I_r$ factor; the nontrivial H^2 and the Wilson holonomy live on the *compact* S_θ^1 factor (Proposition 6.8), whose first cohomology carries the integer $\text{U}(1)_B/\text{U}(1)_{Q_i}$ winding. There is no contradiction: the bundle is topologically a product, trivial along the contractible directions and classified by the circle’s holonomy along the compact one. The objection correctly identifies that the bare phrasing “the contractible bulk” is sloppy — the bulk is contractible only after the circle is excised — and the uniqueness statement is properly read as “unique given the fixed S_θ^1 holonomy class.”

Objection 3 (O-GRIP-3) (the 6D-supergravity specialist)

“You treat $\text{Spin}(1, 5)$ as a fourth internal gauge factor summed into G_{tot} , but the spin connection is fixed by the metric (Levi-Civita); it is not an independent principal connection on equal footing with $\text{SU}(5)$. Treating gravity’s ω as a commuting gauge summand double-counts the geometry.”

Response. This is the sharpest structural objection to TGG, and the answer is that the equal-footing treatment is exactly the constitutional content of AX-2 (GGEP), not an oversight —

but it is a *posture*, and we mark it as such. Under GGEP the spin connection and the gauge connections are facets of one composite primitive, and the Levi-Civita condition that ties ω to the metric is read as the GGEP-closed limit (Proposition 6.4, §8); the standard separated-sector treatment is recovered when that closure is imposed. So the framework does not deny that ω is metric-determined; it claims that the metric and the gauge data are projections of one object, which is precisely the non-standard move the constitution makes. We concede that a reader trained on 6D supergravity will find the equal-footing notation unfamiliar and that, at the level of degrees of freedom, ω is not independent of the metric; the constitutional claim is about the *primitive*, and Proposition 6.4 is where the reconciliation with the metric-determined ω is stated.

Objection 4 (O-GRIP-4) (the category theorist)

“‘Unique up to gauge equivalence’ — in what category, and up to which equivalence? You never specify the homotopy type of the base, never compute $[M_6, BG_{\text{tot}}]$, and never state which moduli you quotient by. The uniqueness is asserted at the level of slogans.”

Response. Fair, and Theorem 10.1 should state its category, which is: principal G_{tot} -bundles with connection over the fixed product geometry M_6 , up to bundle automorphism (gauge transformation) covering the identity, with the S_θ^1 holonomy class held fixed. With the homotopy type fixed — $M_6 \simeq S_\theta^1$ after contracting the non-compact directions — the classifying set $[M_6, BG_{\text{tot}}]$ reduces to the $\pi_1(S_\theta^1)$ data, i.e. the integer holonomies, and “uniqueness” means: given those integers, the bundle-with-connection is determined up to gauge. We grant that the body invokes Kobayashi–Nomizu at a slogan level; the precise statement is the holonomy-fixed one just given, and the uniqueness is uniqueness *relative to* the fixed holonomy class, not an absolute classification claim. We have sharpened the statement of the theorem accordingly.

Objection 5 (O-GRIP-5) (the skeptic of necessary-conditions-as-achievements)

“T1.3 defines $\mathcal{T}_{\text{TGG}}^{\mu\nu}$ as the right-hand side of the Einstein equation, so $\nabla_\mu \mathcal{T}^{\mu\nu} = 0$ is automatic by the Bianchi identity. This is textbook general relativity re-skinned as a ‘Master Conservation Law.’ What is the theorem beyond ‘a stress tensor defined via Einstein’s equations is conserved on-shell?’”

Response. The objection is right that the gravitational half of T1.3 (Theorem 18.4, its TGG prelude) is the contracted Bianchi identity and carries no new content. What is not textbook is the *iff*: the boundary tensor assembles four sectors (matter, gauge, Green–Schwarz inflow, KK), and $\nabla_\mu \mathcal{T}^{\mu\nu} = 0$ holds iff the four boundary Ward currents close *simultaneously* together with the PBI-4 compensator (T2.2). The non-trivial direction is that the four sector divergences are linked — a failure of any one Ward identity shows up as a calculable inflow term, not as a quiet redefinition. So the gravitational conservation is indeed automatic; the content is the four-sector simultaneous-closure condition (Proposition 17.1), and we agree the prose should foreground that rather than the Bianchi half, which the final-review pass does.

Objection 6 (O-GRIP-6) (the string theorist)

“T1.3 asserts $\nabla_\mu T_{\text{KK}}^{\mu\nu} = 0$ because heavy KK modes decouple at boundary scales. De-

coupling does not make their stress contribution separately conserved; integrating them out produces higher-derivative operators and threshold corrections. The clean four-way split into separately-conserved sectors is assumed, not derived.”

Response. Conceded: the separate conservation of the KK sector is a leading-order, decoupling-limit statement, and the higher-derivative and threshold corrections the objection names are real and are exactly the content of open problems M1 and M3 (§57). The honest claim of T1.3 is that *at the boundary scales where the framework makes predictions* the four sectors are separately conserved up to the calculable inflow term; the loop-level completion — where integrating out the KK tower mixes the sectors — is not in this paper. We do not claim the four-way split survives to arbitrary precision; we claim it at leading order and flag the corrections as open. A reader who needs the all-orders statement is correctly pointed to M1/M3 rather than offered a false guarantee.

O CCI, the integer ledger, and the modified-TOV falsifier (Pillar 2)

Objection 7 (O-GRIP-7) (the number theorist)

“T2.1 is advertised as forcing $(1, 3, 5; 7, 35)$, but it takes $N \geq 3$ and $L = 5$ as inputs and then ‘derives’ $N = 3$. Your own remark concedes $L = 5$ is a constitutional commitment and $N \geq 3$ is imported from CP violation. The theorem is ‘given $L = 5$ and $N \geq 3$, the minimal ledger has $N = 3$.’ The empirical content was inserted as a hypothesis.”

Response. The description is accurate and the paper states it (Theorem 18.1 and its remark): $L = 5$ is the smallest-simple-group commitment and $N \geq 3$ is the Kobayashi–Maskawa CP lower bound, both inputs. What T2.1 adds is the *closure*: given those, the MSMF-minimal integer solution of the integrality constraints is the unique tuple $(1, 3, 5; 7, 35)$, and the alternatives $N \in \{4, 6, \dots\}$ fail the parity/positivity filter (Pillar 3) while $N \geq 5$ fails the budget (Pillar 4, conditionally). So T2.1 is not a from-nothing uniqueness theorem; it is the statement that the constitutional inputs plus the integrality constraints have a unique minimal solution. We agree “forces” overstates it when the inputs are empirical, and the corrected reading is “the ledger is the unique MSMF-minimal solution consistent with the smallest-simple-group and CP-lower-bound inputs.”

Objection 8 (O-GRIP-8) (the number theorist)

“The budget $K = 35$ that supposedly forces $N \leq 3$ is itself $M^2 + N^2 + L^2 = 1 + 9 + 25 = 35$ — it already contains $N = 3$. The budget that bounds N is built from the answer. The whole $N = 3$ result is a fixed point of its own inputs.”

Response. This is the single most damaging objection to the paper, and we neither hide it nor claim to have closed it: the closure $K = M^2 + N^2 + L^2 = 35$ does contain $N = 3$, and the BI-6 budget $S_{\text{BI6}}^{\max} = 3\pi$ that bounds the cascade is motivated through $K = 35$. That is precisely why

open problem M10 (§57) is stated in the exact words “derive $S_{\text{BI6}}^{\text{max}} = 3\pi$ independent of $N = 3$ at every intermediate step, so as to rule out any circularity” — the framework documents its own potential circularity and labels closing it the central open task. What survives unconditionally is the *lower* bound $N \geq 3$ from CP violation, which uses none of the budget; the *upper* bound $N \leq 3$ is honestly conditional on M10 (Corollary 18.3 and the T4.1-status remark say so). So $N = 3$ is “ $N \geq 3$ (unconditional) and $N \leq 3$ (conditional on M10’s resolution).” M10’s N_g -independence demand is met by the capacity theorem T-N3-6 of Arisaka (2026, P2-N3) — which consumes neither $K = 35$ nor $J = 7$ — so the upper bound is a theorem conditional on that paper’s coupling clause and integrality theorem L-INT (residue R-INT = M6), and the abstract’s “forces” is read with exactly that conditional attached.

Objection 9 (O-GRIP-9) (the anomaly theorist)

“The coefficient vector $(1, -1, -2/5, -2/5)$ in X_4 is asserted to be fixed by ‘ $\mathcal{H}^2$ normalization,’ with the actual derivation outsourced to A2/A3. In a self-contained A6 the central trace-ratio $5/2$ lock is imported, not shown. A reviewer cannot check the load-bearing $-2/5$ here.”

Response. Correct, and stated: the rank-one collapse of the anomaly matrix and the trace-ratio lock $|a/c| = |a/d| = 5/2$ are proved in GG-A3-BI6 (Arisaka, 2026, A3-BI6) (and the bulk content in GG-A2-GG6 (Arisaka, 2026, A2-GG6)); §15 of the present paper records the result and cites the proof rather than re-deriving it. This is the cohort’s division of labor, not a concealment — A6 is the *formalization* paper for the pillars, and the anomaly computation is A3’s business. We concede the cost the objection names: a reader who wants to verify the $-2/5$ from first principles must consult A3, and A6 is not self-contained on that point. The honest framing, which §16 now carries, is that A6 *uses* the trace-ratio lock as an input from A3 and builds the conservation structure on it; the lock itself is A3’s theorem.

Objection 10 (O-GRIP-10) (the general relativist)

“The modified-TOV ‘astrophysical falsifier’ produces a sub-percent shift in the mass–radius relation. But a sub-percent stress-energy correction is generically mimicked by uncertainties in the nuclear equation of state, which span far more than 1%. How is this distinguishable from ‘general relativity plus a slightly different equation of state’?”

Response. The objection is correct that at the level of a single mass–radius point the OGN-5 correction is degenerate with equation-of-state freedom, and the paper does not call it a clean stand-alone falsifier. The distinguishing claim (Corollary 18.8, Eq. (74)) is structural: the correction enters with a *fixed dimensionless coefficient* κ_{GS} tied to the ledger and the *same* functional dependence across all compact stars, so the falsifiable object is a population-level, equation-of-state-orthogonal ratio rather than a single M – R shift. We concede honestly that whether such an orthogonal ratio is observationally extractable — given that a one-parameter equation-of-state deformation can mimic much of the functional form — is not settled in this paper, and the cosmological/astrophysical falsifiers proper are the territory of GG-A10-Cosmos and GG-A5-GDOS. As stated here, modified TOV is a consistency target, not a near-term kill test, and we have removed the word “falsifier” where it claimed more.

Objection 11 (O-GRIP-11) (the mathematical physicist)

“T2.2’s reverse direction derives the compensator one-form C_μ from the Poincaré lemma applied to the closed boundary 2-form. But Poincaré gives only *local* exactness; you need a *global* smooth C_μ , and the global obstruction is exactly the open M11 monodromy problem. The iff is proved only locally.”

Response. Conceded precisely. The forward direction of T2.2 (Theorem 18.4) is global; the reverse direction constructs C_μ locally via the Poincaré lemma, and the global smooth existence across the codimension-2 instanton strata is the monodromy non-obstruction recorded as open problem M11 (§57, §17.3). So the iff is established locally and globally *modulo* M11. This is the same gap that conditions T4.2, and the paper marks both with the same open problem rather than papering over the local-to-global step. The honest statement of T2.2 is “iff, with the reverse direction global up to the M11 monodromy hypothesis.”

P GFF, Morse theory, and the second law (Pillar 3)**Objection 12 (O-GRIP-12) (the Morse theorist)**

“T3.1 proves strict positivity iff $\dim_{\mathbb{R}} \Sigma \equiv 0 \pmod{4}$, i.e. N odd. But this is equally satisfied by $N_g = 1, 3, 5, 7, \dots$. T3 contributes *zero* upper bound on the number of generations — it is a parity filter, nothing more. Advertising it as one of three legs that force $N_g = 3$ overstates its content.”

Response. Exactly right, and the paper is careful to assign T3 only the parity role: Theorem 26.1 (T3.1) excludes even N via the Hirzebruch-signature obstruction (Proposition 22.14) and contributes *no* upper bound. The three-leg pincer is explicit that the legs do different jobs: CP violation gives $N_g \geq 3$, the signature/parity gives “ N_g odd,” and the budget (Pillar 4, conditional on M10) gives $N_g \leq 3$; only the conjunction yields $N_g = 3$ (Corollary 34.4). If any prose calls T3 a leg that “forces $N = 3$ ” on its own, that is an overstatement we correct: T3 forces *odd*, and odd-plus- $(N \geq 3)$ -plus- $(N \leq 3)$ forces three. We thank the reviewer for the sharpening and have made the parity-only role explicit at §26.

Objection 13 (O-GRIP-13) (the Morse theorist)

“The ‘alternating index sum = Hirzebruch signature’ identity holds only for the equivariant moment-map Morse function under a maximal-torus action on Kähler strata with flat connections. For a generic GFF landscape the index sum is the Euler characteristic, not the signature. You have engineered the Morse function to land on the invariant you need.”

Response. The objection correctly identifies the load-bearing assumption, and the paper states it as a scope condition rather than hiding it: Proposition 22.14 holds for the moment-map height function on the Kähler family-projector strata $\mathbb{CP}^{N-1}$ with gauge connections flat along Σ , and that is the class on which the N_g -counting is done. The substantive claim is that the *physical*

$\mathcal{F}_{\text{GFF}}$ on the family-projector stratum *is* the moment map of the residual torus action — because the stratum is the projective space of family rotations and the free energy is the equivariant height — not an arbitrary Morse function. We concede this identification is an assumption about which functional the physics delivers, and that for a generic, non-torus-invariant landscape the relevant invariant would be the Euler characteristic and the parity argument would not go through. The reader is entitled to ask for a proof that the physical functional is the moment map; §26 states it as the operative hypothesis, and a from-dynamics derivation is part of the open agenda.

Objection 14 (O-GRIP-14) (the topologist)

“Step 3 needs the first Pontryagin class $p_1(\Sigma) \neq 0$, supplied by the $\mathcal{H}_1/\mathcal{H}_2$ theorems. But you admit the $\mathcal{H}_2$ proof leaves the bulk implementation to a separate paper that is *in preparation*. The signature non-triviality — hence the entire parity result — rests on an unwritten paper.”

Response. Conceded, and it is one of the paper’s genuine soft spots. Theorems 27.1 and 27.2 (§27) upgrade the hypotheses $\mathcal{H}_1, \mathcal{H}_2$ of the N=3 program to theorems *within* Gi-4, but the η -invariant non-degeneracy step that underwrites $p_1(\Sigma) \neq 0$ is, as open problem M6 records, sketched here with the detailed bulk implementation deferred to a companion technical note that is not yet public. So the parity result is rigorous modulo M6. We do not claim $\mathcal{H}_2$ is fully proved in print; we claim it is reduced to a stated η -invariant computation, flagged as open. A reader is right to treat the parity leg as conditional on M6 until that computation appears.

Objection 15 (O-GRIP-15) (the open-quantum-systems theorist)

“T3.2 claims $dF/dt \leq 0$ iff CPTP. But CPTP alone gives only monotonicity (Spohn); the clean iff needs an *additional* detailed-balance hypothesis, which you patch into the text. Open systems generically violate detailed balance. State the assumption in the theorem, not in an apology.”

Response. Accepted in full, and the theorem statement (Theorem 26.3/25.3) now carries the hypothesis rather than the footnote: strict monotonicity $dF_{\text{GFF}}/dt \leq 0$ holds for GDOS CPTP dynamics *satisfying detailed balance with respect to F_{GFF}* ; without detailed balance one has only the entropy-production inequality $\dot{S}_{\text{tot}} \geq 0$, and the system can sit in a non-equilibrium steady state. This is exactly the A5 result (its T5 Lyapunov lemma) specialized to the family-projector moduli, and the detailed-balance condition is the same one A5 states. The reviewer is right that “iff CPTP” is too strong; the corrected headline is “if CPTP with detailed balance,” with no converse — Lindblad’s monotonicity needs no detailed balance, so monotonicity cannot certify it — and that condition is now in the theorem. The honesty cost — that the second law here is conditional on a non-generic regularity of the dynamics — is one we state rather than bury.

Objection 16 (O-GRIP-16) (the open-quantum-systems theorist)

“T3.2 is monotonicity from GDOS, but GDOS is A5 and A5’s T5 *is* the boundary second law. So what does A6 add? Either T3.2 is A5’s theorem re-imported (not new

content), or it strengthens A5 with detailed balance (a new assumption, refuting ‘no new axiom’). You cannot have both.”

Response. The dilemma is real and we accept its second horn cleanly: T3.2 *specializes* A5’s general boundary second law ((Arisaka, 2026, A5-GDOS), its Lyapunov theorem) to the stratified family-projector moduli, and in doing so it makes explicit the detailed-balance regularity under which the strict version holds. That regularity is a *condition on the dynamics*, not a new physical axiom about the bulk — it is the standard detailed-balance assumption of non-equilibrium thermodynamics, which A5 already invokes — so “no new physical axiom” survives, but “A6 adds nothing” does not: A6 adds the specialization to $\mathbb{CP}^{N-1}$ and the explicit statement of the operative condition. The honest position is therefore the narrow one: T3.2 is not an independent new law; it is A5’s law restricted to the stratum where the family count lives, with its hypotheses made explicit. We have rewritten §26 to say exactly that.

Objection 17 (O-GRIP-17) (the thermodynamicist)

“The ‘laboratory falsifier’ — equipartition plus the fluctuation–dissipation relation — is standard linear-response theory for any system with a free energy. A coherent violation would falsify ordinary statistical mechanics, not specifically GiFour. This is a prediction of thermodynamics, not of your framework.”

Response. Correct, and we downgrade the claim accordingly. The equipartition/fluctuation–dissipation relation (Corollary 26.6, Eq. (143)) is what *any* system with the GFF free energy at equilibrium must satisfy, and a violation would indeed indict statistical mechanics broadly rather than Gi-4 specifically. Its honest role is not as a discriminating falsifier but as a *consistency requirement*: it certifies that the GFF functional behaves as a genuine free energy, which is a non-trivial check on the construction but not a kill test that singles out the framework. We have relabeled it a consistency condition rather than a falsifier, and the genuinely framework-specific exposure is the integer-ledger structure, not the thermodynamic identity.

Q GRIP, the cascade, and the budget (Pillar 4)

Objection 18 (O-GRIP-18) (the skeptic of conjecture-laden proofs)

“The entire upper bound $N \leq 3$ — the headline result, exported to A9 and P2 — rests on $S_{\text{BI6}}^{\max} = 3\pi$, which your own status remark admits is a named conjecture (M10) taken as input. A 201-page paper whose central numerical claim is an unproven conjecture is a conditional result. The ‘forces $N \leq 3$ ’ language is not warranted.”

Response. We agree the language must carry the conditional, and the theorem does: Theorem 34.1 (T4.1) closes the cascade and yields $N \leq 3$ *conditional on the BI-6 budget* $S_{\text{BI6}}^{\max} = 3\pi$, which the status remark (§34) and open problem M10 mark as an unproven conjecture. So the honest, unconditional content of A6’s GRIP pillar is: the cascade structure, the per-step action πk , the total $\pi N_g(N_g - 1)/2$, and the implication “if the budget is 3π then $N \leq 3$.” The numerical bound itself is conditional. We do not claim a proof of $N \leq 3$; we claim a proof that $N_g \leq 3$

follows from the budget, plus the honest flag that the budget is conjectural. Every downstream use in A9 and P2 inherits that conditional, and those papers cite it as conditional. The reviewer is right that the abstract should not read “forces” without the M10 caveat, and we attach it.

Objection 19 (O-GRIP-19) (the skeptic of conjecture-laden proofs)

“The budget *exactly* equals the $N_g = 3$ cascade cost: $S_{\text{tot}}(\mathbb{CP}^2) = 3\pi = S_{\text{BI6}}^{\text{max}}$. A budget that saturates precisely at the desired answer, motivated only by a ledger that already encodes $N = 3$, is the signature of a tuned constraint. Until M10 is closed, $N \leq 3$ is indistinguishable from ‘we chose the budget to give 3.’”

Response. The exact saturation $S_{\text{tot}}(\mathbb{CP}^2) = 3\pi$ (Proposition 30.11, §34) is real and, as the reviewer says, is the thing a tuned constraint would also produce; the framework’s self-consistency remark explicitly warns that saturation “cannot count as evidence for the budget’s value inside a derivation of $N = 3$.” This is the same circularity as Objection 8 seen from the cascade side, and the answer is the same: it is documented as M10, and the task M10 sets is to derive $S_{\text{BI6}}^{\text{max}} = 3\pi$ *without* using $N_g = 3$ at any step. That task is done: the capacity theorem T-N3-6 of Arisaka (2026, P2-N3) derives $\text{Cap}(N_g) = \pi N_g$ from the stratum’s quantum geometry, consuming neither $K = 35$ nor $J = 7$, conditional on the coupling clause and L-INT (= M6). What we will defend is that the framework does not *pretend* the saturation is independent evidence — it flags it as suspect and quarantines the derivation as open. A reader is entitled to withhold assent to the upper bound until M10 lands, and the paper asks for exactly that posture.

Objection 20 (O-GRIP-20) (the instanton physicist)

“The per-step action $S_{[2k] \rightarrow [2(k-1)]} = \pi k$ is a standard Bogomolnyi / Duistermaat–Heckman result on $\mathbb{CP}^{N_g-1}$ (Witten 1982). The proof is correct but it is not a result *about* generations. What forces the family-projector moduli to be $\mathbb{CP}^{N_g-1}$ with this exact torus height function rather than any other Kähler stratum?”

Response. The objection draws the right line: Theorem 32.1 (the πk result) is standard symplectic geometry on $\mathbb{CP}^{N-1}$, and the physics is entirely in the identification of that space with the family-projector moduli. The justification offered is that the family sector is the projective space of N_g -family rotations modulo phase — i.e. $\mathbb{CP}^{N_g-1}$ by construction of the projector moduli (§31) — and the height function is the equivariant moment map of the residual symmetry, which is why the πk ladder appears. We concede that “the family moduli *are* $\mathbb{CP}^{N_g-1}$ with the moment-map height” is the load-bearing physical input, not a theorem of this paper, and it is the same assumption flagged in Objection 13. The instanton-action theorem is then standard given that input; its novelty is not claimed, only its application. The honest division is: the symplectic result is textbook, the identification with generations is the framework’s posited structure.

Objection 21 (O-GRIP-21) (the mathematical physicist)

“T4.2’s reverse direction requires C_μ to extend smoothly across each codimension-2 instanton stratum, which you admit is established case-by-case for the $\mathbb{CP}^{N_g-1}$ saddles

but lacks a general theorem (M11). So T4.2 is a finite case-check, not a theorem.”

Response. Accepted as stated. Theorem 34.5 (T4.2) establishes runtime gauge-invariance preservation *conditional on* the smooth extension of the PBI-4 compensator across each instanton transition, verified case-by-case for the finitely many $\mathbb{CP}^{N_g-1}$ saddles relevant to the $N_g \leq 3$ application and *not* as a general theorem — the general monodromy non-obstruction is open problem M11 (status remark, §34). So for the Standard-Model application ($N \leq 3$) the needed cases are checked; the general- N statement is open. We do not claim a general theorem; we claim the finite verification sufficient for the physics plus a flagged conjecture for the rest. The reviewer is correct that “theorem” is generous for a finite case-check, and the status remark says exactly that.

Objection 22 (O-GRIP-22) (the dynamical-systems theorist)

“You split GRIP-in (within the fixed landscape) from GRIP-out (which changes it) and quarantine all the hard dynamics into GRIP-out, declared outside the septuple and supplying only initial data. But *why* the system starts at the top DGI saddle — what sets c_{top} — is exactly GRIP-out, which you exclude. The theorem governs only the easy half.”

Response. The objection correctly identifies the scope boundary, and we own it: Theorems 34.1 and 34.5 govern GRIP-In — the descent on a fixed landscape from given initial data c_{top} — and the selection of c_{top} , i.e. how the landscape was built and the system placed at its top, is GRIP-Out, which the paper explicitly places outside the septuple (§31). This is deliberate, not evasive: the cohort assigns the landscape-*building* event (the one-time electroweak GRIP-Out) to the boundary-law and cosmology papers, while A6 formalizes the *traversal*. We concede that this means A6’s theorems do not explain the initial condition, and that the “interesting” question of vacuum selection is GRIP-out and hence out of scope here. The honest claim is that A6 proves the descent dynamics given the landscape, not the genesis of the landscape.

Objection 23 (O-GRIP-23) (the instanton physicist)

“Step 4 of the instanton-action theorem conflates the symplectic *volume* π of one $\mathbb{CP}^1$ with the *action* of a tunneling trajectory. These are different objects; the factor- k cumulative ‘symplectic depth’ is asserted, not computed. Show the action integral, not the Duistermaat–Heckman volume.”

Response. The reviewer is right that the step deserves the explicit Bogomolnyi line integral, and the final-review pass supplies it. The action of the Bogomolnyi-saturating gradient trajectory is $S = \int |\nabla \mathcal{F}_{\text{GFF}}| |\dot{x}| dt = \int d\mathcal{F} = \mathcal{F}(c_+) - \mathcal{F}(c_-)$ (the saturated bound, §32.3, Eq. (166)), and for the moment-map height on $\mathbb{CP}^{N-1}$ the value difference between adjacent fixed points of index $2k$ and $2(k-1)$ is the symplectic area enclosed, which equals πk in Fubini–Study normalization. So the equality of the trajectory action with the symplectic depth is the Bogomolnyi saturation, not a conflation — but we agree the text stated the Duistermaat–Heckman volume and let the reader infer the equality, and it now shows the action integral explicitly. What the reviewer’s instinct does catch is a different thing: the conversion of the depth k to the action πk is a unit,

and the assignment of the pair invariant to the height is an identification that Theorem 24.1 excludes under the transport proposition. The value πk stands as the inflow toll of the gate (**GRIP-ST3**); as the action of the connecting orbit it is $c[\Delta_+(k+1) - \Delta_+(k)]$.

R The weight ladder and the falsifiers

Objection 24 (O-GRIP-24) (the phenomenologist)

“Your geometric gate factor is $e^{-\pi} \simeq 0.043$ (a factor 23), but the observed $\tau_\mu/\tau_\tau \simeq 7.6 \times 10^6$ is delivered almost entirely by the Sargent m^5 phase space ($\sim 10^6$). The geometric contribution is about 0.5% of the log. Calling the generation lifetime ladder a signature of the framework is hype.”

Response. The arithmetic is right and the paper now states it without inflation (§35, §33.2): the cascade supplies only the geometric gate weight $w_{[2k] \rightarrow [2(k-1)]} = e^{-\pi k}$, and the lifetimes are the species-specific phase space, which in fact reverses the naive geometric ordering for the charged leptons. The present version goes further than the reviewer asks: the gate weight is an equilibrium weight, not a rate, so it is not even a small residue in the width unless a mechanism carries it there (**GRIP-ST7**); the robust GRIP-level content is the *existence and generation-ordering* of the weights, read from the mass hierarchy, not a lifetime prediction of any size. The honest claim, which matches (Arisaka, 2026, P2-N3, §7.4) and GG-A9-SM, is that the absolute numbers belong to A9 and the framework owns only the geometric factor. We have removed any phrasing that presented the lifetime numbers as a framework triumph.

Objection 25 (O-GRIP-25) (the phenomenologist)

“Falsifier 4 (F-GRIP-4) is the ordering of gate factors ‘after the known phase-space factors are divided out.’ But those factors are huge and model-dependent, deferred to A9. To extract a factor-23 residual you need a complete SM-lifetime calculation accurate to better than 0.5%. Until A9 exists, this falsifier cannot be executed even in principle.”

Response. Conceded: the gate-ordering falsifier (§33.3) is only executable once the phase-space factors are computed to sub-percent accuracy, and that computation is GG-A9-SM’s, not this paper’s. So Falsifier 4 (F-GRIP-4) is, today, contingent on A9 and is not a stand-alone near-term test. We state the honest version: the robust falsifiable content is structural — a controlled-kinematics decay class in which the Sargent factor is a *uniform* multiplier across gates, so that the residual $e^{-\pi k}$ ordering must survive — and identifying such a class and computing its kinematics is downstream work. A reader is right that, absent A9, the lifetime falsifier is a program, not an executable test, and the paper now says so rather than advertising an immediate kill criterion.

Objection 26 (O-GRIP-26) (the phenomenologist)

“Falsifier 6 (F-GRIP-6), the ‘cascade-action sum rule’ $\sum S = 4\pi$ ‘comparable to the age

of the universe in dimensionless units,’ has no unit conversion and no dimensionful prediction. This is numerology, not a falsifier. Remove it or make it dimensionful.”

Response. Accepted; it is the weakest item in the section and we treat it as such. The sum rule $\sum_i S_{\text{tot}}^{(i)} = \pi N_g(N_g - 1)(N_g + 1)/6 = 4\pi$ at $N_g = 3$ (§33.3) is a closed arithmetic identity about the cascade actions, with no dimensionful bridge to a physical timescale, and the “comparable to the age of the universe” remark is a coincidence with no derivation behind it. The paper labels it “a weak but interesting consistency check,” which is generous; the honest move, taken in the final-review pass, is to demote it from the list of falsifiers to a numerical remark, or to remove the cosmological-timescale comparison entirely. It is not a falsifier and is no longer presented as one.

Objection 27 (O-GRIP-27) (the philosopher of physics)

“Your master move — read every conservation law off one connection — yields *necessary* conditions (anomaly cancellation, Ward closure, Bianchi identities) that any consistent theory must satisfy, then presents satisfying them as a positive achievement. Necessary conditions for consistency are not predictions.”

Response. The distinction is the right one and the paper should make it cleanly (§54.1). Anomaly freedom and Ward closure are indeed necessary conditions that any consistent theory meets, and meeting them is not, by itself, a prediction. What the framework claims as positive content is not “we are consistent” but “the constitution *forces* a specific gauge content and a specific integer ledger, and those force specific numbers” — $\alpha^{-1} = 137$ at the integer level, $M_{\text{lock}} \approx 3.85$ TeV, the kaxion line, the family count — which are the discriminating outputs, and which live downstream in A4/A9/P1/C1. So the honest framing is: the conservation laws are necessary-condition scaffolding; the predictions are the forced integers. We concede that where the prose presents anomaly cancellation *itself* as a triumph it conflates the two, and the corrected emphasis is on what the framework forbids (a non-integer ledger, a fourth family at low budget) rather than on satisfying universal consistency requirements.

Objection 28 (O-GRIP-28) (the working physicist)

“On inspection every falsification channel is degenerate, deferred, or standard physics: modified TOV is equation-of-state-degenerate, equipartition is thermodynamics, the gate-ordering needs an unbuilt A9 calculation, the sum rule is numerology, the cosmology is deferred to A10. What single near-term, GiFour-only measurement could return a result that kills the framework?”

Response. This is the fair bottom line of the falsifier critique, and the honest answer is that A6’s *own* near-term kill tests are weak — the sharp, near-term falsifiers of the program live downstream, and A6 should point to them rather than claim them. The framework-specific, near-term exposures are: the kaxion line at $8.1 \mu\text{eV}$ (≈ 1.95 GHz) in haloscopes ((Arisaka, 2026, C1-Kaxion)), the KK cluster near 3.85 TeV at HL-LHC ((Arisaka, 2026, P1-KK)), and the integer $\alpha^{-1} = 137$ structure ((Arisaka, 2026, A9-SM)); a clean null at the kaxion line, or a fourth light family, would propagate back through the ledger that A6 formalizes. A6’s own

contribution is structural — the rank-one/integer-ledger architecture and the conditional $N \leq 3$ — and its falsifiability is mathematical (the ledger is wrong if the trace data give a different rank) more than experimental. We concede A6 in isolation offers no decisive near-term laboratory kill test, and we now point the reader to the downstream ones rather than overselling the in-paper “falsifiers.”

S T5: constitutional equivalence

Objection 29 (O-GRIP-29) (the logician)

“T5 claims the GiFour system and the Constitution are equivalent in derivational content — mutual derivability. But the statement concedes direction (2) is established only ‘up to two auxiliary propositions whose internal proofs are sketched.’ A bi-implication with one direction admittedly unproven is not an equivalence; it is one theorem plus a conjecture.”

Response. Accepted, and the status note on Theorem 39.1 (T5, §39) says so: direction (1) — the constitution plus BI-6 yield the Gi-4 pillars — is proved; direction (2) — the pillars recover the constitutional axioms — is established up to two auxiliary propositions whose proofs are sketched, and is presented as a structured proof sketch, not a completed proof. So T5 is honestly “one direction proved, the converse sketched,” and the word “equivalence” should carry that qualification, which the status note supplies and the abstract should echo. We do not claim a completed bi-implication; we claim one direction and a sketched converse, and the reviewer is right that “equivalence” unqualified overstates it.

Objection 30 (O-GRIP-30) (the logician)

“‘Equivalent in derivational content’ is never precisely defined. Without specifying the deductive system, the axioms, and what counts as a theorem, mutual derivability is not a checkable claim. As stated it is unfalsifiable hand-waving about two informal vocabularies.”

Response. The objection is correct that §39 gives the equivalence an informal gloss (“anything derivable in one is derivable in the other”) rather than a formal definition over a fixed deductive system. The defensible precise content is narrower and we state it: each constitutional axiom (AX-1–AX-4) appears as a named theorem or proposition of the Gi-4 system (the bridge maps, §40), and each Gi-4 principal theorem is derived from the constitution plus BI-6 (direction (1)); “equivalence” is the conjunction of those two finite lists of named derivations, not a meta-theorem about all possible deductions. So the checkable claim is the explicit correspondence table, not an abstract mutual-derivability over an unspecified logic. We concede the prose reaches for the grander, vaguer statement, and the corrected version is the finite, tabulated correspondence — which is checkable, modulo the direction-(2) sketches.

Objection 31 (O-GRIP-31) (the logician)

“Even the sketch of direction (2) recovers GRIP via a proposition ‘conditional on M10’ and defers GFF to a Stinespring argument that is ‘part of the broader program.’ The equivalence inherits every load-bearing gap — M10, M11 — that the rest of the paper leans on.”

Response. Yes — direction (2) of T5 is no more secure than the pillars it reconstructs, and it inherits exactly their conditionals. Proposition 30.13 (recovering GRIP) is conditional on M10, and the GFF reconstruction (Proposition 22.15) is reduced to a Stinespring dilation argument not completed here. So the equivalence is conditional on the same cascade conditioning that T4 carries — today M11’s general statement, M10 being resolved externally. This is not a separate weakness; it is the same weakness propagating into T5, and the paper marks it rather than presenting T5 as a clean keystone. The honest reading: T5’s forward direction is solid, its converse is a conditional sketch carrying the cascade’s open problems, and the “equivalence” is therefore as conditional as the cascade.

Objection 32 (O-GRIP-32) (the skeptic)

“If T5 is true — the two systems are different presentations of the same content and the GiFour reorganization ‘is not a new theory’ — then nothing in A6 is independently provable that was not already provable from the Constitution plus BI-6. The 201 pages prove no new physics. Your value claim and your equivalence claim are in tension.”

Response. The tension is real and we resolve it by being precise about “new.” T5 says A6 introduces no new *physical content* beyond the constitution plus BI-6 — and that is the intended claim: A6 is the formalization paper, not a new-physics paper. What A6 adds is not new physics but new *rigor and organization*: it gives the pillars explicit theorem statements, proofs (some conditional), and a dual-layer architecture that makes the latent content auditable, in the way that Lagrange and Hamilton added no new mechanics to Newton but made it provable. So the resolution is: “no new physics, genuine new formal content.” We concede the prose sometimes blurs “new theorems” (formal) with “new physics” (content) and that the two must be kept apart; the corrected statement is that A6’s novelty is formalization, and T5 is precisely the assertion that the formalization is faithful.

T The “no new axiom” claim and the promoted hypotheses**Objection 33 (O-GRIP-33) (the philosopher of physics)**

“You insist A6 ‘introduces no additional physical axiom’ and ‘merely makes explicit’ latent content — yet you claim five principal theorems and genuinely new content. A pure formalization cannot also be a source of new theorems. Which is the honest description?”

Response. Both, once “new” is split as in Objection 32: no new *physical axiom*, genuinely new *formal theorems*. The five principal theorems are derivations *from* the existing constitution plus

BI-6; they add no postulate to the framework’s physics, but they are non-trivial mathematical statements (existence/uniqueness of the composite connection, the conservation iff, the parity result, the conditional cascade closure, the equivalence) that were not previously written down or proved. The analogy the paper draws — Newton needed Lagrange and Hamilton — is exactly this: new theorems, no new physics. We accept that the bare phrases “no new axiom” and “new content” read as contradictory, and the paper now consistently says “no new physical axiom; new formal results,” which is the honest description and is what T5 certifies.

Objection 34 (O-GRIP-34) (the philosopher of physics)

“The detailed-balance hypothesis you smuggled into T3.2 *is* a new physical assumption about the open-system dynamics, not present in a generic CPTP map. That alone refutes ‘no new axiom.’ Either retract the claim or prove detailed balance from prior axioms.”

Response. We accept the force of this and draw the line carefully. Detailed balance is a *regularity condition on the dynamics* under which the strict second law holds; it is the standard near-equilibrium assumption of non-equilibrium thermodynamics, and it is already invoked in A5’s boundary second law ((Arisaka, 2026, A5-GDOS)), which A6 specializes. It is therefore not a new *constitutional* axiom of the GG-Theory framework, but it *is* an extra hypothesis on the regime in which T3.2 holds, and the reviewer is right that calling the framework “axiom-free” while relying on it is sloppy. The corrected statement: the four constitutional axioms are unchanged; T3.2’s strict form additionally assumes detailed balance, stated in the theorem. We do not claim to derive detailed balance from the constitution — generic open dynamics violate it — so the honest position is that the strict second law is a conditional result, and “no new axiom” refers to the constitutional axioms, not to every regularity condition a theorem may need.

Objection 35 (O-GRIP-35) (the skeptic)

“You say the hypotheses $\mathcal{H}_1, \mathcal{H}_2$ are ‘promoted to theorems’ in GiFour — then concede A6 does not provide rigorous proofs of them and defers them to a paper in preparation. Promoting a hypothesis to a theorem whose proof is in an unwritten paper is renaming, not promoting.”

Response. Fair, and the wording is corrected. What A6 actually does (Theorems 27.1, 27.2, §27) is *reduce* $\mathcal{H}_1, \mathcal{H}_2$ from free-standing postulates of the N=3 program to consequences of the Gi-4 structure *modulo* a stated η -invariant computation (open problem M6) whose bulk implementation is deferred. “Reduced to a named open computation” is honest; “promoted to a theorem” is not, when the computation is not in print. So the corrected claim is the weaker, true one: within Gi-4, $\mathcal{H}_1$ and $\mathcal{H}_2$ are no longer independent assumptions but follow from the pillar structure once M6 is discharged. We agree “promoted” oversold a reduction-modulo-open-problem, and the text now says “reduced to M6.”

U Cohort interdependence and the risk of circular citation

Objection 36 (O-GRIP-36) (the skeptic of citation loops)

“A6’s T4.1 imports the budget 3π from P2 (Appendix C), while A6’s outward map exports $N \leq 3$ back to P2 and A9 for the hard upper bound on family count. P2 supplies the conjecture A6 needs; A6 supplies the bound P2 cites. This is a citation cycle: neither paper derives the result without the other assuming it. Where is the acyclic foundation?”

Response. The cycle the reviewer draws is real at the level of the *statement* $N \leq 3$, and the acyclic foundation is the shared open problem, named identically in both papers. The unconditional, acyclic content is: A6 proves “if $S_{\text{BI6}}^{\max} = 3\pi$ then $N \leq 3$ ” (Theorem 34.1); P2 proves the same conditional from the family side; and the budget is proved in one place only — the capacity theorem T-N3-6 of Arisaka (2026, P2-N3), conditional on its coupling clause and L-INT (= M6) — while this paper consumes it as input. So there is no circular *proof*: there is one derivation, one residue (M6 = R-INT) that both papers flag, and a conditional implication that each proves independently. The cycle would be vicious only if one paper claimed to *derive* the budget by citing the other; neither does — both cite M10 as open. The honest statement is that $N \leq 3$ rests on the coupling clause and M6 (= R-INT), and the cohort’s job is to close M6, not to launder it between papers. We have made the shared-open-problem status explicit at §F.

Objection 37 (O-GRIP-37) (the skeptic of citation loops)

“T2.1’s Diophantine uniqueness is imported from A2 and A3, the modified-TOV from A5, the $\mathcal{H}^2$ normalization from A2. The four pillar theorems of A6 are, at their load-bearing steps, citations to four other single-author preprints, none re-derived in this document. A reader cannot verify the core results without four unavailable companions.”

Response. The dependency is exactly as described and is the explicit design of the cohort (the inward coherence map, §F): A6 consumes the bulk geometry (A2), the anomaly closure (A3), and the boundary law (A5), and builds the pillar formalization on them. We concede the consequence the reviewer names: A6 is *not* self-contained — its load-bearing inputs are theorems of A2/A3/A5, and verifying them requires those papers. This is the cost of a partitioned program, and the mitigation is that each import is named at its point of use with the specific theorem cited, so the dependency is auditable rather than hidden. We do not claim A6 stands alone; we claim it stands on named, located results of the upstream papers, which a reader must consult. A monograph would re-prove them; a cohort cites them. The reviewer is right to weigh that trade-off, and the inward map is the instrument for doing so.

Objection 38 (O-GRIP-38) (the philosopher of physics)

“The whole cohort cross-validates internally — you call the convergence of the aesthetic and empirical routes ‘one of the strongest internal consistency checks.’ But internal

consistency among co-authored, mutually-citing preprints is not external evidence. The framework risks being a closed deductive island that agrees with itself by construction.”

Response. The epistemic point is correct and we accept it: agreement *within* the cohort is consistency, not confirmation, and the paper should not present internal convergence as empirical support. The genuine external exposure of the whole program is the small set of falsifiable numbers — $\alpha^{-1} = 137$, $M_{\text{lock}} \approx 3.85$ TeV, the kaxion line at $8.1 \mu\text{eV}$, $\Omega_m = 1/\pi$ — which face data independently of how internally coherent the derivation is. Internal consistency checks (such as the two routes to the ledger agreeing) are valuable as *bug-detection* — an inconsistency would falsify the construction — but they are not evidence that the construction is true of the world. We concede that any wording presenting internal convergence as confirmation overclaims, and the corrected framing (§54.1) is that the framework’s truth-claim rests on the downstream empirical numbers, with internal coherence as a necessary but non-confirming precondition.

V Conjecture density at the core: M9, M10, M11

Objection 39 (O-GRIP-39) (the skeptic of conjecture-laden proofs)

“Three named open problems sit directly under the central results: M9 (the norm choice that makes the ledger unique), M10 (the budget that forces $N \leq 3$), M11 (the compensator smoothness that closes the conservation iff). Yet you open the conclusions claiming the framework is ‘mathematically complete, all nine theorems proved.’ That claim is contradicted three subsections later.”

Response. The contradiction is real and the “mathematically complete” phrasing is withdrawn. The accurate statement, which the open-problems triage (§57) already supports, is: the nine dual-layer sub-theorems are proved *conditional on* three named hypotheses — M9 (the MSMF norm), M11 (the compensator monodromy), and, for the upper bound, the coupling clause with M6 (= R-INT) through which M10’s met demand now conditions T4.1 — each underwriting one of the three deliverables (ledger uniqueness, the upper bound, the conservation closure). So the framework is *conditionally* complete: every theorem has a proof modulo a stated, located conjecture, and the conjectures are the honest residue. We should not, and no longer, say “complete” without the conditioning ledger of §57 attached. The reviewer is right that the opening sentence of the conclusions oversold a conditional result, and it is corrected to match the triage that follows it.

Objection 40 (O-GRIP-40) (the number theorist)

“M9 admits the ledger’s uniqueness is proved only under the ℓ^1 norm, with a norm-independence theorem ‘missing.’ If uniqueness is norm-dependent, then $(1, 3, 5; 7, 35)$ is the minimizer of a *chosen* complexity functional, not a forced one. ‘Minimal’ is doing covert work.”

Response. Conceded. The MSMF minimality that selects the ledger is, as M9 states, established for the ℓ^1 complexity norm, and while other natural norms appear to select the same minimizer,

a norm-independence theorem is not in hand. So “minimal” is, strictly, “minimal with respect to the ℓ^1 norm,” and the uniqueness is conditional on that choice. We do not claim the ledger is the minimizer under *every* complexity measure; we claim it under the stated one and flag the independence as open (M9). The reviewer is right that “forced” is too strong while the norm is a choice; the honest reading is “the unique ℓ^1 -minimal admissible ledger,” and whether the choice of norm can be removed is exactly what M9 asks.

Objection 41 (O-GRIP-41) (the skeptic)

“M10’s own phrasing — derive 3π ‘independent of $N = 3$ at every intermediate step, so as to rule out any circularity’ — is an admission that the current derivation is *not* independent of $N = 3$. You have documented your own circularity and labeled it future work.”

Response. That was a fair reading of M10’s demand, and the demand has since been met on its own terms: the capacity theorem T-N3-6 of Arisaka (2026, P2-N3) derives the budget through the stratum’s quantum geometry, with $K = 35$ and $J = 7$ appearing at no step — exactly the N -independence the wording required. The seam-count route this paper documented as circular remains circular; it is simply no longer the route. We regard documenting the circularity precisely — rather than obscuring it — as the correct scholarly posture for a load-bearing gap, but we do not pretend it is anything other than a gap. The upper bound $N \leq 3$ is conditional on the coupling clause and M6 (= R-INT); the formerly documented circularity is dissolved by the external derivation; closing M6 is the remaining task. A reader who concludes “the upper bound is earned modulo one named lemma” has read the paper correctly, and M10 is written so that they will.

W Relationship to string theory, 6D supergravity, holography, and Randall–Sundrum

Objection 42 (O-GRIP-42) (the string theorist)

“You frame GiFour as superior to string theory because $d = 6$ is ‘forced by the magic-dimension ladder’ while string theory ‘chooses’ $d = 10$. But $d = 10$ is *derived* from worldsheet anomaly cancellation, whereas your $d = 6$ rests on the magic-dimension ladder, which is an algebraic preference for normed division algebras, not a consistency requirement. You criticize string theory for a choice you also make.”

Response. The reviewer is right that the magic-dimension ladder is a mathematical organizing principle, not a worldsheet-style consistency derivation, and the comparison should be drawn at that level (§54.5). The honest distinction is not “forced vs chosen” — both frameworks make a structural commitment — but *which rung and why*: string theory takes $d = 10$ because supersymmetry needs the octonionic spinor identity to close, while GG-Theory takes $d = 6$ as the unique magic dimension realizable *without* supersymmetry (the quaternionic rung). Neither

is a pure consistency theorem in the other’s sense. We concede that any wording calling our $d = 6$ “forced” while calling string theory’s $d = 10$ “chosen” applies a double standard, and the corrected comparison states that both rest on a structural principle, the discriminating question being the empirical content each dimension is asked to carry (the integer ledger, here) rather than the dimension count itself.

Objection 43 (O-GRIP-43) (the 6D-supergravity specialist)

“A 6D theory with H_3 , $X_4 = \text{tr } R^2 - \text{Tr } F^2 - \dots$, and a Green–Schwarz inflow term is *exactly* the well-studied 6D (1,0) supergravity anomaly setup (Green–Schwarz–Sagnotti, Sagnotti). Your BI-6 is the standard 6D GS mechanism, and the paper barely cites this literature. The new claims must be located relative to known 6D anomaly results, or they look like rediscovery.”

Response. Accepted, and the location is the same one GG-A3-BI6 (Arisaka, 2026, A3-BI6) draws: the Green–Schwarz mechanism and the modified Bianchi identity are standard, borrowed wholesale, and the engagement with the 6D anomaly literature belongs in A3, where the descent and global completion are proved. A6 *uses* that machinery and does not re-derive it, so the present paper’s thin citation of the 6D-SUGRA literature reflects the division of labor, not a claim of originality for the GS mechanism. What is new — and what must be located relative to the standard results — is the *four-component* X_4 with the integer trace-ratio lock and the constitutional forcing of the gauge content, which is A3’s contribution, not a rediscovery of GS. We concede A6 should signpost the reader to A3 (and through it to the Green–Schwarz–Sagnotti literature) more prominently at §16, and that the GS mechanism itself is in no way claimed as original here.

Objection 44 (O-GRIP-44) (the AdS/CFT specialist)

“You repeatedly invoke ‘holographic descent’ and bulk-to-boundary maps, but the bulk is a warped product with a real boundary (Randall–Sundrum-like), not asymptotically AdS — there is no large- N CFT, no GKP–Witten dictionary, no conformal boundary. Using ‘holographic’ for ordinary KK zero-mode reduction is misleading branding.”

Response. The reviewer is right and the paper should say plainly what the operation is: “holographic descent” here means Kaluza–Klein zero-mode projection plus boundary-term variation onto the four-dimensional boundary of the warped product (§9.2), in the holographic-renormalization-group *sense* in which the warp (radial) coordinate is an energy scale — not AdS/CFT in the GKP–Witten sense. There is no asymptotically-AdS region, no large- N CFT, and no conformal boundary, and we do not claim a holographic duality. The word “holographic” is used for the radial-as-scale intuition and the boundary projection, and we concede that without the explicit disclaimer it reads as a claim of AdS/CFT, which it is not. The clarifying statement — “holographic descent = KK zero-mode projection + boundary variation, with the warp as RG scale” — is now at the first use (§54.7, §54.8).

Objection 45 (O-GRIP-45) (the Randall–Sundrum specialist)

“You claim the hierarchy $E_{\text{GUT}}/E_{\text{Pl}} \sim e^{-35}$ is forced by the ledger $K = 35$ and superior

to RS’s free warp slope. But pinning the warp exponent to the integer 35 from a Diophantine ledger is a choice of boundary condition — exactly the data RS leaves free. Why is e^{-K} with $K = 35$ a derivation and RS’s warp not?”

Response. The honest answer is that the warp exponent is forced *relative to the ledger*, and the ledger is the constitutional commitment — so the comparison with RS is not “derived vs free” in the abstract but “tied to the integer ledger vs continuously tunable.” In RS the warp slope kr_c is a continuous parameter fixed only by matching the observed hierarchy; in GG-Theory the warp depth is the integer $K = 35$ of the OGN-5 ledger (§54.5), so it cannot be slid continuously without breaking the same ledger that fixes $\alpha^{-1} = 137$ and the family count. That is the substantive difference: not that ours descends from nothing, but that it is locked to an integer that simultaneously fixes other observables, whereas RS’s is a free continuous dial. We concede that calling ours a “derivation” while the ledger itself is a constitutional input overstates it; the defensible claim is “integer-locked and over-constrained” versus “continuously free,” and that is the comparison the corrected text draws.

X Method, scope, and specific technical soft spots

Objection 46 (O-GRIP-46) (the skeptic)

“The dual-layer ‘compatibility relations’ $\mathfrak{C}_{\text{TGG}}$, $\mathfrak{C}_{\text{CCI}}$, $\mathfrak{C}_{\text{GFF}}$, $\mathfrak{C}_{\text{GRIP}}$ are each described as ‘a one-line proposition with the proof: ...’ A one-line proposition with the proof folded into the same sentence — especially when the ‘via’ clause is exactly the open M11 — is an assertion, not a proof. The dual-layer architecture multiplies theorem labels without multiplying proven content.”

Response. The criticism of the presentation is fair: the compatibility relations (Eqs. (45)–(176)) link each pillar’s algebraic layer to its operational layer, and three of the four are genuinely one-line consequences of the corresponding pillar theorem, but $\mathfrak{C}_{\text{GRIP}}$ routes through the compensator smoothness and is therefore conditional on M11 — which a one-line statement obscures. We concede that folding “the proof” into the proposition sentence is not a proof and that, for $\mathfrak{C}_{\text{GRIP}}$ especially, the dependence on M11 must be stated in the proposition, not assumed in a clause. The corrected treatment gives each compatibility relation its own short displayed proof (or, for $\mathfrak{C}_{\text{GRIP}}$, an explicit “conditional on M11”), so that the dual-layer labels correspond to checkable content rather than to relabeled assertions. The reviewer is right that label-count is not proof-count.

Objection 47 (O-GRIP-47) (the working physicist)

“Why 201 pages and twelve Parts for content the paper itself calls ‘austere formalization, not a new theory’? The length, the abbreviation glossary, the ‘fourth scientific revolution’ framing, and the Friston free-energy / neuroscience bridge are scope-inflation around a core that reduces to standard 6D GS inflow plus a Diophantine

ledger plus a $\mathbb{CP}^{N_g-1}$ Morse cascade, three of whose key steps are open conjectures.”

Response. There is a real tension between “austere formalization” and 201 pages, and we accept part of the criticism. The length is driven by the dual-layer treatment of four pillars with worked examples and the cohort-standard front and back matter; the core mathematical content is indeed compact — the composite-connection conservation law, the integer ledger (imported), and the $\mathbb{CP}^{N_g-1}$ cascade (conditional, today, on M11’s general statement; M10 resolved externally). The “fourth scientific revolution” framing and the free-energy/neuroscience appendix are motivational and interdisciplinary scope, explicitly marked as such, and a reader who wants only the theorems can take the pillar Parts and skip them. We concede that the framing material should not be mistaken for load-bearing content, and that the paper’s defensible core is exactly the compact list the reviewer gives, three of whose steps are conditional. The honest one-line summary of A6 is: a formal organization of standard inflow plus a Diophantine ledger plus a Morse cascade, with the named conditioning of §57 at the load-bearing steps — and the length is presentation, not additional proof.

Objection 48 (O-GRIP-48) (the differential geometer)

“The chain ‘three quaternionic structures of $\mathbb{H}^2 \Rightarrow$ family triplicity $N = 3$ ’ is asserted, not derived: $\mathbb{H}^2$ has an $SU(2) \cong Sp(1)$ of complex structures (a two-sphere’s worth), not three discrete ones. Mapping ‘ I, J, K ’ to three families is a numerological coincidence ($\Im\mathbb{H}$ is three-dimensional) dressed as structure.”

Response. Conceded, and the paper says it is not a derivation: the quaternionic “triplicity” picture (§8) is a suggestive heuristic, explicitly flagged as “not yet a closed derivation,” and the reviewer’s point that the quaternionic complex structures form a two-sphere, not three discrete objects, is exactly why it cannot be load-bearing. The actual $N = 3$ result does *not* rest on the I, J, K picture — it rests on the CP lower bound, the signature parity, and the (conditional) budget. The quaternionic reading is decoration that gestures at why three is natural; it is not a proof and the paper does not use it as one. We agree it should be clearly cordoned off as heuristic so no reader mistakes the three-dimensionality of $\Im\mathbb{H}$ for a derivation of three generations, and §8 now does so.

Objection 49 (O-GRIP-49) (the representation theorist)

“A footnote flags ‘ $\dim \text{Spin}(1, 5) = 15$, not 12 as a careless count might suggest’ — but $15 = \binom{6}{2}$ is elementary. Flagging 12 as a tempting wrong answer suggests the magic-ladder bookkeeping may rest on similarly ad hoc counts. Where do $K = 35$ and $J = 7$ come from group-theoretically, not just Diophantine fitting?”

Response. The footnote is over-cautious and we will trim it; $\dim \text{Spin}(1, 5) = 15$ needs no defending. On the substantive question: K and J are *not* claimed to be dimensions of a group — they are entries of the OGN-5 integer ledger fixed by the Diophantine cascade ($K = LJ = M^2 + N^2 + L^2 = 35$, $J = 2N + 1 = 7$), with $L = 5$ from $SU(5)$ and $N = 3$ from the trace ratio (§15, and the full derivation in (Arisaka, 2026, A3-BI6)). So they are arithmetic outputs of the ledger, not group dimensions, and the reviewer is right that we should

not let magic-ladder bookkeeping imply otherwise. The honest statement is that K and J are ledger integers with a Diophantine, not a representation-theoretic, origin; any group-dimension coincidence (e.g. $35 = \dim \mathrm{SU}(6)$) is not used and is not claimed. We have removed the misleading footnote and stated the ledger’s arithmetic provenance plainly.

Objection 50 (O-GRIP-50) (the open-quantum-systems theorist)

“You invoke ‘the boundary is CPTP at the foundational level’ as the deep novelty, but a CPTP map requires a specified environment and dilation. You point to a Stinespring-dilation form but never construct the dilation or the environment Hilbert space. ‘The boundary is CPTP’ is a slogan; CPTP-ness is a property of a map you have not written down.”

Response. The construction the reviewer asks for is GG-A5-GDOS’s, not A6’s, and the paper should point there rather than gesture. The GDOS boundary channel — the CPTP map, its Stinespring dilation, the environment (the integrated-out bulk/KK modes), and the GKLS generator — is constructed in (Arisaka, 2026, A5-GDOS) (its open-system theorems); A6 *uses* that channel to define the free-energy flow and does not reconstruct it. So “the boundary is CPTP” is not an A6 slogan standing on nothing; it is an A5 theorem A6 imports. We concede that A6, read alone, asserts CPTP-ness without exhibiting the dilation, and that the reader must go to A5 for the actual map. The honest signpost — “the CPTP channel and its dilation are constructed in A5; here it is used” — now appears at §26, and A6 claims no independent construction of the channel.

Objection 51 (O-GRIP-51) (the mathematical physicist)

“Hopkins–Singer differential cohomology is deployed to make the ‘ $[X_4]$ descends to $p_1(\Sigma)$ ’ step precise, but the decisive η -invariant non-degeneracy lemma is admitted to be a sketch deferred to a paper in preparation (M6). Importing heavy machinery whose decisive lemma is unproven adds apparent rigor without actual rigor.”

Response. Fair, and the same concession as Objection 14 from the machinery side. The Hopkins–Singer differential-cohomology framework (§19) is the correct setting for the integral-class descent, and the framework is used honestly to *state* the non-degeneracy condition precisely — but the decisive η -invariant computation that establishes it on the bulk is open problem M6, sketched here and deferred. So the machinery makes the claim *precise and checkable*, which is a genuine gain, but it does not by itself *prove* the lemma, which is the reviewer’s point. We do not claim the descent step is proved; we claim it is reduced, in the right framework, to a stated η -invariant computation that is open. Apparent rigor and actual rigor coincide only once M6 lands, and the paper marks the gap rather than letting the machinery imply closure.

Objection 52 (O-GRIP-52) (the editor)

“Given all the above — a central result conditional on M10, a conservation iff conditional on M11, a parity leg conditional on M6, a norm-dependent uniqueness (M9), an equivalence with a sketched converse, and load-bearing imports from four companion

papers — what, precisely, does A6 prove *unconditionally and on its own*?

Response. A direct question deserves a direct ledger. Unconditionally and within A6: (i) the existence and holonomy-fixed uniqueness of the composite connection (Theorem 10.1, modulo the stated category); (ii) the gravitational half of the Master Conservation Law as the contracted Bianchi identity, and the *forward* direction of the four-sector iff (Theorem 18.4); (iii) the per-step and total cascade actions πk and $\pi N_g(N_g - 1)/2$ as symplectic facts on $\mathbb{CP}^{N-1}$ (Theorem 32.1, Proposition 30.11); (iv) the conditional implications themselves — “if the budget is 3π then $N \leq 3$,” “if detailed balance then $dF/dt \leq 0$,” “if M11 then the conservation iff closes globally.” What A6 does *not* prove unconditionally: the budget (M10), the compensator monodromy (M11), the parity non-degeneracy (M6), norm-independence (M9), the equivalence converse, and the imported ledger/anomaly results (A2/A3/A5). So A6’s own unconditional content is the connection theorem, the conservation forward-direction, the cascade arithmetic, and a set of clearly-stated conditional implications. The reviewer is right to ask, and the honest answer is that A6 is a conditional formalization whose unconditional core is the structure, not the numbers.

Objection 53 (O-GRIP-53) (the senior community member)

“Stepping back: is GiFour a discovery or a bookkeeping scheme? You have reorganized existing cohort content into four ‘pillars’ and a bridge, added three conjectures, and claimed a unification. Why should the community invest in this organization rather than read A2/A3/A5 directly?”

Response. The fair answer is that Gi-4 is a bookkeeping scheme with a thesis, and its value stands or falls on whether the thesis is right — which is testable. The bookkeeping is the claim that the conservation content of the cohort is carried by exactly four geometric invariants of one composite connection, related by a fixed implication chain $TGG \Rightarrow CCI \Rightarrow GFF \Rightarrow GRIP$; the thesis is that this is the *minimal and complete* set, which T5 attempts (conditionally) to establish via the equivalence to the constitution. If T5’s converse fails, Gi-4 is a convenient but non-canonical repackaging and a reader should indeed go to A2/A3/A5; if it holds, Gi-4 is the canonical statement of what those papers collectively prove, in a form that makes the dependencies and the open problems (M9–M11) visible in one place. We do not claim it is a new discovery about nature — T5 explicitly denies that — only a claim about the *structure* of the existing derivation. The community should invest in it exactly to the extent that an explicit, auditable, conjecture-flagged organization of a large program is worth more than its scattered parts, and no further. That is a modest claim, and it is the one the paper actually makes once the overclaims this section has conceded are removed.

Closing remark on the objections. A framework that claims to read the conservation laws of physics off a single composite connection, and to force the family number by an instanton-action budget, must expect — and should welcome — the hardest scrutiny. We have set out fifty-three objections above, in eleven clusters: the composite connection and TGG (Objections 1–6); the integer ledger and the modified-TOV falsifier (7–11); the Morse and second-law content of GFF (12–17); the cascade and the budget (18–23); the weight ladder and the falsifiers (24–28); the constitutional-

equivalence theorem T5 (29–32); the “no new axiom” claim and the promoted hypotheses (33–35); cohort interdependence and circular citation (36–38); the conjecture density M9–M11 (39–41); the relationship to string theory, supergravity, holography, and Randall–Sundrum (42–45); and method, scope, and the technical soft spots (46–53). The load-bearing conditioning — M11’s general monodromy statement (the conservation closure) and M9’s canonicity question (the norm choice), with M10 resolved externally and M6 discharged on the benchmark — recurs as the honest residue, and we have answered every objection by saying exactly where the framework is unconditional, where it is conditional, and where it depends on a companion paper. The framework’s standing offer to be wrong is the integer ledger and its forced numbers; its standing admission of what it has not yet proved is the pair M9 and M11; and the present section is the record of both.

Part XII

References

Cited works, each with a relevance note linking it to the present paper

Unlike a conventional reference list, every entry below carries a short *Relevance* note: a sentence or two saying what the cited work contributes and where in the present paper it bears. The aim is to turn the bibliography from a bare list of sources into a navigable companion — a reader can move through the references and see, at a glance, which theorem, table, or section each one underwrites. The notes are written for two readers at once. For the newcomer, they trace the intellectual lineage of the geometric backbone: from the foundations of gauge theory, differential cohomology, Morse theory, and open-quantum-systems dynamics, to the twenty-two-paper GG-Theory cohort in which the present paper supplies the Gi-4 Pillars — the geometric backbone on which the GDOS boundary law and its consequences rest. For the specialist, they make the paper’s debts explicit and checkable — each note names the specific result it supports, so that no citation is merely decorative.

The references are organized in two groups — first the author’s GG-Theory papers (Section A below), then all other works (Section B), each in alphabetical order by first author. The cohort papers (GG-A1-Constitution–GG-A11-Origin, with the C- and P-series) are cited under the canonical keys and titles fixed by the cross-paper coherence map of the GG-Theory program.

A. Papers by Katsushi Arisaka (the GG-Theory cohort)

Arisaka, K. (2026). *The GG-Constitution: Four Axioms and Seven Postulates of the GG-Theory Program, From Which GG-6, BI-6, OGN-5, and GDOS Descend as Theorems*. Arisaka-GG-A1-Constitution preprint, available at <https://preprints.arisaka-gg.org/a1-constitution/>. doi:10.5281/zenodo.22063867.

Relevance. The parent constitutional paper of the GG-Theory program. Establishes the four axioms AX-1 LC, AX-2 GGEP, AX-3 OGGE, AX-4 MSMF, together with the seven postulates POS-1–POS-7 (GI, UEP, AF, CP, RG, OGN, MF) and the central Arisaka equation $J_E/Q_B = J_e = J_B = J_{Qi}$ that ties the four boundary currents through one bioenergetic conversion factor. Every theorem in the present paper that names its constitutional source inline cites this paper at the point of invocation; the four Gi-4 pillars TGG/CCI/GFF/GRIP are derived as consequences of the four axioms acting on the 6D bulk. First cited in §1.

Arisaka, K. (2026). *GG-Theory with GG-6: Geometry–Gauge Locking at the 6D Bulk (GG-6) at the Heart of Grand Geometric Theory (GG-Theory)*. Arisaka-GG-A2-GG6 preprint, available at <https://preprints.arisaka-gg.org/a2-gg6/>. doi:10.5281/zenodo.22064298.

Relevance. Establishes the 6D bulk geometry $M_6 = \mathbb{R}^{1,3} \times I_r \times S_\theta^1$ and the role of $\text{Spin}(1, 5)$. The bulk arena on which TGG is constructed in Part II of the present paper. First cited in §1.

Arisaka, K. (2026). *Anomaly Freedom of BI-6: The Modified Bianchi Identity (BI-6) for the Six-Force, Six-Dimensional Bulk (GG-6)*. Arisaka-GG-A3-BI6 preprint, available at <https://preprints.arisaka-gg.org/a3-bi6/>.

[doi:10.5281/zenodo.22064300](https://doi.org/10.5281/zenodo.22064300). **Relevance.** Proves the BI-6 modified Bianchi identity $dH_3 = X_4$ on the 6D bulk and derives the OGN-5 integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$. The differential identity that the algebraic CCI of Part III closes. First cited in §1.

Arisaka, K. (2026). *From Euclid to GG-6: The Realization of Euclid's Elements in Six Dimensions — GU-5, OGN-5, and a Parameter-Free Reconstruction of Modern Physics*. Preprint, available at <https://preprints.arisaka-gg.org/a4-euclid/>.

[doi:10.5281/zenodo.22064302](https://doi.org/10.5281/zenodo.22064302).

Relevance. The historical companion of the GG-6 Trio. Pedagogical commentary on the long chain from Euclid through Newton, Riemann, Maxwell, and Einstein to the GG-Theory stage. Cited here whenever the historical positioning of Gi-4 against its predecessors is invoked, and as one of the companion papers carrying part of the resolution of the cosmological-constant problem (PS-GRIP-12) and the pedagogical/historical synthesis (PS-GRIP-24). First cited in §1.8.

Arisaka, K. (2026). *GDOS: Geometric Dynamics of Open Systems — 6D GG-6 Bulk Projected to 4D Boundary for Quantum Measurement, General Relativity, and QFT*. Arisaka-GG-A5-GDOS preprint, available at <https://preprints.arisaka-gg.org/a5-gdos/>.

[doi:10.5281/zenodo.22064304](https://doi.org/10.5281/zenodo.22064304). **Relevance.** Develops GDOS as the open-system framework for boundary observation. Introduces the boundary projector, the GFF, and the GRIP cascade in their first informal form; the present paper formalizes these as Pillars 3 and 4 of Gi-4. Its Part-VI radiative completion, quantitative-closure, and covariant-kernel subsections supply the derived mobilities, coefficients, fluctuation–dissipation pairing, and the causal runaway-free kernel that Example 4 (§36.4) reads in GRIP vocabulary. First cited in §1.

Arisaka, K. (2026). *The GRIP on GDOS — GRIP: the Geometric Engine of Symmetry Breaking, from the Big Bang to Mind*. Arisaka-GG-A6-GRIP preprint, available at <https://preprints.arisaka-gg.org/a6-grip/>.

[doi:10.5281/zenodo.22064306](https://doi.org/10.5281/zenodo.22064306). **Relevance.** The prior integrated form of the present paper: it introduced GRIP, GFF=GEP+CEP, SIP and the DGI/GSSB/GA cascade as a single dynamical theorem. The present paper reorganizes this content around four explicitly named pillars (TGG, CCI, GFF, GRIP). First cited in §4.1.

Arisaka, K. (2026). *The Qi-4 Pillars — Qi-QUA, Qi-EP, Qi-CON, Qi-HAL: The Quantum-Information Backbone of GDOS*. Arisaka-GG-A7-Qi4 preprint, available at <https://preprints.arisaka-gg.org/a7-qi4/>.

[doi:10.5281/zenodo.22064308](https://doi.org/10.5281/zenodo.22064308).

Relevance. The information-theoretic counterpart of the present paper. Together with the present paper and GG-A5-GDOS, it completes the GDOS trio of foundational papers. First cited in §1.4.

Arisaka, K. (2026). *Quantum Mechanics Derived from GG-Theory as Projection from 6D Bulk to 4D Boundary*. Arisaka-GG-A8-QM preprint, available at <https://preprints.arisaka-gg.org/a8-qm/>.

doi:10.5281/zenodo.22064310.

Relevance. The quantum-mechanical companion of the cohort: it recasts quantum mechanics in the GDOS/OGGEP language whose geometric backbone is the present Gi-4 paper, and carries the measurement-problem carry-over discussed in the Discussion Part. First cited in §4.1.

Arisaka, K. (2026). *The Standard Model as Theorem of GG-Theory: Its 28 Parameters as Observations of One Six-Dimensional Geometry*. Arisaka-GG-A9-SM preprint, available at <https://preprints.arisaka-gg.org/a9-sm/>.

doi:10.5281/zenodo.22064312.

Relevance. Derives the 28 Standard-Model parameters from the OGN-5 ledger in an explicit codebook form. The codebook formulas are re-derived in the present paper as cascade depths of GRIP on the family-projector moduli. First cited in §1.

Arisaka, K. (2026). *The Standard Cosmology as Theorem of GG-Theory: The Cosmic Inventory as Observations of One Six-Dimensional Geometry*. Preprint, available at <https://preprints.arisaka-gg.org/a10-cosmos/>.

doi:10.5281/zenodo.22064314.

Relevance. Derives the cosmological parameters $\Omega_m = 1/\pi$, $\Omega_\Lambda = 1 - 1/\pi$, $\Omega_b = 1/(2\pi^2)$, $\Omega_{DM} = 1/\pi - 1/(2\pi^2)$ from the GG-Constitution and the 6D bulk. Cited here as the companion paper carrying the resolution of cosmological initial conditions (PS-GRIP-10) and the cosmological-constant problem (PS-GRIP-12). First cited in §1.

Arisaka, K. (2026). *The Geometric Inevitability of Cosmos, Life, and Mind: The Self-Referential Closure of the GG-Theory Program*. Arisaka-GG-A11-Origin. Preprint, available at <https://preprints.arisaka-gg.org/a11-origin/>.

doi:10.5281/zenodo.22064319. No result of the present paper depends on it.

Relevance. The closing capstone of the program. Defines the Ground Geometric Attractor (GGA) and proves T-GGA-as-Ground-State — the ground-state refinement of the GRIP cascade adopted into the present paper's runtime vocabulary as Definition 31.4 — together with the GRIP-In/GRIP-Out regime theorems and the eighteen-layer cascade of which the runtime built here is the engine. First cited in §4.1.

Arisaka, K. (2026). *The GG-6 Lock Cluster: A Non-Spatial Kaluza–Klein Resonance near 3.85 TeV as the Standing LHC Falsifier of GG-Theory*. Arisaka-GG-P1-KK preprint, available at <https://preprints.arisaka-gg.org/p1-kk/>.

doi:10.5281/zenodo.22064321.

Relevance. The principal collider falsifier of GG-Theory. Predicts a Kaluza-Klein resonance cluster at the lock scale $M_{\text{lock}} = E_{\text{Pl}}/2e^{35} \simeq 3.849 \text{ TeV}$. In Gi-4 vocabulary, this is the first excitation above the GA of the radial-scale stratum. First cited in §1.1.

Arisaka, K. (2026). *Why Three Generations: A Three-Leg Derivation of $N_g = 3$ from BI-6,*

MSMF, and GDOS Morse Stability. Arisaka-GG-P2-N3. Preprint, available at <https://preprints.arisaka-gg.org/p2-n3/>.
doi:10.5281/zenodo.22064323.

Relevance. The three-leg $N=3$ proof. In the present paper, the hypotheses $\mathcal{H}_1, \mathcal{H}_2$ of that paper are reduced to Theorems 27.1 and 27.2 of Gi-4 modulo the η -invariant computation M6. Its Appendix H carries the Hopkins–Singer pushforward of the BI-6 seam and proves it equal to the boundary’s Berry line bundle (the κ -bridge theorem) on the benchmark background — the rigorous bulk implementation of the descent used in Theorems 27.1 and 27.2, discharging this paper’s residue M6 (= R-INT) there, with the warped and relative refinements declared open. First cited in §1.

Arisaka, K. (2026). *The Electron as the Ground State of Nature: The Six Oldest Questions about the Electron, Closed and Fenced within GG-Theory*. Arisaka-GG-P3-Electron preprint, available at <https://preprints.arisaka-gg.org/p3-electron/>.
doi:10.5281/zenodo.22064325.

Relevance. The dedicated electron paper whose two-sector partition — the phase sector moves the electron, the scale sector transmutes it, only what sheds energy descends — supplies the stratum fence of Example 4 (§36.4): GRIP-In on the orbital stratum descends motion, GRIP-In on the family stratum descends identity, and the same septuple runs on both. First cited in §4.1.

Arisaka, K. (2026). *The Real Geometry Beneath the Complex Numbers of Physics: Positive Frequency is Holomorphy — from Cardano’s Cubics to Twistor Space, via the Sixth Dimension of GG-Theory*. Arisaka-GG-P4-Real preprint, available at <https://preprints.arisaka-gg.org/p4-real/>.
doi:10.5281/zenodo.22064327.

Relevance. Foundations Branch: the origin-of- i member paper, deriving the single imaginary unit of boundary physics from the two real structures of the GG-A2-GG6 arena (the compact phase circle and the causal clock). For the present paper it is the origin of the boundary’s complex structure whose phase conventions the DGI stage of the GRIP pipeline makes dynamical; cited at roster level. First cited in §4.1.

Arisaka, K. (2026). *The Mirrors of Matter: Why Every Particle Has a Twin, Why Nature’s Mirrors Break, and Why the Universe Is Made of Matter*. Arisaka-GG-P5-CPT preprint, available at <https://preprints.arisaka-gg.org/p5-cpt/>.
doi:10.5281/zenodo.22064329.

Relevance. Foundations Branch: the CPT member paper, reading C, P, and T as orientation reversals of the three real bulk factors and carrying the matter–antimatter door; its mirrors and its law-level arrow of time are kin of the irreversible descent that the runtime built here drives one way (the second law of Theorem T3); cited at roster level. First cited in §4.1.

Arisaka, K. (2026). *The Proton as the Ground Geometric Attractor of Matter: The Eleven Questions of the Proton, Closed and Fenced within GG-Theory*. Arisaka-GG-P6-Proton preprint, available at <https://preprints.arisaka-gg.org/p6-proton/>.
doi:10.5281/zenodo.22064331.

Relevance. Predictions Branch: the fourth member paper and the sister of GG-P3-Electron,

running that paper’s ground-attractor method on the baryonic stratum, with baryon number carried as a winding of the compact geometry rather than as a bookkeeping label. Its subject is a selection outcome of the runtime built here — the attractor a GRIP descent settles into once the landscape has stopped changing, which is the fixed-regime branch of §31.8; cited at roster level. First cited in §4.1.

Arisaka, K. (2026). *The Kaxion: Cold Dark Matter as the Non-Spatial Compact-Angle Pseudoscalar of GG-6*. Arisaka-GG-C1-Kaxion preprint, available at <https://preprints.arisaka-gg.org/c1-kaxion/>.

doi:10.5281/zenodo.22064333.

Relevance. The principal cosmological falsifier of GG-Theory. Predicts a $\approx 8.1 \mu\text{eV}$ Kaxion-like dark-matter candidate, the tensor remnant, read here as the slowly-rolling GA of a one-dimensional periodic stratum whose occupant is left open (Remark 11.1). First cited in §1.

Arisaka, K. (2026). *The Hubble Tension in the Light of GG-Theory: The Hubble Constant as One Reading of a Derived Expansion History $H(t)$* . Arisaka-GG-C2-Hubble preprint, available at <https://preprints.arisaka-gg.org/c2-hubble/>.

doi:10.5281/zenodo.22064336.

Relevance. A consequence-side reading of the boundary law: the early- and late-time Hubble constants are the two ends of the scale corridor, a cosmological application of the CCI-fixed X_{GS} developed here. First cited in §4.1.

Arisaka, K. (2026). *Black-Hole Entropy under GG-Theory: The GDOS Reading of the Horizon as an OGGE Port — with the Exact Bekenstein–Hawking $A/4G_N$ Derived from Induced Gravity*. Arisaka-GG-C3-BH preprint, available at <https://preprints.arisaka-gg.org/c3-bh/>.

doi:10.5281/zenodo.22064338.

Relevance. Treats the black-hole horizon as an OGGE registration port; the induced-gravity boundary stress tensor of TGG (Theorem T1.2 here) is the bulk source whose horizon reduction yields the Bekenstein–Hawking area law. First cited in §4.1.

Arisaka, K. (2026). *Symmetry, Conservation, and Geometry: Seven Steps from Noether, Bianchi, to Hopkins–Singer, and the Ledger the Open-System Turn Pays Out*. Arisaka-GG-W1-Symmetry preprint, available at <https://preprints.arisaka-gg.org/w1-symmetry/>.

doi:10.5281/zenodo.22064340.

Relevance. The cohort’s symmetry review, and the source of the sentence on which this paper’s two-fiber reading rests: symmetry breaking as the one-way *reading* of holonomy data — the alternatives held by the fiber that remembers, the selection performed by the fiber that forgets. Its ladder ends at protection without symmetry, which is the register in which the GGA’s labels (charges, windings, parities) are protected here. First cited in §31.8.

Arisaka, K. (2026). *The Circle That Remembers, the Interval That Forgets: GG-Theory and the Sixteen Problems of Particle Physics and Cosmology*. Arisaka-GG-W2-Holonomy preprint, available at <https://preprints.arisaka-gg.org/w2-holonomy/>.

[doi:10.5281/zenodo.22064342](https://doi.org/10.5281/zenodo.22064342).

Relevance. The dedicated treatise on the fork this paper’s engine runs on — holonomy requires closed loops, and an interval has none — proving the group law for circle transport and the semigroup law for interval transport, and deploying the two-fiber anatomy across sixteen standing problems. Its engine section carries the stroke-level mapping (GRIP-In the circle’s, GRIP-Out the interval’s) that Subsection 31.9 completes with the stage-level custody statement. First cited in §31.8.

B. Other references

Abbott, B. P., *et al.* (LIGO Scientific and Virgo Collaborations) (2016). *Observation of gravitational waves from a binary black hole merger*. Physical Review Letters, 116, 061102.

Relevance. The first directly observed inspiral–merger–ringdown waveform. In Example 4 (§36.4) the three phases are read as the GRIP loop drawn end to end: GRIP-In descent on the binary’s stratum, a landscape-reorganizing terminus, and the new stratum’s GRIP-In relaxing to the Kerr attractor. First cited in §36.4.

Adler, S. L. (1969). *Axial-vector vertex in spinor electrodynamics*. Physical Review, 177, 2426–2438.

Relevance. The original derivation of the chiral anomaly in QED, providing the historical foundation for what becomes the four-form anomaly polynomial X_4 of CCI in the present paper. First cited in §13.1.

Aharonov, Y., & Bohm, D. (1959). *Significance of electromagnetic potentials in the quantum theory*. Physical Review, 115, 485–491.

Relevance. The original treatment of holonomy as a physical observable. In the present paper, Wilson-line holonomies of the TGG composite connection on the compact θ -direction (Section 5.5) are the natural generalization of the Aharonov–Bohm phase. First cited in §C.3.

ALEPH, DELPHI, L3, OPAL Collaborations & LEP Electroweak Working Group (2006). *Precision electroweak measurements on the Z resonance*. Physics Reports, 427, 257–454.

Relevance. Establishes the LEP empirical bound $N_\nu = 2.984 \pm 0.008$ on the number of light neutrino species, the most direct experimental confirmation of $N_{\text{gen}} = 3$. This is the empirical input to **T2**’s lower bound $N \geq 3$. First cited in §54.2.

Alvarez-Gaumé, L., & Witten, E. (1984). *Gravitational anomalies*. Nuclear Physics B, 234, 269–330.

Relevance. Foundational treatment of gravitational anomalies and their cancellation. Provides the technical machinery underlying the gravity sector of X_4 in CCI-ST1. First cited in §13.2.

Arisaka, K. (2026). *GG-Theory Foundation Book: Constitution, GG-6 Trio, GDOS Trio, and 4D-Physics Trio Compiled into a Single Reference*. In preparation, available at <https://preprints.arisaka-gg.org/>.

Relevance. The Foundation Book compiles the constitution and its descent chain into a single

accessible reference. Cited here as the eventual textbook synthesis (PS-GRIP-24) covering the cross-domain implications of GG-Theory (life science, neuroscience, complex systems). First cited in §58.1.

Atiyah, M. F. (1982). *Convexity and commuting Hamiltonians*. Bulletin of the London Mathematical Society, 14, 1–15.

Relevance. The convexity theorem and the perfection of moment-map Morse functions; the connectedness of the level sets of a moment-map component, which extends Frankel’s uniqueness of the floor to the Morse–Bott case. Cited for the uniqueness half of the halting corollary. First cited in §26.1.

Bardeen, W. A. (1969). *Anomalous Ward identities in spinor field theories*. Physical Review, 184, 1848–1859.

Relevance. The non-Abelian generalization of the chiral anomaly. Together with Adler (1969), Bell & Jackiw (1969), provides the historical foundation of the four-form X_4 structure in CCI-ST1. First cited in §13.2.

Bell, J. S., & Jackiw, R. (1969). *A PCAC puzzle: $\pi^0 \rightarrow \gamma\gamma$ in the σ model*. Il Nuovo Cimento A, 60, 47–61.

Relevance. The original observation of the axial anomaly in the context of π^0 decay. The historical companion to Adler (1969). First cited in §13.1.

Białynicki-Birula, A. (1973). *Some theorems on actions of algebraic groups*. Annals of Mathematics, 98, 480–497.

Relevance. The cell decomposition attached to the fixed points of a $\mathbb{C}^*$ action on a projective variety; the stable sets of the descent on the family stratum are these cells. First cited in §26.1.

Bianchi, L. (1902). *Sui simboli a quattro indici e sulla curvatura di Riemann*. Rendiconti della Reale Accademia dei Lincei (Series 5), 11, 3–7.

Relevance. The original Bianchi identity $\nabla_{[\mu} R_{\nu\rho]\sigma\tau} = 0$ for the Riemann curvature tensor. A purely geometric statement, derived in 1902 without any reference to dynamics or quantum effects, that is the structural prototype of CCI. First cited in §13.1.

Blumenhagen, R., Körs, B., Lüst, D., & Stieberger, S. (2007). *Four-dimensional string compactifications with D-branes, orientifolds and fluxes*. Physics Reports, 445, 1–193.

Relevance. Comprehensive review of magnetized toroidal compactifications, one of the closed-system attempts to derive $N_g = 3$ discussed in Arisaka (2026, P2-N3, § 12.3). First cited in §54.6.

Boltzmann, L. (1872). *Weitere Studien über das Wärmegleichgewicht unter Gasmolekülen*. Sitzungsberichte der Akademie der Wissenschaften zu Wien, 66, 275–370.

Relevance. The H -theorem and the first statistical-mechanical derivation of the second law of thermodynamics from microscopic dynamics. The structural ancestor of GFF-ST6 (thermodynamic monotonicity) and of the energetic component F_{GEP} of GFF. First cited in §22.1.

Born, M. (1926). *Zur Quantenmechanik der Stoßvorgänge*. Zeitschrift für Physik, 37, 863–867.

Relevance. The Born rule. In the open-system framework of the present paper, the Born rule emerges as the boundary projector selection at the end of the GRIP cascade, rather than being postulated separately. First cited in §54.10.

Bott, R., & Tu, L. W. (1982). *Differential Forms in Algebraic Topology*. Graduate Texts in Mathematics 82, Springer.

Relevance. Standard pedagogical reference for de Rham cohomology, Chern–Weil theory, and characteristic classes used throughout the present paper. The recommended text for graduate students who wish to deepen their understanding of the differential-cohomological machinery underlying the CCI and GFF pillars; cited in Appendix C as the suggested follow-up reading after the pedagogical tour. First cited in §C.

Callan, C. G., & Coleman, S. (1977). *Fate of the false vacuum. II. First quantum corrections*. Physical Review D, 16, 1762–1768.

Relevance. The companion to Coleman (1977) computing the one-loop prefactor of false-vacuum tunneling rates. Cited in Section 33 for the WKB rate formula $\Gamma = Ae^{-S}$ that earlier statements applied to the gates of **GRIP-ST7**, and that the present version says does not apply there: the gates are saddles, not false vacua behind barriers. First cited in §30.1.

Callan, C. G., & Harvey, J. A. (1985). *Anomalies and fermion zero modes on strings and domain walls*. Nuclear Physics B, 250, 427–436.

Relevance. The seminal treatment of anomaly inflow from bulk to boundary. The foundational mechanism that, in modern differential-cohomology form (Hopkins–Singer), drives the descent of X_4 from M_6 onto the boundary moduli in Part IV. First cited in §13.2.

Candelas, P., Horowitz, G. T., Strominger, A., & Witten, E. (1985). *Vacuum configurations for superstrings*. Nuclear Physics B, 258, 46–74.

Relevance. The seminal heterotic-string Calabi–Yau compactification. One of the closed-system family-counting attempts ($|\chi/2| = 3$) compared against in Arisaka (2026, P2-N3, § 12). First cited in §54.6.

Carlen, E. A., & Maas, J. (2014). *An analog of the 2-Wasserstein metric in non-commutative probability under which the fermionic Fokker–Planck equation is gradient flow for the entropy*. Communications in Mathematical Physics, 331, 887–926.

Relevance. Establishes, for the fermionic Fokker–Planck equation, that a GKLS Lindbladian satisfying quantum detailed balance (the GNS-symmetry condition of Kossakowski *et al.*, 1977; the general case is Carlen & Maas, 2017) can be expressed as a gradient flow of quantum relative entropy with respect to a non-commutative Wasserstein-like metric on states. This is the foundational result underlying the gradient-flow representation of the GDOS boundary Lindbladian used in the proof of Theorem 25.3: the boundary moduli-space metric g_Σ is identified with the BKM information metric restricted to the projector manifold. First cited in §25.2.

Carlen, E. A., & Maas, J. (2017). *Gradient flow and entropy inequalities for quantum Markov semigroups with detailed balance*. Journal of Functional Analysis, 273, 1810–1869.

Relevance. The general statement: under the quantum detailed-balance condition of Kossakowski *et al.* (1977), the dissipative part of a GKLS generator is the gradient flow of the relative entropy for a Riemannian metric on states. This is what detailed balance buys in Theorem 25.3 and why T3.2 carries it as a hypothesis. First cited in §25.2.

Coleman, S. (1977). *Fate of the false vacuum: semiclassical theory*. Physical Review D, 15, 2929–2936.

Relevance. The original WKB instanton calculation of false-vacuum decay rates, $\Gamma \propto e^{-S_E}$ where S_E is the Euclidean instanton action. The setting — a local minimum behind a barrier — is what the family stratum’s gates lack, and the reason the cascade exponential $e^{-\pi k}$ of GRIP-ST7 is an equilibrium weight rather than a rate. First cited in §30.1.

Conley, C. (1978). *Isolated Invariant Sets and the Morse Index*. CBMS Regional Conference Series in Mathematics 38, American Mathematical Society, Providence.

Relevance. The fundamental theorem of dynamical systems: every continuous flow on a compact metric space has a complete Lyapunov function, and its Morse decompositions. The theorem the descent of this paper sits inside, and whose hypotheses the second regime leaves (§31.7). First cited in §25.4.

Dirac, P. A. M. (1931). *Quantized singularities in the electromagnetic field*. Proceedings of the Royal Society A, 133, 60–72.

Relevance. The original derivation of magnetic-monopole quantization $\int_{S^2} F/2\pi \in \mathbb{Z}$. This is the historical prototype of the integrality conditions on differential cohomology classes (the CCI integrality $[X_4] \in H^4(M_6; \mathbb{Z})$ in the present paper). Cited in Appendix C as the prototypical physical example of cohomological quantization. First cited in §C.3.

Duistermaat, J. J., & Heckman, G. J. (1982). *On the variation in the cohomology of the symplectic form of the reduced phase space*. Inventiones Mathematicae, 69, 259–268.

Relevance. Exact stationary phase for torus moment maps: the localization of $\int e^{-F/T}$ to the fixed points, which makes the partition function on the family stratum exact, and the normalization that pins the moment-map drop of a circle of weight jump k across a projective line of area π at $k/2$ (Step 5 of Theorem 32.1). First cited in §23.4.

Einstein, A. (1915). *Die Feldgleichungen der Gravitation*. Sitzungsberichte der Preussischen Akademie der Wissenschaften zu Berlin, 844–847.

Relevance. The founding paper of general relativity, recasting gravitation as the curvature of a metric connection on the tangent bundle of spacetime. The structural ancestor of the gravitational component ω of the TGG composite connection. First cited in §6.1.

Fisher, R. A. (1930). *The Genetical Theory of Natural Selection*. Clarendon Press, Oxford.

Relevance. The fundamental theorem of natural selection — the mean fitness rises at the rate of its variance — which is the statement the descent on the family stratum makes in moment coordinates (the replicator remark beside GRIP-ST2). First cited in §30.2.

Frankel, T. (1959). *Fixed points and torsion on Kähler manifolds*. Annals of Mathematics, 70, 1–8.

Relevance. The Morse indices of a moment-map component on a compact Kähler manifold are even, so the function is perfect and a connected stratum has one floor. The theorem behind T3.1’s signature identity and the uniqueness half of the halting corollary. First cited in §26.1.

Franzosa, R. (1988). *The continuation theory for Morse decompositions and connection matrices*. Transactions of the American Mathematical Society, 310, 781–803.

Relevance. Morse decompositions and connection matrices continue under small perturbation of the flow; the sense in which the cascade is robust, and the sense — stability under every deformation within the stratum, because the count is a cohomology class — in which this paper’s claim is stronger. First cited in §30.2.

Freed, D. S., & Hopkins, M. J. (2000). *On Ramond–Ramond fields and K-theory*. Journal of High Energy Physics, 2000(05), 044.

Relevance. Foundational paper on the differential-cohomology classification of higher-form gauge fields and anomaly inflow. Provides the technical machinery underlying the non-degeneracy of the Hopkins–Singer inflow map (Theorem 27.2). First cited in §27.2.

Frigerio, A., & Verri, M. (1982). *Long-time asymptotic properties of dynamical semigroups on W^* -algebras*. Mathematische Zeitschrift, 180, 275–286.

Relevance. Relaxation of every orbit of a quantum dynamical semigroup to a unique faithful stationary state under a commutant condition; the convergence clause of T3.2, which is a relaxation statement and not a consequence of monotonicity. First cited in §26.5.

Friston, K. (2010). *The free-energy principle: a unified brain theory?* Nature Reviews Neuroscience, 11(2), 127–138.

Relevance. The neuroscience free-energy principle, in which the brain minimizes a variational free energy. Provides the prototype for the boundary projection of GFF onto neural information manifolds; GG-Theory subsumes Friston’s principle as the special case where the boundary moduli are restricted to predictive-coding manifolds. First cited in §22.

Friston, K., et al. (2023). *The free-energy principle made simpler but not too simple*. Physics Reports, 1024, 1–29.

Relevance. Modern, mathematically rigorous formulation of the free-energy principle, including its variational and information-geometric structure. Cited in Part IV (Section 25) as the contemporary articulation of the free-energy paradigm that GFF subsumes and generalizes. First cited in §22.

Gell-Mann, M., & Low, F. E. (1954). *Quantum electrodynamics at small distances*. Physical Review, 95, 1300–1312.

Relevance. The β -function formalism in QED, the first explicit statement of the running coupling and the seed of the renormalization-group flow. The structural ancestor of GRIP-ST1 (DGI). First cited in §30.1.

Goresky, M., & MacPherson, R. (1988). *Stratified Morse Theory*. Springer-Verlag, Berlin.

- Relevance.** Standard reference for Morse theory on Whitney stratified spaces. The mathematical framework underlying Appendix E and the stratified GFF structure of Part IV. First cited in §E.
- Gorini, V., Kossakowski, A., & Sudarshan, E. C. G. (1976). *Completely positive dynamical semigroups of N -level systems*. Journal of Mathematical Physics, 17, 821–825.
- Relevance.** The GKLS theorem characterizing the most general CPTP generator of open-system quantum dynamics. Provides the mathematical prototype for the boundary-projection structure that OGGEF elevates to foundational status. First cited in §A.
- Green, M. B., & Schwarz, J. H. (1984). *Anomaly cancellations in supersymmetric $D = 10$ gauge theory and superstring theory*. Physics Letters B, 149, 117–122.
- Relevance.** The original Green–Schwarz anomaly-cancellation mechanism — and, in the GG-Theory reading, the technical pioneering of AX-2 GGEF: the geometrization of gauge fields on the same footing as gravity. BI-6 is the GGEF completion of this 1984 construction, extended from the heterotic ten-dimensional setting to the 6D bulk and the four-factor gauge group G_{tot} of GG-Theory; the full supportive treatment of the GS legacy lives in Arisaka (2026, A3-BI6, § 8.7, “BI-6 as the GGEF completion of GS: the geometrization reading”), which the present paper cites rather than re-derives. The Chern–Simons construction of H_3 from TGG data (Proposition 11.3) is the natural Gi-4 adaptation of the GS Chern–Simons primitive. First cited in §13.1.
- Helgason, S. (1979). *Differential Geometry, Lie Groups, and Symmetric Spaces*. Academic Press.
- Relevance.** Standard reference for the Fubini–Study metric on $\mathbb{CP}^{N-1}$ and the structure of Spin groups over normed division algebras. Used throughout Part II. First cited in §8.3.
- Hinton, G. E., & van Camp, D. (1993). *Keeping the neural networks simple by minimizing the description length of the weights*. In *Proceedings of the 6th Annual Conference on Computational Learning Theory (COLT '93)*, pp. 5–13.
- Relevance.** The original variational free-energy formulation in machine learning, $F_{\text{var}} = D_{\text{KL}}(q||p) - \log p(o)$. This is the prototype of Friston’s free-energy principle and provides the precise mathematical bridge between **GFF-ST6** and Bayesian inference, as discussed in Appendix I. First cited in §I.
- Hirzebruch, F. (1966). *Topological Methods in Algebraic Geometry* (3rd ed.). Springer-Verlag, Berlin.
- Relevance.** Foundational reference for the Hirzebruch signature theorem and the L-polynomial computation. Used throughout Part IV (GFF Morse-positivity from CCI) and the explicit $\sigma(\mathbb{CP}^{N-1})$ computation. First cited in §22.4.
- Hofbauer, J., & Sigmund, K. (1998). *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge.
- Relevance.** The replicator equation and its convergence theory; the form the descent on the family stratum takes in moment coordinates. First cited in §30.2.
- ’t Hooft, G. (1993). *Dimensional reduction in quantum gravity*. arXiv:gr-qc/9310026.
- Relevance.** The original holographic-principle proposal: the fundamental degrees of freedom of

- a $(d + 1)$ -dimensional gravitational theory project onto a d -dimensional boundary. Conceptual ancestor of OGGEF. First cited in §30.1.
- Hopkins, M. J., & Singer, I. M. (2005). *Quadratic functions in geometry, topology, and M-theory*. Journal of Differential Geometry, 70, 329–452.
Relevance. Foundational paper on differential cohomology and its application to anomaly inflow. The technical machinery underlying Theorems 27.1 and 27.2 of Part IV. First cited in §13.1.
- Hosotani, Y. (1983). *Dynamical mass generation by compact extra dimensions*. Physics Letters B, 126, 309–313.
Relevance. The original Wilson-line / Hosotani mechanism for electroweak symmetry breaking. In Gi-4 vocabulary, this is the GRIP cascade on the Higgs moduli stratum Σ_H . First cited in §11.2.
- Hosotani, Y. (2005). *Gauge-Higgs unification and the Hosotani mechanism*. Progress of Theoretical Physics Supplement, 172, 313–326.
Relevance. Modern review of gauge-Higgs unification. The 6D realization of Hosotani (1983) that operates on the compact θ -direction of M_6 . First cited in §1.
- Hulse, R. A., & Taylor, J. H. (1975). *Discovery of a pulsar in a binary system*. The Astrophysical Journal, 195, L51–L53.
Relevance. The binary pulsar PSR B1913+16, whose orbital decay has tracked the quadrupole prediction for five decades to a few parts in 10^3 — the empirical anchor of the inspiral leg of Example 4’s GRIP-In descent (§36.4). First cited in §36.4.
- Hurley, M. (1995). *Chain recurrence, semiflows, and gradients*. Journal of Dynamics and Differential Equations, 7, 437–456.
Relevance. Chain recurrence for semiflows on arbitrary metric spaces, characterized through attractors, and the gradient criterion: for a gradient-like semiflow the chain-recurrent set is the critical set. The form of Conley’s theory the boundary semigroup needs. First cited in §25.4.
- Ibáñez, L. E., Nilles, H. P., & Quevedo, F. (1987). *Orbifolds and Wilson lines*. Physics Letters B, 187, 25–32.
Relevance. The $\mathbb{Z}_3$ -orbifold construction with three twisted sectors, the most natural prior model giving exactly three generations. Compared against in Arisaka (2026, P2-N3, § 12.3). First cited in §54.6.
- Ibáñez, L. E., & Uranga, A. M. (2012). *String Theory and Particle Physics: An Introduction to String Phenomenology*. Cambridge University Press.
Relevance. Standard reference for D-brane and intersecting-brane family counting. Provides the broader phenomenological context for Section 44.3 (the Gi-4 three-leg N=3 derivation). First cited in §54.6.
- Jaynes, E. T. (1957). *Information theory and statistical mechanics*. Physical Review, 106(4), 620–630.
Relevance. The maximum-entropy principle in statistical inference, equivalent to free-energy minimization in equilibrium statistical mechanics. Cited in Part IV (Section 22) as one of the universal projections of the GFF framework. First cited in §22.

Kadanoff, L. P. (1966). *Scaling laws for Ising models near T_c* . Physics, 2, 263–272.

Relevance. The block-spin coarse-graining picture for statistical-mechanical systems near criticality; the conceptual bridge between Gell-Mann–Low’s β -function and Wilson’s full renormalization-group framework. First cited in §30.1.

Kanno, S. (2010). *Anomaly cancellation conditions in extended models*. Progress of Theoretical Physics, 124, 1085–1100.

Relevance. Reviews anomaly-cancellation constraints on family number in BSM extensions. Family number is constrained but not uniquely fixed by closed-system anomaly cancellation alone. First cited in §18.3.

Kobayashi, M., & Maskawa, T. (1973). *CP-violation in the renormalizable theory of weak interaction*. Progress of Theoretical Physics, 49, 652–657.

Relevance. The original observation that $N \geq 3$ is required for CP violation. The historical precursor to leg L1 of the $N=3$ proof; the lower bound used in Theorem T2. First cited in §44.3.

Kobayashi, S., & Nomizu, K. (1969). *Foundations of Differential Geometry, Vol. II*. Wiley-Interscience, New York.

Relevance. Standard reference for Chern–Weil theory and the Pontryagin / Chern class identifications used throughout Parts II and III. First cited in §6.1.

Kobayashi, T., Oba, T., & Shinjo, T. (2020). *Modular flavor symmetry and three generations*. Journal of High Energy Physics, 2020(11), 199.

Relevance. A modern attempt to derive $N = 3$ from modular flavor symmetry. Representative of the prior approaches that obtain three generations conditionally on a chosen symmetry input. First cited in §54.6.

Kossakowski, A., Frigerio, A., Gorini, V., & Verri, M. (1977). *Quantum detailed balance and KMS condition*. Communications in Mathematical Physics, 57, 97–110.

Relevance. The quantum detailed-balance condition in the form the gradient-flow representation needs (self-adjointness of the generator in the GNS inner product of the stationary state), and the theorem that a weakly coupled bath in a KMS state supplies it. The hypothesis (2) of T3.2. First cited in §25.2.

Kvalheim, M. D., Gustafson, P., & Koditschek, D. E. (2021). *Conley’s fundamental theorem for a class of hybrid systems*. SIAM Journal on Applied Dynamical Systems, 20, 784–825.

Relevance. The fundamental theorem for hybrid systems with a closed guard set and a continuous reset map, proved by gluing into one compact space; the only extension of Conley’s theorem under which the second regime of the runtime is a theorem, at the price that the reset fires before the descent halts (§31.7). First cited in §31.6.

Lewis, A. D. (2025). *The fundamental theorem of dynamical systems: all at once and all in the same place*. arXiv:2508.10525.

Relevance. A unified construction of complete Lyapunov functions for flows and semiflows

in continuous and discrete time; the semiflow form of Conley’s theorem, which is the one a forward-only boundary dynamics needs. First cited in §25.4.

Lindblad, G. (1975). *Completely positive maps and entropy inequalities*. Communications in Mathematical Physics, 40(2), 147–151.

Relevance. Lindblad’s monotonicity theorem: relative entropy is non-increasing under any CPTP map, and hence the relative entropy to a stationary state descends under any CPTP semigroup. This is the unconditional half of Theorem **T3.2**: it needs no detailed balance, and for that reason it cannot certify detailed balance either, which is why T3.2 is stated as a conditional and not an equivalence. First cited in §22.1.

Lindblad, G. (1976). *On the generators of quantum dynamical semigroups*. Communications in Mathematical Physics, 48, 119–130.

Relevance. The Lindblad master equation, derived independently of Gorini, Kossakowski & Sudarshan (1976). Provides the technical foundation for CPTP open-system dynamics, the closed-system structure that the OGGEF-elevated boundary projector generalizes. First cited in §25.2.

Łojasiewicz, S. (1984). *Sur les trajectoires du gradient d’une fonction analytique*. In Seminari di Geometria 1982–1983, Università di Bologna, 115–117.

Relevance. Bounded trajectories of an analytic gradient flow converge to a single point. What the halting corollary would need where a floor is a manifold of positive dimension inside a larger stratum, and does not need on the one-dimensional Wilson-line stratum. First cited in §25.4.

Maldacena, J. (1997). *The large N limit of superconformal field theories and supergravity*. Advances in Theoretical and Mathematical Physics, 2, 231–252.

Relevance. The original AdS/CFT correspondence. The prototype of bulk-to-boundary descent that OGGEF/GFF/GRIP elevates to a foundational principle. First cited in §54.7.

Maxwell, J. C. (1865). *A dynamical theory of the electromagnetic field*. Philosophical Transactions of the Royal Society of London, 155, 459–512.

Relevance. The founding paper of classical electromagnetism and, in the modern geometric reading, the first gauge theory in physics: the four-potential A_μ is defined only up to a gradient, and the field strength $F_{\mu\nu}$ is the curvature of this connection on a U(1) principal bundle. This is the structural ancestor of TGG and of CCI (in its $dF = 0$ form). First cited in §6.1.

Miller, M. C., Lamb, F. K., Dittmann, A. J., et al. (2021). *The radius of PSR J0740+6620 from NICER and XMM-Newton data*. Astrophysical Journal Letters, 918, L28.

Relevance. NICER measurement of neutron-star radius for the heaviest precisely-measured pulsar. Provides the empirical baseline against which the modified-TOV correction $\Delta_{\text{GS}}(r)$ of CCI-ST7 is to be tested. First cited in §17.4.

Milnor, J. (1963). *Morse Theory*. Annals of Mathematics Studies 51, Princeton University Press.

Relevance. Foundational reference for Morse theory. The height-function-on- $\mathbb{CP}^{N-1}$ Morse polynomial is the textbook example that drives Lemmas in Parts IV–V. First cited in §22.2.

Mittnenzweig, M., & Mielke, A. (2017). *An entropic gradient structure for Lindblad equations and couplings of quantum systems to macroscopic models*. Journal of Statistical Physics, 167, 205–233.

Relevance. Generalizes the Carlen–Maas gradient-flow representation to a broad class of Lindblad master equations: any GKLS evolution with a detailed-balanced fixed point is a gradient flow of relative entropy under an appropriate information-geometric metric. This is the technical underpinning for the boundary-projection step (Step 2) of the proof of Theorem 25.3: it ensures that the boundary moduli-space gradient-flow representation $[\dot{\Pi}] = -g_{\Sigma}^{-1} \nabla \mathcal{F}_{\text{GFF}}$ is well-defined for the GDOS Lindbladian whenever its OGGEF-inherited detailed-balance condition holds. First cited in §25.2.

von Neumann, J. (1932). *Mathematische Grundlagen der Quantenmechanik*. Springer-Verlag, Berlin.

Relevance. The foundational mathematical formalization of quantum mechanics as unitary evolution on Hilbert space. In the methodological reflection of Section 60.4, this is the moment when closed-system QM was formalized; the open-system aspects of measurement were already present but were not elevated to foundational status. First cited in §1.

Noether, E. (1918). *Invariante Variationsprobleme*. Nachrichten der Königlichen Gesellschaft der Wissenschaften zu Göttingen, Math.-Phys. Klasse, 235–257.

Relevance. The symmetry–conservation theorem. The foundational identity tying continuous symmetries of an action to conserved currents. The dynamical face of CCI traces directly back to this paper. First cited in §13.1.

Norton, D. E. (1995). *The fundamental theorem of dynamical systems*. Commentationes Mathematicae Universitatis Carolinae, 36, 585–597.

Relevance. The name, the standard statement and the standard proof of Conley’s theorem. First cited in §25.4.

NOvA Collaboration (2022). *Improved measurement of neutrino oscillation parameters by the NOvA experiment*. Physical Review D, 106, 032004.

Relevance. Empirical evidence for non-trivial PMNS mixing and a non-zero leptonic CP phase. Independent support for $N \geq 3$ in the leptonic sector. First cited in §44.3.

Onsager, L. (1931). *Reciprocal relations in irreversible processes. I*. Physical Review, 37, 405–426.

Relevance. The reciprocal relations $L_{ij} = L_{ji}$ for the linear response of irreversible systems. The first quantitative result on near-equilibrium open-system dynamics; the structural ancestor of GFF-ST7 (Onsager reciprocity). First cited in §22.1.

Particle Data Group (2024). *Review of Particle Physics*. Progress of Theoretical and Experimental Physics, 2024, 083C01.

Relevance. Standard reference for measured Standard-Model parameters: charged-lepton masses, quark masses, mixing angles, CP phases, and the Jarlskog invariant $J_{\text{CKM}} = (3.18 \pm 0.15) \times 10^{-5}$. The empirical baseline against which the GG-Theory codebook is checked. First cited in §56.

Patrão, M. (2011). *Existence of complete Lyapunov functions for semiflows on separable metric spaces*. Far East Journal of Dynamical Systems, 17, 49–54.

Relevance. The complete Lyapunov function for continuous-time semiflows; the form of Conley’s theorem the boundary dynamics, a semigroup and not a group, needs. First cited in §25.4.

Planck Collaboration (2020). *Planck 2018 results. VI. Cosmological parameters*. Astronomy & Astrophysics, 641, A6.

Relevance. Cosmological measurement $N_{\text{eff}} = 2.99 \pm 0.17$, an independent confirmation of $N_\nu = 3$ complementing ALEPH *et al.* (2006). First cited in §54.2.

Randall, L., & Sundrum, R. (1999). *A large mass hierarchy from a small extra dimension*. Physical Review Letters, 83(17), 3370–3373.

Relevance. The original extra-dimension scenario. Provides the prototype for the scale-profile identification $T_{\text{GFF}} \sim M_{\text{KK}}$ used in Section 25 (the Tolman-like identification of scale with effective temperature). First cited in §25.1.

de Rham, G. (1931). *Sur l’analyse situs des variétés à n dimensions*. Journal de Mathématiques Pures et Appliquées, 10, 115–200.

Relevance. The original de Rham theorem identifying real singular cohomology with cohomology of closed differential forms. This is the foundational result underlying every step of the differential-cohomology refinement used in Parts III, IV, and Appendix C. First cited in §C.2.

Riley, T. E., Watts, A. L., Ray, P. S., et al. (2021). *A NICER view of the massive pulsar PSR J0740+6620 informed by radio timing and XMM-Newton spectroscopy*. Astrophysical Journal Letters, 918, L27.

Relevance. Companion NICER mass–radius determination for the heaviest precisely-measured pulsar. Together with Miller *et al.* (2021), defines the empirical envelope inside which the modified-TOV correction $\Delta_{\text{GS}}(r)$ must lie if CCI is to remain falsifiable but not yet falsified. First cited in §17.4.

Sands, M. (1970). *The Physics of Electron Storage Rings: An Introduction*. SLAC Report SLAC-R-121, Stanford Linear Accelerator Center.

Relevance. The canonical storage-ring treatment of radiation damping, quantum excitation, and equilibrium emittance. Example 4 (§36.4) identifies the equilibrium emittance as the orbital stratum’s geometric attractor and the measured constant $C_q = (55/32\sqrt{3}) \hbar/mc$ as the coefficient-level realization of **GFF-ST7**(2). First cited in §36.4.

Schwarz, M. (1993). *Morse Homology*. Birkhäuser Progress in Mathematics 111.

Relevance. Modern reference for transversality of gradient flows between Morse critical points. Underlies the cascade-trajectory analysis of Part V. First cited in §30.4.

Spohn, H. (1978). *Entropy production for quantum dynamical semigroups*. Journal of Mathematical Physics, 19, 1227–1230.

Relevance. Spohn’s entropy-production theorem for quantum dynamical semigroups: the entropy production $-dD(\rho_t \parallel \rho_*)/dt$ is non-negative for any CPTP semigroup with a stationary state ρ_* ;

- it vanishes at ρ_* and may vanish elsewhere. Spohn proves no equivalence between monotonicity and detailed balance, and T3.2 no longer attributes one to him: detailed balance enters T3.2 as a hypothesis, for the gradient-flow structure, not as a consequence of monotonicity. First cited in §26.5.
- Susskind, L. (1995). *The world as a hologram*. Journal of Mathematical Physics, 36, 6377–6396.
Relevance. The holographic principle in its string-theory formulation. One of the foundational steps toward the open-system / boundary-projection perspective elevated to foundational status by GDOS. First cited in §1.3.
- T2K Collaboration (2020). *Constraint on the matter–antimatter symmetry-violating phase in neutrino oscillations*. Nature, 580, 339–344.
Relevance. Empirical evidence of leptonic CP violation, providing independent support for the lower bound $N \geq 3$ in the leptonic sector. First cited in §44.3.
- Ward, J. C. (1950). *An identity in quantum electrodynamics*. Physical Review, 78, 182.
Relevance. The Ward identity for QED, expressing the gauge-invariance of the effective action as an identity among Green’s functions. The quantum-field-theoretic reformulation of Noether’s theorem and the QFT progenitor of CCI. First cited in §13.1.
- Wilson, K. G. (1974). *Confinement of quarks*. Physical Review D, 10, 2445–2459.
Relevance. The seminal treatment of Wilson loops in lattice gauge theory. Provides the physical interpretation of the Wilson-line order parameters of TGG (Definition 7.2). First cited in §30.1.
- Witten, E. (1982). *Supersymmetry and Morse theory*. Journal of Differential Geometry, 17, 661–692.
Relevance. The foundational paper relating supersymmetric quantum mechanics to Morse theory. Provides the Bogomolnyi-saturating instanton-action computation on $\mathbb{CP}^{N-1}$ that Theorem 32.1 reproduces; Witten’s Morse-complex weights are entries of the connection matrix, which on an admissible stratum vanishes, so they are not the weights of GRIP-ST7. First cited in §30.4.
- Witten, E. (1996). *On flux quantization in M-theory and the effective action*. Journal of Geometry and Physics, 22, 1–13.
Relevance. Foundational application of differential cohomology to anomaly inflow in M-theory. Provides the prototype for the non-degeneracy of the Hopkins–Singer descent (Theorem 27.2). First cited in §1.1.
- Witten, E. (1985). *Global anomalies in string theory*. In: Symposium on anomalies, geometry, topology (Argonne, IL, 1985), 61–99.
Relevance. Reviews global-anomaly constraints and their relation to family number. One of the precursors to the BI-6 anomaly polynomial. First cited in §F.
- Yang, C. N., & Mills, R. L. (1954). *Conservation of isotopic spin and isotopic gauge invariance*. Physical Review, 96, 191–195.
Relevance. The non-abelian generalization of Maxwell’s gauge theory; the introduction of SU(2) gauge fields and the non-linear field strength $F_{\mu\nu}^a$. The structural ancestor of the boundary gauge component A of the TGG composite connection. First cited in §6.1.

Zeh, H. D. (1970). *On the interpretation of measurement in quantum theory*. Foundations of Physics, 1, 69–76.

Relevance. The first articulation of decoherence: the mechanism by which open-system entanglement with the environment produces effective classicality from quantum dynamics. An early foundational step toward the open-system measurement perspective. First cited in §1.4.

Zurek, W. H. (1981). *Pointer basis of quantum apparatus: Into what mixture does the wave packet collapse?* Physical Review D, 24, 1516–1525.

Relevance. The pointer-basis selection mechanism in decoherence theory. Together with Zeh (1970), provides the technical machinery underlying the open-system measurement perspective. First cited in §54.4.