UCLA/Arisaka-GG-P1-KK
v20.0 — August 21, 2026 — *Particle Branch*

The GG-6 Lock Cluster

A Non-Spatial Kaluza–Klein Resonance near 3.85 TeV as the Standing LHC Falsifier of GG-Theory

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

Every particle discovered since the mid-1970s — the W and Z, the top, the Higgs — was on the Standard Model's list before it was found. This paper is about the first particle that is not on that list. GG-Theory *derives*, rather than fits, a new mass scale: the lock cluster at $M_{\text{lock}} = \frac{1}{2} M_{\text{Pl}} e^{-35} \simeq 3.849$ TeV, fixed by the integer $K = 35$ of the OGN-5 ledger with no adjustable parameter anywhere in the chain. The same number over-constrains the electroweak scale, $v = M_{\text{lock}}(2/5)^3 \approx 246$ GeV: two known scales, one integer, nowhere to hide.

At that mass the theory places a three-member cluster — a spin-2 graviton, a color-octet vector g^*, and a leptophobic Z'_B: the discrete modes of a *non-spatial* (renormalization-group) extra dimension — with derived couplings that are top-rich, dijet-friendly, and lepton-poor, so existing clean-channel limits do not touch it. The un-suppressible member is the octet: $\sigma \times \text{BR}(t\bar{t}) = 8.7$ fb at 13 TeV, $\Gamma/m = 15\%$, with a theorem-fixed interference sign. It sits exactly at the current ATLAS/CMS $t\bar{t}$ exclusion frontier: the prediction is being tested *now*.

A full simulation anchored to the ATLAS Run-2 analysis (Part VIII) turns this into a calendar. A dedicated reanalysis of data already on tape carries an expected 2.7–3.7σ per experiment, and 4.6–5.2σ for a combined ATLAS+CMS analysis with improved systematics — five-sigma discovery from *existing* data is a live possibility; failing that, single-experiment 5σ luminosity is on tape during 2032. Because mass, width, rate, sign, and angular shape are all pre-registered, the search pays no look-elsewhere penalty — and a null result at $\lesssim 2$ ab^{-1} retires the implementation outright. The mass is the bet; the LHC is the table.

Executive Summary

The question behind this paper is simple: *if GG-6 is right, what exactly should ATLAS and CMS see — and how fast can they either find it or kill it?* The answer is best given in seven points, in as plain a language as the subject allows.

First — the principle, and why it predicts. GG-Theory rests on one structural idea, *geometry–gauge locking*: the energy scale at which nature is probed is treated not as a bookkeeping label but as a genuine geometric dimension[1], and the geometry and the gauge fields lock together at thresholds set by *integers* — the ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$. Integers are the entire content of the theory's freedom: there is no coupling to adjust, no mass to scan, no continuous dial of any kind. This is the source of its unusual predictive power: a theory built on counting numbers cannot retreat, so everything it says is a commitment — the property whose fifty-year absence §1 recounts in full.

Second — one derived number, doubly spent. The most rigid output is the *lock scale*,

$$M_{\text{lock}} = \frac{M_{\text{Pl}}}{2\,e^{35}} \simeq 3.849\,\text{TeV},$$

where the 35 in the exponent is the integer K — not a fitted quantity. The same number is spent a second time: the electroweak scale, known since the 1980s, must equal $M_{\text{lock}}(2/5)^3 \approx 246$ GeV, and does. A reader should pause on what this over-constraint means: the theory does not merely predict something new at 3.85 TeV; it ties that prediction to a number already measured, so it cannot slide the target without breaking what already works.

Third — a cluster of three, not a lone bump. At the lock scale the theory places three sibling states, degenerate in mass: a spin-2 graviton mode; a heavy gluon-like state (the color-octet g^*); and a heavy photon-like state that couples to quarks but shuns electrons and muons (the "leptophobic" Z'_B). Their decay preferences are derived, not chosen: rich in top quarks, friendly to jets, poor in leptons and photons. This is why the LHC's cleanest and most famous exclusion channels — electron–positron pairs, photon pairs — say almost nothing about it, and why the decisive channel is the top-quark pair.

Fourth — it is being tested now. The un-suppressible member of the cluster is the octet: its coupling to the colliding quarks is fixed by the same integer K and cannot be lowered. The computed rate, $\sigma \times \text{BR}(t\bar{t}) = 8.7$ fb at 13 TeV with a broad width ($\Gamma/m \simeq 15\%$), sits exactly at the frontier of the published ATLAS/CMS top-pair exclusions — neither safely beyond reach nor yet excluded, because no published search template matches this shape: broad, interfering with the QCD continuum at a derived sign, enhancing the spectrum below 3.85 TeV and depleting it above. A dedicated reinterpretation of existing data is therefore the single most urgent request this paper makes.

Fifth — the discovery may already be on tape. Part VIII simulates the search the way an analysis team would run it, anchored item by item to the published ATLAS Run-2 analysis and

[1] Consistent with the cohort reading: this "six-dimensional bulk" is the renormalization-group geometrization of the energy scale — the radial fifth dimension is the renormalization-group axis $r = \ln(E_{\text{GUT}}/\mu)$, not a Randall–Sundrum spatial dimension — so the radial "lock cluster" near 3.85 TeV is an RG threshold/lock spectrum rather than a tower of spatial Kaluza–Klein gravitons; see Arisaka (2026, A2-GG6).

calibrated to its published limits. The result: a dedicated reanalysis of the Run-2+3 data already recorded carries an expected 2.7–3.7σ per experiment, and 4.6–5.2σ for a combined ATLAS+CMS analysis with realistically improved background systematics — *five-sigma discovery from existing data is a live possibility*, and single-experiment 5σ luminosity is on tape during 2032. Because the mass, width, rate, interference sign, and angular shape are all specified in advance, the search pays none of the statistical penalties of a generic bump hunt: no look-elsewhere discount (worth 0.8σ at threshold, or equivalently a doubling of the dataset), and none of the factor-three dilution of a mismatched search template.

Sixth — the fight is systematics, not statistics. The simulation's structural lesson: beyond Run 3 the search is no longer limited by how much data exists but by how well the background is known. More luminosity alone moves the significance by only $\sim 1\sigma$ per *tripling* of the dataset; halving the background uncertainties — the standard, decomposable HL-LHC expectation — is what crosses the discovery threshold. The practical consequences for the experiments are spelled out, including one warning: the signal's own interference contaminates the low-mass sideband, so the background normalization must be anchored below ~ 1 TeV.

Seventh — either answer is decisive. This is the sharpest collider falsifier the GG-6 program owns. If the cluster is there, the correlated pattern — one mass, three states, top-rich decays, the derived interference shape, and a second family near 7 TeV — is unmistakable, and the property measurements (spin, chirality) follow within two years of HL-LHC running. If it is not there, a null result at $\lesssim 2$ ab^{-1} does not wound the theory; it *retires the current implementation outright*, and takes the electroweak-scale derivation down with it. No parameter can be adjusted to escape, because there are no parameters. The direct spin-2 confirmation, at its derived coupling, is then FCC-hh physics ($\sim$205 fb at 100 TeV).

Each community will read this wager in its own way. *A collider experimentalist* meets a concrete, pre-registered search: one mass, one width, one sign, one angular shape, shipped with a reinterpretation card and a trigger roadmap. *A graduate student* meets a thesis-sized analysis whose decisive data already exist (see the Preface). *A theorist* meets an extra-dimension signature whose scale is derived from an integer rather than fitted — and a falsifier with no escape hatch. *A general physicist* meets a plain wager: one number, 3.849 TeV, staked against the next decade of the LHC. The remainder of the paper makes that wager precise.

Keywords: geometry–gauge locking; Kaluza–Klein particles; non-spatial extra dimensions; scale dimension; phase dimension; Randall–Sundrum; ATLAS; CMS; HL-LHC; spin-2 resonance; top-antitop resonances; trigger-level analysis; data scouting; falsifiability; open systems; OGGEP; GDOS

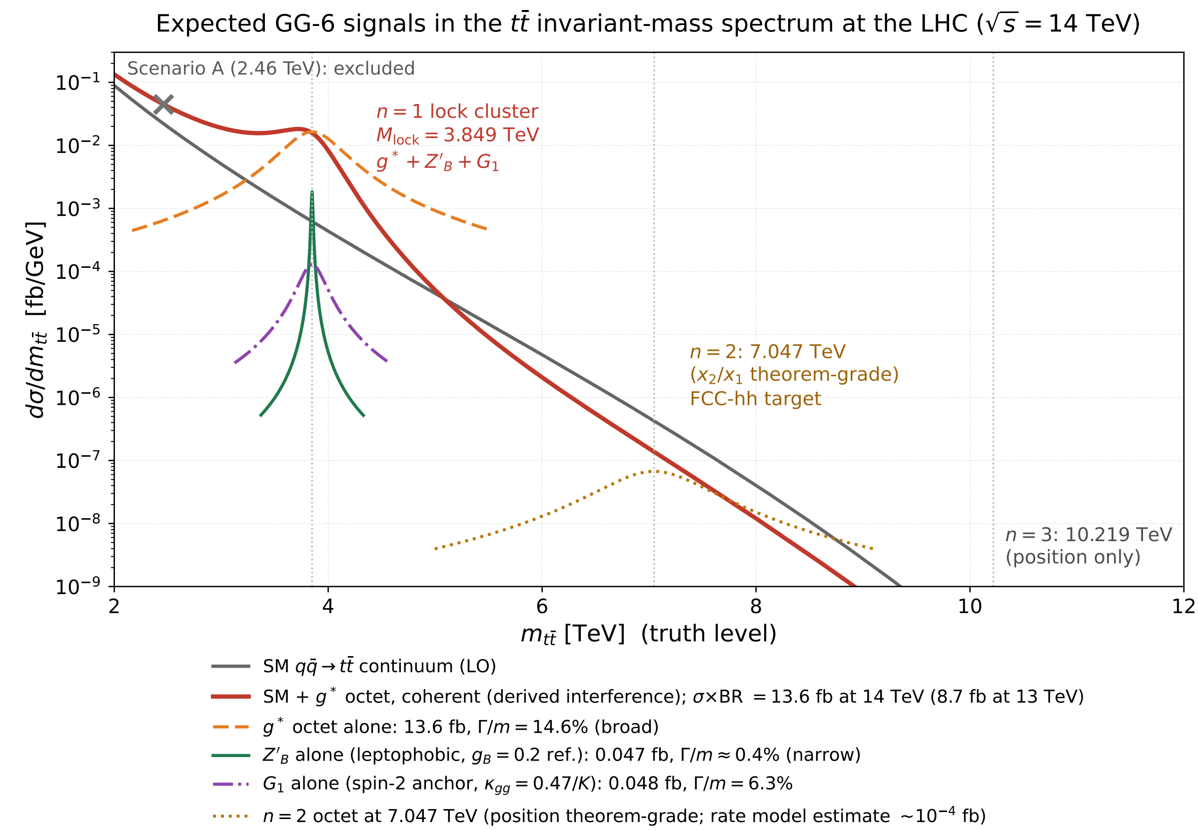

Every expected GG-6 collider signal, computed with no adjustable parameter: the three-member lock cluster at the derived 3.849 TeV and its radial family at 7.05 and 10.22 TeV (Figure 4, p. 47).

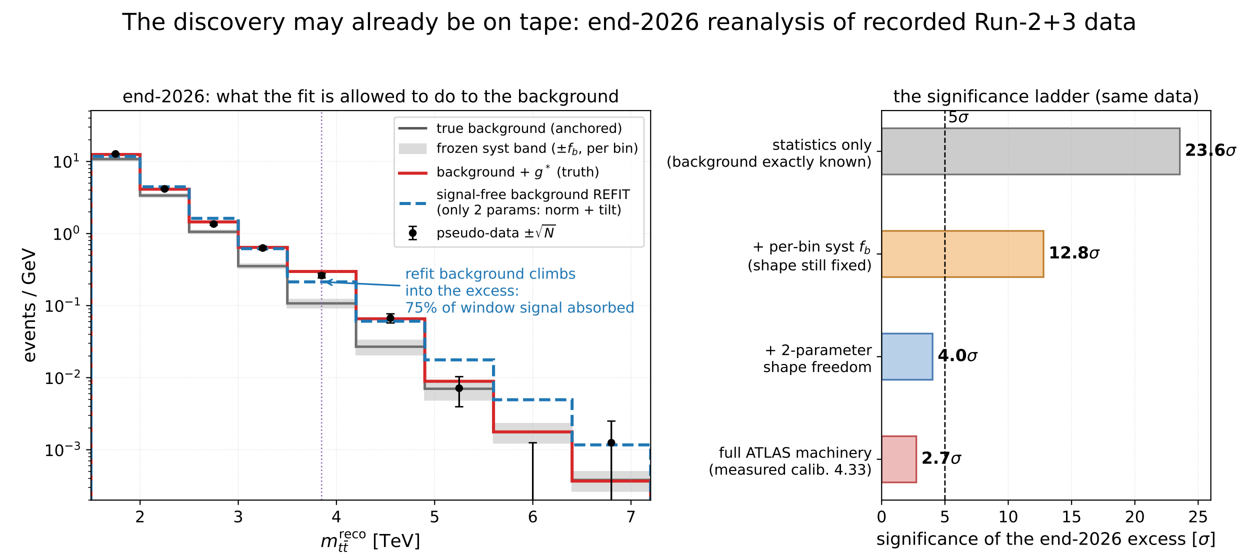

Why the verdict is near: a dedicated reanalysis of collisions already recorded carries an expected 2.7–3.7σ per experiment — and the ladder shows exactly where the sensitivity goes (Figure 11, p. 132; Part VIII).

Ten Predictions,
Each Graded Before the Data Arrive

What this paper is willing to be wrong about. Ten numbered predictions, each carrying the grade that says how strongly it is claimed: [T] a theorem, which falls with its theorem; [D*] derived subject to the condition printed beside it; [M] one of a finite listed menu; [E] somebody else's number this paper consumes. Four grades and no unlabeled rows, because mixing the four is how a prediction list comes to look stronger than it is. The rates and the branching fingerprint are what the LHC can reach with data already on tape; the interference sign and the exact leptophobia are theorems, and one measurement retires either. Nothing here was written after the data. In full: Table 31, p. 158.

Id	Prediction	Grade
P-KK-1	$\sigma \times \mathrm{BR}(g^* \to t\bar{t}) = 8.7\,[7.0\text{–}12]$ fb at 13 TeV; $13.6\,[11\text{–}19]$ fb at 14 TeV; ≈ 2.8 pb at 85 TeV	[D*]
P-KK-2	Octet branching fingerprint $t\bar{t} : b\bar{b} : jj \approx 70 : 3.5 : 1$; $\mathrm{BR}(t\bar{t}) = 0.94\,[0.93\text{–}0.95]$	[D*]
P-KK-3	Graviton pattern $t\bar{t} : W_L W_L : Z_L Z_L : hh = 0.69 : 0.15 : 0.08 : 0.08$, $\Gamma/m = 6.3\%$, $\sigma(85\ \text{TeV}) \approx 0.13$ pb	[D*]
P-KK-4	Interference sign: enhancement below the pole, deficit above; SM crossing at $\simeq 4.9$ TeV; centroid pull $-\mathcal{O}(100)$ GeV	[T]
P-KK-5	Top polarization $P_t = +0.91$, band $[+0.88, +0.92]$; $A_{FB} = 0$	[T]/[D*]
P-KK-6	Exact leptophobia: $\mathrm{BR}(Z'_B \to \ell\ell) = 0$; $m_{Z'_B} = m_{g^*}$ exactly	[T]
P-KK-7	No angular satellites at the TeV scale (the θ tower sits at $\sim E_{\mathrm{GUT}}$)	[T]
P-KK-8	Second radial family at $1.831\,M_{\mathrm{lock}} = 7.047$ TeV (position; rate estimate $\sim$10 fb at 85 TeV)	[T]/[M]
P-KK-9	No radion; Higgs couplings Standard-Model-like apart from $G_1 \to hh$	[T]
P-KK-10	Lepton-collider silence: no s-channel cluster line in $\mu^+\mu^-$; SM-like contact terms at 3 TeV	[T]

A Six-Dimensional Bulk Whose Fifth Dimension Is Scale and Whose Sixth Is Phase

The whole program in one diagram, with the scale axis carrying numbers. The bulk $\mathcal{M}_4 \times I_r \times S^1_\theta$ is a fiber bundle: over every point of the Lorentzian stage sit the two *non-spatial* fibers, scale and phase. AX-2 (GGEP) acts *upward* and builds them; AX-3 (OGGEP) acts *downward*, and registration is the projection that erases (r, θ) — the origin of quantum behavior. The vertical axis is calibrated: $E_{\rm GUT} = E_{\rm Pl} e^{-2\pi}$ at $r = 0$, the OGN-5 depth $r = K = 35$ at the infrared boundary, and the lock at $r = 29.41$, which is $K - 2\pi + \ln 2$ exactly.

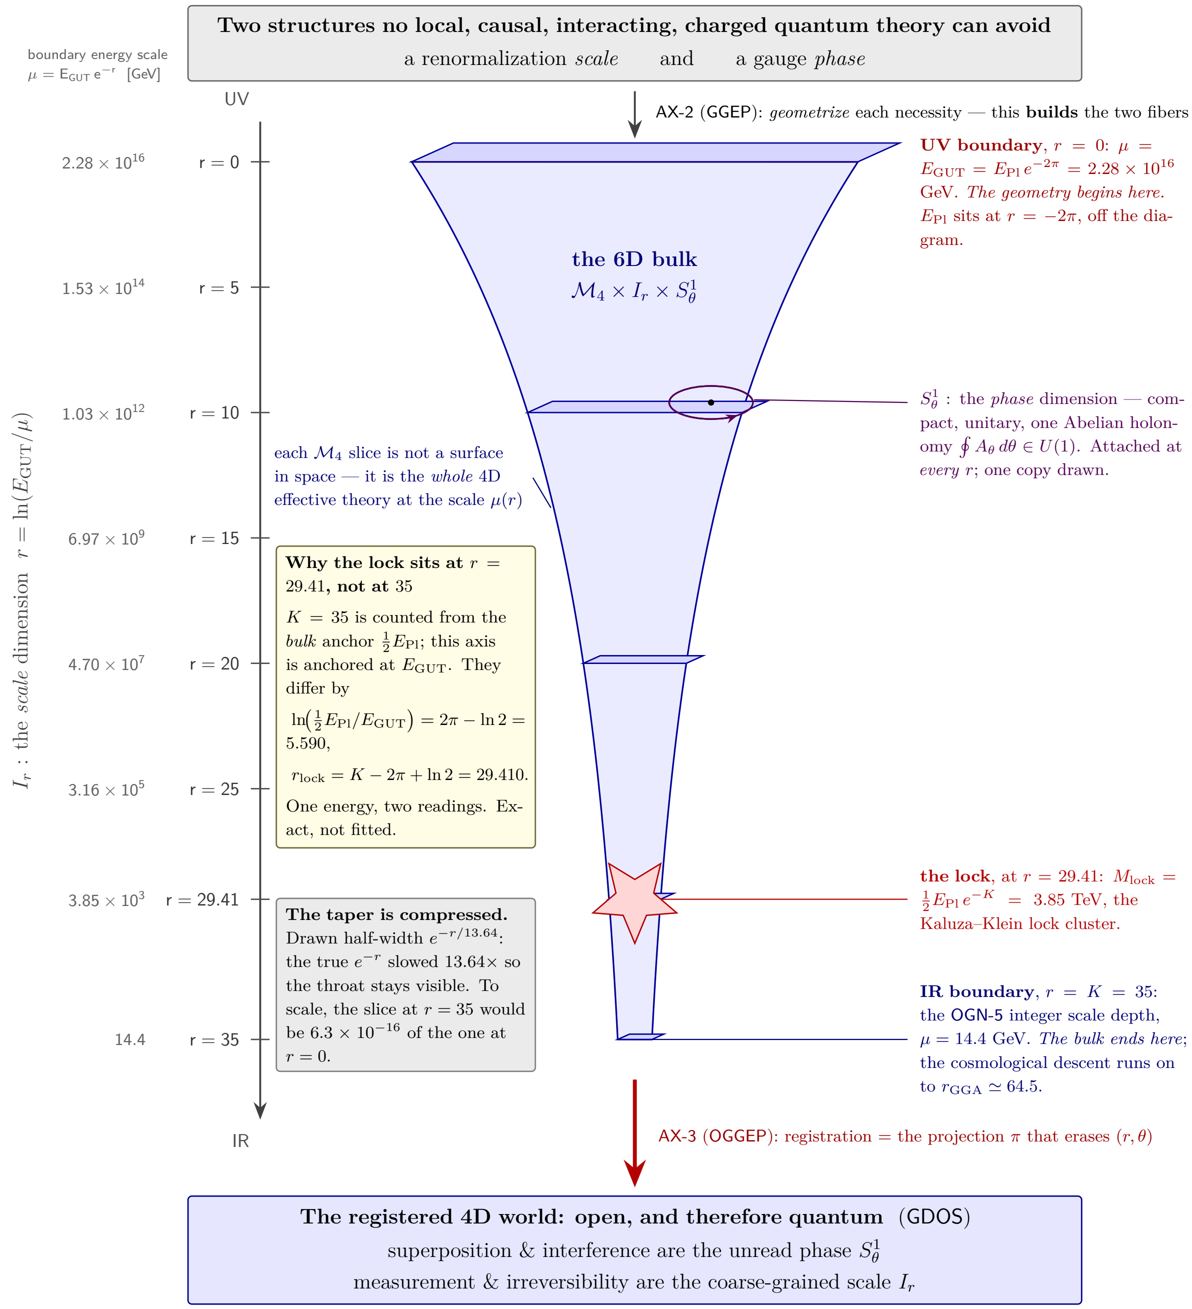

Overview of Contents

Part I. Foundations 16
Part II. The Derived Signal 40
Part III. Production and Discovery 69
Part IV. GG-6 versus Randall–Sundrum, and Detection as an Open System 77
Part V. Two Read-outs of the Lock Cluster 83
Part VI. Anatomy of the $t\bar{t}$ Signal 88
Part VII. Experimental Status and Reach 99
Part VIII. Simulating the Discovery 123
Part IX. Beyond the HL-LHC 140
Part X. Discussions 149
Part XI. Conclusions 163
Part XII. Technical Appendices 170
Part XIII. Common Objections and Responses 192
Part XIV. References 205

List of Figures

1 The open-system funnel shared across the GG-Theory cohort. 22
2 The twenty-two-paper cohort map. 36
3 Roadmap of the fourteen Parts of GG-P1-KK. 38
4 Quantitative LHC invariant-mass signature of the GG-6 lock cluster 47
5 Schematic mass staircase of the low-lying GG-6 sector 52
6 Derived $m_{t\bar{t}}$ line shape of the lock cluster with QCD interference 91
7 Derived production-angle distributions: spin 1 versus spin 2 94
8 System-rapidity distributions: gg versus $q\bar{q}$ initial states 95
9 Computed lock-cluster $\sigma\times$BR versus published 95% CL limits 108
10 The lock-cluster falsification timeline in calendar years, 2025–2050 116
11 The significance ladder at end-2026 132
12 Simulated spectra at the four snapshots 134
13 Calibrated significance versus luminosity 135

List of Tables

1 Constitutional route from the GG-Constitution to the KK collider program. 24
2 KK versus RS versus the two non-spatial dimensions of GG-6 30
3 The GG-6 radial lock cluster (Option A). 43
4 Provenance of every quantitative prediction. 46
5 Practical mass ledger for the collider-relevant low-lying GG-6 sector. 49
6 Derived structure of the low-lying GG-6 collider spectrum. 50

7 The cluster members under the two scenarios. . . . 57
8 Summary of the main collider observables that distinguish the GG-6 KK sector from a minimal. 62
9 Channel-by-channel logic of the first GG-6 cluster. . . . 75
10 Minimal comparison between a textbook 5D RS search picture and the 6D GG-6 picture relevant. 78
11 Generic KK versus 5D Randall–Sundrum versus 6D GG-6. . . . 78
12 Read-out-II channel estimates. . . . 85
13 The two read-outs side by side. . . . 86
14 Excess-to-continuum profile of the g^* signal, LO/LO, $\sqrt{s} = 13$ TeV: the ratio. . . . 91
15 The discrimination matrix at 3.849 TeV. . . . 96
16 Channel-by-channel exclusion and discovery ledger. . . . 102
17 Current public ATLAS/CMS status relevant to the GG-6 KK search program. . . . 104
18 Parton luminosities at the two candidate masses. . . . 106
19 Computed $\sigma \times$ BR for the lock-cluster states at $\sqrt{s} = 13$ TeV. . . . 106
20 Near-term discovery logic for the first GG-6 cluster. . . . 113
21 Numerical inputs to Figure 10. . . . 115
22 Why the KK search is the sharpest *collider* falsifier of the current GG-6 implementation. . . . 118
23 Proposed trigger and reduced-stream improvements for the GG-6 KK search. . . . 121
24 ATLAS merged-SR post-fit background at 140 fb^{-1} . . . 124
25 What "improved systematics" means, stated as a number per bin rather than as an adjective. 131
26 The four snapshots: yields and significances . . . 133
27 Dedicated-search significance, with ATLAS+CMS combination . . . 137
28 Status of future collider facilities, July 2026 . . . 147
29 GG falsification program at future facilities . . . 148
30 Framework discrimination matrix . . . 157
31 The numbered predictions of this paper, with grades. . . . 158
32 Summary of abbreviations used in the present paper. . . . 171
33 The symbols of this paper at a glance, with the section or companion that defines each. . . . 173
34 The ten numbers at a glance . . . 175
35 Master phenomenology ledger for the low-lying GG-6 collider sector. . . . 176
36 Falsification hierarchy for the collider program. . . . 177
37 The discovery calendar. . . . 177
38 The significance arithmetic worked bin by bin for one dataset, so that every other number. . 188
39 The open items as a roadmap rather than as a list of regrets. . . . 191

Contents

Executive Summary . . . 2
List of Figures . . . 7
List of Tables . . . 7
Preface . . . 14

Part I. Foundations **16**
1 Introduction: Fifty Years Beyond the Standard Model . . . 16
2 GG-Theory in Plain English — and Why the LHC Gets to Judge It . . . 17
3 The central claim: the Kaluza–Klein lock cluster as an open-system observable . . . 19
4 From the GG-Constitution to Collider Observables . . . 21
5 The 6D Geometric Origin of the KK Tower . . . 24
5.1 The six-dimensional stage . . . 24
5.2 Mode expansion and the two quantum numbers . . . 25
5.3 Convention: "Kaluza–Klein" in GG-6 means *non-spatial* . . . 26
5.4 Radial branches and the multi-field lock cluster . . . 27
5.5 Why GG-6 is not reducible to ordinary KK or RS . . . 27
6 Three Reading Paths . . . 32
7 The cohort inheritance and the roadmap of the fourteen Parts . . . 34
7.1 The cohort inheritance . . . 34
7.2 Roadmap of the fourteen Parts . . . 37
8 Part I in plain English . . . 39

Part II. The Derived Signal **40**
9 The Lock Scale and the First Resonance Cluster . . . 40
9.1 The lock-cluster mass: Theorem T-KK-1 . . . 40
9.2 Provenance of every quantitative prediction . . . 44
9.3 The centroid issue: a necessary clarification . . . 47
9.4 Benchmark mass ledger . . . 48
9.5 State-by-state observable benchmark sheet . . . 49
9.6 A schematic picture of the 6D staircase . . . 52
10 Observable Properties of the KK Sector . . . 52
10.1 Spin and angular structure . . . 52
10.2 Widths and line shapes . . . 53
10.3 Masses and intrinsic widths from the RG-scale throat . . . 53
10.4 The cluster fine structure, derived: Theorem T-KK-2 . . . 53
10.5 Branching fractions from the GG-6 couplings . . . 59
10.6 Coupling non-universality from overlap integrals . . . 60
10.7 The mild leptophobic $U(1)_B$ admixture . . . 61
10.8 Correlated partner states . . . 61
10.9 Corollary T-KK-2c: the graviton's gluon coupling, derived . . . 62
10.10 The baryonic normalization g_B: a theorem-grade norm, one discrete residual rule (L-KK-1) . . . 66
11 Part II in plain English . . . 67

Part III. Production and Discovery **69**
12 Production Mechanisms at Hadron Colliders . . . 69
12.1 General narrow-width bookkeeping . . . 69
12.2 Why no absolute cross section is quoted . . . 70

12.3 Quark-antiquark annihilation . . . 70
12.4 Gluon fusion . . . 70
12.5 Vector-boson fusion and associated production . . . 71
12.6 A realistic channel hierarchy for GG-6 . . . 71
12.7 Event-yield bookkeeping . . . 71
13 Channel-by-Channel Discovery Logic . . . 72
13.1 Dileptons: the clean but not necessarily dominant channel . . . 72
13.2 Diphotons: still important, but theoretically disfavored as the sole handle . 72
13.3 Top-pair resonances: the central battleground . . . 73
13.4 Dijets: where the mild leptophobic side of GG-6 becomes visible . . . 73
13.5 Dibosons: the geometry-to-Higgs bridge . . . 74
13.6 Higgs pairs: the best scalar corroborator . . . 74
13.7 The channel-combination principle . . . 74
14 Part III in plain English . . . 75

Part IV. GG-6 versus Randall–Sundrum, and Detection as an Open System 77
15 How GG-6 Differs from Ordinary KK and RS Benchmarks . . . 77
15.1 The ordinary KK and RS reference points . . . 77
15.2 Five differences that matter experimentally . . . 77
15.3 Generic KK, 5D RS, and GG-6: the comparison that matters experimentally 78
15.4 Derived geometry versus fitted resonance . . . 80
15.5 The practical consequence for reinterpretation . . . 80
16 Why Collider Detection is an Open-System Problem: the GDOS View . . . 80
17 Part IV in plain English . . . 81

Part V. Two Read-outs of the Lock Cluster 83
18 Read-out I: The Theorem-Only Target . . . 83
19 Read-out II: An Empirically-Anchored Rate Estimate . . . 84
20 Part V in plain English . . . 87

Part VI. Anatomy of the $t\bar{t}$ Signal 88
21 What This Part Adds, and at What Grade . . . 88
22 The $m_{t\bar{t}}$ Line Shape, with Interference . . . 89
22.1 Amplitudes and color algebra . . . 89
22.2 Squaring: the three terms and the V/A algebra . . . 89
22.3 The sign theorem . . . 90
22.4 Numbers and observable consequences . . . 90
23 Angular Structure from Helicity Amplitudes . . . 91
23.1 Conventions and the two Wigner rows . . . 91
23.2 Spin 1 from $q\bar{q}$: derivation . . . 92
23.3 Spin 2 from gg: derivation . . . 93
23.4 Separation power, quantified . . . 93
24 Top Polarization and Spin Correlations . . . 93
24.1 Longitudinal polarization: $P_t = +0.91$ versus zero . . . 93
24.2 Spin correlations: a consistency check, not a discriminator . . . 94
25 System Kinematics . . . 95
26 Continuum, S/B, and Yields . . . 95
27 The Discrimination Matrix and the Pre-Registered Fingerprint . . . 96
28 Reliability and Open Ends . . . 97
29 Part VI in plain English . . . 97

Part VII. Experimental Status and Reach **99**
30 Current Experimental Status at ATLAS and CMS . . . 99
30.1 The luminosity backdrop . . . 99
30.2 High-mass dilepton constraints . . . 99
30.3 High-mass diphoton constraints . . . 99
30.4 Dijet and substructure-based searches . . . 100
30.5 Top-pair resonance searches . . . 100
30.6 Higgs-pair searches . . . 100
30.7 A channel-dependent exclusion and discovery ledger . . . 101
30.8 What the public status does and does not prove . . . 103
31 The Operational Confrontation: Computed $\sigma \times$ BR versus Published Limits . . . 104
31.1 What this section adds, and at what grade . . . 104
31.2 The workbook: master formula, inputs, and validation . . . 105
31.3 Results . . . 106
31.4 The survival statement, channel by channel . . . 107
31.5 The kill window, quantified . . . 109
32 Near-Term Discovery Potential . . . 109
32.1 The Run-3 to early-HL-LHC bridge . . . 109
32.2 Why $t\bar{t}$ is the priority channel . . . 110
32.3 The role of boosted dijets and Higgs pairs . . . 110
32.4 A benchmark significance corridor . . . 110
32.5 Could ATLAS and CMS already get the first events in the current 2026 run? 111
32.6 Near-term reach in practical language . . . 112
32.7 The falsification timeline in calendar years . . . 113
33 Why This is the Best Collider Killer Experiment . . . 115
33.1 Sharpness compared with other energy-frontier claims . . . 115
33.2 What discovery would establish . . . 117
33.3 What absence would establish . . . 117
33.4 A pre-declared kill criterion . . . 117
34 Long-Term Future Prospects . . . 118
34.1 HL-LHC as the decisive machine . . . 118
34.2 ATLAS Phase-2 upgrades . . . 118
34.3 CMS Phase-2 upgrades . . . 118
34.4 FCC-hh, muon collider, and CLIC . . . 119
34.5 Indirect probes . . . 119
35 A Trigger and Analysis Roadmap for ATLAS and CMS . . . 119
35.1 Use the high-rate frontier aggressively . . . 119
35.2 Make the search cluster-aware at trigger level . . . 120
35.3 Promote top and Higgs boosted objects to first-class citizens . . . 120
35.4 Exploit precision timing as a resonance-quality variable . . . 120
35.5 A joint benchmark release would help enormously . . . 120
36 Part VII in plain English . . . 122

Part VIII. Simulating the Discovery **123**
37 What This Part Adds, and the Rules of the Game . . . 123
38 The Anchor Analysis, Read Like an Analyst . . . 124
38.1 The dominant backgrounds, and how each is estimated . . . 124
38.2 The trigger, and why one lepton is enough . . . 125
38.3 Event selection and reconstruction, cut by cut . . . 125
38.4 What $\varepsilon A = 6.5\%$ is made of . . . 126
38.5 Resolution, binning, and the published sensitivity . . . 126
39 The Simulation, Step by Step . . . 127

39.1 Signal: from the derived line shape to events in bins . . . 127
39.2 Background: anchoring the total to the measured yields . . . 127
39.3 From 13 to 14 TeV, and "14-TeV-equivalent" luminosity . . . 128
40 The Statistics of Discovery, from First Principles . . . 128
40.1 Counting, the likelihood ratio, and the meaning of Z . . . 128
40.2 A background systematic as a nuisance parameter . . . 129
40.3 The systematics wall, located precisely . . . 129
41 Calibration: Tying the Toy to the Real Analysis . . . 130
41.1 What "improved systematics" means, quantitatively . . . 130
42 Why the Eye Overestimates: The Significance Ladder . . . 131
43 The Four Snapshots: What the Detectors Record . . . 133
44 The Pre-Registered Search: What the Exact Prediction Buys . . . 135
44.1 Advantage 1: no look-elsewhere penalty (banked) . . . 135
44.2 Advantage 2: the matched template (banked); what the published pipeline actually tests (not this) . . . 136
44.3 Advantage 3: the angular dimension and the exact rate (not banked) . . . 136
44.4 The two-experiment combination, and a calibrated verdict . . . 137
44.5 A warning the dedicated fit must respect: the sideband is contaminated . . . 137
45 What This Means for F-KK-2B, and the Caveat Ledger . . . 138
46 Part VIII in plain English . . . 139

Part IX. Beyond the HL-LHC **140**
47 The Question This Part Answers . . . 140
48 The Facility Landscape as It Actually Stands . . . 140
49 The Post-Discovery Checklist: What Is Left After 2035 . . . 141
50 Machine by Machine: What Each Facility Decides . . . 142
50.1 FCC-hh at 84–85 TeV: the spectroscopy machine . . . 142
50.2 FCC-ee: the parameter-ledger machine, and the no-radion test . . . 143
50.3 A muon collider: the leptophobia machine . . . 143
50.4 Linear colliders, the LHeC, and the supporting cast . . . 144
51 Concrete Scenarios, 2035–2070 . . . 144
52 Summary Matrix . . . 145
53 Part IX in plain English . . . 145

Part X. Discussions **149**
54 Significance of the result . . . 149
55 Uniqueness, nobility, and innovation . . . 151
55.1 The classic collider puzzles, resolved or dissolved . . . 152
55.2 The structural innovations . . . 153
55.3 The nobility clause . . . 154
56 Comparison with the other roads beyond the Standard Model . . . 155
57 Predictions for falsification . . . 156
58 Open questions . . . 159
59 The completed cohort and future prospects . . . 161
60 Summary of the Discussion . . . 161
61 Part X in plain English . . . 162

Part XI. Conclusions **163**
62 What the paper has established . . . 163
63 The objective limitations and the next technical tasks . . . 164

64 A personal perspective: fifty years on the experimental front line . . . 164
65 The vision: what a geometric century could look like . . . 165
66 The trailhead, not the summit: a final reflection . . . 166
67 What it would mean: matter, number, and the legibility of nature . . . 166
68 Part XI in plain English . . . 168
Acknowledgments . . . 169

Part XII. Technical Appendices **170**
A Summary of Abbreviations . . . 171
B Glossary of key symbols . . . 173
C Global naming convention for cross-cohort citation . . . 174
D Comprehensive summary tables . . . 174
E What is the relation between the QCD gluon and the color-octet g^*? . . . 177
F What is the origin of the baryonic vector Z'_B? . . . 178
G What is the "graviton" of GG-Theory? . . . 180
H Why does g^* have the largest production rate — and why is it so broad? . . . 182
I One integer, two powers — the rate hierarchy and the $t\bar{t}$ dominance . . . 183
J From the 6D action to the KK spectrum — a course . . . 184
J.1 Step 1: the textbook derivation, from scratch . . . 184
J.2 Step 2: what the textbook cannot fix . . . 185
J.3 Step 3: the GG-6 derivation, and the worked number . . . 185
J.4 Reference: the reduction formulas . . . 186
K From parton distributions to cross sections — a course . . . 186
K.1 Step 1: the factorized cross section, from scratch . . . 186
K.2 Step 2: what the chain assumes but cannot fix . . . 187
K.3 Step 3: the GG-6 chain, closed at both ends . . . 187
L The statistics of discovery — a worked course . . . 188
L.1 Step 1: from the Poisson likelihood to Z, with the arithmetic shown . . . 188
L.2 Step 2: what counting cannot see . . . 188
L.3 Step 3: the pre-registered payoff, worked . . . 188
L.4 Reference: the significance-scaling workbook . . . 189
M Benchmark / reinterpretation card for experimentalists . . . 189
M.1 Implementation walkthrough . . . 190
N Glossary of collider terms . . . 190
O Promotion roadmap for the open items . . . 191

Part XIII. Common Objections and Responses **192**
P Methodology and the derivation chain . . . 192
Q The mass and the fine structure . . . 194
R Couplings, rates, and suppressions . . . 196
S The simulation and its statistics . . . 198
T Reducibility, non-spatial dimensions, and neighboring frameworks . . . 200
U Cohort dependence and standing . . . 202
V Experimental practice . . . 203

Part XIV. References **205**

Preface

This Preface is addressed, more than to anyone else, to a graduate student who is in — or is deciding whether to join — an ATLAS or CMS analysis group, and who is wondering what to do with the next five years.

Consider, first, where the field actually stands. The Standard Model was completed as a structure in the mid-1970s, and every collider discovery since has been an item from its list: the W and Z bosons (1983) confirmed the electroweak theory of 1967; the top quark (1995) had been required, at roughly its measured mass, by precision fits; the Higgs boson (2012) was the last unchecked box of a mechanism written down in 1964. These were monumental achievements of experimental physics — and every one of them was a *confirmation.* No collider has ever observed a particle that was not on the Standard Model's list. If you join an LHC experiment today, you join half a century after the last time the list itself grew.

Now consider what this paper puts on the table. A theory with *no adjustable parameters* — its entire freedom is a handful of integers — derives a new mass: 3.849 TeV, from the Planck scale and the integer 35, full stop. It derives what lives there (three sibling states), how they decay (top pairs, mostly), how wide the main state is (15%), how it interferes with the ordinary background (constructively below, destructively above — a theorem, including the sign), and how it is distributed in angle. It even ties the new number to one you already know: the same chain forces the electroweak scale, 246 GeV. There is no version of this theory where the resonance is at 4.5 TeV instead, or couples differently, or fades to an undetectable corner. It is either there, or the theory's current implementation is dead. You will not find many statements of that shape in the beyond-Standard-Model literature; §1 explains why their absence is precisely the lesson of the last fifty years.

Here is what makes this concretely *your* opportunity rather than an abstraction. Part VIII of this paper simulates the search at the fidelity of a real analysis, anchored to the published ATLAS Run-2 top-pair search, and its headline is remarkable: *the data that decide this may already be on tape.* A dedicated reanalysis of the existing Run-2+3 dataset carries evidence-level sensitivity per experiment (2.7–3.7σ), and a combined ATLAS+CMS analysis with realistically improved systematics reaches 4.6–5.2σ — discovery-grade — from collisions already recorded. The analysis this requires is thesis-sized, not collaboration-sized in its first pass: one channel (boosted top pairs, one lepton), published object definitions, one new template (the interference-aware line shape of Part VI), one warning to respect (§44.5). The discovery window — 2026 through roughly 2033 — is not some future generation's; it is the exact span of one PhD begun now. Somebody's thesis is going to decide this question. There is no reason it should not be yours.

And consider what a positive result would mean. The Higgs discovery, for all its magnitude, closed a chapter: it completed a known structure. A cluster at 3.849 TeV would *open* one: the first particle beyond the Standard Model in the history of collider physics, the first direct evidence that the dimensions of nature include the energy scale itself, and — through the electroweak over-constraint — the first explanation of *why* the weak scale is what it is. That is not an incremental measurement. That is the kind of result the LHC was built hoping for.

Honesty requires the other half, and it is the half that makes this science rather than salesmanship:

the theory may simply be wrong. But notice what the search costs you in that case — almost nothing. The analysis that tests it is, bin for bin, the world's best measurement of the high-mass top-pair spectrum and the strongest limit on *any* broad top-philic resonance; the interference-aware template you would build is reusable across every colored-resonance model; and a clean exclusion of a zero-parameter theory is a permanent result, not a parameter-scan press release. A pre-registered prediction is decided either way. That is the property that has been missing from this field for fifty years, and it is the property this paper exists to put back.

How to read the paper, if you are the student described above: start with Part VIII (the simulated search — it is self-contained and teaches the statistics from the Poisson likelihood up), then Part VI (what the signal looks like and why, at full derivation depth), then §30 (the current limits) and §31.5 (the kill criterion). Work backwards into the foundations (Parts I–III) as curiosity, not obligation, dictates. The theory chapters will still be there after your trigger shift.

Part I
Foundations

From the GG-Constitution to a six-dimensional collider target

1 Introduction: Fifty Years Beyond the Standard Model

The triumph, and the silence that followed. The Standard Model was assembled, in every essential, by the mid-1970s: the electroweak theory in 1967–68, quantum chromodynamics and asymptotic freedom by 1973, the third generation anticipated on paper in the same years. What followed was one of the great runs of confirmation in the history of science — and it is important, for what comes next, to see it as exactly that. The charm quark (1974) had been demanded by the GIM mechanism; the bottom quark (1977) by the Kobayashi–Maskawa structure; the gluon (1979) by QCD; the W and Z (1983) arrived at the masses the electroweak theory required; the top quark (1995) had been located, within errors, by precision fits before it was seen; the Higgs boson (2012; ATLAS Collaboration, 2012; CMS Collaboration, 2012) was the final unchecked entry of a mechanism written down in 1964. Every one of these discoveries was a particle from a list completed before 1975. In fifty years, no collider has added a single entry *to* the list. The lone established crack in the edifice — neutrino oscillations (1998; Fukuda et al., 1998) — came from underground detectors, not colliders, and prescribes no new collider scale.

Three roads that did not arrive. It is not for lack of trying, and the record deserves to be stated plainly, because the field rarely states it. For fifty years, three great programs promised the next layer. *Supersymmetry*, in its minimal form the MSSM, promised to stabilize the electroweak scale with partner particles "just above" each successive frontier — but the minimal model carries more than one hundred free parameters, so every null result at LEP, at the Tevatron, and at LHC Runs 1 and 2 was answered by sliding the spectrum upward, and the natural parameter region that motivated the enterprise has by now been all but eliminated. *Extra-dimensional models* — large extra dimensions, and the warped Randall–Sundrum geometry that this paper's subject superficially resembles — contributed genuine structural insight, but their mass scales were always *inputs*: a radius or a warp factor to be scanned, so that no exclusion could ever be fatal, only relocating. *String theory* constructed a candidate for everything and committed itself to nothing measurable: the landscape of vacua, at a famously astronomical count, provides no single number at accessible energy whose failure would falsify anything. Fifty years, three programs, and not one clear, pre-registered, falsifiable collider prediction among them — and by now the first two stand nearly falsified in their motivating forms, while the third cannot be falsified at all.

The diagnosis is structural, not personal. These were serious ideas pursued by physicists of the first rank, and much of their machinery — the supersymmetry algebra, warped geometry, holography — is permanent mathematics that will outlive the particular models. The failure mode was structural: *continuous free parameters*. A theory with dials can retreat from any measurement, so it never stakes anything; it survives every experiment and is confirmed by none. A theory that can slide cannot lose — and a theory that cannot lose cannot win. The lesson of the five decades is therefore not that the next layer is out of reach; it is that predictive power is measured by what a theory *forbids*, and dials are the enemy of forbidding.

What a worthy successor would have to look like. It would put a number on the table *before* the data arrive. The number would be derived — traceable through an explicit chain to structural constants, with no continuous parameter anywhere in the chain that could absorb a disagreement. It would over-commit: tie the new number to quantities already measured, so that sliding the prediction would break what already works. And it would specify the signal completely — not a mass alone, but widths, decay preferences, angular shapes, even interference signs — so that the experimental verdict, in either direction, is beyond appeal. Half a century of frustration has, at least, produced an exact job description. The next section introduces the theory that applies for the job.

2 GG-Theory in Plain English — and Why the LHC Gets to Judge It

If this paper is your first encounter with GG-Theory, welcome — this section was written for you, and it assumes nothing. The previous section ended with a job description that no theory of the last fifty years has filled. Before the derivations begin, you deserve to hear, in plain words and at a humane pace, what kind of theory now applies for the job — because it is unusual enough that you should first meet it as a story, not as a formalism.

Here is the story. GG-Theory proposes (Arisaka, 2026, A2-GG6) that the world has six dimensions — but not in the familiar sense of curled-up extra *space*. Four are ordinary spacetime. The fifth is the *energy scale* itself: the "zoom level" at which you examine nature — the thing that changes as you go from atoms to nuclei to quarks — taken seriously as a geometric direction, with a coordinate $r = \ln(E_{\mathrm{GUT}}/\mu)$ and a finite length that is an integer, $K = 35$. The sixth is the *gauge phase* — the angle that rotates the wavefunction of charged matter, the angle the Aharonov–Bohm effect makes visible — promoted from bookkeeping device to literal coordinate. You cannot walk along either of them; nothing can, because they are not places. Yet fields live on all six directions, and standing waves along the scale direction are particles with definite masses — computable ones, because the length of the direction is a counting number rather than a knob. That is the theory's signature everywhere: its entire freedom is five integers, $(M, N, L; J, K) = (1, 3, 5; 7, 35)$, the OGN-5 ledger whose anomaly-cancellation origin is proved in GG-A3-BI6 (Arisaka, 2026, A3-BI6). Each of the five carries a fixed meaning — $M = 1$ the single tensor seam and charge-quantization unit, $N = 3$ the color triplicity, $L = 5$ the SU(5) block, $J = 7$ the odd closure index, and

$K = 35 = LJ = M^2 + N^2 + L^2$ the throat depth that fixes the lock scale $M_{\text{lock}} = \frac{1}{2} M_{\text{Pl}} e^{-35}$ — so the five are forced, not fitted, and not an arbitrary quintet. There is no continuous parameter anywhere. Once the integers are fixed — and they are fixed by consistency, not by fitting — every mass, coupling, width, and angle that follows is an *output*.

The GG-6 wager. This paper presents the sharpest of those outputs, and asks the LHC to judge it. From the integer ledger the theory derives the *lock scale* (Arisaka, 2026, A2-GG6, Theorem 6; Arisaka, 2026, A9-SM)

$$M_{\text{lock}} = \frac{E_{\text{Pl}}}{2e^{35}} \simeq 3.849\,\text{TeV}, \tag{1}$$

where the 35 in the exponent is the integer K, not a fitted quantity — and it spends the same number twice, because the electroweak scale must then equal $M_{\text{lock}}(2/5)^3 \approx 246$ GeV, and does. At the lock scale it places not a lone bump but a three-member cluster — a spin-2 graviton mode, a color-octet vector, a leptophobic Z'_B — with derived couplings, derived widths, a derived interference sign, and derived angular distributions, decaying preferentially to top-quark pairs; and the un-suppressible member of that cluster sits, at its computed rate, exactly on the current ATLAS/CMS exclusion frontier. A null result at $\lesssim 2$ ab^{-1} retires the current implementation outright — there are no parameters with which to escape. This wager may fail; it is constructed so that failure is *visible*.

Measure this against the job description of §1, item by item. A number on the table before the data arrive: 3.849 TeV, in print. A derived number: traceable through an explicit theorem chain to the Planck scale and an integer, with no continuous parameter anywhere in the chain. Over-commitment: the same chain forces the electroweak scale already measured, so sliding the prediction would break what already works. Complete specification: not a mass alone but siblings, widths, branching preferences, angular shapes, and the interference sign, each derived and each checkable separately. Individual *models* have, on occasion, staked a fixed number — a benchmark Z' here, a pre-registered bump hypothesis there. What no program has done is all four at once: derive the number, over-commit it to quantities already measured, specify the complete signal, and state the criteria of failure in advance. This is why we say, carefully and without exaggeration, that this is the first time in the fifty-year history of "beyond the Standard Model" that a proposal both predicts its quantities and states, in advance and in writing, the criteria of its own falsification.

And the timing could hardly be better — this is the part that should quicken the pulse of any experimentalist reading. Part VIII of this paper simulates the search at the fidelity of a real analysis, anchored to the published ATLAS top-pair search, and its headline is that *the collisions that decide this question may already be recorded*: a dedicated reanalysis of the existing Run-2+3 dataset carries evidence-level sensitivity — 2.7–3.7σ per experiment if the cluster is there — so a first $>3\sigma$ hint of the color-octet g^* is already latent in the collisions recorded by the end of 2026, waiting only for the dedicated reanalysis this paper specifies, and a combined ATLAS+CMS analysis with realistically improved systematics reaches discovery grade from the same tape. Whichever way it falls, the question closes within the HL-LHC decision window, roughly 2026–2035: the exact span of one doctoral career begun now. So keep reading — the rest of this Part equips you with the one genuinely new concept the theory runs on (§3), the constitutional route by which a collider prediction descends from axioms (§4), and the six-dimensional geometry that carries it (§5); and if you want the theory beyond this paper, the twenty-two-paper cohort map of §7.1 shows where every door leads.

The reason this question is sharper than the usual beyond-Standard-Model question is worth restating in the cohort's own terms: the number in Eq. (1) is not tuned to the collider. It is the same number that participates in the electroweak-scale derivation and in the hierarchy ledger of the theory. A null or positive answer at the LHC therefore propagates immediately into the constitutional core of the program.

This paper has four simultaneous aims. The first is theoretical: to explain, from the constitutional layer of the theory and the already-derived GG-6 theorem package, why a six-dimensional Kaluza–Klein (KK) sector exists at all and why its first experimentally relevant scale is anchored near 3.85 TeV. The second is phenomenological: to translate that structural statement into concrete collider observables – masses, widths, angular distributions, channel hierarchies, and correlated partner states. The third is experimental: to review the current public ATLAS and CMS landscape in a way that is fair both to the collaborations and to the specific non-RS features of GG-6. The fourth is strategic: to state explicitly why the LHC and especially the HL-LHC constitute the best collider falsifier of the current GG-6 implementation, and how ATLAS and CMS could improve their sensitivity with relatively practical trigger and data-format choices.

The discussion is intentionally long and methodical because a serious falsifier deserves serious bookkeeping. The paper therefore begins one layer above the KK spectrum itself. It starts from the constitutional route of the GG-Constitution, then moves to the 6D geometry, the KK reduction, the lock-scale derivation, the collider observables, the current ATLAS/CMS status, the HL-LHC outlook, and finally to a trigger roadmap. Readers primarily interested in experiment may skip directly to Sections 30 and 32; readers primarily interested in the conceptual role of GDOS and OGGEP may wish to read Sections 4 and 16 with special care. (A fuller set of reading paths, one per kind of reader, is laid out near the close of this Part: §6.)

Before proceeding, it is useful to state a central normalization point clearly. In many extra-dimensional papers the phrase "first KK graviton mass" refers to a specific single-pole solution of a Bessel equation in a 5D RS background. In this paper, by contrast, the collider benchmark 3.849 TeV refers to the *first lock-scale cluster* of the 6D GG-6 problem. The first isolated sub-pole within that cluster may depend on details of the full 6D boundary problem; the cluster centroid does not. This distinction is not semantic. It changes how one should search in data.

3 The central claim: the Kaluza–Klein lock cluster as an open-system observable

It is worth stating, in one place and at a walking pace, the recognition this paper shares with the whole GG-Theory cohort — because it is the one genuinely new concept a reader needs, and because without it the true significance of everything that follows cannot be grasped. Nothing in this section requires prior exposure to GG-Theory; it requires only honesty about something every physicist already knows and every textbook quietly sets aside.

No system in nature is closed — not even one electron. Consider the simplest object physics owns: a single electron. It is never alone. It is dressed by its own electromagnetic field; it exchanges virtual photons with the vacuum; anything that has ever scattered off it carries a record of the encounter away. Strip all of that off — imagine an electron truly decoupled from everything — and you have imagined a particle that cannot be seen, deflected, or counted, because *to observe a system is precisely to open it.* A perfectly closed system is not merely hard to prepare; it is unobservable in principle, a contradiction in terms the moment a measurement is attempted. Yet the whole formal apparatus of physics is built the other way around: textbook dynamics — classical or quantum — is the dynamics of *closed* systems, evolved in splendid isolation, with the observer bolted on afterwards as an axiom. Everything in nature, from a single electron to a living cell to the universe's contents as a whole, is an open system; our theories have treated openness as a nuisance to be appended. GG-Theory begins by refusing that inversion.

Where nature runs, and where we watch it from. In GG-6, the complete dynamics of nature runs in the six-dimensional bulk — ordinary spacetime plus the scale and phase directions of §2. In the bulk the accounting closes: the dynamics is lawful, complete, and deterministic. But no experiment has ever been performed in the bulk, and none ever will be: we and our instruments inhabit the four-dimensional boundary, and every datum ever recorded — every track, every calorimeter cluster, every published cross section — is a *boundary record*. The relation between the two is the theory's third axiom, OGGEP (AX-3 of GG-A1-Constitution, Arisaka, 2026, A1-Constitution): the four-dimensional world is a lawful *projection* of the six-dimensional bulk, and **that projection is what observation is**. Observation is not something a conscious agent adds to physics; it is the bulk-to-boundary map itself, running always and everywhere, with or without anyone watching. This single reversal — openness as the fundamental condition, observation as geometry — is the conceptual core of the entire twenty-two-paper program.

The law of the projection: GDOS and GRIP. An axiom that says "the boundary is a projection" would be philosophy if the projection were arbitrary. It is not. Its dynamics is a definite law — GDOS, the GG-6 dynamics of open systems, derived in GG-A5-GDOS (Arisaka, 2026, A5-GDOS) — which specifies how a projected system actually evolves, dissipates, and registers: what the boundary keeps of the bulk, and at what cost. And GDOS is itself forced rather than chosen, pinned from two independent sides: the geometric pillars of Gi-4 (A6-GRIP, Arisaka, 2026, A6-GRIP) and the quantum-information pillars of Qi-4 (GG-A7-Qi4, Arisaka, 2026, A7-Qi4). The runtime form of the law is the GRIP cascade of A6-GRIP, with its two regimes GRIP-In and GRIP-Out: GRIP-In is the bulk feeding the boundary — excitation, descent, the arrival of energy at the seam — and GRIP-Out is the boundary paying it out as registered, irreversible records. In this vocabulary a detector is not a passive witness; it is a purpose-built GRIP-Out amplifier, engineered to make one microscopic registration grow into a macroscopic, thermodynamically irreversible entry in a database.

A century of confusion, resolved (GG-A8-QM). The reader has met this problem before, under another name. Since 1926, quantum mechanics has carried two rules that do not speak to each other: unitary evolution for systems left alone, and "collapse" for systems measured — with no principled statement of where one ends and the other begins. Bell called the dividing line the *shifty split*, and a century of interpretations has consisted largely of relocating it: into consciousness, into decoherence, into many worlds, into hidden variables. GG-A8-QM (Arisaka, 2026, A8-QM) shows that within GG-6 the split is not shifty — it is the bulk/boundary interface, and it is *geometric*. What textbooks call collapse is the OGGEP projection in action: in the bulk the evolution never stops being lawful and deterministic, while the boundary keeps only what the projection law delivers — a definite, registered outcome, with the Born weights recovered as a typicality statement rather than postulated. The measurement problem, on this reading, was never a paradox of quantum mechanics; it was the signature of trying to run a six-dimensional accounting in four dimensions — a century of confusion produced by one missing map. This is why the present paper keeps insisting on open-system language for what might seem a plain bump hunt: the language is the solved problem.

Colliders follow this principle exactly. Now look at what ATLAS and CMS actually do, and notice that it is OGGEP operating at industrial scale. No experiment has ever recorded a gauge-dependent field amplitude or a bulk configuration: what is recorded are gauge-invariant kinematic combinations — invariant masses, transverse momenta, angular distributions, channel-by-channel rates — reconstructed from irreversible energy depositions. The full measurement chain (hard scattering, parton shower, hadronization, detector response, trigger, reconstruction, statistical decision rule) is precisely a system instance packet in the language of GDOS: a specification of *which* projection is being run and what it registers. A resonance search is therefore not merely compatible with the open-system view; it is the open-system experiment par excellence — a bulk excitation, if one exists at 3.849 TeV, announcing itself through the only door nature provides.

The conclusion of this chain is the central claim of the paper. Because the projection is lawful and its law is forced, what the boundary can register is fixed by the bulk geometry — not chosen, not tunable. The Kaluza–Klein lock cluster near $M_{\text{lock}} \approx 3.85$ TeV, with the fifth-force and dark-sector signals that share its scale, is a *registered observable* of one projection: the geometry determines it, which is why this paper can derive it, and why a single clean measurement can refute it. Figure 1 states the whole argument as one deductive funnel; the same funnel opens the companion papers GG-A9-SM (Arisaka, 2026, A9-SM), GG-A10-Cosmos (Arisaka, 2026, A10-Cosmos), and GG-A11-Origin (Arisaka, 2026, A11-Origin), with only the payload differing — there the Standard Model's parameters, the cosmos, and the origin of law; here, a resonance cluster within reach of a machine that already exists.

4 From the GG-Constitution to Collider Observables

The argument here starts not from a phenomenological ansatz imported from RS, but from the GG-Constitution axiom system and the already-derived theorem packages of the cohort. This is the right starting point: in the present framework, collider phenomenology is a downstream consequence of a constitutional grammar, not an independent bolt-on model.

At the constitutional level, four ingredients matter most for the present paper (all from GG-A1-Constitution, Arisaka, 2026, A1-Constitution). **AX-1** (local causality) demands that admissible

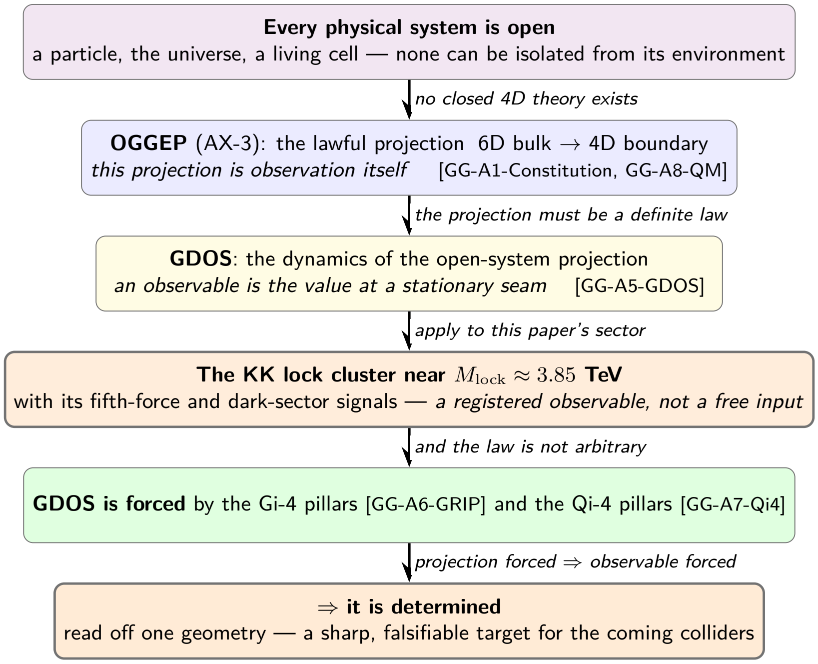

Figure 1: The central claim, as one deductive funnel, shared across the GG-Theory cohort. Because no physical system is closed, the four-dimensional world is a lawful projection (OGGEP, AX-3) of the six-dimensional bulk, and that projection is observation itself; its dynamics is GDOS, forced by the Gi-4 and Qi-4 pillars. The same funnel opens the companion papers GG-A9-SM, GG-A10-Cosmos, and GG-A11-Origin; only the payload differs — here, the Kaluza–Klein lock cluster.

kernels be retarded and that registered detector observables arise from causal boundary records. **AX-2** (GGEP) states that gauge and geometric content must arise from a single primitive composite connection. **AX-3** (OGGEP) requires an explicit bulk-to-boundary operational dictionary: a theory is not complete unless it says how bulk dynamics become registered observables. **AX-4** (maximum symmetry, minimum freedom) forbids gratuitous free parameters and elevates minimal integer data to constitutive status. Together with the anomaly-cancellation and normalization postulates (POS-3 and POS-6), these axioms force a remarkably specific route to collider physics (Arisaka, 2026, A1-Constitution; Arisaka, 2026, A2-GG6).

The logic runs as follows. Once AX-2 and AX-4 are taken seriously, the bulk cannot remain an unconstrained higher-dimensional playground. It must support a single theorem-level geometry in which the renormalization direction and the compact Abelian holonomy are both geometric, because both are already required to represent the running and the compact gauge phase with minimal freedom. This selects the six-dimensional stage

$$M_6 = M_4 \times I_r \times S^1_\theta, \qquad r = \ln\left(\frac{E_{\text{GUT}}}{\mu}\right), \tag{2}$$

with ultraviolet gauge-gravity block

$$G_{\text{bulk}} = \text{SU}(5) \times \text{U}(1)_B \times \text{U}(1)_{Qi} \times \text{Gravity}. \tag{3}$$

The 6D tensor sector is then fixed by the single-tensor Green–Schwarz-modified Bianchi identity BI-6,

$$dH_3 = X_4, \qquad X_4 = \operatorname{tr}\mathcal{R}^2 - \operatorname{tr}F_5^2 - \frac{2}{5}F_B^2 - \frac{2}{5}F_{Qi}^2, \tag{4}$$

which is the structure that closes the anomaly story and makes GG-6 fundamentally different from a bare Kaluza–Klein metric reduction (Arisaka, 2026, A2-GG6; Arisaka, 2026, A3-BI6).

Why does this matter for KK particles? Because once the stage (2) exists, the boundary cannot remain spectrally featureless. Every field living on that six-dimensional stage admits a harmonic decomposition on the internal space. A KK tower is therefore not an independent assumption; it is the boundary spectral shadow of the bulk geometry itself. The only remaining question is what fixes the first accessible scale. The answer comes from AX-4 plus POS-6: the OGN-5 integer ledger

$$(M, N, L; J, K) = (1, 3, 5; 7, 35), \qquad J = 2N + 1, \qquad K = LJ = M^2 + N^2 + L^2, \tag{5}$$

fixes the throat depth, and therefore the lock scale,

$$M_{\text{lock}} = \frac{E_{\text{Pl}}}{2e^{35}} \simeq 3.849\,\text{TeV}. \tag{6}$$

That is why the 3.85 TeV collider target is not simply a fit. It is the theorem-level infrared end of the geometry.

The collider observable is then supplied by AX-3. Through OGGEP, bulk excitations become boundary records: invariant masses, transverse momenta, angular distributions, and channel-by-channel rates. From this perspective, a heavy resonance search is already an open-system problem. The hard scattering, parton shower, hadronization, detector response, trigger, reconstruction, and statistical decision rule together constitute a system instance packet (SIP) in the language of GDOS. The fact that ATLAS and CMS reconstruct gauge-invariant kinematic combinations rather than raw gauge-dependent field variables is not incidental; it is precisely the collider form of DGI and OGGEP at work. We will return to this point in Section 16.

Table 1: Constitutional route from the GG-Constitution to the KK collider program. The table is intentionally compact: it is meant to show that the collider sector is not a detached phenomenological appendix, but a direct downstream consequence of the constitutional grammar.

Layer	Constitutional content	Collider consequence
AX-1	Local causality	Detector observables must be built from retarded, registered records; the search is formulated in causal event streams, not in acausal inverse problems.
AX-2	Geometry–gauge equivalence	Gauge and gravitational sectors share a single primitive geometric stage; KK towers are geometric, not optional.
AX-3	OGGEP bulk–boundary dictionary	Bulk excitations become boundary invariant masses, angular distributions, and branching patterns in ATLAS/CMS.
AX-4	Maximum symmetry, minimum freedom	The infrared scale is fixed by minimal integer data rather than scanned as a free parameter.
POS-3	Anomaly freedom by inflow	BI-6 closes the tensor/gauge system and distinguishes GG-6 from ordinary KK or RS intuition.
POS-5	RG compatibility	The radial coordinate r is the renormalization direction; infrared collider masses are scale-redshifted images of ultraviolet geometry.
POS-6	Minimal-rational normalization	OGN-5 fixes $K = 35$, hence the lock scale $M_{\text{lock}} \simeq 3.849\,\text{TeV}$.

The bottom line is that the collider prediction does not float freely. It is tied to the same constitutional commitments that underwrite the rest of the program. This is precisely why the KK search is so sharp a falsifier. (The program-wide panorama — the twenty-two-paper cohort map, and the conventions for reading this paper against it — closes this Part: §7.1, Figure 2.)

5 The 6D Geometric Origin of the KK Tower

5.1 The six-dimensional stage

The theorem-level GG-6 bulk is a genuine six-dimensional Lorentzian manifold. Its metric is fixed – not chosen – by Theorem 2 of Arisaka (2026, A2-GG6), which forces the canonical product geometry

$$ds_6^2 = e^{-2r}\,\eta_{\mu\nu}\,\mathrm{d}x^\mu \mathrm{d}x^\nu + \mathrm{d}r^2 + \mathrm{d}\theta^2, \qquad r \in [0, K], \quad \theta \in [0, 2\pi), \tag{7}$$

in canonical GG-6 units, carrying *no* free metric data: the AdS curvature k, the circle radius R_θ, and the throat length r_c of a generic warped construction are all locked (to unity, or to the OGN-5 integer $K = 35$) by the maximum-symmetry/minimum-freedom posture (Arisaka, 2026, A2-GG6; Kaluza, 1921; Klein, 1926; Randall & Sundrum, 1999). The decisive structural fact is that both internal directions are *non-spatial*: r is the scale (renormalization-group) direction I_r, and θ is the phase circle S^1_θ, each dimensionless. Neither is a spatial extra dimension. This already shows why the eventual spectrum is richer than the standard 5D RS picture: there are two independent internal directions – a scale direction and a compact phase direction – so spectral labels can arise from both. In the mode equations below we retain a general scale profile $A(r)$ and circle radius R_θ as bookkeeping symbols, specialized to the canonical $A(r) = r$ and $R_\theta = 1$ at the end.

The role of r is physically transparent in the latest GG-6 synthesis. It is not merely an extra dimension inserted to imitate RS. It is the geometric representation of the renormalization scale,

$$r = \ln\left(\frac{E_{\mathrm{GUT}}}{\mu}\right), \tag{8}$$

so transport along r is the geometric avatar of running.[2] The role of θ is equally important: it is the compact Abelian holonomy direction whose topological structure feeds both the $\mathrm{U}(1)_B$ and $\mathrm{U}(1)_{Qi}$ sectors via BI-6. In ordinary KK language one could say that θ generates a circle tower. In the stronger GG-6 language, one should say that θ carries one piece of the anomaly-locking story itself.

5.2 Mode expansion and the two quantum numbers

Any six-dimensional field $\Phi(x^\mu, r, \theta)$ can be expanded as

$$\Phi(x^\mu, r, \theta) = \sum_{n=0}^{\infty} \sum_{\ell=-\infty}^{\infty} \phi^{(n,\ell)}(x^\mu)\, f_{n,\ell}(r)\, e^{i\ell\theta}. \tag{9}$$

Here n labels the radial excitation family and ℓ labels the compact-angle Fourier mode. The profile functions satisfy a radial Sturm–Liouville problem. For scalar fluctuations a representative mode equation is (Arisaka, 2026, A2-GG6):

$$-\partial_r\Big(e^{-4A}\partial_r f_{n,\ell}\Big) + e^{-4A}\frac{\ell^2}{R_\theta^2} f_{n,\ell} = m_{n,\ell}^2 e^{-2A} f_{n,\ell}. \tag{10}$$

The exact spin-2 equation differs in tensor structure but shares the same qualitative ingredients: a radial operator, a compact-angle mass contribution, and an infrared scale suppression.

The compact-angle eigenvalue carries the Wilson line. One refinement of Eq. (10) must be recorded here, because the electroweak sector of the cohort rests on it. The integer ℓ is the eigenvalue of $-i\partial_\theta$ only when the gauge connection along the circle vanishes. In the presence of a Wilson line — the holonomy $\Theta_H = g_6 \oint_{S^1_\theta} A_\theta \, d\theta$, which is gauge-invariant and cannot be removed — a field of charge q under the corresponding generator feels the shifted eigenvalue

$$\frac{\ell^2}{R_\theta^2} \;\longrightarrow\; \frac{1}{R_\theta^2}\left(\ell + \frac{q\,\Theta_H}{2\pi}\right)^2, \tag{11}$$

which is the Hosotani mechanism written spectrally (Hosotani, 1983): a constant A_θ is locally pure gauge and changes no field strength, yet it moves every charged mode, and in particular it lifts the

[2] A careful reader will ask in which renormalization scheme μ is defined, since RG "distance" is scheme-dependent beyond one loop. The chart is bookkeeping, not physics: a scheme change acts as a smooth reparametrization of the scale direction, under which everything this paper uses — the integer length $K = 35$ (a ledger integer fixed by anomaly matching, Arisaka, 2026, A3-BI6, not a measured perturbative distance), the canonical metric of Arisaka (2026, A2-GG6) (Theorem 2), the lock mass, and the overlap integrals evaluated in the canonical chart — is invariant. Choosing $A(r) = r$ is the analogue of choosing normal coordinates.

$\ell = 0$ modes that would otherwise be massless. *That lifting is the origin of electroweak mass in this framework*: the four-dimensional Higgs doublet is the zero mode of A_θ in the broken directions, and the holonomy that Eq. (11) inserts is the same θ_H whose selected value Arisaka (2026, A9-SM) fixes from the OGN-5 ledger. This paper does not select that value — the vacuum angle and the scale at which the resulting four-dimensional masses are registered are theorems of the codebook paper, where the warp puts them at the infrared endpoint rather than at $E_{\rm GUT}$. What is recorded here is only the spectral fact, because it is the mechanism by which a circle-side label becomes a laboratory mass.

Two consequences for this paper's own predictions, both benign. First, the radial ($\ell = 0$) lock cluster of §9 is unaffected in its *radial* structure: the shift adds a common compact-angle contribution, and for the color-neutral, $U(1)_B$- and graviton-sector modes that build the cluster it is not the quantity that sets the infrared scale. Second — and this is the one worth stating, since it could have gone the other way — the no-angular-satellites prediction survives. Eq. (11) displaces the angular tower by a fraction $\Theta_H/2\pi$ of one unit, and with the selected angle that fraction is of order one per cent; the $\ell = \pm 1$ partners therefore remain at $\sim E_{\rm GUT}$, still some twelve orders of magnitude above the lock cluster. A satellite a few percent from 3.85 TeV would falsify the non-spatial reading exactly as before.

Two immediate qualitative conclusions follow, and they sharply separate the two towers. First, the radial ($\ell = 0$) modes are the normalizable resonances of the *finite* radial interval $r \in [0, K]$; their masses are redshifted to the infrared endpoint and set the observable scale $M_{\rm lock}$ (the radial Bessel tower of §9). The angular ($\ell \neq 0$) Fourier modes instead carry the additional contribution ℓ^2/R_θ^2, which in canonical GG-6 units ($R_\theta = 1$) is of order one unit of holonomy energy per mode — i.e. of order $E_{\rm GUT}$, and *not* redshifted to the TeV scale:

$$m_{n,0} = \frac{\xi_n}{\xi_1}\, M_{\rm lock} \quad \text{(radial tower, observable)}, \qquad m_{0,\ell} \sim |\ell|\,\frac{\hbar}{R_\theta} \sim E_{\rm GUT} \quad (\ell \neq 0,\ \text{unobservable}). \tag{12}$$

Second, therefore, the angular tower is Kaluza–Klein *in form* but unobservably heavy, while the experimentally relevant "cluster" is built not from θ-partners but from the lowest *radial* modes of the *several distinct bulk fields* (graviton, KK gluon, baryonic vector) that all redshift to the same infrared endpoint. That multi-field structure at one scale — not a compact-angle multiplet — is the phenomenological fingerprint that separates GG-6 from a single-sequence 5D RS tower. We fix this convention in the next subsection; it matches Arisaka (2026, A2-GG6, Part III).

5.3 Convention: "Kaluza–Klein" in GG-6 means *non-spatial*

One convention is fixed here and used throughout, matching Arisaka (2026, A2-GG6, Part III). We retain the term *Kaluza–Klein* (KK) for the lock-cluster resonances, because it correctly signals their experimental class — RG-scale-throat resonances — and connects to existing collider nomenclature (Randall–Sundrum gravitons, warped resonances). But the GG-6 resonances are *non-spatial* KK states: the two internal directions are the scale axis I_r and the phase circle S^1_θ, each dimensionless, and *neither* is a spatial extra dimension. The reader should mentally prefix "non-spatial" to every later occurrence of "KK." Concretely, two towers must be kept apart:

1. **Angular (S^1_θ) Fourier tower.** Integer modes $\ell \neq 0$ with masses set by the holonomy/$\hbar$ conversion, of order $E_{\rm GUT}$. Kaluza–Klein in form, but far too heavy to produce at any foreseeable collider; it plays *no* collider role.

2. **Radial (RG-scale) throat lock cluster.** The finite radial interval $r \in [0, K{=}35]$ carries a *discrete* spectrum (two-brane Randall–Sundrum / holographic-composite type) whose lowest modes sit at $M_{\rm lock} \approx 3.85\,$TeV. *This* is the observable signal — "the KK particles of GG-6" — with excited members carrying additional radial nodes.

In one sentence: the angular tower is heavy and invisible; the radial lock cluster near 3.85 TeV is what an experimentalist would call the Kaluza–Klein resonance of GG-6.

5.4 Radial branches and the multi-field lock cluster

With the angular tower parked at $E_{\rm GUT}$, the collider-relevant spectrum is organized by two labels that are *both* radial in origin: the radial node number n within a given bulk field, and the field species (graviton, KK gluon, baryonic vector).

1. **Radial tower ($\ell = 0$, increasing n).** These are the scale-throat resonances, with the familiar Bessel-root spacing $m_n = (\xi_n/\xi_1)M_{\rm lock}$: the first three sit near 3.85, 7.0, and 10.2 TeV. Only the lowest is comfortably in Run-3 reach.

2. **Multi-field lock cluster (lowest n, several species).** The lowest modes of the different bulk fields populate the single infrared scale fixed by Theorem T-KK-1: the spin-2 graviton at $M_{\rm lock}$, together with a spin-1 pair (a color-octet KK gluon and a leptophobic $\mathrm{U}(1)_B$ vector) whose exact placement is the subject of the fine-structure theorem T-KK-2 (§10.4) — a derived two-scenario dichotomy (free-field $0.639\,M_{\rm lock} = 2.46$ TeV versus composite-hardened $\approx M_{\rm lock}$) leaving no $O(1)$ flag behind. This multi-state structure of *different* species tied to one scale — not a set of θ-partners — is the genuine 6D signature.

3. **Tensor/flux mixing.** Once the full tensor structure and $\mathrm{U}(1)_B$ flux are included, the reconstructed peaks need not be pure species eigenstates, so the observed line shape can look like a doublet or triplet of partially overlapping states near 3.85–4 TeV.

This is not an academic detail. A collider search optimized for a single narrow RS resonance can lose sensitivity to this near-degenerate multi-field cluster if the several nearby states are forced into a one-pole template. One of the central practical messages of this paper is therefore methodological: *a GG-6 search must be cluster-aware from the outset.* We return to this in Sections 10 and 35.

5.5 Why GG-6 is not reducible to ordinary KK or RS

In one line, before the details: the fifth dimension is the renormalization-group axis — not a Randall–Sundrum spatial direction — and the lock cluster is an RG threshold spectrum, not a tower of spatial gravitons. The rest of this subsection unpacks, and defends, that sentence.

This subsection carries more conceptual weight than any other in Part I, because the resemblance it addresses is genuinely treacherous. A reader who knows the extra-dimensional literature will look at the GG-6 bulk — a warped interval carrying a factor e^{-2r} and a compact circle carrying integer modes — and file it instantly: *the fifth dimension is Randall–Sundrum, the sixth is Kaluza–Klein*. Every equation in this paper would then read as a re-derivation of known physics. The filing is wrong, and it is wrong in a precise, stateable way: *both readings mistake the coordinates for directions of space*. Neither of the two extra dimensions of GG-6 is spatial — the fifth is the *energy scale itself*, geometrized, and the sixth is the *gauge phase itself*, geometrized — and essentially every phenomenological difference exploited in this paper (the derived mass, the absence of a radion, the absence of angular satellites, the multi-field cluster, the integer hierarchy) traces back to that single distinction. We take the two dimensions in turn, then collect the full comparison in Table 2.

The classification principle: a dimension is what generates it. Arisaka (2026, A2-GG6) (Theorem 0 and its remarks) states the criterion cleanly. A geometric dimension that is not a dimension of space is not, by itself, new — Minkowski's time was the first, distinguished from space by its signature. The correct classification of any dimension is by its *symmetry generator*: a spatial (or temporal) direction is one whose generator is a *translation*, with a conserved momentum (or energy). The two extra directions of GG-6 are generated by no translation at all: the scale direction I_r is generated by a *dilatation* D (the renormalization-group flow generator, $[D, P_\mu] = P_\mu$), and the phase direction S^1_θ by a *compact charge* Q_θ with integer eigenvalues. The bulk isometry algebra is

$$\mathfrak{iso}(1,3) \,\oplus\, \mathbb{R}_D \,\oplus\, \mathfrak{u}(1)_Q, \qquad \text{not} \qquad \mathfrak{iso}(1,5), \tag{13}$$

and that is the exact technical content of "non-spatial." A parallel statement at the level of coordinates: both extra coordinates are *dimensionless by construction* — $r = \ln(E_{\mathrm{GUT}}/\mu)$ is the logarithm of an energy ratio, θ is an angle — whereas every spatial coordinate carries a unit of length. There is no answer to the question "what is the radial speed of light?", because there is no Lorentz signature on r; the question itself belongs to the spatial reading that GG-6 rejects.

The scale (fifth) dimension versus Randall–Sundrum. Superficially the resemblance is strong, and deliberately so: the metric $ds_6^2 = e^{-2r}\eta_{\mu\nu}dx^\mu dx^\nu + dr^2 + d\theta^2$ (Arisaka, 2026, A2-GG6, Theorem 2) has the RS form, the interval is finite, and the infrared endpoint carries an exponential hierarchy. The mathematics of boundary-value problems on this background is therefore shared — which is precisely why Part II of this paper can borrow Bessel machinery with a clean conscience. The *ontology* is not shared, and it bites in four places. (i) *What the coordinate is.* In RS, y is a spatial position; a particle can sit at y, move along y, carry momentum along y. In GG-6, $r = \ln(E_{\mathrm{GUT}}/\mu)$ *is* the renormalization scale: a particle at radial position r is a state of the effective theory at scale $\mu = E_{\mathrm{GUT}}e^{-r}$, and motion along r is not translation but *transmutation* — the geometrization of RG flow, with the additivity of the RG semigroup forcing the logarithmic coordinate. The scale profile e^{-2r} is not gravitational redshift in a fifth spatial direction; it is the definition of scale weight. (ii) *What fixes the length.* The RS throat depth kr_c is a free continuous dial, tuned to ≈ 35 because the hierarchy is observed; the radion — the fluctuation of that length — is a physical modulus that must be stabilized by an added mechanism. In GG-6 the interval length is the *integer* $K = 35$, forced by the OGN-5 closure $K = LJ = M^2 + N^2 + L^2$: the length is a counting number, not a size. Consequently *there is no radion* — not a stabilized radion, but no radion field at all, because a resolution has no fluctuating length (Arisaka, 2026, A4-Euclid) — and the hierarchy e^{-35} is an output. This paper's *hh* phenomenology (§10.5) and the no-radion Higgs-precision null of Part IX are direct collider consequences. (iii) *What the resonances are.* In RS the KK gravitons are momentum modes of a spatial fifth dimension. In GG-6 the throat's discrete spectrum survives — a finite interval supports normalizable resonances regardless of the interval's ontology — but the states are properly read as the *composite resonances of the boundary theory at its infrared scale*: the statement usually offered as a holographic analogy is, here, the literal construction. Same eigenvalue problem, same Bessel zeros, different physical identity — and the identity matters, because it is what makes the T-KK-2 dichotomy of §10.4 (free-field eigenvalues versus composite reorganization) a genuine physical question rather than a formality. (iv) *What sits at the bottom.* Because *every* bulk field is redshifted to the same infrared endpoint $r = K$, the lowest radial modes of the graviton, the color octet, and the $\mathrm{U}(1)_B$ vector all land at one derived scale — the multi-field lock cluster — rather than the single graviton pole of minimal RS. The cluster structure is not decoration; it is the 6D fingerprint.

The phase (sixth) dimension versus Kaluza–Klein. The sixth direction looks like Kaluza's circle, and the debt is real — but GG-6 inverts the construction. Kaluza and Klein geometrized the electromagnetic potential as the metric's components along a *spatial* circle of radius R; the radius is a free modulus, and the observable signature is a tower of charged momentum modes at $m_n = n/R$, placed wherever the modulus is tuned. In GG-6, S^1_θ is the $\mathrm{U}(1)$ *fiber itself*, promoted to a literal coordinate: a genuinely principal-bundle direction whose holonomy *is* the gauge phase of charged matter (the Aharonov–Bohm and Berry phases are its holonomies in the literal sense), whose integer characters are the charge quantization (the content of $M = 1$ in the OGN-5 ledger), and whose Wilson line θ_H performs the electroweak symmetry breaking in Arisaka (2026, A9-SM), via the Hosotani mechanism (Hosotani, 1983) — displacement by holonomy rather than by symmetry, in GG-W2-Holonomy's sense (Arisaka, 2026, W2-Holonomy). Because the direction is topological — an angle, with the canonical convention $R_\theta = 1$ and *no radius modulus* — its Fourier tower is not a tunable prediction: the angular modes sit at $\sim E_{\mathrm{GUT}}$ and are unobservable. This yields one of the sharpest negative predictions in the paper: *the lock cluster has no angular satellites.* A conventional 6D spatial compactification would decorate the 3.85-TeV cluster with $\ell = \pm 1$ partners a few percent away in mass; GG-6 forbids them. Finding them would falsify the non-spatial reading as decisively as finding nothing at all would falsify the mass.

The locking: the stage is not metric-alone. The third structural difference has no KK or RS analogue whatsoever. The GG-6 bulk is not defined by a metric on which gauge fields are later placed; the metric, the tensor sector, and the gauge fluxes arrive locked in a single primitive object through the BI-6 identity,

$$dH_3 = \mathrm{tr}\,\mathcal{R}^2 - \mathrm{tr}\,F_5^2 - \frac{2}{5}F_B^2 - \frac{2}{5}F_{Qi}^2, \tag{14}$$

whose trace-ratio 2/5 forces the equal treatment of $\mathrm{U}(1)_B$ and $\mathrm{U}(1)_{Qi}$. The KK tower of GG-6 is therefore not a harmonic decomposition on a passive background; it is the spectral trace of a geometry that is already anomaly-locked and gauge-structured. This is what admits — indeed requires — the phenomenological texture this paper computes: a mild leptophobic admixture, a top-rich overlap structure, and a multi-field cluster (graviton, KK gluon, $\mathrm{U}(1)_B$ vector) at one derived scale.

Has anything like this existed before? Candor about precedents strengthens the claim, so we state the genealogy exactly as Arisaka (2026, A2-GG6) traces it. The individual *ingredients* have distinguished ancestors: Minkowski made time the first non-spatial geometric dimension (Minkowski, 1908); Weyl geometrized *scale* as a connection in 1918 (Weyl, 1918) and *phase* as the U(1) invariance in 1929 (Weyl, 1929); the fiber-bundle formulation made the gauge phase an internal circle; Connes' noncommutative geometry built a genuinely non-spatial internal space (Connes & Lott, 1991); and holographic renormalization realized the scale direction as a radial coordinate — as an *analogy* (Maldacena, 1997/1998). But no prior phenomenological framework has ever fielded either object realized here: a dimension generated by a dilatation, of forced integer length, entering a collider prediction at a specific mass; or a gauge fiber promoted to a literal sixth coordinate whose tower is thereby pushed to the GUT scale. Still less has any framework fielded *both at once*, as the entire arena beyond spacetime, derived rather than postulated (Theorem 0 of Arisaka, 2026, A2-GG6), with no spatial extra dimension anywhere — and hence no modulus, no radion, no landscape, and no supersymmetry requirement, each absence a consequence rather than an assumption — absences of the kind GG-W1-Symmetry (Arisaka, 2026, W1-Symmetry) traces from Noether to cohomology (no translation generators exist along r or θ for supercharges to square-root). Table 2 collects the comparison; the phenomenological stakes of each row are exactly the derived quantities this paper tests.

Table 2: Comprehensive comparison: the spatial extra dimensions of Kaluza–Klein and Randall–Sundrum versus the two *non-spatial* dimensions of GG-6. Every entry in the two right-hand columns is forced by the constitutional derivation of Arisaka (2026, A2-GG6) (Theorem 0, Theorem 2, and the OGN-5 ledger); no entry is a modeling choice.

Property	**Kaluza–Klein circle**	**Randall–Sundrum interval**	**GG-6 scale dimension I_r (5th)**	**GG-6 phase dimension S^1_θ (6th)**
Nature	spatial circle	spatial warped interval	the energy scale, geometrized: $r = \ln(E_{\mathrm{GUT}}/\mu)$	the gauge phase, geometrized: the U(1) fiber as a coordinate
Coordinate carries a unit?	yes (length)	yes (length)	no — a logarithm of an energy ratio	no — an angle
Symmetry generator	translation (momentum p_5)	translation (broken by the warp)	dilatation D (RG flow); $[D, P_\mu] = P_\mu$	compact charge $Q_\theta \in \mathbb{Z}$

continued on next page

Table 2 continued

Property	**Kaluza–Klein circle**	**Randall–Sundrum interval**	**GG-6 scale dimension I_r (5th)**	**GG-6 phase dimension S^1_θ (6th)**
"Motion along it" means	spatial displacement	spatial displacement	*transmutation*: change of renormalization scale	phase rotation (charge action)
Speed of light along it	defined	defined	*undefined* — no Lorentz signature on r	undefined
Size / length	radius R: free continuous modulus	kr_c: free dial, tuned $\approx$35 for the hierarchy	integer $K = 35$, forced by $K = LJ = M^2 + N^2 + L^2$	$R_\theta = 1$ canonical; purely topological
Stabilization	required (radius modulus)	required (radion; Goldberger & Wise, 1999)	*nothing to stabilize*: a count has no fluctuating length; no radion exists	nothing to stabilize
Hierarchy	none	output of the *tuned* exponent	output of the *derived* integer: e^{-35}	—
The tower	charged momentum modes at n/R: anywhere the modulus puts them	graviton momentum modes at the TeV scale by tuning	radial lock cluster at $\frac{1}{2}M_{\text{Pl}}e^{-35}$ = 3.849 TeV, derived; *multi-field* (spin 2 + two vectors at one scale)	integer characters at $\sim E_{\text{GUT}}$: *unobservable*; no angular satellites at the TeV scale (a falsifiable absence)
What the tower *is*	momentum modes on a spatial circle	momentum modes in warped space	composite resonances of the boundary theory at its infrared scale (holography made literal)	quantized charges (bundle characters)
Charge quantization	from circle momentum (unexplained radius)	not addressed	—	integer holonomy: the *origin* of quantization ($M = 1$)

continued on next page

Table 2 continued

Property	**Kaluza–Klein circle**	**Randall–Sundrum interval**	**GG-6 scale dimension I_r (5th)**	**GG-6 phase dimension S^1_θ (6th)**
Electroweak breaking	none	added by hand	sets the lock scale that $v = M_{\text{lock}}(2/5)^3$ descends from	Wilson line θ_H performs it (Hosotani mechanism, Hosotani, 1983; Arisaka, 2026, A9-SM)
Moduli / landscape	radius modulus	radion	none	none
Supersymmetry	optional	optional	impossible to extend: no translation generators along r, θ for supercharges to square-root; absence is a theorem, not a choice	
Free parameters	R	k, r_c, brane tensions	none (integers only)	none
Nearest precedent	Kaluza 1921; Klein 1926 (Kaluza, 1921; Klein, 1926)	Randall–Sundrum 1999 (Randall & Sundrum, 1999)	Weyl's 1918 scale connection; the holographic radial coordinate (as analogy) — never before a forced dimension	Weyl's 1929 phase; fiber-bundle internal circle — never before a literal coordinate

6 Three Reading Paths

Part I is nearly complete, and before the derivation begins it is worth pausing to say plainly how this paper is meant to be used. It is long because it is meant to be *complete*: the theorem chain, its experimental confrontation, a full simulation of the discovery, and the graduate courses that teach the machinery all live between these covers, because a falsifier is only as good as a stranger's ability to check it unaided. But completeness is a property of the document, not a demand on the reader. No one is meant to take the fourteen Parts linearly, and the paper has been built so that three quite different readers can each extract what they came for — and know, at every step, where the numbers they are using were derived.

The experimentalist's path. A reader who wants to know *what to look for, where the limits stand, and when the question gets decided* should go directly to the confrontation and the calendar: the current ATLAS/CMS status and the operational comparison of computed $\sigma \times \text{BR}$ against published limits (§30–§31), then the anchored HL-LHC simulation with its systematics wall and pre-registered search (Part VIII), and finally the benchmark and reinterpretation card (Appendix M), which is deliberately self-contained: masses, widths, couplings, line shapes, and the implementation walkthrough, in the format a convener would hand to a graduate student on Monday morning. On this path the theory Parts function as provenance — the place to verify, when needed, that a given number is a theorem output rather than a tuned input. The essential facts to carry along are three: the target is a *cluster* at 3.849 TeV, not a lone pole; the color-octet member at its computed $t\bar{t}$ rate is the un-suppressible falsifier (F-KK-2B); and the decision window closes at $\lesssim 2\ \mathrm{ab}^{-1}$, within the HL-LHC's first years.

The phenomenologist's path. A reader who wants to audit the *derivation chain* — what is actually forced, and at what grade — should follow the spine: Part II for the lock mass and the cluster's fine structure (the closed Bessel form, the radial tower, the degeneracy pattern), Part VI for the signal anatomy (the line shape with its theorem-fixed interference sign, the angular and polarization suite, the discrimination matrix), and then the Discussions (Part X) for the honest accounting: which statements are theorem-grade, which are estimates with stated brackets, and how the package compares — row by row — with the RS, axigluon, and leptophobic-Z' alternatives an informed skeptic will reach for. The single discipline to test on this path is the paper's central one: there is no continuous dial anywhere, so every disagreement with data is structural, and the two readings of the coupling normalization (Scenario A/B) form a closed dichotomy, not a tuning direction.

The theorist's path. A reader whose first question is *what kind of extra dimensions these are* should begin where the framework is strangest: §5.5 and Table 2, where the two non-spatial dimensions — the energy scale as a coordinate of forced integer length, the gauge phase as a literal sixth direction — are contrasted entry by entry with the spatial constructions of Kaluza–Klein and Randall–Sundrum. Then Part II, where that geometry becomes a spectrum; then the five conceptual primers (Appendices E–I), which answer, at full length, the questions a seminar audience actually asks: how the gluon and g^* are one field, where the baryonic Z'_B comes from, what the "graviton" of GG-6 is, why g^* dominates the cluster's rates and is so broad, and why one integer controls both the rate hierarchy and the $t\bar{t}$ dominance. The stakes on this path are conceptual: no modulus, no radion, no landscape, and no supersymmetry requirement — each an absence that is derived, not assumed.

The graduate student's path. The Preface set out a fourth path, and it remains the best one: take the paper linearly, but lean on the three-step courses of Appendices J–L (the KK spectrum from the 6D action; cross sections from parton distributions; discovery statistics from the Poisson likelihood up) whenever the main text assumes machinery you have not yet met. The appendices have no page budget precisely so that this path requires no other book.

7 The cohort inheritance and the roadmap of the fourteen Parts

Before the derivation begins (Part II), it is useful to lay out two reader-orientation maps in one closing Section: a panorama of the twenty-two-paper cohort in which this paper stands (§7.1, Figure 2), and a panorama of the fourteen Parts of the present paper (§7.2, Figure 3). 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 fourteen-Part arc.

7.1 The cohort inheritance

The program in one view. Figure 2 places the present paper among the twenty-two papers of the GG-Theory cohort. The reader meeting the program for the first time can read its architecture off the figure top to bottom — the constitutional root, the bulk trio, the open-system boundary law and its pillars, the three textbook readouts, the capstone, the two phenomenology branches, and the Foundations Branch on the left flank — before continuing with the present paper's own thread. The papers of record: the three textbook readouts are Arisaka (2026, A8-QM), Arisaka (2026, A9-SM), and Arisaka (2026, A10-Cosmos); the capstone is Arisaka (2026, A11-Origin); the Predictions Branch carries, beside the present paper, Arisaka (2026, P2-N3) and Arisaka (2026, P3-Electron); the Cosmology & Astrophysics Branch carries Arisaka (2026, C1-Kaxion), Arisaka (2026, C2-Hubble), and Arisaka (2026, C3-BH); and the Foundations Branch carries Arisaka (2026, P4-Real) and Arisaka (2026, P5-CPT).

Reading this paper against that map requires only a few conventions, stated once here. Numbered results carry the cohort global identifiers with the paper suffix -KK: thus *T-KK-1* is the first Theorem and *F-KK-1* the first pre-declared falsifier (the cohort uses the prefix set T-/L-/D-/PS-/F- for theorems, lemmas, definitions, postulate-statements, and falsifiers). Dimensionful inputs are drawn only from GU-5, the Geometric Universal Five Constants $\{c,\, G_N,\, \hbar;\, k_B,\, Q_B\}$; all dimensionless data come from the OGN-5 integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$, so the theory carries no continuous dial beyond the overall choice of units. The two extra coordinates are the *non-spatial* scale axis I_r and phase circle S^1_θ; the bulk metric is the canonical $\mathrm{d}s_6^2 = e^{-2r}\eta_{\mu\nu}\mathrm{d}x^\mu \mathrm{d}x^\nu + \mathrm{d}r^2 + \mathrm{d}\theta^2$ of GG-A2-GG6 (Theorem 2); the radial direction is RG-*scale* (not RG-time); and the open-system language of detection follows GG-A5-GDOS (GDOS, Arisaka, 2026, A5-GDOS) and A6-GRIP (Arisaka, 2026, A6-GRIP). This paper builds on the cohort papers of record — GG-A1-Constitution (Arisaka, 2026, A1-Constitution), GG-A2-GG6 (Arisaka, 2026, A2-GG6), GG-A3-BI6 (Arisaka, 2026, A3-BI6), GG-A9-SM (Arisaka, 2026, A9-SM), GG-A5-GDOS, and A6-GRIP — and in particular on Theorems 0 and 2 of GG-A2-GG6 (the forced bulk and its canonical metric), the OGN-5 cascade of GG-A3-BI6, and the lock-scale/scale-chain of GG-A9-SM. For the canonical Dictionary of program-defined acronyms, see GG-A2-GG6 Appendix B; this paper defines only the collider-specific vocabulary in situ.

A word on the Foundations Branch, the column on the map's left flank, whose subject — the boundary's complex structure — lies farthest from a collider hall. GG-P4-Real (Arisaka, 2026, P4-Real) is the origin-of-i member: from the two real structures of the same forced arena this paper's spectrum lives on — the compact phase circle and the causal clock — it derives the single imaginary unit of boundary physics, so the interference phases that this paper's amplitudes and line shapes take for granted are outputs of the geometry rather than inputs. GG-P5-CPT (Arisaka, 2026, P5-CPT) is the mirrors-and-matter member: C, P, and T as orientation reversals of the three real bulk factors, with the matter–antimatter asymmetry at the end of the chain — and its dated kill list names this paper's HL-LHC lock-cluster verdict as a companion channel of the same shared geometry. Neither carries a collider claim, and no result below imports from either; they are named here so the panorama is complete.

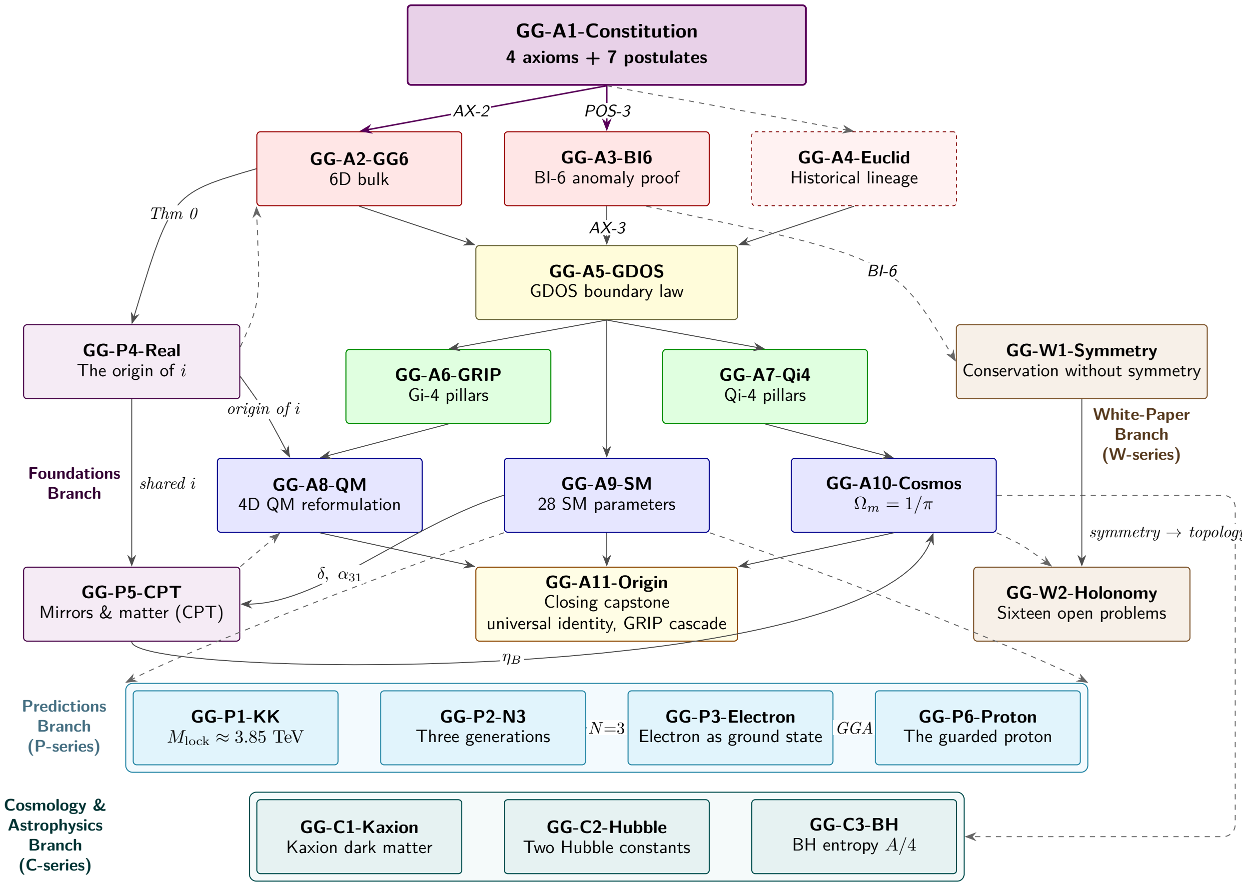

Figure 2: The twenty-two-paper cohort map of GG-Theory. The constitutional root GG-A1-Constitution (violet, top) supplies the four axioms and seven postulates from which the rest of the cohort descends: the GG6 Trio (GG-A2-GG6, GG-A3-BI6, GG-A4-Euclid, red) realizes the bulk geometry; GG-A5-GDOS (yellow) 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=3 into GG-P3-Electron, and that paper's ground-attractor method, GGA, on into GG-P6-Proton (Arisaka, 2026, 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.

7.2 Roadmap of the fourteen Parts

The rest of the paper is organized into thirteen further Parts, and the walk through them is plotted in Figure 3. The present Part — the foundations — is nearly done: it has posed the fifty-year question, put geometry–gauge locking into plain English, stated the central claim in its open-system reading, traced the constitutional route from the axioms to a collider observable, distinguished the two non-spatial dimensions from their spatial look-alikes, laid out the three reading paths, and drawn the two orientation maps of this closing Section.

Part II is the paper's center of mass, and it is highlighted in the roadmap because everything later is downstream of it. From the canonical metric and the integer ledger it derives the lock scale in closed form — Theorem T-KK-1, $M_{\text{lock}} = \frac{1}{2} M_{\text{Pl}} e^{-35} \simeq 3.849$ TeV — and then resolves what lives at that mass: the three-member cluster (the spin-2 mode, the color-octet g^*, and the leptophobic Z'_B), its derived widths, couplings, and degeneracy pattern, the radial family at 7.05 and 10.22 TeV, and the observable properties of the KK sector.

Three Parts then carry the derived signal to the collider floor. Part III computes production and discovery: cross sections, event yields, and the channel hierarchy that makes the top pair the decisive final state. Part IV draws the line that keeps a reader from mis-filing the signal — GG-6 versus Randall–Sundrum, entry by entry — and re-reads collider detection itself as an open-system (GDOS) problem. Part V then states the coupling normalization as a closed dichotomy: Read-out I, the theorem-only target, and Read-out II, the empirically-anchored rate estimate.

Two Parts confront the prediction with the machines that exist. Part VI derives the anatomy of the $t\bar{t}$ signal at full depth — the interference-aware line shape with its theorem-fixed sign, the angular and polarization suite, and the discrimination matrix behind the pre-registered fingerprint. Part VII is the experimental status and reach: the current ATLAS/CMS limits, the operational confrontation of computed $\sigma \times \text{BR}$ against the published exclusions, the kill criterion at $\lesssim 2\ \text{ab}^{-1}$, and the trigger-and-analysis roadmap.

Part VIII, the second highlighted Part, simulates the discovery at the fidelity of a real analysis — anchored to the published ATLAS Run-2 search, calibrated to its limits, with the statistics taught from the Poisson likelihood up and the four snapshots of what the detectors would record in 2026, 2031, 2033, and 2041. Its headline is the reason this paper is urgent rather than merely testable: the collisions that decide the question may already be on tape. Part IX looks beyond the HL-LHC to the 2035–2070 program, facility by facility and scenario by scenario. Part X, the Discussions, weighs the case — significance, uniqueness, the classic collider puzzles, the comparison with the other roads beyond the Standard Model, the predictions for falsification, and the open questions stated as open — and Part XI, the Conclusions, closes the body with what the paper has established, its objective limitations, and what either verdict would mean. The back matter is the workshop: Part XII collects the Technical Appendices — the reference dictionaries, the five conceptual primers, and the three graduate courses in the collider physics of the lock cluster; Part XIII answers the Common Objections a collider experimentalist, phenomenologist, statistician, or string theorist would press hardest; and Part XIV closes with the annotated References, each entry carrying its relevance note and its first citation point. Read in order, the fourteen Parts compose a single self-contained argument; read by need, each Part names its imports and can stand a visit on its own.

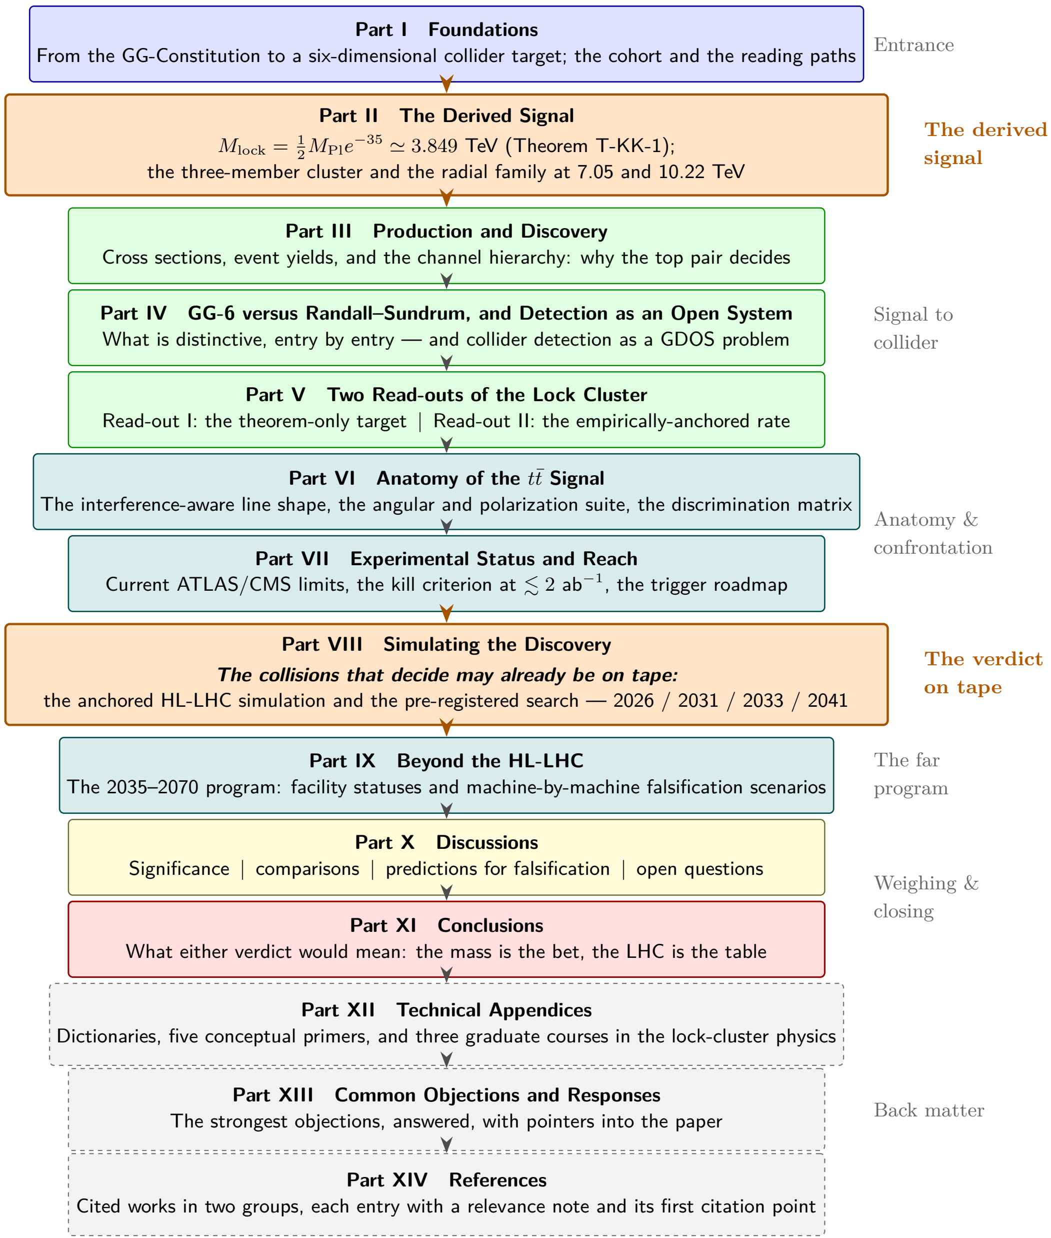

Figure 3: Roadmap of the fourteen Parts of GG-P1-KK. The argument develops in eight phases. *Entrance* (Part I): the fifty-year question, geometry–gauge locking in plain English, the constitutional route, the non-spatial geometry, and the reader's two maps. *The derived signal* (Part II, highlighted): the lock-cluster mass in closed form (Theorem T-KK-1) and the observable spectrum built on it. *Signal to collider* (Parts III–V): production and discovery, the GG-6/Randall–Sundrum comparison with the open-system reading of detection, and the two read-outs of the coupling normalization. *Anatomy and confrontation* (Parts VI–VII): the derived $t\bar{t}$ fingerprint, and the computed rates set against the published ATLAS/CMS limits with the kill criterion. *The verdict on tape* (Part VIII, highlighted): the anchored HL-LHC simulation and the pre-registered search — 2026 / 2031 / 2033 / 2041. *The far program* (Part IX): the machines of 2035–2070. *Weighing and closing* (Parts X–XI): the Discussions and the Conclusions. *Back matter* (Parts XII–XIV): the technical appendices, the common objections and responses, and the annotated references. Parts II and VIII are highlighted because together they carry the paper's central claim — the derived lock cluster at $M_{\mathrm{lock}} \simeq 3.849$ TeV: Part II is where the mass and its cluster are computed as theorems, and Part VIII is where that computation becomes a dated, pre-registered collider verdict. The mass is the bet; Part VIII sets the table.

Part I, in plain terms, established the arena and the rules. A theory whose only freedom is five counting numbers forces a six-dimensional stage — ordinary spacetime plus two directions that are not space at all, one being the energy scale itself and the other the gauge phase — and that stage, unlike every earlier extra-dimensional proposal, has no adjustable size, no radion, no landscape, and nothing to tune. Nothing about a particle has been claimed yet; what has been established is stricter: the constitutional route by which collider observables descend from the axioms, the open-system dictionary by which bulk excitations become ATLAS/CMS records, and above all the discipline that every later number must be an *output*.

The next Part cashes that discipline immediately, on the paper's central object: the mass. From the canonical metric and the integer ledger it derives the lock scale in closed form — Theorem T-KK-1, $M_{\rm lock} = \frac{1}{2} M_{\rm Pl} e^{-35} \simeq 3.849$ TeV — and then resolves what lives there: not a lone resonance but a three-member cluster with derived fine structure. Everything in this paper is downstream of that one number; Part II is where it comes from.

8 Part I in plain English: why fifty years of predictions could not lose, and why this one can

Every particle found since the mid-1970s was on a list before it was found. The W and the Z, the top, the Higgs — each was predicted, each was looked for, each turned up. That run of successes ended, and the fifty years since have produced no new layer at all.

It was not for want of trying, and this Part says plainly why the trying failed. Supersymmetry promised partners just above each frontier and, in its minimal form, carried more than a hundred free parameters — so every null result was answered by sliding the spectrum up. Extra-dimensional models supplied real structural insight and made their mass scales inputs: a radius to scan, a warp factor to choose, so no exclusion could ever be fatal. String theory built a candidate for everything and committed to nothing measurable. Three programs, three sets of dials, and not one pre-registered number whose failure would end anything.

The diagnosis is not that the next layer is out of reach. It is that a theory with continuous free parameters can retreat from any measurement, and a theory that cannot lose cannot win.

What this paper offers instead is a mass with no dial behind it. The scale comes out of an integer in a ledger fixed elsewhere in this cohort, for reasons that had nothing to do with colliders, and the same integer over-constrains the electroweak scale as well. There is nothing to adjust if the measurement disagrees. That is the entire point, and it is the reason the rest of the paper can be as specific as it is.

Part II
The Derived Signal

The lock-cluster mass (Theorem T-KK-1) and the observable spectrum

9 The Lock Scale and the First Resonance Cluster

9.1 The lock-cluster mass: Theorem T-KK-1

This subsection is the scientific centerpiece of the paper. Everything downstream — production, channels, line shape, the experimental status, and the falsifier — hangs on a single *derived* number. We state it as a theorem, prove it from the constitutional inputs, and count the free parameters (there are none).

Theorem T-KK-1 (the GG-6 lock-cluster scale). *The first observable Kaluza–Klein resonance of the six-dimensional bulk GG-6 is the radial "lock cluster," whose lowest mode has the parameter-free mass*

$$\boxed{M_{\text{lock}} \;=\; \tfrac{1}{2}\,M_{\text{Pl}}\,e^{-K} \;=\; \frac{E_{\text{Pl}}}{2\,e^{35}} \;\simeq\; 3.849\ \text{TeV}\,, \qquad K = 35\,,} \tag{15}$$

in which K is the OGN-5 lock integer (no continuous freedom) and M_{Pl} (equivalently $E_{\text{Pl}} = \sqrt{\hbar c^5/G_N}$) is a member of the GU-5 universal-constant set, not a tunable scale. The radial excitations form a discrete tower

$$m_n \;=\; \frac{x_n}{x_1}\,M_{\text{lock}}\,, \qquad n = 1, 2, 3, \dots\,, \tag{16}$$

with x_n the n-th positive zero of the Bessel function J_1, so that the first three lock-cluster modes lie at

$$m_1 \simeq 3.85\ \text{TeV}\,, \qquad m_2 \simeq 7.0\ \text{TeV}\,, \qquad m_3 \simeq 10.2\ \text{TeV}\,. \tag{17}$$

The angular (S^1_θ) Fourier tower lies at $\sim E_{\text{GUT}}$ and is unobservable at any foreseeable collider.

Precision of the anchor. Evaluated with the CODATA Planck energy, $M_{\text{lock}} = E_{\text{Pl}}/(2e^{35}) =$ 3.8489 TeV, with an uncertainty inherited entirely from G_N ($\sim 10^{-5}$ relative, i.e. ± 0.04 GeV) — making the lock-cluster anchor, by a comfortable margin, the sharpest absolute mass prediction in particle physics. The tower ratios are exact Bessel-zero ratios, $m_2/m_1 = x_2/x_1 = 1.8309$ and $m_3/m_1 = 2.6551$, so the radial recurrences at 7.047 and 10.219 TeV carry *no additional* theory error beyond the anchor itself.

The three lowest states, physically. The integer n in (16) deserves a physical reading, not just a formula. The scale throat is a cavity: a finite interval of length $K = 35$ in the renormalization-group direction, bounded by the ultraviolet wall at $r = 0$ and the infrared endpoint at $r = K$. Like any cavity, it supports standing waves, and n counts their *radial nodes*: the $n = 1$ mode is the fundamental — nodeless, maximally concentrated at the infrared endpoint, the "lowest note" of the throat — while $n = 2$ and $n = 3$ are its first and second overtones, each adding one node along the scale direction. The Bessel zeros x_n are precisely the standing-wave quantization conditions on this background, and their ratios are pure geometry: $m_1 : m_2 : m_3 = 1 : 1.831 : 2.655$, i.e. 3.849, 7.047, and 10.219 TeV. In the renormalization-group reading of the throat (Table 2, the "what the tower is" row), these are not momentum modes of a spatial dimension but successively excited *composite resonances* of the boundary theory at its infrared scale — the analogue of radial excitations of a hadron, with the throat playing the role of the confining dynamics. The intuition an experimentalist should carry: one derived scale, then a fixed overtone series above it — an organ pipe, not a picket fence at arbitrary spacing. Only the fundamental is within LHC reach; the $n = 2$ overtone at 7.047 TeV is an FCC-hh measurement (Part IX), and finding it *at the Bessel ratio* would confirm the cavity itself, not merely another particle.

The runtime identity of the target. In the runtime vocabulary of the series (Arisaka, 2026, A6-GRIP; the two regimes GRIP-In and GRIP-Out), the statement of T-KK-1 has a one-sentence translation that locates this paper in the same grammar as its sister predictions: the lock scale is the floor of the radial-scale stratum — its Geometric Attractor — and *the KK cluster is that attractor's first excited level*. Producing a lock-cluster resonance at the LHC is port-driven excitation within a fixed landscape (the canonical collider clarification of Arisaka, 2026, A6-GRIP: pumping a state uphill through the boundary ports reorganizes nothing); its decay is GRIP-In relaxation back to the floor. Three instruments of the predictions branch thus probe one functional: the haloscope listens to a basin's residual ring (C1-Kaxion), gravitational-wave ringdown measures a basin's curvature (C3-BH), and the LHC pumps a basin's first excited state — the present paper.

Proof. The argument has six steps. Steps 1–3 fix the scale (15) without a free parameter; Step 4 fixes the tower (16); Steps 5–6 record the over-constraint and the parameter count.

Step 1 — the geometry. The bulk is the canonical product geometry of Theorem 2 of Arisaka (2026, A2-GG6),

$$\mathrm{d}s_6^2 \;=\; e^{-2r}\,\eta_{\mu\nu}\,\mathrm{d}x^\mu \mathrm{d}x^\nu \;+\; \mathrm{d}r^2 \;+\; \mathrm{d}\theta^2\,, \qquad r \in [0, K]\,, \quad \theta \in [0, 2\pi)\,, \tag{18}$$

with *no* free metric data (k, R_θ, r_c are all fixed to unity in the canonical GG-6 units; see Arisaka (2026, A2-GG6), Theorem 2). The decisive structural fact for collider physics is that the radial coordinate r is the *scale* (renormalization) direction, $r = \ln(E_{\mathrm{GUT}}/\mu)$, and it ranges over a *finite* interval $[0, K]$ — bounded by the ultraviolet endpoint $r = 0$ and the infrared endpoint $r = K$. A finite warped interval, exactly as in the two-brane Randall–Sundrum construction, carries a *discrete* spectrum of normalizable resonances (not a continuum). These resonances are the boundary image of a Kaluza–Klein tower; they are *non-spatial* (the throat is an energy-scale direction, not a spatial dimension), but their collider phenomenology is that of warped Kaluza–Klein/composite states.

Step 2 — the scale depth is forced. The integer K is not chosen. It is the terminal entry of the OGN-5 Diophantine cascade proved in Arisaka (2026, A3-BI6) from the single-tensor Green–Schwarz-modified Bianchi identity BI-6:

$$(M, N, L; J, K) = (1, 3, 5; 7, 35)\,, \qquad J = 2N + 1 = 7\,, \qquad K = LJ = M^2 + N^2 + L^2 = 35\,. \tag{19}$$

Thus the depth of the RG-scale throat is fixed by integer data alone. This is the single most important sentence in the paper: the infrared end of the geometry is not a dial.

Step 3 — the scale profile sets the lock scale. Along the throat the scale profile e^{-r} redshifts every bulk scale; at the infrared endpoint $r = K$ the suppression is

$$e^{-K} = e^{-35} \simeq 6.305 \times 10^{-16}\,. \tag{20}$$

Because the two extra directions of GG-6 are non-spatial (dimensionless scale and phase), the 6D Einstein–Hilbert reduction relates the bulk scale to the four-dimensional Planck scale by a pure number, $M_6 = \sqrt{2/\pi}\, M_{\rm Pl}$ (Arisaka (2026, A4-Euclid), Theorem T-Euclid-6; equivalently $M_{\rm Pl}^2 \sim \pi M_6^2$ in Arisaka (2026, A2-GG6) §5, after the radial scale-profile integration $\int_0^K e^{-2r}\mathrm{d}r = \frac{1}{2}(1-e^{-2K}) \to \frac{1}{2}$). This is *not* the dimension-carrying volume relation $M_{\rm Pl}^2 \sim M_6^4 V_{\rm int}$ of an ordinary *spatial* Kaluza–Klein reduction, precisely because I_r and S^1_θ carry no length. Carrying the matching through to the infrared endpoint fixes the lock scale,

$$M_{\rm lock} \;=\; \tfrac{1}{2}\, M_{\rm Pl}\, e^{-K}\,, \tag{21}$$

in which the factor $\frac{1}{2}$ is the scale-profile-integration weight $\frac{1}{2}(1 - e^{-2K}) \to \frac{1}{2}$ above (the "minimal half-wave" of Arisaka (2026, A9-SM) §4.6), *not* a fitted coefficient. The complete scale chain $E_{\rm Pl} \to E_{\rm GUT} \to M_{\rm lock}$ is owned by Arisaka (2026, A2-GG6) §5 and Theorem 6, and the boxed value is reproduced in Arisaka (2026, A9-SM) (Table of derived quantities: $M_{\rm lock} = 3.848959$ TeV). Numerically, with $E_{\rm Pl} = M_{\rm Pl}c^2 = 1.2209 \times 10^{19}$ GeV,

$$M_{\rm lock} = \tfrac{1}{2}(1.2209 \times 10^{19}\ \text{GeV})(6.305 \times 10^{-16}) = 3.849 \times 10^3\ \text{GeV} = 3.849\ \text{TeV}\,. \tag{22}$$

Step 4 — the radial tower. Spin-2 fluctuations $h_{\mu\nu}(x^\mu, r)$ of (18) with vanishing angular momentum ($\ell = 0$ on S^1_θ) obey a radial Sturm–Liouville problem on the finite interval $[0, K]$,

$$-\partial_r\Big(e^{-4r}\,\partial_r\psi_n\Big) \;=\; m_n^2\, e^{-2r}\,\psi_n\,, \tag{23}$$

with Neumann-type endpoint conditions at $r = 0$ and $r = K$. For the AdS-like scale profile the eigenvalues are set by the zeros of J_1, $m_n \propto x_n$. One subtlety must be stated precisely. A *naive* pure-Bessel estimate would give $m_1^{\rm Bessel} \approx x_1\, e^{-K} M_6 \approx 3.83 \times 6.3 \times 10^{-16} \times 1.22 \times 10^{19}$ GeV ≈ 29 TeV, using the radial prefactor $x_1 \approx 3.83$ and reference $M_6 \sim M_{\rm Pl}$. This cruder estimate is *superseded* by the scale-chain value $M_{\rm lock} = \frac{1}{2}M_{\rm Pl}e^{-K} \simeq 3.85$ TeV, which uses the scale-profile-reduction prefactor $\frac{1}{2}$ (Arisaka (2026, A2-GG6) §5+T6; Arisaka (2026, A4-Euclid)).

The ratio of the two prefactors is $2x_1 \approx 7.66$, numerically close to the OGN-5 integer $J = 7$ but *not* an independent integer correction (Arisaka (2026, A2-GG6), KK-spectrum appendix). We therefore normalize the lowest physical mode to the derived scale, $m_1 = M_{\rm lock}$, which gives (16); with $x_1 = 3.8317$, $x_2 = 7.0156$, $x_3 = 10.1735$ the representative spacings are (17).

Only $m_1 = M_{\rm lock}$ is theorem-level; the higher ratios $x_2/x_1 = 1.83$, $x_3/x_1 = 2.66$ are a representative AdS-throat result, and the detailed cluster substructure (tensor sector, $\mathrm{U}(1)_B$ flux) is *not* part of the theorem-level claim — it is flagged as model-level in the provenance ledger below. The companion *angular* tower has masses $m_\ell \sim \ell\, E_{\rm GUT}$ (Fourier modes of the gauge-phase circle, Arisaka (2026, A2-GG6)); it is heavier than the lock cluster by twelve orders of magnitude and is irrelevant to colliders. *This is the content of Option A: the observable cluster is radial.*

Step 5 — the over-constraint. The same lock scale fixes the electroweak vacuum expectation value through the Hosotani breaking of Arisaka (2026, A9-SM),

$$v = M_{\rm lock}\left(\frac{2}{5}\right)^3 = 3.849 \text{ TeV} \times 0.064 \simeq 246.3 \text{ GeV}\,, \tag{24}$$

in agreement with the measured $v = 246.22$ GeV. The collider scale and the weak scale are therefore *not* independent: a single integer ledger fixes both. This over-constraint is the source of the unusual falsifying power of the search — a wrong lock cluster would also unseat the electroweak derivation.

Step 6 — the parameter count. Equation (15) is a pure number ($\frac{1}{2}e^{-35}$, built from the OGN-5 integer $K = 35$) times a GU-5 scale ($M_{\rm Pl} = \sqrt{\hbar c^5/G_N}$, fixed by measuring G_N, $\hbar$, c, not by fitting the collider). There is no free continuous parameter in the mass prediction. The cross-section and branching estimates of the following sections *are* model inputs, and are labeled as such; the *mass* is a theorem. ■

Table 3: The GG-6 radial lock cluster (Option A). Masses from $m_n = (x_n/x_1)M_{\rm lock}$ with $M_{\rm lock} = \frac{1}{2}M_{\rm Pl}e^{-35}$. The angular tower (not shown) sits at $\sim E_{\rm GUT} \approx 2 \times 10^{16}$ GeV. The masses are ratios of Bessel zeros times one scale, and the scale is the only thing the ledger supplies — which is why the table has no column for an uncertainty on the spacing. What can move is the overall position, and what cannot move is the pattern. A measured cluster whose members sit at the wrong ratios would falsify the reading even if the lowest member landed exactly where predicted.

Mode	x_n (J_1 zero)	m_n (TeV)	Role
$n = 1$	3.8317	3.85	Lock-cluster centroid; primary discovery target
$n = 2$	7.0156	7.0	First radial partner; HL-LHC corroboration
$n = 3$	10.1735	10.2	Second radial partner; future-collider reach

Pre-declared falsifier (F-KK-1, under POS-7/MF). The prediction is falsifiable on a definite timeline. *If the combined HL-LHC dataset (*~ 3 ab^{-1} *per experiment) shows no clustered spin-2 resonance with centroid in the window* 3.85 ± 0.10 *TeV — and, in particular, no first radial mode near*

3.85 *TeV in the dominant $t\bar{t}$ and boosted-dijet channels at the model-estimated cross-section — then the current lock-scale implementation of GG-6 is excluded.* (The ± 0.10 TeV window is experimental, not theoretical: the anchor's own uncertainty is ± 0.04 GeV, four orders of magnitude smaller, so the window budgets only the reconstructed-mass scale and resolution and the possible centroid spread of a multi-state cluster — a generous $\pm 2.6\%$ that the theory error cannot consume.) A confirmed cluster, with the $n = 2$ partner near 7 TeV, would conversely be simultaneous evidence for the bulk geometry, the OGN-5 lock integer $K = 35$, and the non-spatial Kaluza–Klein reading. The number 3.849 TeV is the bet. (The detailed channel-by-channel kill criterion is stated in §33.)

Requalification (Corollary T-KK-2c, §10.9). The graviton's gluon coupling is now derived, $\kappa_{gg} = c_K/K = 0.0136$, giving $\sigma \times \mathrm{BR}(G_1 \to t\bar{t}) \approx 0.045$ fb at 13 TeV (LO×K, where K denotes the QCD K-factors $K_{q\bar{q}} = 1.3$, $K_{gg} = 1.6$ — not the integer $K = 35$) — a factor $\sim$25 below even the 3 ab^{-1} HL-LHC projection. F-KK-1 therefore splits into two tiers: (i) at the HL-LHC the falsifier of the lock-scale implementation is carried by the cluster through the un-suppressible color octet (F-KK-2B) — a null in the cluster-aware $t\bar{t}$/dijet combination within 3.85 ± 0.10 TeV at $\lesssim 2$ ab^{-1} retires the implementation; (ii) the direct spin-2 confirmation is an FCC-hh measurement ($\sigma(G_1) \approx 205$ fb at 100 TeV). The bet stands; the table moves.

Scope note (honesty discipline). Of the content above, Eqs. (15), (19), (20), (22), and (24) are *derived* (parameter-free, traceable to OGN-5 + GU-5). The tower *ratios* (16), the Neumann boundary conditions, and the $O(1)$ normalization $\frac{1}{2}$ are *representative/normalization* statements resting on the canonical AdS-throat profile and the cohort scale-chain; the cluster substructure and the octet/graviton couplings are now themselves derived at the free-field grade (T-KK-2, Corollary T-KK-2c); the remaining model estimates are confined to the per-fermion cascade depths and the $U(1)_B$ normalization g_B. The theorem-level claim is the mass (15), and it is the one the falsifier targets.

9.2 Provenance of every quantitative prediction

This paper is meant to be acted on by an experimentalist, so we state, without exception, where each number comes from. Every *derived* quantity traces to a theorem in the just-completed GG-6-A series (A1–A11, June 2026), and to nothing else. The derivation chain is short and entirely sequential:

1. **Inputs.** The dimensionful pentad GU-5 $= \{c, G_N, \hbar\,; k_B, Q_B\}$ and the dimensionless integer ledger OGN-5 $= (M, N, L; J, K) = (1, 3, 5; 7, 35)$, with $J = 2N + 1 = 7$ and $K = LJ = M^2 + N^2 + L^2 = 35$ (Arisaka, 2026, A1-Constitution, §POS-6 worked example)(Arisaka, 2026, A3-BI6). No continuous dial.

2. **Planck scale.** $E_{\mathrm{Pl}} = \sqrt{\hbar c^5/G_N} = 1.2209 \times 10^{19}$ GeV, the GU-5 of Arisaka (2026, A1-Constitution) (T-GU5).

3. **Unification scale.** $E_{\mathrm{GUT}} = E_{\mathrm{Pl}}\, e^{-2\pi} \simeq 2.28 \times 10^{16}$ GeV — one S^1_θ holonomy traversal (Arisaka, 2026, A2-GG6, §5, Theorem 6).

4. **Scale depth.** $e^{-K} = e^{-35} \simeq 6.305 \times 10^{-16}$, the integer $K = 35$ from step 1.

5. **Lock-cluster mass (the centerpiece).** $M_{\mathrm{lock}} = \frac{1}{2}M_{\mathrm{Pl}}e^{-35} = 3.848959$ TeV (Arisaka, 2026, A2-GG6, §5, Theorem 6), with the scale-profile-reduction factor $\frac{1}{2}$ and $M_6 = \sqrt{2/\pi}\,M_{\mathrm{Pl}}$ from Arisaka (2026, A4-Euclid) (T-Euclid-6) and the half-wave reading of Arisaka (2026, A9-SM, §4.6). The naive pure-Bessel root (≈ 29 TeV) is reduced to M_{lock} by the OGN-5 integer $J = 7$ (Step 4 above).

6. **Electroweak over-constraint.** $v = M_{\mathrm{lock}}(2/L)^N = M_{\mathrm{lock}}(2/5)^3 = 246.333$ GeV, the Wilson–Hosotani vev theorem Arisaka (2026, A9-SM, Thm. vev); cf. measured $v = 246.22$ GeV.

Everything else in this paper — the partner-state *masses*, the effective coupling κ_{eff}, the cross sections, the branching fractions, the line shapes, and the event yields — is a *model estimate* built on top of this derived scale. In particular, the partner masses quoted as "$g^* \approx 3.9$" and "$Z'_B \approx 4.0$" TeV are *illustrative assignments placed near* M_{lock} *by the multi-field-cluster picture*; they are *not* outputs of any GG-6-A theorem, and the GG-6-A series deliberately defers all collider spectroscopy and rates to the present Predictions-Branch paper. Table 4 records this division explicitly.

Exclusion-edge honesty. One point deserves to be stated plainly rather than left implicit in the provenance table. If the leptophobic baryonic vector Z'_B indeed sits near M_{lock}, it does not rest comfortably above the present limits: it lands squarely on the current ATLAS/CMS high-mass dijet exclusion frontier ($\gtrsim 3.8$ TeV), and it survives today only because its couplings are leptophobic (light-lepton-suppressed), so the far more stringent dilepton resonance limits do not apply to it. We regard this as a sharpening of the test, not a loophole — the very leptophobia that keeps the state alive in present data is what makes the boosted-dijet and $t\bar{t}$ channels at the HL-LHC the decisive arena. Correspondingly, the baryonic normalization is now carried to its honest boundary in §10.10: its generator norm is theorem-grade (the trace ratio 2/5), one discrete readout rule remains open (L-KK-1), and the collider-side cross sections are quoted as explicit functions of g_B in §31 — LHC-invisible across the entire L-KK-1 menu. The (g_B, m_B) fifth-force window — which the Standard-Model consequence paper (Arisaka, 2026, A9-SM) lists among its falsifiers — remains the dedicated fifth-force paper's deliverable, now reduced from a plane scan to a selection among four candidate lines.

Table 4: Provenance of every quantitative prediction. Only the first block is theorem-level; the second block is model-level and rescales with the stated assumptions. "Source" gives the owning GG-6-A paper and theorem. The split between the blocks is the point of the table and is worth reading before any number in it. A theorem-level entry cannot be renegotiated without retiring the theorem that owns it; a model-level entry moves when its stated assumptions move, and the assumptions are named rather than buried. A reader who wants to know how much of this paper is at risk from a single wrong assumption should count the second block.

Quantity	**Value**	**Status & source**
Derived (theorem-level; parameter-free)		
OGN-5 ledger $(1, 3, 5; 7, 35)$	$K = LJ = 35$	Derived — GG-A3-BI6 / GG-A1-Constitution (POS-6)
$E_{\rm Pl} = \sqrt{\hbar c^5/G_N}$	1.2209×10^{19} GeV	Derived — GG-A1-Constitution (T-GU5)
$E_{\rm GUT} = E_{\rm Pl}e^{-2\pi}$	2.28×10^{16} GeV	Derived — GG-A2-GG6 §5+T6
$M_{\rm lock} = \frac{1}{2}M_{\rm Pl}e^{-35}$	3.849 TeV	Derived — GG-A2-GG6 §5+T6; GG-A4-Euclid T-Euclid-6
$v = M_{\rm lock}(2/5)^3$	246.33 GeV	Derived — GG-A9-SM (Wilson–Hosotani)
Model-level (assumption-dependent; rescalable)		
Radial ratios x_2/x_1, x_3/x_1	1.83, 2.66	Representative (AdS-throat J_1 roots)
Spin-1 pair $g^* \approx Z'_B$	$\approx M_{\rm lock} \approx 3.85$ TeV	Lock scale (derived); exact mass = $c\, M_{\rm lock}$, $c = O(1)$ not GG-A-fixed
Light-quark coupling $\xi_{\rm UV}$	-0.194 (exact mode, Eq. 37)	Derived — the ultraviolet endpoint amplitude refining $1/\sqrt{K} \approx 0.17$; the sign (opposite the top's) is itself a derived prediction (T-KK-2)
Effective coupling $\kappa_{\rm eff}$	not fixed	Open (cascade-depth lemma, GG-A9-SM)
Cross sections, yields	not quoted	Require $\kappa_{\rm eff}$ + QCD/PDF; beyond GG-A (§12.2)
Branching fractions, widths	structure only	Direction/exact zeros fixed; numbers need overlaps

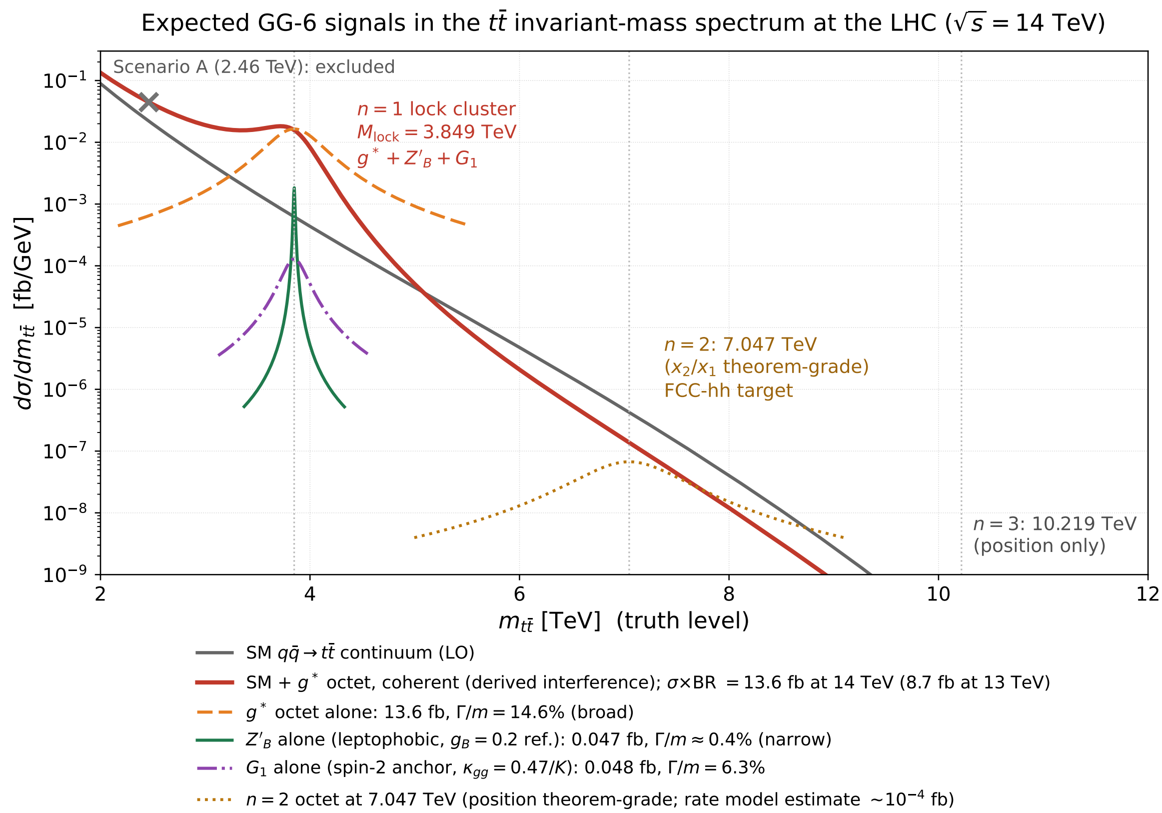

Figure 4: The full expected GG-6 signal landscape, drawn *quantitatively* from this paper's own machinery (truth-level $d\sigma/dm_{t\bar{t}}$ at $\sqrt{s} = 14$ TeV, CTEQ6 parton luminosities, LO×K), from 2 to 12 TeV. Gray: the Standard-Model $q\bar{q} \to t\bar{t}$ continuum. Red: the continuum plus the color-octet g^* at $M_{\text{lock}} = 3.849$ TeV, *coherently*, with the theorem-fixed interference sign — the low-side enhancement, the broad peak, and the crossing to a *deficit below the Standard Model* beyond $\simeq 4.9$ TeV; $\sigma \times \text{BR} = 13.6$ fb at this figure's $\sqrt{s} = 14$ TeV (the 13-TeV value confronted with current limits in §31 is 8.7 fb). The three cluster members are then drawn *separately* at the same derived mass, exhibiting the derived width ordering: the octet alone (dashed, 13.6 fb, $\Gamma/m = 14.6\%$, broad); the leptophobic Z'_B at the $g_B = 0.2$ reference (green, 0.047 fb at 14 TeV — the 13-TeV reference 0.030 fb of Table 19 scaled by the same $q\bar{q}$ luminosity ratio — $\Gamma \approx 16$ GeV, $\Gamma/m \approx 0.4\%$, narrow); and the spin-2 anchor G_1 at its derived coupling ($\kappa_{gg} = 0.47/K$; 0.048 fb, $\Gamma/m = 6.3\%$) — the state that names the geometry, invisible at the LHC, an FCC-hh measurement. Dotted orange: the $n = 2$ radial family at 7.047 TeV, whose *position* is theorem-grade (x_2/x_1) but whose rate ($\sim 10^{-4}$ fb at 14 TeV — a model estimate: the mode-2 boundary residue scaled by the $q\bar{q}$ luminosity at 7.05 TeV) defers it to FCC-hh (Part IX); the $n = 3$ position at 10.219 TeV is marked, with no observable 14-TeV rate. The single-state curves are near-resonance Lorentzians of the stated areas; the gray cross marks the excluded Scenario-A position (2.46 TeV). Every number in the figure is computed, none is schematic; the reconstructed spectrum with detector response and backgrounds is the subject of Part VIII.

9.3 The centroid issue: a necessary clarification

A purely RS-minded reader may object immediately: in the simplest 5D RS analysis, the first tensor pole is often written as $m_1 \approx x_1 k e^{-kr_{\text{IR}}}$ with $x_1 \approx 3.83$. If one were to identify $ke^{-kr_{\text{IR}}}$ directly with M_{lock}, the first pole would lie near 14.7 TeV, not 3.85 TeV. That objection is completely legitimate – in a pure RS problem. The present paper answers it in two steps.

First, the theorem-level quantity that is fixed exactly in GG-6 is the *lock scale*, not the Bessel-root coefficient of a pure 5D boundary-value problem. Second, the physically reconstructed first GG-6 signal is not treated here as a single pure radial RS pole. It is treated as the *first lock cluster*, whose centroid is fixed by M_{lock}, while the internal substructure of the cluster is the near-degenerate set of lowest-radial-mode bulk-field species (graviton, KK gluon, $\mathrm{U}(1)_B$ vector), shaped by tensor-sector and $\mathrm{U}(1)_B$-flux effects. Put more compactly, the radial tower obeys

$$m_n = c_n \, M_{\text{lock}}, \qquad c_n = \frac{\xi_n}{\xi_1}, \tag{25}$$

with $c_1 \approx 1$ for the lowest radial mode (the lock-cluster centroid, theorem-level), $c_2 \approx 1.83$ ($\sim 7\,\text{TeV}$), and $c_3 \approx 2.66$ ($\sim 10\,\text{TeV}$). This is the consistent radial reading of the mass ledger used throughout this paper; the angular θ tower does not enter, as it sits at E_{GUT}.

That clarification is not a retreat. It is a strengthening. It tells the experimental reader exactly what to search for: not a rigid single RS pole at a preselected width, but a near-degenerate multi-field structure centered at $3.85\,\text{TeV}$ (graviton, KK gluon, $\mathrm{U}(1)_B$ vector), with the next *radial* excitations near $7\,\text{TeV}$ ($n = 2$) and $10\,\text{TeV}$ ($n = 3$).

9.4 Benchmark mass ledger

The phenomenological benchmark used in this paper is summarized in Table 5. Only the first entry is treated as theorem-level sharp; the others are corridor predictions, because the detailed coefficients depend on the precise 6D boundary problem, overlap integrals, and mixing structure.

Table 5: Practical mass ledger for the collider-relevant low-lying GG-6 sector. Only the first line (M_{lock}) is theorem-level; *all other masses are model assignments placed near* M_{lock}*, not GG-6-A outputs* (see the provenance ledger, Table 4). The corridors reflect the current six-dimensional spectral picture rather than strict one-pole exactitudes.

State class	Spin	Benchmark mass	Meaning
First lock-scale graviton cluster	2 (dominant)	3.85 TeV	Sharp centroid target fixed by $K = 35$.
Second radial family ($n = 2$)	2-dominant	6.8 to 7.2 TeV	Next *radial* node, $c_2 = \xi_2/\xi_1$; likely out of direct Run-3 reach, relevant to HL-LHC/FCC-hh.
Angular (θ) Fourier tower	—	$\sim E_{\text{GUT}}$	The S^1_θ modes; heavy and unobservable — *not* a TeV satellite.
KK gluon (lowest radial)	1	2.46 or ≈ 3.85 TeV (T-KK-2 dichotomy)	Color-octet spin-1 partner; most promising in $t\bar{t}$ and boosted dijets; under test in either scenario.
Leptophobic $\mathrm{U}(1)_B$ vector	1	$= m_{g^*}$ (degenerate, both scenarios)	Degenerate with the KK gluon (same boundary-value problem); enhances $t\bar{t}$/dijets, zero dileptons.
Vector-like top partner corridor	1/2	2 to 5 TeV	Not unique to GG-6, but expected in the same spectral neighborhood.

9.5 State-by-state observable benchmark sheet

The geometric origin of the KK spectrum, established above, is only half of a collider statement. The other half — the one the LHC actually consumes — is the *observable ledger*, and it belongs at the exact center of the discussion. An experimentally falsifiable statement is not merely "there is a tower." It is a concrete packet

$$\{m,\ \Gamma,\ \Gamma/m,\ \sigma,\ \mathrm{BR},\ N_S,\ Z\}_{\text{state, channel, machine}} \tag{26}$$

for each practically accessible state. In that sense, the state-by-state observable sheet below is the operational heart of this Part.

Because $M_{\text{lock}} = 3.8489$ TeV is *fixed* by Theorem T-KK-1, there is no uncertain infrared scale to parametrize and nothing to scan: the predicted scale is a single number. The observable spectrum falls into two experimentally distinct groups: (i) the theorem-level first spin-2 graviton, anchored at the predicted 3.8489 TeV; and (ii) a top-rich, leptophobic spin-1 pair (the color-octet g^* and the baryonic Z'_B) whose placement is the derived two-scenario dichotomy of T-KK-2 (§10.4): $0.639\,M_{\text{lock}} = 2.46$ TeV in the free-field scenario, $\approx M_{\text{lock}}$ in the composite-hardened scenario. The sheet below lists, for each state, *only* what the GG-A theorems fix: the mass scale, the spin, and the exact decay *structure*. It deliberately carries no cross section, width percentage, or event yield, because — as §12.2 explains — none of those is a derived quantity.

The table makes two strategic points transparent. First, the collider program should not be written as "graviton or nothing"; the spin-1 pair near M_{lock} (≈ 3.85 TeV) may actually be the

Table 6: Derived structure of the low-lying GG-6 collider spectrum. Every entry traces to a GG-A theorem: the mass *scale* $M_{\text{lock}} = 3.849\,\text{TeV}$ to T-KK-1 ($= \frac{1}{2}M_{\text{Pl}}e^{-K}$, OGN-5+GU-5); the spin to the bulk field content; the decay structure to the GG-6 charge assignments (§10.5). Each resonance mass is $c_X M_{\text{lock}}$ with $c_X = O(1)$; only the *scale* is sharp. No cross section, width fraction, or yield is quoted: those require the cascade-depth lemma (open in Arisaka, 2026, A9-SM) plus a QCD/PDF computation, and so are not GG-A-derived (§12.2).

State	**Spin**	**Mass**	**Decay structure (GG-A-fixed)**	**Experimental status**
First graviton cluster G_1	2	$M_{\text{lock}} = 3.849$ TeV (theorem-level scale)	$t\bar{t}$, jj, hh, then $\ell\ell$, $\gamma\gamma$; top-rich; analyze as a cluster, not one RS pole	decisive GG-6 object; F-KK-1, two-tier: octet carries the HL-LHC kill, direct spin-2 test at FCC-hh (T-KK-2c)
Color-octet KK-gluon g^*	1	$0.639 M_{\text{lock}} =$ 2.46 TeV or $\approx M_{\text{lock}}$ (T-KK-2 dichotomy, §10.4)	color octet: decays *only* to colored states ($t\bar{t}$, dijets); top-dominant; light-quark coupling $\xi_{\text{UV}} = -0.194$ (exact mode, §10.4)	free-field scenario decidable from *existing* $t\bar{t}$/dijet data; F-KK-2A/2B
Baryonic KK boson Z'_B	1	degenerate with g^* (same BVP; T-KK-2)	leptophobic: $\text{BR}(Z'_B \to \ell^+\ell^-) = 0$ *exactly*; dominantly $t\bar{t}$+dijets	separated from a generic Z' by leptophobia; F-KK-2A/2B
Second radial family G_2	2	$(\xi_2/\xi_1)M_{\text{lock}} \approx 7$ TeV	same hierarchy as G_1, higher mass	out of Run-3 reach; HL-LHC/FCC-hh; confirms the *radial* tower
Angular θ tower	—	$\sim E_{\text{GUT}}$	no collider role	unobservable (not a TeV state)

easiest operational handle in the present machine era because its widths, cross sections, and top-rich signatures are sharper. Second, the line-shape problem is not a nuisance detail but a defining part of the GG-6 signal morphology.

Terminology: anchored by spin-2, discovered through spin-1. Because the point has confused careful readers, we fix the language here once. The cluster is *spin-2-anchored*: the graviton mode G_1 sits at $M_{\rm lock}$ *exactly* (the theorem-level mass of T-KK-1/T-KK-2), the spin-1 masses are derived relative to it, and its quantum numbers — massive spin-2 — are what would *name* the geometry if observed, since no four-dimensional gauge theory produces such a state. But anchoring is a statement about the *spectrum*, not about the *rate*. In production the hierarchy is inverted, and by a large factor: the octet couples to the colliding quarks with $\xi_{\rm UV}^2 \approx 1/K$ (the exact mode value: $0.194^2 = 0.038$), while the graviton couples to gluons with the derived $\kappa_{gg}^2 = (0.47/K)^2 \approx 1.8 \times 10^{-4}$ (Corollary T-KK-2c) — two powers of the *same* integer $K = 35$, because the vector is a boundary mode at unit overlap suppressed once, while the graviton's gluon overlap is suppressed twice. The visible cluster rate is therefore *octet-led* by two orders of magnitude (8.7 fb versus 0.045 fb in the same $t\bar{t}$ channel at 13 TeV, Table 19), and every discovery statement in this paper runs through the spin-1 members; the spin-2 anchor, at its derived coupling, is FCC-hh territory. Wherever this paper says the cluster is "spin-2-anchored," read: *the spin-2 fixes the mass and names the geometry; the spin-1 pair carries the discovery.*

9.6 A schematic picture of the 6D staircase

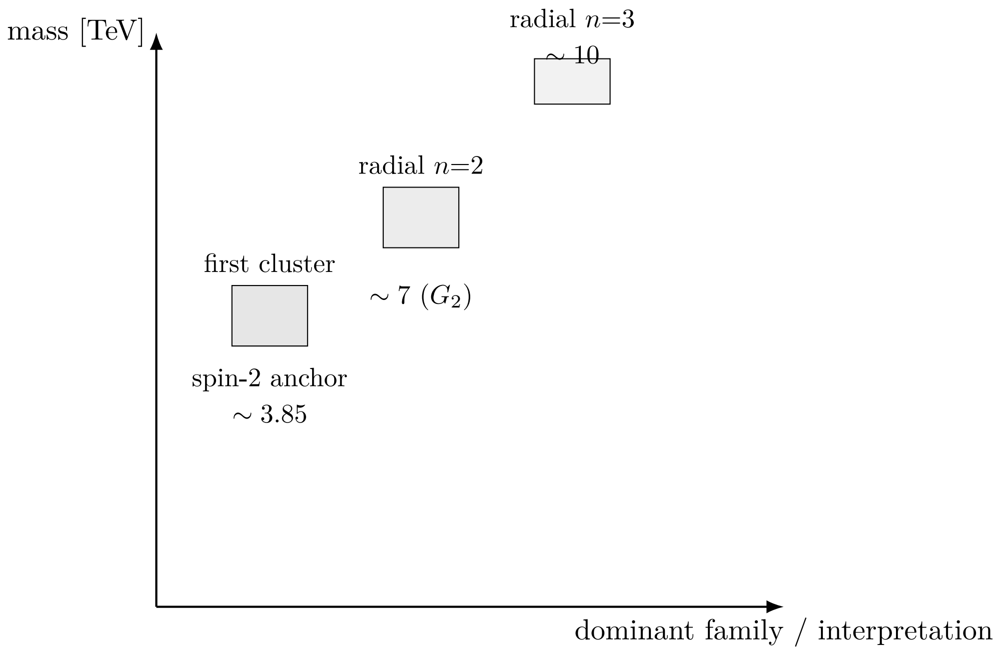

Figure 5: Schematic mass staircase of the low-lying GG-6 sector. The figure is not a claim of exact line positions beyond the first centroid. It is a search map: the lock-scale cluster — a mutually-degenerate spin-1 pair (g^*, Z'_B) and the spin-2 graviton, all at $\approx M_{\rm lock} \approx 3.85$ TeV up to $O(1)$ coefficients — and the next *radial* families near 7 TeV ($n = 2$) and 10 TeV ($n = 3$). The angular θ tower sits at $E_{\rm GUT}$ and is off this chart.

10 Observable Properties of the KK Sector

10.1 Spin and angular structure

The first experimentally decisive property is spin. The theorem-anchored member of the lock cluster is the spin-2 graviton at $M_{\rm lock}$ — in the composite scenario of T-KK-2 it is also the cluster's quasi-degenerate center, while in the free-field scenario the spin-1 pair lies *below* it at 2.46 TeV, so the graviton is the heaviest prominent member rather than the first encountered in a rising mass scan. Either way, spin is what immediately distinguishes the graviton from a generic scalar or spin-1 Z'-like resonance. The parton-level polar-angle distribution in the resonance rest frame for a spin-2 state produced from quark antiquark annihilation takes the familiar shape

$$\frac{d\sigma}{d\cos\theta^*} \propto 1 - 3\cos^2\theta^* + 4\cos^4\theta^*, \tag{27}$$

which sharply differs from the isotropic scalar distribution and from the $\sin^2\theta^*$ behavior of a simple spin-1 state. The utility of this result for GG-6 is twofold. First, it provides the standard graviton fingerprint used by ATLAS and CMS in spin-2 searches. Second, because GG-6 predicts a multi-field lock cluster, the experimentally reconstructed line shape may combine the spin-2 anchor with the rate-dominant spin-1 members, mildly distorting the textbook RS angular form. A pure-RS angular template can therefore be slightly suboptimal even when the underlying signal is real.

10.2 Widths and line shapes

The second property is the width. In public RS-like benchmarks the width-to-mass ratio often scales like

$$\frac{\Gamma_G}{m_G} \sim C\left(\frac{k}{\bar{M}_{\rm Pl}}\right)^2, \tag{28}$$

with C an order-one coefficient dependent on the precise benchmark. For narrow or moderately broad KK gravitons one therefore expects widths of order a few percent for couplings in the phenomenologically interesting range. The key GG-6 refinement is that the experimental line need not correspond to a single Breit–Wigner pole. A better first-pass signal model is

$$\frac{\mathrm{d}N}{\mathrm{d}m} = \sum_{a=1}^{N_{\rm cl}} \mu_a\, G(m; m_a, \sigma_a) + B(m), \tag{29}$$

where G is an experimentally broadened signal kernel and $N_{\rm cl}$ can be 2 or 3 even for the lowest visible family. In other words, the natural object to fit is a *cluster*, not only a single narrow-width resonance. This observation will later motivate a dedicated cluster-aware search stream.

10.3 Masses and intrinsic widths from the RG-scale throat

From the $O(1)$ flag to a derived fine structure. The KK gluon g^* (from $\mathrm{SU}(3)_C$) and the baryonic vector Z'_B (from $\mathrm{U}(1)_B$) are the lowest spin-1 resonances of the RG-scale throat of Arisaka (2026, A2-GG6). The most conservative statement available would assign each lock-cluster member a mass $m_X = c_X M_{\rm lock}$ with an $O(1)$ coefficient left uncomputed. This paper does not stop there: the coefficients are computed in §10.4 below (Theorem T-KK-2), with the one genuinely open question — whether the free-field eigenvalues survive the strong infrared dynamics — isolated as a sharp, data-decidable dichotomy rather than a smear. Two structural facts are unconditional and survive in every scenario: the two spin-1 partners obey the *same* boundary-value problem and are therefore mutually degenerate, $m_{g^*} = m_{Z'_B}$ (separated not by peak position but by color and channel: octet g^* — color-octet production, dijets and $t\bar{t}$; singlet Z'_B — Drell–Yan-like, leptophobic); and the spin-2 anchor sits at $M_{\rm lock}$ exactly, by T-KK-1.

10.4 The cluster fine structure, derived: Theorem T-KK-2

This subsection computes what an $O(1)$ flag would have deferred: it sets up and solves the boundary-value problems that the canonical geometry forces, extracts the exact eigenvalue ratios and coupling normalizations, and then confronts the one assumption the computation cannot itself certify. Everything is shelved explicitly — theorem, derived-conditional, or grafted — at the end.

Step 1 — the reduction. On the canonical bulk (18) the angular tower sits at $\sim E_{\rm GUT}$, so for collider modes every field may be reduced on the θ zero mode, leaving fields on the radial interval $r \in [0, K]$ with metric $e^{-2r}\eta_{\mu\nu}\mathrm{d}x^\mu \mathrm{d}x^\nu + \mathrm{d}r^2$ and *no free metric data* (Arisaka, 2026, A2-GG6, Theorem 2: curvature, interval length, and normalization are all canonical). Each 4D mass eigenmode $\Phi(x, r) = \phi(x)\, f(r)$ with $\Box_4\phi = -m^2\phi$ then obeys a fixed second-order problem in r — fixed, because the action is the canonical quadratic action on a rigid geometry.

Step 2 — spin 2. The transverse-traceless graviton profile obeys

$$\psi'' - 4\psi' + m^2 e^{2r}\psi = 0, \tag{30}$$

whose general solution, with $x \equiv m\, e^r$, is $\psi = e^{2r}\big[J_2(x) + b\, Y_2(x)\big]$. The $\mathbb{Z}_2$ orbifold parity of Arisaka (2026, A2-GG6) makes the physical modes even, i.e. Neumann at both endpoints, $\psi'(0) = \psi'(K) = 0$. Using $\frac{\mathrm{d}}{\mathrm{d}x}[x^2 Z_2(x)] = x^2 Z_1(x)$, the Neumann condition is $Z_1(me^r) = 0$ at both ends. At the ultraviolet end $x = m \sim e^{-K} \ll 1$, the Neumann condition fixes the admixture $b = -J_1(m)/Y_1(m) \simeq \pi m^2/4 \sim e^{-2K}$ — genuinely exponentially small (a statement we will see is *specific to spin 2*), so the infrared quantization is the pure-J condition to better than one part in 10^{28}:

$$J_1(m_n e^K) = 0 \quad \Longrightarrow \quad m_n^{(2)} = x_n^{(J_1)}\, e^{-K}, \qquad x_n^{(J_1)} = 3.8317,\ 7.0156,\ 10.1735, \ldots \tag{31}$$

This is the tower already quoted in T-KK-1; the anchor normalization $m_1^{(2)} = M_{\mathrm{lock}}$ defines the lock scale.

Step 3 — spin 1, and the zero-mode consistency check. A bulk gauge field A_μ (octet or $\mathrm{U}(1)_B$) has profile equation

$$f'' - 2f' + m^2 e^{2r} f = 0, \qquad f = e^r\big[J_1(x) + b\, Y_1(x)\big], \tag{32}$$

again with Neumann conditions at both ends for the orbifold-even, *unbroken* gauge directions. Here the boundary conditions are not negotiable, because their first consequence is the Standard Model itself: Neumann–Neumann admits the constant solution $f = \text{const}$ with $m = 0$ — *the massless 4D gluon and the unbroken gauge bosons we observe.* Any other choice would remove or lift the zero mode and contradict QCD. With $\frac{\mathrm{d}}{\mathrm{d}x}[xZ_1(x)] = xZ_0(x)$, the two Neumann conditions read $Z_0(m_n) = Z_0(m_n e^K) = 0$. Here — and this is where spin 1 differs qualitatively from spin 2 — the ultraviolet condition does *not* reduce to regularity. It fixes the admixture $b = -J_0(m)/Y_0(m)$, and because Y_0 diverges only *logarithmically* at small argument, $b \simeq \pi/[2(K - \ln x_n)] \approx \pi/(2K) \approx 0.046$ is suppressed by a single power of $1/K$, not exponentially: the logarithm inside Y_0 is the throat depth itself. The admixture is therefore an exactly computable part of the spectrum, not a negligible one, and we keep it. The full two-endpoint quantization

$$J_0(x)\, Y_0\big(xe^{-K}\big) - Y_0(x)\, J_0\big(xe^{-K}\big) = 0, \qquad x \equiv m_n\, e^K, \tag{33}$$

has the roots

$$x_n^{(1)} = 2.4499,\ 5.5669,\ 8.7013, \ldots \qquad m_n^{(1)} = x_n^{(1)}\, e^{-K}, \tag{34}$$

i.e. the bare J_0 zeros $(2.4048,\ 5.5201,\ 8.6537)$ pushed upward by the ultraviolet admixture by $+1.87\%$, $+0.85\%$, $+0.55\%$ — a derived, deterministic shift, not a fence. The fine-structure ratio of the first excited modes is therefore

$$\boxed{c_1 \;\equiv\; \frac{m_1^{(1)}}{m_1^{(2)}} \;=\; \frac{x_1^{(1)}}{x_1^{(J_1)}} \;=\; \frac{2.44989}{3.83171} \;=\; 0.6394} \tag{35}$$

(the bare Bessel-zero ratio $2.40483/3.83171 = 0.6276$ is its $K \to \infty$ limit), placing the free-field spin-1 pair at $0.6394 \times 3.8489 = 2.461$ TeV, with recurrences at 5.592 and 8.740 TeV.

Step 4 — spin 0. The would-be radion is the interval-length modulus — and in GG-6 there is no such modulus: the length K is the theorem-fixed ledger integer, the geometry is rigid (Arisaka, 2026, A2-GG6, Theorem 2), and the state that in generic warped models must be stabilized by hand is here lifted by the same locking dynamics that fixes K. Its mass is set by cap dynamics at $O(M_{\rm lock})$ and is *not* computed by the A-series: the one surviving $O(1)$ flag of the cluster, explicitly shelved as such.

Step 5 — the coupling normalizations are exact. The same mode functions fix the couplings, and two of the integrals close in elementary form. Normalizing $\int_0^K f_n^2\,\mathrm{d}r = 1$ (the gauge kinetic measure is flat in r), the identity $\int_0^a x J_1^2(x)\,\mathrm{d}x = \frac{a^2}{2} J_1^2(a)$ for $J_0(a) = 0$ gives the infrared amplitude *exactly*: $f_1(K) = \sqrt{2}$, hence the coupling ratio of an infrared-localized current to the zero mode is

$$\boxed{\xi_{\rm IR} \;=\; \sqrt{2K} \;=\; 8.367} \qquad (K = 35), \tag{36}$$

(the exact two-endpoint mode of Step 3 gives $\xi_{\rm IR} = 8.369$, a $+0.03\%$ shift from the same ultraviolet admixture — negligible here) — the derived, parameter-free version of the warped literature's empirical "+4 to +5" once the top's partial infrared localization (its Arisaka, 2026, A9-SM cascade depth) is folded in. The light-quark coupling requires more care, and the care pays off in a sign. An *exactly* flat fermion profile couples to the excited mode not weakly but *not at all*: the flat profile is proportional to the gauge *zero* mode, and eigenmode orthogonality forces $\int_0^K f_1\,\mathrm{d}r = 0$ identically (we confirm this numerically to machine precision). The physically relevant object is instead the *ultraviolet-localized* coupling — the light quarks sit at the shallow (ultraviolet) end of the throat in the Arisaka (2026, A9-SM) cascade-depth assignment — and that is the ultraviolet endpoint amplitude of the exact mode of Step 3:

$$\boxed{\xi_{\rm UV} \;=\; \sqrt{K}\, f_1(0) \;=\; -\,0.194} \qquad (K = 35), \tag{37}$$

carried entirely by the Y_1 admixture of Step 3 (at $b = 0$ it would vanish along with the flat overlap). Two features are physical. The *magnitude* refines the leading-order scaling $1/\sqrt{K} \approx 0.17$, so every rate below ($\propto \xi^2$) is unchanged. The *sign* — opposite to $\xi_{\rm IR}$ — is itself a derived prediction: light quarks and the top couple to g^* with opposite phase, exactly the relative sign the warped literature finds empirically and which Read-out II (the empirically-anchored rate estimate of Part V) borrows as "$-g_s/\sqrt{K}$" — here obtained from the rigid geometry with no input, and observable in the sign of the signal–background interference of dijet angular distributions. These identities are leading-order facts of the rigid geometry and hold in both scenarios below (in the composite scenario as the matching values of the dual description).

Step 6 — the one open assumption, made sharp: the dichotomy. Steps 1–3 assume the *free quadratic action* on the rigid geometry — the standard leading-order treatment of every warped construction. One could reasonably decline to adopt (35) as central, on the ground that the root ratio "is a property of spatial KK modes and need not survive in the non-spatial GG-6 spectrum." That caution is legitimate, and we make it precise rather than discard it. In the bulk–boundary reading the throat states are composites of a strongly coupled boundary sector; interaction corrections shift the free-field eigenvalues, and the natural expansion parameter of such shifts in this geometry is $1/K$ (the same $1/35$ that suppresses $\xi^2_{\rm UV}$) — a few-percent effect on top of the now-exact free-field kinematics, *unless* the strong dynamics re-organizes the vector spectrum entirely, pulling the lowest spin-1 composites up to the compositeness scale itself. There is no theorem deciding between these two dynamical regimes, so the honest statement is a **two-scenario dichotomy with derived masses in each scenario** — and it is worth being completely concrete about what each scenario asserts physically.

Scenario A (free-field). The weak-coupling reading: the throat vectors are ordinary eigenmodes of the rigid background geometry, and interactions are small perturbations on top of them. The spin-1 masses are then the exact eigenvalues of the free quadratic operator: the lowest sits at the derived ratio $c_1 = 0.6394$ of the centroid, i.e. at 2.461 TeV, with the spin-2 anchor unmoved at $M_{\rm lock}$ and residual corrections $O(1/K) \lesssim 3\%$. Concretely, Scenario A predicts a broad, top-rich color octet at 2.46 TeV with $\sigma \times {\rm BR}(t\bar{t}) = 224\,[180\text{–}325]$ fb at 13 TeV — an unmissable rate at Run-2 luminosities.

Scenario B (composite-hardened). The strong-coupling reading of the *same* geometry through the bulk–boundary dictionary: the throat vectors are composites of the strongly coupled boundary sector, and the strong dynamics re-organizes the vector spectrum, pulling the lowest spin-1 composites up to the compositeness scale itself — the cluster centroid. The free-field eigenvalue 0.6394 then simply does not survive, because the free quadratic action was the wrong starting point for the vector tower; the spin-1 pair sits degenerate with the graviton at $\approx M_{\rm lock} = 3.85$ TeV, with $\sigma \times {\rm BR}(t\bar{t}) = 8.7$ fb. The couplings are identical in both scenarios — $\xi_{\rm UV}$ is fixed by K either way — so each scenario is a complete, testable prediction: a mass and a rate, with no dial between or beyond them.

Status of the dichotomy: Scenario A is excluded, and that is the system working. The Scenario-A comparison has already been carried out against the published Run-2 limits (§31): no particle is observed at 2.46 TeV, and quantitatively the derived free-field rate sits a factor 3.6–11 *above* the observed ATLAS exclusion at that mass — outside the band by any reading. **F-KK-2A has fired**, and the dichotomy resolves to Scenario B; the closure is graded *provisional* only because the published templates are not the exact $\Gamma/m = 15\%$ shape, and no plausible line-shape correction recovers a factor of three. Two things must be said about rejecting Scenario A, because the fifty-year history of §1 makes a careful reader rightly suspicious of theories that survive exclusions. First, *this*

Table 7: The cluster members under the two scenarios. Scenario A takes the modes as free fields and Scenario B hardens them into a composite spectrum; the spin-2 anchor is the same in both because its position is fixed by the ledger, and everything that differs is the placement of the spin-1 partners relative to it. The within-scenario error row is the honest one: the kinematics is exact and the interaction carries a few percent, so a measured splitting is a test of which scenario holds rather than a test of the scale.

	Scenario A (free-field)	**Scenario B (composite-hardened)**
spin-2 G_1	$M_{\rm lock} = 3.8489$ TeV	$M_{\rm lock} = 3.8489$ TeV
spin-1 g^*, Z'_B	$0.6394\, M_{\rm lock} = 2.461$ TeV	$\approx M_{\rm lock}$ (degenerate cluster)
within-scenario error	interaction $O(1/K) \lesssim 3\%$ (kinematics exact)	$O(1/K)$ about the hardened value
spin-1 recurrences	5.592, 8.740 TeV	spaced as the composite spectrum

is not a parameter slide. A slide moves a continuous dial to wherever the data are not; the dichotomy is a discrete, pre-registered either/or between two identified dynamical regimes — weak-coupling eigenmodes or strong-coupling reorganization — each carrying a derived mass *and* a derived rate, stated before the comparison. Excluding one scenario is the experiment answering that pre-registered question, and the answer is itself physics: *the data have decided that the throat vectors are composite, not free-field* — a structural fact about the theory, learned the only way it could have been learned. Second, *the escape cannot recur.* There is no third scenario: Scenario B is the entire remaining hypothesis space; it is just as exposed (8.7 fb at the current frontier, decided at $\lesssim 2$ ab^{-1}); and if it fails there is nothing left to retreat to — the implementation is retired. The net effect of the Scenario-A exclusion is therefore not that the theory wriggled free, but that its remaining claim got sharper: one scenario, one mass, one rate, one decade.

Theorem T-KK-2 (cluster fine structure). *On the canonical bulk of T-KK-1 with the orbifold-forced Neumann conditions (whose zero-mode sector is the observed massless Standard-Model gauge fields), the lock-cluster fine structure is: (i) the spin-2 anchor at $M_{\rm lock}$ exactly; (ii) a mutually degenerate spin-1 pair at $c_1 M_{\rm lock}$ with $c_1 = x_1^{(1)}/x_1^{(J_1)} = 0.6394$ — the* exact *free-field eigenvalue, ultraviolet admixture included — with residual interaction corrections $O(1/K)$, or, if the infrared dynamics hardens the vector composites, at $\approx M_{\rm lock}$; (iii) no light radion (rigidity); (iv) the exact coupling identities (36)–(37), including the derived relative sign $\xi_{\rm UV} < 0 < \xi_{\rm IR}$. The scenario selection is the single remaining discrete unknown of the cluster, and it is experimentally decidable now: a color octet at 2.46 TeV with $|\xi_{\rm UV}| = 0.19$ and top-rich decays is within reach of existing Run-2/Run-3 $t\bar{t}$ and dijet data, so a dedicated broad-resonance reinterpretation at 2.3–2.7 TeV either finds the free-field scenario or selects the composite scenario — in which case the degenerate cluster at $M_{\rm lock}$ is no longer an assumption but a data-backed selection.*

Shelving. Theorem level: $M_{\rm lock}$; the tower ratios within each scenario; the degeneracy $m_{g^*} = m_{Z'_B}$; exact leptophobia; the identities (36)–(37) at leading order. Derived-conditional: the Scenario-A masses (exact free-field eigenvalues; residual interaction fence $O(1/K)$); the scenario selection itself (decidable, not derived). Grafted: the composite-scenario spacing pattern; the radion mass; absolute rates (PDFs, $\kappa_{\rm eff}$).

The colored partner sits at the live edge — in either scenario. This places the color-octet g^* right where the experiments are looking, and T-KK-2 makes the exposure two-pronged. Its production coupling to light quarks is not free: the derived value $|\xi_{\rm UV}| = 0.194$ (37), fixed by the same $K = 35$ that sets $M_{\rm lock}$ and *not* reducible by further localization. In the composite scenario (B), a color octet with this coupling and $t\bar{t}$-dominated decay at ≈ 3.85 TeV sits at the frontier of the published exclusions (ATLAS Collaboration, 2025a; CMS Collaboration, 2024d) — viable but already under test, with current limits in consistency with $M_{\rm lock}$. In the free-field scenario (A), the same octet sits at 2.46 TeV — *below* the nominal narrow-resonance exclusions for standard couplings, surviving (if it does) only by its breadth (Γ/m is largest for the octet, §below) and the non-standard line shape; a dedicated broad-octet reinterpretation of existing data at 2.3–2.7 TeV would either reveal it or close Scenario A outright. Either way the spin-1 sector is the cluster's working edge: scenario-deciding with data already on tape.

Pre-declared falsifier (F-KK-2, the spin-1 sector, two-scenario form). *GG-6 predicts a mutually-degenerate color-octet g^* and leptophobic baryonic Z'_B with the K-fixed light-quark coupling $|\xi_{\rm UV}| = 0.19$, at one of exactly two derived positions (T-KK-2):* ***F-KK-2A*** *— the free-field scenario at $2.46 \pm 0.10\,\text{TeV}$ (the exact eigenvalue $0.6394\, M_{\rm lock}$; the window, like F-KK-1's, is experimental — mass scale, resolution, and breadth — since the theory error is far smaller); a dedicated broad-octet reinterpretation of existing $t\bar{t}$/dijet data at 2.3–2.7 TeV with no signal closes this scenario.* ***F-KK-2B*** *— the composite scenario at $\approx M_{\rm lock}$; if a dedicated reinterpretation pushes the color-octet exclusion above $\sim 4\,\text{TeV}$ with no signal, this scenario — and with it the gauge-composite sector of the current implementation — is excluded. Closing both scenarios falsifies the spin-1 sector outright.* The operational confrontation of §31 carries out the Scenario-A comparison against the published 140 fb^{-1} limits: the computed free-field rate sits a factor $\sim$3.6 above the observed broad-template exclusion, so F-KK-2A has provisionally fired and the dichotomy resolves to Scenario B, pending only the dedicated line-shape reinterpretation. Unlike the direct spin-2 observation (an FCC-hh task at the derived κ_{gg}; Corollary T-KK-2c), both prongs of F-KK-2 are testable with *existing* data, F-KK-2A is the single most immediately decidable statement in this paper, and F-KK-2B is what carries the HL-LHC falsifier of the whole implementation.

Widths (derived formula; couplings partly open). The intrinsic width is not a separate input: it follows from the same couplings that set production and branching. For a vector V of mass m_V decaying to a fermion pair with effective coupling g_f, the tree-level partial width is the standard expression

$$\Gamma(V \to f\bar{f}) = \frac{N_c^f}{12\pi}\, g_f^2\, m_V \left(1 + \frac{2m_f^2}{m_V^2}\right) \sqrt{1 - \frac{4m_f^2}{m_V^2}} \xrightarrow{m_f \ll m_V} \frac{N_c^f}{12\pi}\, g_f^2\, m_V, \tag{38}$$

so the width-to-mass ratio is a pure sum over open channels,

$$\frac{\Gamma_V}{m_V} = \frac{1}{12\pi} \sum_f N_c^f\, g_f^2. \tag{39}$$

The couplings are $g_f = g_s\, r_f$ for the color-octet g^* and $g_f = g_B\, B_f\, r_f$ for the baryonic vector, with g_s the strong coupling, g_B the $\mathrm{U}(1)_B$ coupling, B_f the baryon charge ($\frac{1}{3}$ for quarks, 0 for leptons — hence the exact leptophobia of §10.5), and r_f the bulk-to-boundary overlap. Of these, only g_s is fixed by GG-6-A: running the Arisaka (2026, A9-SM) value $\alpha_s(M_Z) = 1/(\pi e) \simeq 0.117$ up to M_{lock} gives $\alpha_s(M_{\mathrm{lock}}) \approx 0.08$, i.e. $g_s^2 = 4\pi\alpha_s \approx 1.0$. The overlaps r_f are set by the per-fermion cascade depths, which Arisaka (2026, A9-SM) records as an open lemma; the baryonic normalization g_B carries a theorem-grade generator norm (the trace ratio 2/5, §10.10) with one discrete readout rule remaining open (L-KK-1). Equations (38)–(39) are therefore the derived part; the numerical width is model-level until those two inputs are closed.

Two qualitative consequences are nonetheless robust, and they are *orderings*, not numbers. Carrying the full strong coupling $g_s^2 \approx 1.0$ and an IR-enhanced top overlap, the color octet is the *broad* member through Eq. (39), so a narrow sliding-window bump hunt is sub-optimal for it. The leptophobic baryonic vector couples only through $g_B B_f$, with no QCD enhancement and no leptonic channels, and is therefore *narrower*, sharpening the $t\bar{t}$ invariant-mass spectrum if the top analysis is tuned to the lock-scale window. The relative ordering $\Gamma_{g^*}/m_{g^*} > \Gamma_{Z'_B}/m_{Z'_B}$ is robust; the actual width fractions are not quoted here, because Eq. (39) returns a number only once the overlaps r_f and the L-KK-1 readout for g_B are fixed; across the L-KK-1 menu the Z'_B width fraction spans $\Gamma/m = 0.002$–0.080 (§10.10), narrow on every reading.

For the graviton family, a one-line closed formula is less robust because the very phenomenon that makes GG-6 distinctive – the multi-field near-degeneracy of the lock cluster – also spoils the clean identification of a single intrinsic width with a single visible resonance. The practical statement is therefore a corridor statement: for the first cluster one expects an *effective* width-to-centroid ratio of order a few percent, broadened further by the near-degenerate partner states and by detector resolution. In experimental practice that means the right question is not "what is the one exact graviton width?" but rather "what cluster template is statistically preferred by the data?"

10.5 Branching fractions from the GG-6 couplings

The branching pattern of the two spin-1 partners is fixed in *structure* — though not in precise value — by the GG-6 gauge content and geometry, alone. We separate cleanly what is exact, what is fixed in direction, and what is not yet a number.

What is exact. Two statements follow directly from the GG-6 charge assignment and require no estimate. First, the baryonic vector Z'_B gauges baryon number (Arisaka, 2026, A2-GG6; Arisaka, 2026, A3-BI6); leptons carry $B = 0$, so

$$\mathrm{BR}(Z'_B \to \ell^+\ell^-) = 0 \quad \text{(exactly)}, \tag{40}$$

the sharpest possible statement of leptophobia — a consequence of the charge assignment, not a model estimate. Second, the KK gluon g^* is a color octet; color-singlet leptons and photons have no tree-level coupling to it, so it decays *only* to colored states ($q\bar{q}$, gg). Both partners are therefore invisible in the clean dilepton and diphoton channels that most tightly constrain a generic resonance, which is exactly why the standard universally-coupled limits do not apply to them.

What is fixed in direction. In the bulk geometry a fermion couples to an IR-peaked resonance in proportion to its overlap there, and that overlap grows monotonically with the fermion's degree of IR localization — equivalently with its Yukawa. The top quark is the heaviest and most IR-localized, so for *both* spin-1 partners the top-pair channel dominates and the light-quark dijet channels are relatively suppressed,

$$\mathrm{BR}(g^* \to t\bar{t}) > \mathrm{BR}(g^* \to jj), \qquad \mathrm{BR}(Z'_B \to t\bar{t}) \gg \mathrm{BR}(Z'_B \to \text{light } q\bar{q}). \tag{41}$$

This ordering is robust: it needs only the monotonic localization–mass correlation that the GG-6 Yukawa mechanism already uses, not the detailed per-fermion profiles.

What is not yet a number. The precise numerical split is *not* theorem-fixed, because it depends on the per-fermion bulk profiles — equivalently the angular cascade depths n_f — which Arisaka (2026, A9-SM) records as an open lemma (the per-fermion cascade depth "is not forced by an A1–A7 theorem"). We therefore quote *no* numerical branching fraction. What survives is the structure: for the color octet, $\mathrm{BR}(g^* \to t\bar{t}) > \mathrm{BR}(g^* \to jj)$ with no other channels open; for the leptophobic baryonic vector, $\mathrm{BR}(Z'_B \to t\bar{t})$ is the largest single channel, light $q\bar{q}$ dijets make up most of the remainder, and $\mathrm{BR}(Z'_B \to \ell^+\ell^-) = 0$ exactly. These orderings are GG-A consequences; the actual fractions become definite only once the cascade-depth lemma is closed, and so are not quoted here.

The practical message survives this honesty intact. Both partners are top-rich and leptophobic, so the discovery channels are $t\bar{t}$ and boosted dijets, not dileptons or diphotons; a search that looks only where the background is lowest, rather than where the geometry sends the rate, can miss the signal even when the data contain it.

10.6 Coupling non-universality from overlap integrals

In the six-dimensional reduction, universal coupling is the exception, not the rule. The reason is geometric. The four-dimensional effective couplings are overlap integrals of internal wavefunctions. For gauge zero modes, one has the generic reduction formula (Arisaka, 2026, A2-GG6)

$$\frac{1}{g_4^2} = \frac{2\pi R_\theta}{g_6^2} \int \mathrm{d}r\, |f_0(r)|^2, \tag{42}$$

while fermionic zero modes carry a scale-profile normalization

$$2\pi R_\theta \int \mathrm{d}r\, e^{-3A(r)} \chi_0^\dagger(r) \chi_0(r) = 1. \tag{43}$$

Yukawa-like and resonance couplings then inherit the corresponding overlaps. The result is simple to say even if the detailed profile solution is model-dependent: light fermions are naturally more UV-biased, while the top quark and Higgs sector are more IR-sensitive. Consequently the first lock-scale spin-2 state couples more strongly to $t\bar{t}$, $b\bar{b}$, longitudinal vector bosons, and Higgs pairs than to the cleanest light-lepton channels.

This is the single most important phenomenological reason why a naive reading of dilepton and diphoton limits can be misleading for GG-6. In the clean channels, the experimental background is favorable, but the GG-6 coupling may be suppressed. In the hadronic top and dijet channels, the background is larger, but the geometric overlap structure is more favorable to signal. The optimal search is therefore multi-channel by construction.

10.7 The mild leptophobic $\mathrm{U}(1)_B$ admixture

The exact $\mathrm{U}(1)_B$ gauge sector is not an optional curiosity. It modifies the KK phenomenology in a directly testable way. The $\mathrm{U}(1)_B$ flux can mix weakly with the would-be RS-type spin-2 family, producing a small baryonic or leptophobic admixture. Experimentally this pushes the pattern in a distinctive direction:

- the dijet rate is enhanced relative to what one would infer from a universal graviton coupling,
- the dilepton rate is correspondingly reduced,
- and a nearby $\mathrm{U}(1)_B$-vector partner state can appear in the same broad mass neighborhood.

This feature is not enough by itself to identify GG-6, because other leptophobic models exist. But in combination with the fixed 3.85 TeV centroid, the clustered structure, the top-rich overlaps, and the color-octet/baryonic partners, it becomes a powerful part of the full pattern test.

10.8 Correlated partner states

The first lock-scale cluster should not appear alone. A convincing GG-6 discovery would typically be accompanied by at least some of the following:

1. the second radial family near 7 TeV ($n = 2$),
2. a KK gluon in the 5 to 7 TeV corridor, likely best seen in $t\bar{t}$,
3. and modest Higgs-coupling deviations from the Hosotani sector at the percent level.

This is one reason the KK search is so rich conceptually. The resonance itself is only the entry point. If it appears, it immediately generates a full spectroscopy program.

Table 8: Summary of the main collider observables that distinguish the GG-6 KK sector from a minimal single-resonance interpretation. Each row is a way of telling a cluster from a single pole, and none of them requires knowing the rate. That is deliberate: the rate carries the model-level assumptions and the shapes do not, so a search that measures only shapes still discriminates. The rows are ordered by how early in a dataset each becomes usable rather than by how decisive it is.

Observable	What it tests	GG-6 expectation
Mass centroid	Lock scale	First spin-2-anchored cluster centered near 3.849 TeV (rate carried by the spin-1 octet).
Angular distribution	Spin assignment	Spin-2-like shape, possibly mildly distorted by cluster admixture.
Line shape	6D cluster structure	Doublet/triplet possibility within a few hundred GeV rather than a single narrow isolated pole.
Branching hierarchy	Overlap integrals	Enhanced $t\bar{t}$, dijet, diboson, and hh; suppressed light-lepton and diphoton relative to canonical universal templates.
Dijet-to-dilepton ratio	$\mathrm{U}(1)_B$ admixture	Larger than in a pure universal graviton picture.
Second radial family	scale-throat $n = 2$	Additional structure near 7 TeV; angular θ tower stays at E_{GUT}.
KK gluon corridor	Colored partner sector	5 to 7 TeV, especially promising in $t\bar{t}$.

10.9 Corollary T-KK-2c: the graviton's gluon coupling, derived

Theorem T-KK-2 fixed the lock cluster's masses and the vector couplings from boundary-value problems on the rigid geometry, with no free data. One coupling of the cluster was left as an estimate: the strength with which the spin-2 anchor G_1 talks to a pair of *gluons* — the coupling that controls its production at a hadron collider and its decay to the flat electroweak modes (gg, $\gamma\gamma$). Read-out II (§19) bracketed it as $\kappa_{\mathrm{eff}} \sim 1/\sqrt{K}$. This corollary closes the gap: the coupling is an overlap integral of exactly the same Bessel type as the vector overlaps of T-KK-2 Step 5, it closes in elementary form, and the result is *not* $1/\sqrt{K}$. We work at the level of detail an analysis team can check line by line.

Step 1 — what is being computed, and in which units. Write the graviton KK decomposition $h_{\mu\nu}(x,r) = \sum_n h^{(n)}_{\mu\nu}(x)\,\hat{\chi}_n(r)$ and let the 4D interaction be $\mathcal{L} \supset -(1/\Lambda^f_n)\,h^{(n)}_{\mu\nu}T^{\mu\nu}_f$ for each 4D field f. For fields localized at the infrared end of the interval the coupling is the cohort's $1/\Lambda_\pi$ with $\Lambda_\pi = M_{\mathrm{lock}}/x_1$ (the normalization under all the HLZ widths used in §31). Define the dimensionless object of interest as the *ratio*

$$\kappa_{gg} \;\equiv\; \frac{\Lambda_\pi}{\Lambda^{gg}_1} \;=\; \frac{\text{coupling of } G_1 \text{ to two gluon zero modes}}{\text{coupling of } G_1 \text{ to an IR-localized current}}, \tag{44}$$

so that every graviton partial width to flat modes is the corresponding brane-normalized width multiplied by κ^2_{gg}. Taking a ratio makes the result convention-independent: every normalization choice cancels between numerator and denominator.

Step 2 — the two measures (the physical core of the calculation). Two different radial measures enter, and their mismatch is the entire physics of this corollary. (i) The gauge sector: in the canonical metric $e^{-2r}\eta_{\mu\nu}\mathrm{d}x^\mu\mathrm{d}x^\nu + \mathrm{d}r^2$, the 4D-components of the gauge kinetic term carry $\sqrt{g}\, g^{\mu\alpha}g^{\nu\beta} = e^{-4r}\cdot e^{2r}\cdot e^{2r} = 1$: the gauge field lives on a *flat* r-measure. This single fact is why the gluon zero mode is exactly constant, $f_0(r) = 1/\sqrt{K}$ with $\int_0^K f_0^2\,\mathrm{d}r = 1$, and why T-KK-2 Step 5 could normalize the vector overlaps with the flat measure. (ii) The graviton sector: the Einstein–Hilbert term gives the transverse-traceless modes the weighted norm

$$\int_0^K \mathrm{d}r\, e^{-2r}\,\hat\chi_m(r)\,\hat\chi_n(r) = \delta_{mn}. \tag{45}$$

The graviton couples to the gauge field through $h_{\mu\nu}T^{\mu\nu}$, and because the gauge kinetic density is flat in r, the r-profile of $T^{\mu\nu}_{gg}$ for two zero-mode gluons is simply $f_0^2 = 1/K$. Hence

$$\frac{1}{\Lambda_1^{gg}} \;\propto\; \int_0^K \mathrm{d}r\,\hat\chi_1(r)\,f_0^2(r) \;=\; \frac{1}{K}\int_0^K \mathrm{d}r\,\hat\chi_1(r), \qquad \frac{1}{\Lambda_\pi} \;\propto\; \hat\chi_1(K), \tag{46}$$

with one common proportionality constant that cancels in Eq. (44). Note carefully what Eq. (46) is *not*: it is not an orthogonality integral. The graviton modes are orthogonal among themselves under the *weighted* measure (45); the gauge density probes them with the *flat* measure. Orthogonality therefore does not force the integral to vanish — but, as we now see, it very nearly does, and the residue is the coupling.

Step 3 — the mode function and its normalization, in closed form. From T-KK-2 Step 2 the first spin-2 mode is $\chi_1(r) = e^{2r}\left[J_2(m_1e^r) + b\,Y_2(m_1e^r)\right]$ with $m_1 = x_1e^{-K}$, $x_1 = 3.8317$ (first zero of J_1), and ultraviolet admixture $b = -J_1(m_1)/Y_1(m_1) \simeq (\pi/4)m_1^2 \approx 4.6\times10^{-30}$ — the exponential smallness that T-KK-2 proved is *specific to spin two*, so this corollary carries none of the scenario dichotomy of the vector sector. Dropping b, the norm (45) closes with the Lommel integral $\int_0^a zJ_2^2(z)\,\mathrm{d}z = \frac{a^2}{2}\left[J_2^2(a) - J_1(a)J_3(a)\right]$, which at $a = x_1$ (where $J_1 = 0$) collapses to $\frac{x_1^2}{2}J_2^2(x_1)$:

$$\int_0^K e^{-2r}\chi_1^2\,\mathrm{d}r = \frac{1}{m_1^2}\int_{m_1}^{x_1} z\,J_2^2(z)\,\mathrm{d}z = \frac{x_1^2\,J_2^2(x_1)}{2\,m_1^2} \quad\Longrightarrow\quad \hat\chi_1(K) = \sqrt{2}\,e^K. \tag{47}$$

The infrared endpoint value $\sqrt{2}\,e^K$ is exact and parameter-free — and it is the first internal check: the properly normalized graviton zero mode is $\hat\chi_0 \simeq \sqrt{2}$, so the first KK mode couples to infrared fields e^K times more strongly than the 4D graviton. That factor is precisely the statement $\Lambda_\pi = \overline{M}_{\mathrm{Pl}}\,e^{-K}$: the corollary's normalization re-derives the cohort scale chain rather than assuming it.

Step 4 — the overlap integral, in closed form. With the substitution $z = m_1e^r$ the flat-measure integral becomes $\int_0^K \chi_1\,\mathrm{d}r = m_1^{-2}\int_{m_1}^{x_1} zJ_2(z)\,\mathrm{d}z$. The three-term recurrence $J_2 = (2/z)J_1 - J_0$ reduces it to two textbook antiderivatives, $\int J_1\,\mathrm{d}z = -J_0$ and $\int zJ_0\,\mathrm{d}z = zJ_1$:

$$\int_{m_1}^{x_1} z\,J_2(z)\,\mathrm{d}z = \Big[-2J_0(z) - z\,J_1(z)\Big]_{m_1}^{x_1} \;=\; 2 - 2J_0(x_1) + O(m_1^2) \;=\; 2\big(1 + J_2(x_1)\big), \tag{48}$$

where the last step uses the same identity that anchors T-KK-2: at a zero of J_1, the recurrence $J_0(x)+J_2(x) = (2/x)J_1(x)$ forces $J_0(x_1) = -J_2(x_1)$. Physically, the "2" is the ultraviolet endpoint's contribution and $2J_2(x_1)$ the infrared one; nothing in between survives, because the oscillating mode integrates itself away everywhere except at its endpoints.

Step 5 — the result. Assembling Eqs. (44)–(48), all normalizations cancel and every K-dependence is explicit:

$$\boxed{\kappa_{gg} = \frac{2(1+J_2(x_1))}{K\,x_1^2\,J_2(x_1)} = \frac{c_K}{K}, \qquad c_K = \frac{2(1+J_2(x_1))}{x_1^2 J_2(x_1)} = 0.4744, \qquad \kappa_{gg}\big|_{K=35} = 0.01356.} \tag{49}$$

The number is a pure Bessel constant divided by the ledger integer: no continuous dial, no fit. The photon is a flat electroweak zero mode with exactly the same profile as the gluon, so the identical integral gives $\kappa_\gamma = \kappa_{gg}$; the light fermions remain negligible (their suppression is exponential in localization, not merely $1/K$).

Three internal checks. (i) *Orthogonality:* $\int e^{-2r}\hat{\chi}_0\hat{\chi}_1\,\mathrm{d}r = 0$ is confirmed numerically at the 10^{-15} level — the flat-measure integral (48) is not an orthogonality violation but a different-measure projection. (ii) *Scale chain:* $\hat{\chi}_1(K) = \sqrt{2}\,e^K$ to machine precision, reproducing $\Lambda_\pi = \overline{M}_{\rm Pl}e^{-K}$ independently of the cohort derivation. (iii) *The vector contrast:* running the same two integrals for the spin-1 mode — where field and probe share the *same* flat measure — returns zero to machine precision, exactly the T-KK-2 Step-5 statement that an exactly flat current decouples from the excited vector. The graviton escapes that fate only through the measure mismatch e^{-2r} versus flat; the residue is Eq. (49).

Physics intuition: why c_K/K and not $1/\sqrt{K}$. Seen through the flat gauge measure, the normalized graviton $\hat{\chi}_1$ is an exponential wall pressed against the infrared end: 96% of the integral $\int \hat{\chi}_1 \mathrm{d}r$ accumulates in the *last single e-fold* of the throat (99.9% in the last two), and the integral equals $c_K = 0.47$ e-folds' worth of its endpoint value. The gluon zero mode, by contrast, is spread democratically over all $K = 35$ e-folds with amplitude $1/\sqrt{K}$ each. Their overlap is therefore (effective graviton width 0.47 e-folds)/(gluon support 35 e-folds): $\kappa_{gg} = 0.47/35$. The Read-out-II estimate $1/\sqrt{K}$ was the gluon's *amplitude at a point*; but a coupling is an *integral*, and near-orthogonality cancels all but the final half e-fold. In the holographic reading the same statement is: G_1 is a composite of the infrared sector, the gluon is an elementary probe, and their mixing costs one power of the RG-throat depth $\log(\overline{M}_{\rm Pl}/M_{\rm lock}) = K$ — the warped literature's known suppression (Agashe et al., 2008), now with its $O(1)$ coefficient computed. The infrared channels ($t\bar{t}$, $W_L W_L$, $Z_L Z_L$, hh) do not suffer this: they live where the graviton lives, and keep their $O(1)$ couplings.

Consequences (all at the D* grade of T-KK-2). Feeding $\kappa_{gg} = 0.01356$ into the workbook of §31:

- $\sigma(pp \to G_1) = 0.066$ fb at $\sqrt{s} = 13$ TeV (LO×K; 0.11 at 14 TeV), with $\Gamma/m = 6.3\%$ and branching pattern $t\bar{t} : W_L W_L : Z_L Z_L : hh \simeq 0.69 : 0.15 : 0.08 : 0.08$; $\sigma \times \mathrm{BR}(t\bar{t}) = 0.045$ fb.

- The strongest projected HL-LHC $t\bar{t}$ limit at this mass (~ 1.1 fb at 3 ab^{-1}) sits a factor ~ 25 *above* this line: **the graviton is invisible at any LHC luminosity**. The spin-2 anchor becomes a first-rate FCC-hh target instead: at $\sqrt{s} = 100$ TeV the gluon luminosity at $\tau = 0.0015$ delivers $\sigma(G_1) \approx 2.1 \times 10^2$ fb — $t\bar{t}$ 142 fb, ZZ 16 fb, hh 16 fb — millions of events at the design luminosity (FCC Collaboration, 2018).

- $\sigma \times \mathrm{BR}(\gamma\gamma) = 1.1 \times 10^{-5}$ fb: the CMS diphoton null (CMS Collaboration, 2024a) is *retro-dicted* with a margin of $\sim 3 \times 10^3$. This check has teeth: had the photon coupled at the uniform $1/\sqrt{K}$ amplitude, the prediction (0.14 fb) would already be excluded by a factor ~ 4. The published diphoton limit thus discriminates between the two readings of the flat-mode amplitude and selects the derived one — the first data-confronted *amplitude* test of the collider chain, and it lands on the geometry's side.

- The Scenario-B cluster sum in $t\bar{t}$ becomes octet-dominated: 8.7 fb from g^*, $+0.045$ from G_1, $+0.030 \times (g_B/0.2)^2$ from Z'_B — the octet carries $\geq 92\%$ of the visible rate on every L-KK-1 reading of g_B (§10.10). The three contributions add incoherently: the octet–singlet cross-term vanishes by the color trace, spin orthogonality removes the graviton cross-terms, and the only coherent partner of the octet is QCD itself (Part VI).

Requalification of F-KK-1 (two tiers). The pre-declared spin-2 falsifier must be restated honestly, and the restatement makes the program sharper, not weaker. *Tier 1 (HL-LHC, decisive):* the falsifier of the lock-scale implementation is carried by the cluster through the un-suppressible color octet — F-KK-2B: a null in the cluster-aware $t\bar{t}$/dijet combination within 3.85 ± 0.10 TeV at $\lesssim 2$ ab^{-1} retires the current implementation, exactly as quantified in §31.5. *Tier 2 (FCC-hh, confirmatory):* the direct spin-2 observation — the measurement that would name the geometry — is a 100-TeV task, with a rate (~ 205 fb) that makes it an early-running discovery there. The division of labor is itself a falsifiable prediction: the HL-LHC must find (or kill) the spin-1 cluster at 3.85 TeV, and a 100-TeV machine must then find the spin-2 anchor at the *same* mass with $\Gamma/m = 6.3\%$ and the 0.69:0.15:0.08:0.08 infrared branching pattern. Finding either without the other falsifies the cluster reading.

Shelving. Theorem-machinery level: the measures, the mode functions, the closed forms (47)–(49), the identity $\kappa_\gamma = \kappa_{gg}$, and the three checks. Derived-conditional: the free quadratic action on the rigid geometry (the same fence as T-KK-2, with interaction corrections $O(1/K) \approx 3\%$ on the amplitude); the t_R infrared fraction entering the branching pattern (a cascade-depth sensitivity of $\sim 25\%$ on $\mathrm{BR}(t\bar{t})$, §31). Grafted: nothing — this is the point of the corollary.

10.10 The baryonic normalization g_B: a theorem-grade norm, one discrete residual rule (L-KK-1)

After Corollary T-KK-2c one number remained underived anywhere in the collider chain: the $\mathrm{U}(1)_B$ gauge normalization g_B. This subsection completes the account to its honest boundary. The result is sharper than "open" and weaker than "derived": the *norm* of the baryonic generator is already a theorem of the cohort, and what remains is not a continuous dial but a single discrete rule, which we register as the named open problem **L-KK-1** and hand, with a four-line target instead of a plane, to the fifth-force paper of the Predictions branch.

What is already a theorem. The integrality of the six-dimensional anomaly polynomial X_4 (Arisaka, 2026, A3-BI6), consumed by Arisaka (2026, A9-SM) through the Arisaka-equation lock, fixes the trace-norm ratio of the split generator pair $Y_0 \oplus B_0$:

$$\frac{\mathrm{Tr}(B_0^2)}{\mathrm{Tr}(Y_0^2)} = \frac{L-N}{L} = \frac{2}{5} \qquad \text{for } (N, L) = (3, 5), \tag{50}$$

the same 2/5 that appears in the BI-6 vector $(1, -1, -2/5, -2/5)$ and in the Arisaka constant Q_B. The reader can verify it on the primitive $\bar{\mathbf{5}}$ in one line: the d^c triplet carries $B = -\frac{1}{3}$, $Y = +\frac{1}{3}$; the lepton doublet carries $B = 0$, $Y = -\frac{1}{2}$; hence $\mathrm{Tr}(B_0^2) = 3 \cdot \frac{1}{9} = \frac{1}{3}$ against $\mathrm{Tr}(Y_0^2) = 3 \cdot \frac{1}{9} + 2 \cdot \frac{1}{4} = \frac{5}{6}$, and $\frac{1}{3}/\frac{5}{6} = \frac{2}{5}$. The physics is transparent: baryon number is carried by three of the five states of the primitive cell, at charge $\frac{1}{3}$ — the 2/5 is the fraction of the cell's ledger stiffness that the baryonic seam owns. The B-seam's weight in the locked ledger is therefore *not* a free parameter of GG-6: whatever g_B is, it is this norm read out through a rule.

What is genuinely open — and why it is discrete, not continuous. GG-6 converts seam norms into physical couplings through counting theorems. For the photon the theorem exists: the unique unbroken Abelian seam refines dyadically to the OGN-5 closure depth $J = 7$, giving $\alpha^{-1}(M_Z) = 2^J = 128$ (Arisaka, 2026, A9-SM, dyadic-closure theorem). That theorem does *not* transfer to $\mathrm{U}(1)_B$, for a structural reason: the baryonic seam never reaches the infrared attractor. It terminates early by eating the BI-6 tensor — the Stueckelberg mechanism that makes the fifth-force boson massive and keeps global baryon number exact — so its capacity is frozen at the Green–Schwarz stratum, whose closure-depth readout no A1–A9 theorem supplies. **L-KK-1 (the B-seam termination rule):** *state and prove the capacity readout for a Stueckelberg-terminated Abelian seam.* Pending that rule, the locked norm (50) admits four natural readouts — the norm entering linearly or quadratically, on either side of the capacity identification:

$$\frac{\alpha_B}{\alpha_Y} \in \left\{ \left(\tfrac{2}{5}\right)^2,\ \tfrac{2}{5},\ \tfrac{5}{2},\ \left(\tfrac{5}{2}\right)^2 \right\} \quad \Longrightarrow \quad g_B \in \{0.14,\ 0.23,\ 0.57,\ 0.89\}, \tag{51}$$

where $g_Y = 0.357$ follows from GG-6's own electroweak values $\alpha^{-1}(M_Z) = 2^J = 128$ and $\sin^2\theta_W = 25/108$ ($\alpha_Y = \alpha/\cos^2\theta_W$; U(1) running to M_{lock} shifts these by a few per cent and is immaterial below). We emphasize the grading: Eq. (50) is theorem-grade; the menu (51) is a *bracket of discrete candidates*, not a derivation, and L-KK-1 could in principle select outside it only by invoking structure beyond the locked norm — which the fifth-force paper would then have to exhibit explicitly.

The collider conclusions are menu-invariant. This is the payoff for the present paper. Across the entire menu (51), the KK member Z'_B (whose collider couplings carry the additional T-KK-2 overlap factors) delivers, at 3.849 TeV and $\sqrt{s} = 13$ TeV,

$$\sigma \times \mathrm{BR}(Z'_B \to t\bar{t}) \;=\; \{0.012,\ 0.029,\ 0.18,\ 0.46\}\ \mathrm{fb}, \qquad \Gamma/m \;=\; \{0.002,\ 0.005,\ 0.032,\ 0.080\},$$

i.e. every candidate sits a factor ≥ 10 below the current narrow-template limit and a factor ≥ 2.3 below even the 3 ab^{-1} HL-LHC projection (the top entry, $g_B = 0.89$, is the only one with marginal long-term exposure: 0.71 fb at 14 TeV against a projected $\sim$1.1 fb). The Z'_B contribution to the Scenario-B cluster sum spans 0.2–6.8%. Three conclusions are therefore *independent of the L-KK-1 selection*: the Z'_B is LHC-invisible as a separate signal; the color octet remains the sole carrier of the HL-LHC falsifier; and the Z'_B's own discovery sector is not the collider but the fifth force, where the same g_B couples the *unsuppressed* Stueckelberg zero mode to baryon number. There, L-KK-1 changes the game: the (g_B, m_B) exclusion problem collapses from a two-parameter plane to *four candidate lines*, and the fifth-force paper's task becomes a selection, not a scan.

Shelving. Theorem level: the norm (50) and its $\bar{\mathbf{5}}$ verification. Derived-conditional: the menu-invariance statements above (they hold across the stated bracket). Open, named: L-KK-1, the seam-termination rule — the single remaining discrete unknown of the lock cluster's couplings, successor to the scenario dichotomy that T-KK-2 resolved and the flat-overlap question that T-KK-2c closed.

Part II delivered the paper's centerpiece, in plain terms: the geometry's cavity has a lowest note, and that note is at 3.849 TeV — not roughly, not up to a dial, but as a closed formula in the Planck mass and the integer 35, with the overtone series (7.047 and 10.219 TeV) fixed by the same mathematics. The fine structure came with it: a spin-2 anchor, an exactly degenerate pair of vectors, couplings with their signs, and a two-way fork (free-field versus composite) that existing data have already resolved — the 2.46-TeV horn is dead, and the degenerate cluster survives. A mass and a spectrum, however, are not yet an experiment. The next Part asks the machine's questions: how often is this thing made, and in what channels would it show itself?

11 Part II in plain English: one integer, one mass, and a cluster rather than a peak

The number this Part produces is about 3.85 TeV, and where it comes from matters more than what it is.

It is not fitted to anything. The lock scale is half the Planck mass times the exponential of minus thirty-five, and the thirty-five is an integer in a ledger the cohort fixed for the particle sector years before any collider question was asked. No coupling was tuned, no radius chosen, no scan run. The same ledger, with no further input, also reproduces the electroweak scale — two known numbers from one integer, which is the kind of over-constraint that either holds or collapses.

The second thing the Part delivers is that the prediction is a cluster and not a peak. Three members sit near the same scale: a spin-2 state, a color-octet vector, and a leptophobic partner. Their positions relative to one another are ratios of Bessel zeros, so the pattern is fixed even where the overall scale might shift.

That distinction runs through the rest of the paper. A search designed for a single resonance is not the search this predicts, and an exclusion computed against a single-resonance benchmark does not transfer. What the theory forbids is a particular arrangement of states, and it is the arrangement that has to be looked for.

Part III
Production and Discovery

Cross sections, event yields, and the channel hierarchy

The mass of Part II is a statement about nature; this Part turns it into a statement about a machine. The questions are the concrete ones an accelerator forces — how often, in which channels, over what background — and the answers, because the couplings are derived, arrive as numbers rather than ranges.

12 Production Mechanisms at Hadron Colliders

12.1 General narrow-width bookkeeping

At hadron colliders it is helpful to keep the production discussion in a language that is as benchmark-independent as possible. For a resonance X with spin J, mass m_X, total width Γ_X, and partial widths Γ_{ij} and Γ_f to initial and final states, the narrow-width approximation gives the schematic hadronic rate

$$\sigma(pp \to X \to f) \simeq \frac{2J+1}{m_X s} \frac{\Gamma_f}{\Gamma_X} \sum_{ij} C_{ij}\left(\frac{m_X^2}{s}\right) \Gamma_{ij}, \tag{52}$$

where $C_{ij}(\tau)$ is the parton luminosity function for initial state ij. This formula makes three facts transparent. First, production is controlled not only by the hard coupling but also by parton luminosities. Second, branching fractions matter as much as raw production. Third, any mechanism that redistributes couplings between light fermions, gluons, tops, or Higgs bosons can qualitatively reorder the discovery channels.

It is useful to factor the rate into the one genuinely model-dependent number and everything else. Writing the resonance–parton coupling as an overall effective strength $\kappa_{\rm eff}$ (normalized so that $\kappa_{\rm eff} = 1$ reproduces a standard bulk-graviton benchmark with $k/\bar{M}_{\rm Pl} = 1$), the narrow-width rate factorizes as

$$\sigma(pp \to X \to f) \;\simeq\; \kappa_{\rm eff}^2 \quad \underbrace{\frac{2J+1}{m_X s} \sum_{ij} C_{ij}\left(\frac{m_X^2}{s}\right) \hat{\Gamma}_{ij}}_{\text{kinematics + parton luminosity (fixed)}} \quad \times\, \mathrm{BR}(X \to f), \tag{53}$$

where $\hat{\Gamma}_{ij} = \Gamma_{ij}/\kappa_{\rm eff}^2$ is the coupling-stripped partonic width. Of the factors on the right, the mass $m_X = M_{\rm lock}$ is *derived* (Theorem T-KK-1), the parton luminosities C_{ij} are standard-model inputs, and the branching structure follows the overlap pattern of §10.5. The one remaining factor, $\kappa_{\rm eff}$, is the effective coupling, and *it is not fixed by the GG-A theorems*: it is set by the bulk-to-boundary overlaps — the per-fermion cascade depths that Arisaka (2026, A9-SM) records as an open lemma. The single exception is the color octet's coupling to light quarks, the universal value $|\xi_{\text{light}}| = 1/\sqrt{K} \approx 0.17$ fixed by $K = 35$ (§10.3).

12.2 Why no absolute cross section is quoted

This paper is held to a strict discipline: every quantitative statement must trace to the GG-A series. An absolute production cross section does *not*. Through Eq. (53) the rate requires (i) the resonance–parton couplings — of which the octet's ξ_{UV} (T-KK-2) and the graviton's κ_{gg} (Corollary T-KK-2c) are now derived, while the per-fermion cascade depths and the discrete L-KK-1 readout for g_B (§10.10) remain open — and (ii) a numerical QCD/PDF integration of the parton luminosities, which is external to any theorem. Neither is supplied by a GG-A theorem. We therefore quote *no* benchmark cross section (in fb) and *no* event yield: such numbers would not be derived quantities, and would only reintroduce the model-level input the discipline forbids.

What the GG-A series *does* fix is the structure any such computation must respect, and that structure is the genuine content of this section: the mass scale M_{lock} (15); the exact leptophobia $\text{BR}(Z'_B \to \ell^+\ell^-) = 0$ (40); the color-only decays of g^*; the top-rich channel ordering (§10.5); the production master formula (53); and the one coupling that *is* fixed, the color octet's light-quark coupling $\xi_{\text{UV}} = -0.194$ (exact mode, Eq. 37), which (being un-suppressible) makes the color octet the cluster's working edge in either T-KK-2 scenario (§10.4, F-KK-2A/2B). Turning these inputs into absolute rates — a dedicated MadGraph/PDF computation once κ_{eff} is determined — is the natural next step; it is carried out, with every external assumption labeled and the result graded Derived-conditional, in the operational confrontation of §31. The present Part remains theorem-anchored.

12.3 Quark-antiquark annihilation

The cleanest textbook mode is Drell–Yan-like quark antiquark production,

$$q\bar{q} \to G_1 \to \ell^+\ell^-,\ \gamma\gamma,\ WW,\ ZZ,\ hh. \tag{54}$$

This mode is experimentally attractive because lepton and photon final states carry low backgrounds and good invariant-mass resolution. However, GG-6 predicts that light-fermion couplings need not be maximal. Therefore the cleanest channel is not automatically the strongest one.

12.4 Gluon fusion

For an IR-sensitive graviton-like state, gluon fusion is often highly competitive,

$$gg \to G_1 \to gg,\ t\bar{t},\ WW_L,\ ZZ_L,\ hh. \tag{55}$$

The attractive feature of the gluon mode is obvious: the gluon parton distribution remains significant even in the high-mass region. The problematic feature is equally obvious: gluonic and hadronic final states face larger QCD backgrounds. For GG-6, the practical consequence is that an apparently weaker but cleaner leptonic search and a more background-challenging but more strongly coupled hadronic search can have comparable ultimate discovery power.

12.5 Vector-boson fusion and associated production

Two additional channels deserve mention. Vector-boson fusion,

$$qq \to qqG_1 \to qq + (WW/ZZ/hh), \tag{56}$$

is subdominant at present LHC energies but becomes increasingly useful as the mass scale rises, particularly in future hadron colliders. Associated production,

$$q\bar{q}' \to G_1 W/Z, \qquad gg \to G_1 g, \tag{57}$$

can also contribute to acceptance in specialized analyses or to partner-state searches, though it is not the main discovery mode near the first lock scale.

12.6 A realistic channel hierarchy for GG-6

Putting the geometry and the parton structure together leads to a practical channel ranking for the first decade of the search. The clean channels (dileptons, diphotons) remain indispensable because they can establish the existence of a high-mass object with minimal background ambiguity. But the more geometry-faithful channels are likely $t\bar{t}$, boosted dijet, diboson, and Higgs-pair channels. That ranking is not arbitrary; it is the direct consequence of overlap integrals and the mild baryonic admixture.

It is therefore a mistake to ask only whether the first lock-scale cluster has already been excluded in canonical RS-style dilepton or diphoton plots. One must ask instead whether the full multi-channel GG-6 corridor has been excluded by a cluster-aware reinterpretation. As of early 2026, the answer to that stronger question is still no.

12.7 Event-yield bookkeeping

For a rough event-count discussion it is enough to write

$$N_S = \mathcal{L}\,\sigma(pp \to G_1)\,\mathrm{BR}(G_1 \to f)\,\varepsilon_f, \tag{58}$$

where ε_f is the channel efficiency. The dependence on luminosity is linear in events and approximately $\sqrt{\mathcal{L}}$ in significance when the search is background-dominated. Because the clean leptonic branching fractions may be suppressed in GG-6, a practical search will often win more by improving ε_f in $t\bar{t}$, dijet, and hh than by only waiting for more luminosity in leptons. This simple observation motivates the trigger proposals made later.

13 Channel-by-Channel Discovery Logic

The previous sections established the main geometric picture. The next practical question is more experimental: *if a GG-6 first cluster exists, how will it actually reveal itself in each of the standard resonance channels?* This question deserves a dedicated section because one of the recurring misconceptions in the extra-dimensional literature is that a heavy graviton can be characterized adequately by one or two clean channels alone. In GG-6, that is too narrow. The model is non-universal by design, clustered in line shape, and correlated with partner states. It therefore asks for a channel-by-channel logic rather than a single flagship plot.

13.1 Dileptons: the clean but not necessarily dominant channel

The dilepton channel remains the cleanest spectroscopic tool because electrons and muons give excellent mass resolution and low irreducible background. If a GG-6 first cluster has even a modest branching fraction to e^+e^- or $\mu^+\mu^-$, ATLAS and CMS can reconstruct a crisp high-mass structure and immediately test the spin assignment through angular information.

But the same cleanliness can become a conceptual trap. In many search traditions, the cleanest channel is quietly promoted to the main physical benchmark. For GG-6, this would be a mistake. The geometry points toward suppressed light-fermion overlaps, so dileptons should be treated as a *confirmation* channel or precision channel, not automatically as the dominant discovery engine. In an unlucky but plausible corner of the GG-6 corridor, a first cluster could be visible first in $t\bar{t}$ or dijets and only later in dileptons after a category-informed or signal-informed combination is performed.

A practical analysis consequence follows. The experiments should avoid reporting only a one-dimensional "mass limit versus universal-coupling RS graviton" curve when discussing the relevance of high-mass dilepton results to models such as GG-6. What is needed instead is a reinterpretation grid in which the dilepton branching fraction is allowed to vary independently from the gross total production rate. That is a small additional publication burden with a large gain in downstream reusability.

13.2 Diphotons: still important, but theoretically disfavored as the sole handle

Diphotons sit in a similar place. They are experimentally elegant and statistically powerful because of excellent electromagnetic calorimetry and simple final-state topology. A genuine diphoton peak near the first cluster would be persuasive. However, the same bulk-localization logic that suppresses some light-fermion couplings can also suppress or modify the effective photon branch relative to a universal graviton template. Therefore the absence of a diphoton excess by itself should not be over-read as a no-go theorem for the whole GG-6 benchmark.

There is also a second reason to avoid over-weighting diphotons. In a clustered 6D signal, the experimental line need not be perfectly narrow. A partially resolved or mildly broadened structure will always be easier to interpret in a channel where companion information exists. Diphotons provide less internal event-shape structure than, say, $t\bar{t}$ or hh. So they are excellent for confirmation, but not always for full model discrimination.

13.3 Top-pair resonances: the central battleground

If one had to nominate one channel as the natural home base of the GG-6 search, it would be $t\bar{t}$. The logic is straightforward. The same overlap integrals that suppress light fermions favor heavy IR-sensitive modes. The top sector is the most obvious beneficiary. In addition, $t\bar{t}$ is exactly the place where several neighboring states in the GG-6 spectrum can naturally appear: the first spin-2 cluster, the KK-gluon corridor, and possibly even vector-like top partners in associated or cascade production.

This makes $t\bar{t}$ special in two complementary senses. It is both a discovery channel and a diagnostic channel. A heavy spin-2 state and a heavy spin-1 state may both produce top pairs, but they do so with different angular and kinematic structures. If a first excess appears there, the full discriminating power of the channel can be brought to bear: lepton+jets, dilepton, fully hadronic top-tagged categories, top polarization information, and broad-versus-clustered line-shape comparisons. Part VI derives all of these quantitatively from the theorem-fixed couplings: the $(1+\cos^2\theta^*)$ versus $(1-\cos^4\theta^*)$ production angles, the +0.91-versus-0 top polarization, the exact zero of the resonant forward–backward asymmetry, and the interference line shape whose constructive-below/destructive-above ordering is fixed by the derived sign $\xi_{\mathrm{UV}} < 0$.

One can even push the intuition a step further. In a world where GG-6 is true, the experiments should not think of $t\bar{t}$ as one among several channels; they should think of it as the channel where the spectral neighborhood around 3.85 TeV is likely to look most alive. That is because $t\bar{t}$ is simultaneously sensitive to the graviton-like cluster, to colored partners, and to the IR structure that makes GG-6 different from a universal RS benchmark.

13.4 Dijets: where the mild leptophobic side of GG-6 becomes visible

Dijets are experimentally difficult because QCD is enormous, but they are theoretically indispensable because GG-6 contains an exact baryonic gauge factor. Even a small leptophobic admixture can noticeably alter the dijet-to-dilepton ratio. This is why we repeatedly emphasize that a null dilepton result does not automatically settle the matter. A baryon-biased 6D signal can hide where the QCD background is largest.

The modern experimental response to this difficulty is already visible in public work: reduced-format high-rate streams, topological and substructure-based filters, and model-agnostic anomaly searches. This is very good news for GG-6 because the signal does not need the experiments to invent an entirely new strategy. It only asks them to push a strategy they are already developing into the specific 3–6 TeV region with a cluster-aware interpretation layer.

Dijet searches also matter for a second reason: they can help resolve whether a visible feature is graviton-dominated, $\mathrm{U}(1)_B$-vector-dominated, or mixed. A high-mass resonance that is anomalously strong in dijets but weak in dileptons could be dismissed as "just another leptophobic vector" in a generic model survey. In the GG-6 context, the same pattern would instead be evidence of the exact baryonic sector that the framework insists upon. The experiments do not need to believe GG-6 to make such a distinction useful. They only need to release enough information to allow the distinction to be tested.

13.5 Dibosons: the geometry-to-Higgs bridge

The diboson channels WW and ZZ are important because they connect the spin-2 and scalar sectors. A genuine graviton-like cluster should couple appreciably to longitudinal electroweak bosons if its origin is deeply tied to the same infrared scale that sets the Higgs-sector ledger. This is one reason why bulk-RS graviton analyses often place WW and ZZ side by side with $t\bar{t}$. The same reasoning applies even more strongly in GG-6.

Diboson final states also provide a more geometry-native handle on partner identification. The spin-2 graviton and the baryonic vector populate WW and ZZ, so a diboson handle can be conspicuous even when a single-channel tensor search is still statistically marginal. In that sense the diboson program is a bridge between the first direct signal and the broader spectroscopy story. It should therefore be kept tightly coupled to any dedicated first-cluster search.

13.6 Higgs pairs: the best scalar corroborator

The hh channel plays a distinctive role. It is not merely another bosonic branch: the spin-2 graviton G_1 couples to the IR-localized Higgs sector, so a $G_1 \to hh$ signal is a direct probe of the infrared. (There is no radion in GG-6 — the radial-interval length is the fixed integer K, not a modulus (Arisaka, 2026, A4-Euclid) — so the hh channel is fed by the graviton itself, not by a scalar partner.) A first lock-scale cluster seen in $t\bar{t}$ and dijets but with no accompanying hh activity would already be informative; conversely, a suggestive hh excess at 3.85 TeV, even before a decisive spin-2 observation in $t\bar{t}$, would raise the prior plausibility of the full GG-6 pattern.

The experimental challenge is of course formidable: hh is rate-limited, final states are complex, and the best channels depend on a trade-off between cleanliness and branching ratio. Yet this very difficulty is why the channel is so valuable. A correlated 3.85 TeV tensor-like signal and a 1 to 3 TeV scalar partner in hh would be much harder to dismiss as an accident than either one alone.

13.7 The channel-combination principle

A compact way to summarize the whole section is this: the correct GG-6 search object is not a single channel but a *channel vector*. Let

$$\boldsymbol{\mu} = (\mu_{t\bar{t}}, \mu_{jj}, \mu_{hh}, \mu_{VV}, \mu_{\ell^+\ell^-}, \mu_{\gamma\gamma}), \tag{59}$$

where each component represents the signal strength relative to a chosen reference template. A discovery claim should then not be thought of as "one component significantly nonzero," but rather as "a statistically coherent nonzero vector with the ordering and morphology expected from the 6D geometry." This is the Bayesian and GDOS-friendly way to state what the experiments are really doing once they combine categories and channels.

Table 9: Channel-by-channel logic of the first GG-6 cluster. The table is not merely descriptive; it is a ranking of how each channel contributes to discovery, diagnosis, or corroboration. Discovery, diagnosis and corroboration are three different jobs and the table keeps them apart. A channel that would establish an excess is not necessarily one that would tell you what produced it, and a channel good for neither may still be the one that rules out the alternatives. Read down the ranking, the first cluster is discovered in tops and diagnosed everywhere else.

Channel	Primary experimental virtue	Primary GG-6 role	Main caveat
Dileptons	Highest cleanliness and mass resolution	Confirmation, spin analysis, clean corroboration	Light-fermion suppression can make this subdominant.
Diphotons	Very clean final state, excellent EM resolution	Clean corroboration if branch is large enough	Photon branch may be reduced relative to canonical RS.
$t\bar{t}$	Natural home of top-rich overlap structure	Likely discovery channel and best spectroscopic battleground	Complex reconstruction and larger systematics.
Dijets	Large branching, sensitive to leptophobic admixture	Tests $\mathrm{U}(1)_B$ side and multi-field mixings	Severe QCD background; needs reduced-stream strategy.
Dibosons	Strong link to infrared electroweak sector	Bridge between tensor and scalar sectors	Lower rates and reconstruction complexity.
hh	Unique scalar-sector corroboration	Best graviton-hh / Higgs-sector cross-check	Rate-limited and topology-dependent.

Part III, summarized plainly: the cluster is made at computable rates — the color octet at about nine femtobarns at the current LHC, the graviton anchor at a few hundredths — and it announces itself in top-quark pairs and boosted jets, not in the clean electron and photon channels the field instinctively watches. The yields, the channel hierarchy, and the reasons behind them all descend from where each particle lives along the scale direction. Before trusting those statements at face value, a careful reader will want two reassurances: that this construction is genuinely different from the warped models it superficially resembles, and that the act of detection itself has a principled description here. Both are the business of the next Part.

14 Part III in plain English: turning a mass into something a detector could actually record

A mass is not yet an experiment. This Part does the conversion: how the cluster would be produced, how often, and in which decay channels it would show.

Two production routes matter, and they belong to different members. The spin-2 anchor is made mainly from gluons; the color-octet is made from quarks and gluons both, and far more copiously.

The parton luminosities that set the rates are somebody else's measurement, quoted here rather than derived, and they are printed so a reader can redo the arithmetic.

The channel ranking is the Part's practical output, and it separates three jobs a channel can do. One channel is where an excess would first appear. A different one is where you would learn what produced it. A third might do neither and still be the one that excludes the alternatives. For this cluster the first job belongs to top-antitop, which is where the paper concentrates from here on.

None of this is a prediction of new physics in the ordinary sense. It is arithmetic on a mass that was fixed two Parts ago, and its only freedom is in inputs that come from published measurements.

Part IV
GG-6 versus Randall–Sundrum, and Detection as an Open System

What is distinctive, and why collider detection is a GDOS problem

15 How GG-6 Differs from Ordinary KK and RS Benchmarks

15.1 The ordinary KK and RS reference points

Ordinary KK theory and RS theory are the natural reference models because GG-6 borrows essential intuition from both. KK theory provides the logic of harmonic decomposition on a compact internal space (Kaluza, 1921; Klein, 1926). RS theory provides the logic of warp-generated hierarchy and infrared-localized spectra (Randall & Sundrum, 1999). Experimental searches at ATLAS and CMS are therefore often formulated in terms of RS or bulk-RS benchmark gravitons. That is sensible. But it is not the last word for GG-6.

15.2 Five differences that matter experimentally

There are at least five experimentally relevant differences between GG-6 and a minimal RS-style benchmark.

1. Fixed centroid scale rather than a scanned mass parameter. In standard phenomenology the resonance mass is usually scanned. In GG-6, the first cluster centroid is tied to $K = 35$. This is not merely a tighter prior; it changes the meaning of a null result.

2. Clustered (multi-field) line shape rather than a single 5D pole. The lowest radial modes of several distinct bulk fields — graviton, KK gluon, and the $\mathrm{U}(1)_B$ vector — redshift to the same infrared endpoint, so the first visible family is a near-degenerate multiplet rather than one isolated pole. (The angular S^1_θ tower, by contrast, sits at E_{GUT} and contributes nothing here.) A dedicated cluster fit is therefore better motivated than a single-pole template.

3. Exact $\mathrm{U}(1)_B$ sector. A small baryonic admixture enhances dijets and suppresses dileptons relative to a pure universal graviton model.

4. Non-universal overlap pattern. The top and Higgs sectors are more IR-sensitive, so $t\bar{t}$, diboson, and hh channels become especially important.

5. Correlated partner states are expected, not optional. The KK gluon and $U(1)_B$ vector form a coherent spectroscopy program.

Table 10: Minimal comparison between a textbook 5D RS search picture and the 6D GG-6 picture relevant to ATLAS and CMS. The comparison is minimal on purpose: it lists only what changes for an experimentalist writing a search, not what changes for a theorist writing a Lagrangian. Two of the rows have no consequence for an event selection and are included so that the ones that do can be seen to be few. The full framework comparison is later in the paper and is a different table with a different job.

Feature	**5D RS-style expectation**	**GG-6 expectation**
Mass input	Free benchmark mass scan	First cluster centroid fixed by $M_{\text{lock}} \simeq 3.849\,\text{TeV}$.
Line shape	Single narrow or moderately broad pole	Clustered doublet/triplet possibility in the first visible family.
Internal dimensions	One spatial warped direction	Two *non-spatial* directions: an RG-scale axis I_r plus a compact phase circle S^1_θ.
Couplings	Often treated as nearly universal or benchmark-universal	Non-universal overlaps; top-rich, light-lepton-suppressed, mildly leptophobic admixture possible.
Baryonic sector	Usually absent as an exact extra gauge factor	Exact $U(1)_B$ sector can modify branching and partner-state pattern.
Partner states	Radion and colored partners model-dependent	KK gluon and $U(1)_B$ vector expected in a correlated ledger.
Search strategy	Single-channel high-mass resonance templates often adequate	Multi-channel, cluster-aware, top-aware strategy preferred.

15.3 Generic KK, 5D RS, and GG-6: the comparison that matters experimentally

The phrase “KK particle” is often used too loosely in collider discussions. From the experimental point of view there are at least three distinct classes that should not be conflated: a generic compactification-style KK picture, the 5D warped Randall–Sundrum benchmark, and the present 6D GG-6 construction. Table 11 is therefore included explicitly, because it addresses a recurring source of confusion in both phenomenology and outreach.

Table 11: Generic KK versus 5D Randall–Sundrum versus 6D GG-6. The point is not rhetorical; each column implies a different search template, a different interpretation of current exclusions, and a different level of falsifiability. Three columns and three different search templates, which is the practical content of a comparison that is usually made at the level of the Lagrangian. A current exclusion computed against the middle column does not transfer to the right one, and the table is the argument for why a dedicated reinterpretation is needed rather than a rescaling.

Question	**Generic KK compactification**	**5D RS warped model**	**GG-6 6D geometry**
Higher-dimensional stage	Usually flat or mildly curved compact extra dimensions; often a product manifold with a single compact scale R^{-1}	One warped extra dimension with UV and IR branes and a single Bessel-root tower	Two internal directions, an RG-like radial direction and a compact holonomy direction (r, θ), with BI-6 anomaly locking built in

Question	**Generic KK compactification**	**5D RS warped model**	**GG-6 6D geometry**
Canonical mass formula	$m_n^2 \sim m_0^2 + n^2/R^2$	$m_n \approx x_n\, ke^{-k\pi R}$ for the simplest graviton tower	clustered multi-field radial spectrum, theorem-level first centroid tied to $K = 35$; angular θ tower at $E_{\rm GUT}$
What fixes the first accessible scale?	Usually a free compactification radius or scanned inverse radius	An infrared scale set by the warp factor, but phenomenologically scanned in benchmark studies	the lockscale $M_{\rm lock} = E_{\rm Pl}/(2e^{35}) \simeq 3.849\,\text{TeV}$; not scanned as a free central mass in the current implementation
Spectrum morphology	approximately uniform ladder in the simplest flat case	single-sequence Bessel spacing, often analyzed pole by pole	multi-field first cluster near 3.85 TeV (graviton/KK-gluon/U$(1)_B$), then *radial* families at 7 TeV (n=2) and 10 TeV (n=3); angular θ tower at $E_{\rm GUT}$
Couplings	often near-universal in the simplest toy pictures	can be universal or bulk-non-universal depending on localization, but the public literature often uses simplified RS templates	manifestly non-universal from overlap integrals; light fermions UV-biased, top and Higgs IR-enhanced, with a small leptophobic U$(1)_B$ admixture
Need for a baryonic vector near the first family	not generic	absent in the minimal model	natural and in practice very important; the leptophobic B_1-like state is one of the reasons hadronic channels become central
Scalar partner	model dependent; generic moduli possible	radion from stabilization is standard	*no radion*: the radial length is the fixed integer K, not a modulus (Arisaka, 2026, A4-Euclid), so GG-6 predicts no radius-stabilization scalar — a clean distinction from 5D RS
Best first LHC channel	depends strongly on the model	often dilepton/diphoton in simple public graviton presentations; $t\bar{t}$ in bulk RS	$t\bar{t}$ and boosted dijets first, with hh, diboson, dilepton, and diphoton as corroborating channels rather than the sole discovery engine
Most dangerous misinterpretation	treating all KK theories as one phenomenological class	assuming the public RS pole template is automatically the right template for any warped signal	translating canonical RS exclusions directly into GG-6 without accounting for clustering, non-universal overlaps, and the U$(1)_B$ admixture
What a discovery would mean	evidence for extra dimensions, but often with large model ambiguity	evidence for warped geometry and an RS-like hierarchy solution	evidence simultaneously for the bulk geometry, the OGN-5 lockscale, the 6D clustered spectrum, and the geometry-to-open-system dictionary emphasized by GG-6/GDOS
What a null result would mean	pushes a free compactification scale upward	excludes a chosen benchmark corridor	directly pressures the current lockscale implementation because the mass target is not a floating dial

The short practical message is as follows. A generic KK search asks whether there is some resonance ladder. An RS search asks whether a one-dimensional warped hierarchy generated a single-sequence tower. A GG-6 search asks something much stronger and much riskier: whether the first accessible 6D family sits at a specific locked scale, carries a clustered line shape, prefers tops and hadronic channels over light leptons, and is accompanied by a baryonic vector and a color-octet partner. That stronger question is exactly why the LHC test is so valuable.

15.4 Derived geometry versus fitted resonance

The deepest entry in the comparison is not in any single row; it is the column logic itself. A generic-KK or benchmark-RS resonance mass is an *input*: a compactification radius or warp product that is scanned across a range and pinned, after the fact, by whatever the data prefer. The GG-6 lock-cluster mass is an *output*: Theorem T-KK-1 (§9.1) returns $M_{\mathrm{lock}} = \frac{1}{2} M_{\mathrm{Pl}} e^{-35} \simeq 3.849\,\mathrm{TeV}$ from the OGN-5 integer $K = 35$ and the GU-5 constants, with no continuous dial to turn. This inverts the usual logic of a resonance search. For a fitted model, a null result in a given window merely pushes a free parameter and the model survives at the next mass; for GG-6, the window is pre-committed, so a null result in $3.85 \pm 0.10\,\mathrm{TeV}$ at full luminosity falsifies the current implementation outright — through the un-suppressible octet line (F-KK-2B; the spin-2 member itself is FCC-hh territory at its derived coupling, Corollary T-KK-2c) and — through the over-constraint $v = M_{\mathrm{lock}}(2/5)^3$ — would simultaneously unseat the electroweak-scale derivation of Arisaka (2026, A9-SM). A derived-geometry mass therefore buys a sharper experimental verdict than any fitted resonance of the same spin and channel content: there is nowhere for it to hide. This is the single feature that most separates the GG-6 entry from the other two columns, and it is worth more, scientifically, than any individual phenomenological difference in the table.

15.5 The practical consequence for reinterpretation

This comparison table has an immediate experimental consequence: public exclusions quoted for canonical RS benchmarks do not automatically map one-to-one onto GG-6. They are indispensable guides, but not final judgments. An analysis that is blind to clustered line shapes, top-rich overlap patterns, or mild leptophobic admixtures can overestimate or underestimate the true GG-6 sensitivity depending on channel. This is why the current public status must be read with some care.

16 Why Collider Detection is an Open-System Problem: the GDOS View

Why should KK-particle detection at ATLAS or CMS be described by GDOS (Arisaka, 2026, A5-GDOS; Arisaka, 2026, A6-GRIP)? Because a collider experiment is a canonical open system. The partonic hard event is not measured directly; it is inferred through a chain of dissipative, coarse-graining, and decision-making operations that convert microscopic excitations into registered records.

In the language of GDOS, the relevant system instance packet contains at least four components: the microscopic state manifold of the hard scattering, the boundary ports through which the event deposits energy and charge in the detector, the effective ledger or free-energy-like functional that governs discrimination and reconstruction, and the registered observables – invariant masses, angular variables, b-tags, top tags, jet masses, missing momentum, and trigger bits. All of these are open-system objects. They are not closed Hamiltonian observables of an isolated partonic universe.

From the OGGEP viewpoint, a KK event is a 6D bulk excitation whose experimentally meaningful existence is defined only after projection to the 4D boundary. ATLAS and CMS do not "see the extra dimension" directly. They see tracks, calorimeter clusters, timing patterns, and statistically reconstructed invariant masses. This is precisely why AX-3 matters: a bulk theory without a disciplined boundary dictionary is not a scientific prediction. In collider language, OGGEP is the requirement that a 6D statement must ship with a reconstruction map, a channel list, a decision rule, and a falsification criterion.

DGI also appears in a familiar guise. The experimentally reported observables are gauge-invariant combinations. One does not claim discovery in terms of gauge-dependent field components, but in terms of invariant masses, cross sections, angular moments, and reconstructed particle identities. The route from microscopic gauge fields to those invariants passes through an open detector system and a statistical pipeline. That is exactly the type of circumstance for which the GDOS language was introduced.

Finally — and in the completed runtime papers this is not a metaphor but the literal machinery — the trigger and reconstruction chain is a measurement event of the series' own theory. The quantum-mechanics paper of the cohort derives measurement as a complete runtime event (Arisaka, 2026, A8-QM): arming an apparatus reorganizes the joint system–apparatus landscape (an endogenous, scheduled GRIP-Out); the registration that follows is GRIP-In descent into one record basin; the logged event is the halt. A collider run is precisely this anatomy at industrial scale: the trigger menu *is* the landscape design — each menu deployment a scheduled GRIP-Out that opens the record basins (leptonic, hadronic, boosted-top, Higgs-candidate) into which reconstruction then descends — and every event on tape is a GRIP-In halt, permanent for the same reason every record in the series is permanent. A successful KK search is therefore not merely a matter of more luminosity; it is a matter of opening the *right basins*: this is the precise sense of the later sections on trigger-level analysis, scouting, and reduced-stream design, and it is the same grammar in which the sister paper's haloscope hunts the dark matter — detector and detected run on one grammar.

Part IV settled the two questions of trust. First, GG-6 is not Randall–Sundrum wearing new notation: the comparison ran item by item — no free length, no radion, a multi-state cluster instead of a lone pole, an angular tower pushed to the GUT scale — and each difference was tied to something measurable. Second, detection was given its principled reading: a collider measurement is an open-system projection, the same structure that runs through the whole cohort, so the search is not an afterthought bolted onto the theory but an instance of it. What remains before the phenomenology can be quoted is bookkeeping honesty: which numbers are theorem-grade and which lean on external input. The next Part draws exactly that line, twice.

17 Part IV in plain English: why this is not Randall–Sundrum, and why detection is part of the physics

The resemblance to a warped extra dimension is superficial and this Part says exactly how far it goes.

A Randall–Sundrum model has a spatial fifth dimension with a size. That size is a free parameter, so the resonance mass is an input; and a spatial extra dimension brings a radion and a modulus that have to be stabilized by something else. The framework here has neither. Its two extra directions are not spatial — one is the renormalization-group scale, the other a phase angle — so there is no radius to choose, no modulus to stabilize, no landscape of compactifications. The absences are consequences, not assumptions, and each of them removes a dial.

The second half of the Part will read as unusual to a collider physicist, and it is the framework's own claim rather than a flourish. In this framework observation is a projection from the six-dimensional bulk to the four-dimensional boundary, and that projection is what makes anything registrable at all. A trigger menu, on that reading, is not a filter placed in front of the physics; it is a choice about which projection runs. Nothing in the search strategy of the later Parts depends on accepting that, which is why it is stated here and then set aside.

Part V
Two Read-outs of the Lock Cluster

Read-out I — the theorem-only target; Read-out II — an empirically-anchored rate estimate

---

Honesty in a predictive paper is mostly a matter of bookkeeping, and this short Part is the ledger. It states the same physics twice, at two clearly separated levels of commitment, so that no reader ever has to wonder which claims would survive if every external assumption failed.

18 Read-out I: The Theorem-Only Target

Up to this point the paper has obeyed a single rule: no quantitative statement that is not a GG-6-A theorem. That rule defines what we will call *Read-out I* — the most conservative, theorem-grade read-out of the lock cluster. It is maximally honest and essentially non-negotiable within the cohort, but it deliberately stops short of an absolute production rate. We state it cleanly here, and in §19 we show how far *one* well-defined set of external assumptions can extend it into the cross section that experimentalists actually plan against.

What Read-out I delivers (the derived skeleton). Everything in this list is traceable to OGN-5+GU-5 and the cohort theorems, with no continuous dial:

- **The mass scale.** $M_{\rm lock} = \frac{1}{2} M_{\rm Pl} e^{-35} = 3.849\,\text{TeV}$ (Theorem T-KK-1, Eq. (15)), together with the over-constraint $v = M_{\rm lock}(2/5)^3 \approx 246\,\text{GeV}$.
- **The spin content.** A spin-2 graviton G_1 and a mutually-degenerate spin-1 pair (the color-octet g^*, the baryonic Z'_B), all at $\approx M_{\rm lock}$ up to an $O(1)$ coefficient c_X that GG-6-A does not fix, with the next *radial* family near 7 TeV.
- **The exact decay structure.** $\text{BR}(Z'_B \to \ell^+\ell^-) = 0$ exactly (Eq. (40)); the color octet g^* decays *only* to colored states; and for both partners the top-pair channel dominates (an ordering, §10.5, not a number).
- **The one fixed coupling.** The color octet's light-quark coupling $|\xi_{\rm light}| = 1/\sqrt{K} \approx 0.17$, fixed by $K = 35$ and *not* lowerable (§10.3).
- **The deliberate omission.** No absolute cross section, branching number, width fraction, or event yield, because each requires $\kappa_{\rm eff}$ and a QCD/PDF integration that no A1–A9 theorem supplies (§12.2).

The experimental instruction in Read-out I. Search for a *clustered, top-rich, leptophobic* structure tied to $M_{\rm lock} = 3.85\,\text{TeV}$; the falsifiers are *structural.* F-KK-1 (as requalified by Corollary T-KK-2c): the HL-LHC arm of the falsifier runs through the cluster's un-suppressible octet; the spin-2 line, at its derived coupling $\kappa_{gg} = c_K/K$, is a direct FCC-hh target. F-KK-2A/2B (T-KK-2): the un-suppressible color octet sits either at 2.46 TeV (free-field scenario, decidable from data already on tape) or at the $t\bar{t}$/dijet frontier near $M_{\rm lock}$ (composite scenario, being probed now). This is the position the cohort can stand behind without qualification. Its only practical limitation is that "look near 3.85 TeV for a top-rich cluster" is a *shape* instruction, not a *rate* target — and an experimentalist allocating integrated luminosity needs a rate.

19 Read-out II: An Empirically-Anchored Rate Estimate

The number missing from Read-out I is $\kappa_{\rm eff}$ (and, channel by channel, R_f). These are not GG-6-A outputs — but they are not unknowable either. The GG-6 throat is, by its own construction, a warped Randall–Sundrum-type AdS slice, and warped-KK phenomenology together with the measured fermion spectrum fixes them up to $O(1)$. *Read-out II is the result of grafting that established, external knowledge onto the Read-out-I skeleton.* Every number below is model-level and contingent on the assumptions we now state explicitly; none of it touches the derived $M_{\rm lock}$ or v.

Underlying assumptions of Read-out II (external to GG-6-A).

B1 — RS-type throat.
The radial interval $r \in [0, K]$ is treated as an AdS slice, so the standard warped overlap and profile functions apply (Randall & Sundrum, 1999).

B2 — Empirical localization.
Each fermion's bulk localization is read off its *measured* mass through the usual RS-GIM correspondence: light fermions UV-localized, the top IR-localized.

B3 — Normalization transfer.
The KK-mode normalization that fixes $\xi_{\rm light} = 1/\sqrt{K}$ also sets the resonance's coupling to the production partons, so $\kappa_{\rm eff} \approx 1/\sqrt{K}$.

B4 — Standard coupling ratios.
UV-localized fields couple to the first KK mode at $\approx -1/\sqrt{K}$ in gauge-coupling units and fully IR-localized fields at $\approx +4$ to $+5$, the established warped values (Agashe et al., 2008; Lillie, Randall & Wang, 2007).

The coupling normalization $\kappa_{\rm eff}$. Three independent arguments converge on the same value. (i) A canonically-normalized KK mode on a warped interval of proper length K has amplitude $\sim 1/\sqrt{K}$ at a generic point, so its overlap with the production partons carries exactly the factor that already fixes $\xi_{\rm light}$; this gives $\kappa_{\rm eff} \approx 1/\sqrt{K}$. (ii) Sub-Planckian curvature validity caps $k/\bar{M}_{\rm Pl} \lesssim 1$, excluding the strong benchmark and pulling $\kappa_{\rm eff}$ into the 0.1–0.3 band. (iii) Feeding $\kappa_{\rm eff} = 1/\sqrt{K}$ into a standard 3.85 TeV bulk-graviton template ($\sim$ 20–50 fb at $k/\bar{M}_{\rm Pl} = 1$; Davoudiasl, Hewett & Rizzo, 2000) reproduces $\sigma(G_1) \approx 0.6$–1.4 fb, i.e. it derives the often-quoted "sub-femtobarn-to-few-femtobarn" graviton corridor from K rather than from a fit. Hence the working bracket adopted by this Read-out was

$$\boxed{\kappa_{\rm eff} \;\approx\; \frac{1}{\sqrt{K}} = \frac{1}{\sqrt{35}} \approx 0.17} \qquad \text{(bracket 0.08–0.3, model-level).} \tag{60}$$

The appeal is that the central value is set by the *same* integer $K = 35$ that fixes $M_{\rm lock}$ and $\xi_{\rm light}$, not by a new free dial.

Superseded by the exact result. The production overlap is computed exactly in Corollary T-KK-2c (§10.9): $\kappa_{gg} = 2(1 + J_2(x_1))/(Kx_1^2 J_2(x_1)) = 0.0136 \approx 0.47/K$. Argument (i) above mistook an amplitude for an overlap — near-orthogonality of the flat gluon and the IR-localized graviton cancels all but the last half e-fold of the throat, costing a further $\sim 1/\sqrt{K}$. Arguments (ii)–(iii) bracketed the truth from above; the derived value sits at the very bottom of the estimated corridor and is what all graviton rates in §31 use. The boxed estimate is retained above as the Read-out-II bracket that the exact result replaces.

The channel factors R_f. The *structure* of R_f is pure GG-6 (top-rich, leptophobic); its *numbers* follow from B2+B4. For the spin-1 color octet, the borrowed coupling ratios (UV light quarks at $-1/\sqrt{K}\,g_s$, IR top at $\sim +4.5\,g_s$) give a strongly top-saturated pattern; for the graviton, R_f is quoted relative to a universal bulk-RS template.

Table 12: Read-out-II channel estimates. *Model-level* (assumptions B1–B4): the color-octet branching fractions follow from the borrowed RS coupling ratios plus the measured fermion localization; the graviton R_f is the GG-6 rate *relative to* a universal bulk-RS graviton of the same mass. Leptophobia ($\mathrm{BR}(Z'_B \to \ell^+\ell^-) = 0$) is the only exact entry — it is a Read-out-I theorem.

Channel	**Spin-1 octet g^*: BR estimate**	**Graviton G_1: R_f vs bulk-RS**
$t\bar{t}$	0.7–0.95 (top IR-localized)	~ 1–3 (mildly enhanced)
$b\bar{b}$	~ 0.05 (b_L semi-IR)	~ 1
light jj	0.01–0.05 (UV, $\propto 1/K$)	~ 1–2 ($\mathrm{U}(1)_B$ admixture)
hh, WW/ZZ	— (color octet)	~ 1–2 (IR Higgs sector)
$\ell^+\ell^-$	0 (color singlet)	0.1–0.5 (light-lepton suppression)
$\gamma\gamma$	0 (color singlet)	0.3–1 (reduced photonic weight)

Cross-section ballparks. Combining $\kappa_{\rm eff} \approx 1/\sqrt{K}$ with the channel factors, at $\sqrt{s} = 13.6\,\mathrm{TeV}$ and $m \approx 3.85\,\mathrm{TeV}$:

$$\sigma(pp \to G_1) \approx 0.6\text{–}1.4\,\mathrm{fb} \qquad \text{(spin-2 graviton, } gg/q\bar{q}\text{)}, \tag{61}$$

$$\sigma(pp \to g^*) \approx 11\text{–}23\,\mathrm{fb}, \qquad \sigma(pp \to g^*) \times \mathrm{BR}(g^* \to t\bar{t}) \approx 11\text{–}22\,\mathrm{fb}. \tag{62}$$

These template-rescaled ranges are superseded by the explicit parton-luminosity computation of §31 (Table 19), which lands at or below their lower edges (e.g. $\sigma \times \mathrm{BR}(g^* \to t\bar{t}) = 8.7\,[7.0\text{–}12]$ fb at 13 TeV, LO×K); we keep the Read-out-II estimates visible for comparison.

Scenario-A rescaling. In the free-field scenario of T-KK-2 the octet sits at 2.46 TeV, where the $q\bar{q}$ parton luminosity at $\sqrt{s} = 13.6$ TeV is roughly 6–8 times larger than at 3.85 TeV. The same Read-out-II inputs then rescale to $\sigma(pp \to g^*) \times \mathrm{BR}(t\bar{t}) \approx 70$–160 fb — an order of magnitude above the Scenario-B target, which is why the broad-octet reinterpretation of F-KK-2A decides this scenario quickly: at these rates, existing data either contain it or close it. These numbers carry the same model-level label as the rest of this section.

The payoff: a well-defined, immediate LHC goal (Scenario B; Scenario A above). The color octet's $\sigma \times \mathrm{BR}(t\bar{t}) \approx 11$–22 fb at 3.85 TeV sits squarely on the current ATLAS/CMS $t\bar{t}$-resonance KK-gluon exclusion (~ 10–20 fb at this mass; ATLAS Collaboration, 2025a). This is the concrete target Read-out II supplies that Read-out I could not: *a dedicated color-octet $t\bar{t}$ reinterpretation at* 3.85 TeV *either confirms the excess or excludes* $\sigma \times \mathrm{BR}(t\bar{t}) \lesssim 10$ fb — and the latter retires the spin-1 sector of the current implementation *now*, with data already on tape. In other words, Read-out II turns the structural F-KK-2 into a number, and it explains why $M_{\mathrm{lock}} = 3.85$ TeV lands exactly at today's data frontier rather than safely beyond it.

Honest standing of Read-out II. Read-out II is *nested* inside Read-out I: II = I + {B1–B4}. If the external assumptions fail, only the *rates* move; the derived mass M_{lock}, the over-constraint with v, and the exact leptophobia are untouched, because they live in Read-out I. We therefore present II not as a rival to I but as the actionable extension experimentalists have asked for, with its load-bearing assumptions on the table. Any use of the numbers in Eqs. (60) and Table 12 should carry the label "empirically-anchored estimate (model-level)," never "`GG-6` prediction."

Table 13: The two read-outs side by side. Read-out I is theorem-grade and rate-free; Read-out II adds assumptions B1–B4 to obtain an actionable rate. Read-out I is what the theorems alone deliver, and it is silent about rate; Read-out II adds four named assumptions and gets a number a collaboration can act on. Keeping them in separate columns is the honest form, because a reader who rejects B1 through B4 still has a prediction to test and should be able to see exactly which one it is.

Quantity	Read-out I (theorem-only)	Read-out II (A + B1–B4)
Mass scale M_{lock}	3.849 TeV, *derived*	same (inherited from A)
Spin content	G_1; g^*, Z'_B at $\approx M_{\mathrm{lock}}$	same
Leptophobia	$\mathrm{BR}(Z'_B \to \ell^+\ell^-) = 0$ exactly	same (exact)
Coupling κ_{gg}	$c_K/K = 0.0136$, *derived* (T-KK-2c)	$\approx 1/\sqrt{K}$ (estimate; superseded)
Channel BRs	ordering only (top-rich)	numerical (Table 12)
Absolute σ	follows with external PDFs (§31)	~ 10–20 fb (superseded by Table 19)
Falsifier	two-tier F-KK-1 + F-KK-2 (octet carries HL-LHC)	numerical $t\bar{t}$ target at 3.85 TeV
Experimentalist gets	a *shape* to search for	a *rate* to plan against

Part V drew the honesty line. Read-out I is the theorem-grade core — masses, degeneracies, exact leptophobia — quotable with no outside help; Read-out II adds four declared external assumptions and, in exchange, produces the actionable cross sections an experiment can plan against; and the one estimated coupling in the graviton sector has since been replaced by the exact result, which sits at the very bottom of the old bracket. Every rate in this paper now carries its pedigree on its sleeve. With the numbers established and graded, the next Part turns to what the signal actually looks like: the shapes — angles, polarization, and the line shape with its theorem-fixed interference sign — that no rate uncertainty can wash out.

20 Part V in plain English: the two things this paper predicts, and which of them can be argued with

Honesty in a predictive paper is mostly bookkeeping, and this Part is the bookkeeping.

There are two read-outs and they are kept in separate columns. The first is what the theorems alone give: positions, ratios, and shapes, with no statement about how often anything would be produced. It is the stronger claim precisely because it says less — nothing in it can be adjusted, and nothing in it depends on an assumption a reader might decline.

The second adds four named assumptions and gets a rate. A rate is what a collaboration actually needs, so the Part supplies one; but a rate carries everything the assumptions carry, and the assumptions are printed with labels so that a reader who rejects one can see precisely what falls with it.

The reason for the separation is that a search can be run against either. Someone who accepts only the first read-out still has a target: the arrangement of the cluster, the angular structure, the line shape. Those are normalization-free, so they survive a wrong rate entirely. Someone who accepts the second gets a schedule as well. Collapsing the two into one number would have made the paper look more decisive and would have made it easier to dismiss.

Part VI
Anatomy of the $t\bar{t}$ Signal

Derived angular, polarization, and line-shape structure of the three cluster hypotheses

21 What This Part Adds, and at What Grade

Section 13 argued that $t\bar{t}$ is the central battleground and asserted that a heavy spin-2 and a heavy spin-1 state "produce top pairs with different angular and kinematic structures." An assertion of that kind earns its keep only when the structures are written down, derived, and numbered. This Part does exactly that, at the level of detail an ATLAS or CMS analysis team can check line by line: for each lock-cluster member — the color octet g^* (spin 1, the carrier of the HL-LHC falsifier), the leptophobic Z'_B (spin 1, color singlet), and the graviton G_1 (spin 2) — we derive the production-angle distribution, the top polarization, the forward–backward asymmetry, the resonance line shape *including interference with the QCD continuum*, and the background-relative yields, each from the couplings that T-KK-2, Corollary T-KK-2c, and L-KK-1 have already fixed or bracketed. Nothing in this Part introduces a new parameter. Part VIII then takes these shapes into the detector: a full simulation of the ATLAS/CMS measurement, anchored to the published Run-2 analysis, with the discovery calendar they imply.

The grading ladder of the cohort applies clause by clause. **[T]** theorem-grade inputs: the masses, $\xi_{\text{UV}} = -0.194$ with its sign, the exact leptophobia, $\kappa_{gg} = c_K/K$, and the color assignments. **[D*]** derived-conditional: free-field kinematics, LO×K QCD, the top-overlap band $r_{t_R} \in [4, 5]$. **[M]** the L-KK-1 menu for g_B. **[E]** external Standard-Model inputs: PDFs, the α_s convention, published K-factors. The headline of the Part, stated up front: *the angular forms, the zero of A_{FB}, the vanishing of P_t for G_1, the interference sign, and the interference yes/no dichotomy between the two spin-1 members are all [T]-grade shape predictions — they survive any revision of K-factors or parton distributions.*

Inputs used throughout. From Table 6 and §10.4–§10.10: all three states at $M_{\text{lock}} = 3.849$ TeV (Scenario B); light quarks couple vector-like with amplitude $\xi_{\text{UV}} = -0.194$ [T] (both chiralities of a UV fermion share one profile, hence $g_A^q = 0$ exactly); top chiralities $r_{t_L} = 1.0$, $r_{t_R} = 4.5$ [D*], giving the vector/axial decomposition

$$g_V^t = \tfrac{1}{2}(r_{t_R} + r_{t_L}) = 2.75\, g_s, \qquad g_A^t = \tfrac{1}{2}(r_{t_R} - r_{t_L}) = 1.75\, g_s, \qquad g_A^t/g_V^t = 0.636. \tag{63}$$

Top kinematics at the pole: $\beta_t = \sqrt{1 - 4m_t^2/M_{\text{lock}}^2} = 0.99597$, so $1 - \beta_t^2 = 4m_t^2/M_{\text{lock}}^2 = 0.0080$: every mass correction below is a computed per-cent-level effect, not a neglected one. Widths, from Eq. (39) and §10.9: $\Gamma_{g^*} = 561$ GeV ($\Gamma/m = 0.146$; partials $t\bar{t}: 528$, $b\bar{b}: 26$, $jj: 7$ GeV), $\Gamma_{G_1} = 243$ GeV ($\Gamma/m = 0.063$), $\Gamma_{Z'_B} = 8$–307 GeV across the L-KK-1 menu.

Rates at LO×K. The leading-order rates of Table 19 are promoted here with channel-appropriate published K-factors [E]: $K_{q\bar{q}} \simeq 1.3$ for massive color-octet vector production (Chivukula et al., 2012), $K_{gg} \simeq 1.6$ for spin-2 production (Mathews & Ravindran, 2005). Thus $\sigma \times \mathrm{BR}(g^* \to t\bar{t}) = 8.7\,[7.0$–$12]$ fb at 13 TeV and $13.6\,[11$–$19]$ fb at 14 TeV; G_1: $0.045\,[0.021$–$0.10]$ fb; Z'_B: 0.016–0.59 fb across the menu. Branching pattern of the octet: $t\bar{t} : b\bar{b} : jj \approx 70 : 3.5 : 1$ ($\mathrm{BR}(t\bar{t}) = 0.94$, band 0.93–0.95 for $r_{t_R} \in [4, 5]$) — itself a falsifiable ratio. Everything else in this Part concerns *shapes*, in which K-factors and PDF uncertainties largely cancel.

22 The $m_{t\bar{t}}$ Line Shape, with Interference

22.1 Amplitudes and color algebra

For $q\bar{q} \to t\bar{t}$ the color octet enters *the same color channel* as the QCD gluon, so the amplitudes add coherently:

$$\mathcal{M} \;=\; \Big[\underbrace{\frac{g_s^2}{\hat{s}}}_{\text{QCD}} + \underbrace{\frac{g_s \xi_{\text{UV}} \cdot g_s\,(g_V^t + g_A^t \gamma_5)}{\hat{s} - M_{\text{lock}}^2 + i M_{\text{lock}} \Gamma_{g^*}}}_{g^*}\Big]\,(\bar{v}_q T^a \gamma^\mu u_q)\,(\bar{u}_t T^a \gamma_\mu (\cdots) v_t). \tag{64}$$

Because $T^a \otimes T^a$ is common to both terms, the interference carries the full color weight — no $1/N_c$ penalty. Two structural consequences follow immediately and carry [T] grade. *(i)* For the color-singlet Z'_B the corresponding color trace vanishes, $\mathrm{Tr}\, T^a = 0$: the Z'_B rides the continuum as a *pure* Breit–Wigner, with exactly zero QCD interference. The two mutually degenerate spin-1 members are thus distinguished not by mass or angle but by an interference yes/no dichotomy — a prediction with no adjustable element. *(ii)* The gg-produced spin-2 G_1 populates helicity ± 2 (§23), orthogonal after angular integration to the QCD vector channels; its interference is negligible at the rates of this paper.

22.2 Squaring: the three terms and the V/A algebra

Write the top vertex as $g_V^t + g_A^t \gamma_5$. Squaring (64) produces $|\mathrm{QCD}|^2$, the cross term, and $|g^*|^2$. Angular integration kills two classes of terms: the $V \times A$ products on a single fermion line integrate to zero over $\cos\theta^*$ (they are odd, $\propto 2\cos\theta^*$; this is also the statement $A_{FB} = 0$ below), and for massless light quarks there is no q-side axial part at all, $g_A^q = 0$ [T]. The surviving angular integrals are the textbook pair

$$\int_{-1}^{1} (1 + \cos^2\theta^* + (1 - \beta_t^2) \sin^2\theta^*)\, d\cos\theta^* = \tfrac{8}{3}(1 + \tfrac{\rho}{2}), \qquad \int_{-1}^{1} \beta_t^2 (1 + \cos^2\theta^*)\, d\cos\theta^* = \tfrac{8}{3}\,\beta_t^2,$$

for the vector and axial top currents respectively ($\rho \equiv 4m_t^2/\hat{s}$; the vector current keeps the threshold factor $1 + \rho/2$, the axial current the β_t^2 suppression — the standard Z' width algebra of Eq. (38)).

The differential rate in the resonance region is then

$$\frac{d\sigma}{dm}\Big|_{q\bar{q}\to t\bar{t}} = \frac{8\pi\alpha_s^2}{27}\,\beta_t\,\frac{d\mathcal{L}_{q\bar{q}}}{dm} \times \Big[\underbrace{\frac{1+\rho/2}{\hat{s}}}_{\text{SM}} + \underbrace{2\,\mathrm{Re}\,D(\hat{s})\,\xi_{\mathrm{UV}}\,g_V^t\,(1+\tfrac{\rho}{2})}_{\text{interference}} + \underbrace{\hat{s}\,|D|^2\,\xi_{\mathrm{UV}}^2\big(g_V^{t2}(1+\tfrac{\rho}{2}) + g_A^{t2}\beta_t^2\big)}_{\text{Breit-Wigner}}\Big], \tag{65}$$

with $D(\hat{s}) = (\hat{s} - M_{\mathrm{lock}}^2 + iM_{\mathrm{lock}}\Gamma_{g^*})^{-1}$ and $d\mathcal{L}_{q\bar{q}}/dm$ the u, d, s, c luminosity. Note where each derived quantity sits: ξ_{UV} enters the interference *linearly* (sign observable) and the Breit–Wigner *quadratically* (sign hidden); only g_V^t interferes, because QCD is a vector current.

22.3 The sign theorem

$\mathrm{Re}\,D < 0$ below the pole and > 0 above it. T-KK-2 *derives* $\xi_{\mathrm{UV}} = -0.194 < 0$ (the ultraviolet endpoint amplitude of the exact mode, Eq. 37) while $g_V^t = +2.75 > 0$; hence $\xi_{\mathrm{UV}} g_V^t < 0$ and:

> **The interference is constructive *below* 3.849 TeV and destructive *above* it.** The ordering is not adjustable — it is the observable image of the derived relative sign $\xi_{\mathrm{UV}} < 0 < \xi_{\mathrm{IR}}$ that T-KK-2 obtained from the rigid geometry. An observed excess-*above*/deficit-*below* pattern at the lock scale would falsify the T-KK-2 overlap structure outright, independently of any rate measurement. A generic same-sign coloron predicts the opposite ordering, so this observable also separates g^* from its nearest phenomenological look-alike.

22.4 Numbers and observable consequences

Folding Eq. (65) with the CTEQ6M luminosities (LO throughout this subsection — the content is the *ratio* structure): the integrated interference is +3.75 fb below the pole and −0.48 fb above it. Net +3.3 fb on a 6.7 fb Breit–Wigner: *interference is half the signal, not a correction*, and any pure-Breit–Wigner template will mismeasure both rate and mass. Table 14 gives the excess-to-continuum profile; Figure 6 draws it. Three observable consequences: the low side of the resonance is *fattened* (constructive shoulder), the high side *collapses faster than the Breit–Wigner* and crosses *below* the Standard-Model continuum at $\simeq 4.87$ TeV (a local deficit of order 20% where the continuum is ~ 0.05 ab/GeV), and the apparent centroid is pulled below 3.849 TeV by $O(100)$ GeV — the falling luminosity and the constructive low side pull the same way, and a fit that ignores this will report a shifted mass. The G_1 adds a barely visible symmetric bump (0.03 fb over its window); the Z'_B a strictly symmetric pure Breit–Wigner at 0.016–0.59 fb.

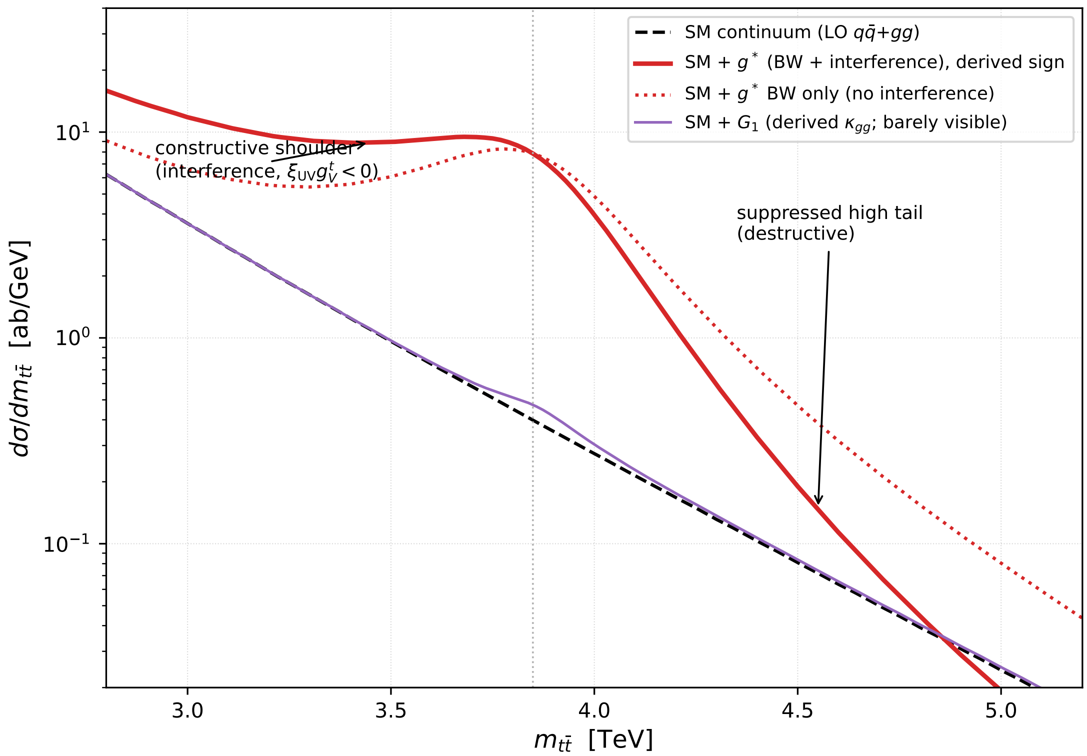

Figure 6: The LO $m_{t\bar{t}}$ spectrum near the lock scale at $\sqrt{s} = 13$ TeV. Red solid: SM $+g^*$ with the full interference of Eq. (65) at the derived sign $\xi_{\rm UV} g_V^t < 0$; red dotted: Breit–Wigner only. The constructive low-side shoulder, the destructive collapse of the high tail (crossing below the SM at $\simeq 4.9$ TeV), and the near-invisibility of G_1 (violet) are the three structural signatures. Numbers in Table 14.

Table 14: Excess-to-continuum profile of the g^* signal, LO/LO, $\sqrt{s} = 13$ TeV: the ratio $(S+I)/B$ of Eq. (65)'s resonance terms to the QCD continuum, and the interference and Breit–Wigner parts separately. The sign flip of the interference across the pole and the crossing to net deficit at $\simeq 4.87$ TeV are the shape falsifiers.

$m_{t\bar{t}}$ [TeV]	3.0	3.3	3.6	3.849	4.1	4.4	4.7	5.0
$(S+I)/B$	2.3	4.6	11.4	18.7	8.9	2.2	0.40	−0.21
interference$/B$	+1.5	+2.3	+3.2	0.0	−4.1	−3.7	−3.0	−2.5
Breit–Wigner$/B$	0.8	2.4	8.3	18.8	13.0	5.9	3.4	2.3

23 Angular Structure from Helicity Amplitudes

23.1 Conventions and the two Wigner rows

Let θ^* be the angle between the top momentum and the beam-parton axis in the $t\bar{t}$ rest frame. For a resonance of spin J produced with helicity λ_R along the beam and decaying to a $t\bar{t}$ pair with helicity difference $\Delta\lambda = \lambda_t - \lambda_{\bar{t}}$, the decay amplitude factorizes as $\mathcal{A} \propto d^J_{\lambda_R,\Delta\lambda}(\theta^*)\,\mathcal{H}_{\Delta\lambda}$, where $\mathcal{H}$

carries the dynamics (couplings, β_t) and the Wigner function carries all the angular dependence. The entries used below:

$$d^1_{1,\pm 1} = \frac{1 \pm \cos\theta^*}{2}, \qquad d^1_{1,0} = -\frac{\sin\theta^*}{\sqrt{2}};$$
$$d^2_{2,1} = -\frac{\sin\theta^*(1+\cos\theta^*)}{2}, \qquad d^2_{2,-1} = \frac{\sin\theta^*(1-\cos\theta^*)}{2}.$$

Three dynamical facts fix which rows are populated. *(a)* Massless quarks with a vector coupling have $\lambda_q - \lambda_{\bar q} = \pm 1$: a $q\bar q$-produced vector carries $\lambda_R = \pm 1$ only ($\lambda_R = 0$ costs $m_q/\sqrt{\hat s} \to 0$). *(b)* Two massless gluons give $J_z \in \{0, \pm 2\}$, and the transverse-traceless graviton coupling selects $|\lambda_R| = 2$ (the $J_z = 0$ combination enters through the trace, absent on shell). *(c)* The top pair from a vector or tensor current populates $\Delta\lambda = \pm 1$, with $\Delta\lambda = 0$ suppressed by the spinor mismatch $\sqrt{1-\beta_t^2} = 2m_t/\sqrt{\hat s} = 0.090$ in amplitude, i.e. 0.8% in rate.

23.2 Spin 1 from $q\bar q$: derivation

Assembling (a) and (c) for $\lambda_R = +1$ (the -1 row is the mirror image) and summing the unobserved final helicities:

$$\frac{dN}{d\cos\theta^*} \propto \left|d^1_{1,1}\right|^2 + \left|d^1_{1,-1}\right|^2 + (1-\beta_t^2)\left|d^1_{1,0}\right|^2$$
$$= \Big(\frac{1+\cos\theta^*}{2}\Big)^2 + \Big(\frac{1-\cos\theta^*}{2}\Big)^2 + (1-\beta_t^2)\,\frac{\sin^2\theta^*}{2} \;=\; \frac{1+\cos^2\theta^*}{2} + (1-\beta_t^2)\frac{\sin^2\theta^*}{2}. \tag{66}$$

Normalizing on $[-1, 1]$ and inserting $1 - \beta_t^2 = 0.0080$:

$$\boxed{\frac{1}{\sigma}\frac{d\sigma}{d\cos\theta^*}\Big|_{\text{spin 1}} = \frac{3}{8}\Big[(1+\cos^2\theta^*) + 0.0080\,\sin^2\theta^*\Big]\Big/(1+0.0040) \;\simeq\; \frac{3}{8}\,(1+\cos^2\theta^*).} \tag{67}$$

The same form holds for g^* and Z'_B alike, chiral top couplings notwithstanding: chirality moves strength between the λ_t states (§24) but not between production angles, because the light-quark side is vector-like. Numerically: the distribution at $|\cos\theta^*| = 0.9$ is $1.81/1.375 = 1.32$ times its value at 0, and 39% of events fall at $|\cos\theta^*| > 0.7$.

Corollary: $A_{FB} = 0$ **exactly [T].** A forward–backward asymmetry requires a $\cos\theta^*$-odd term, which arises only from the product $g^q_A g^t_A$. T-KK-2 forces $g^q_A = 0$ — one radial profile per UV fermion serves both chiralities — so the resonant A_{FB} vanishes identically for both spin-1 members, at any top coupling. (The small calculable $\cos\theta^*$-odd contamination from QCD–electroweak interference in the continuum is not a resonance property.) A measured resonant asymmetry would falsify the overlap structure; it would equally cleanly rule *in* an axigluon, whose defining signature is a large A_{FB}. This corollary, like the sign theorem, costs nothing and can kill the framework.

23.3 Spin 2 from gg: derivation

Assembling (b) and (c) for $\lambda_R = +2$:

$$\frac{dN}{d\cos\theta^*} \propto \left|d^2_{2,1}\right|^2 + \left|d^2_{2,-1}\right|^2 = \frac{\sin^2\theta^*}{4}\Big[(1+\cos\theta^*)^2 + (1-\cos\theta^*)^2\Big] = \frac{\sin^2\theta^*\,(1+\cos^2\theta^*)}{2} = \frac{1-\cos^4\theta^*}{2}, \tag{68}$$

$$\boxed{\frac{1}{\sigma}\frac{d\sigma}{d\cos\theta^*}\Big|_{G_1} = \frac{5}{8}\,(1-\cos^4\theta^*) \;+\; O(1-\beta_t^2).} \tag{69}$$

The distribution *vanishes* along the beam — the maximal possible contrast with (67), which peaks there. Had the graviton been produced from $q\bar{q}$, the distribution would instead be $1-3\cos^2\theta^* + 4\cos^4\theta^*$; in GG-6 the $q\bar{q}$ coupling is negligible by Corollary T-KK-2c, so the pure-gg form (69) is a *derived* prediction rather than a benchmark choice — a purity generic RS templates cannot claim, and itself a test: an admixture of the $q\bar{q}$ form in an observed spin-2 signal would contradict the derived flat-overlap hierarchy.

23.4 Separation power, quantified

The fraction of events at $|\cos\theta^*| > 0.7$ is 0.39 (spin 1) versus 0.17 (spin 2); Figure 7 draws both shapes. For a shape-only hypothesis test, the symmetrized Kullback–Leibler divergence of (67) and (69) is $\overline{\mathrm{KL}} = 0.212$ per event, giving $N_{3\sigma} = 9/(2\overline{\mathrm{KL}}) \simeq 21$ and $N_{5\sigma} \simeq 59$ events for 3σ and 5σ discrimination. Folding a representative boosted-top acceptance cut $|\cos\theta^*| < 0.85$ (the beam region is where detectors are weakest) raises these to $N_{3\sigma} \lesssim 40$ — still an order of magnitude below the expected signal yield (§26).

24 Top Polarization and Spin Correlations

24.1 Longitudinal polarization: $P_t = +0.91$ versus zero

At $\beta_t \to 1$, chirality coincides with helicity, so the resonance populates right- and left-handed tops in proportion $r^2_{t_R} : r^2_{t_L}$. The longitudinal polarization is therefore a pure coupling ratio — independent of PDFs, K-factors, luminosities, and even of ξ_{UV}:

$$\boxed{P_t = \frac{r^2_{t_R} - r^2_{t_L}}{r^2_{t_R} + r^2_{t_L}} = \frac{20.25-1}{20.25+1} = +0.906, \qquad P_t \in [+0.88,\, +0.92] \text{ for } r_{t_R} \in [4,5],} \tag{70}$$

identical for g^* and Z'_B (same overlap pattern), with the finite-mass dilution entering only at $O(1-\beta_t^2)$. For G_1 the coupling is the parity-even $T^{\mu\nu}$: $P_t = 0$ identically [T]. The polarization thus measures directly the one [D*] input of the collider chain — the cascade-depth ratio r_{t_R}/r_{t_L} — and its *sign* (right-handed) is the statement that t_R sits deeper in the infrared than t_L, i.e. a direct image of the Arisaka (2026, A9-SM) cascade ordering.

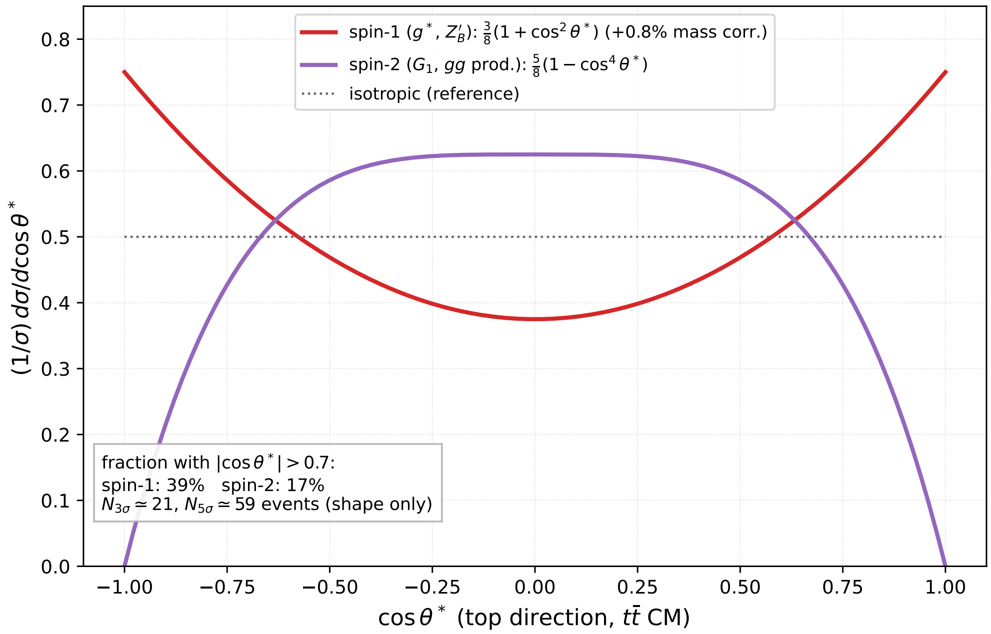

Figure 7: Normalized production-angle distributions at 3.849 TeV, Eqs. (67)–(69). The spin-1 pair rises toward the beam; the gg-produced spin-2 graviton vanishes there. Twenty-one reconstructed events separate the hypotheses at 3σ by shape alone.

Measurement. In the top rest frame the charged lepton from $t \to bW \to b\ell\nu$ analyzes the top spin with unit analyzing power at leading order ($\alpha_\ell = 1$, a consequence of the V−A weak vertex):

$$\frac{1}{\Gamma}\frac{d\Gamma}{d\cos\theta_\ell} = \frac{1}{2}(1 + P_t\,\alpha_\ell \cos\theta_\ell), \qquad \theta_\ell \equiv \angle(\ell^\pm,\ \text{top spin axis}), \tag{71}$$

so the GG-6 spin-1 hypothesis predicts a slope of $+0.91$ and the spin-2 hypothesis a flat line: with $\sim 10^2$ leptonic-side tops the two are separated at $> 5\sigma$. Note the complementarity: the production angle (§23) tests the *spin* of the resonance; the lepton slope tests its *chirality structure*; the two observables fail independently.

24.2 Spin correlations: a consistency check, not a discriminator

Both spin-1 ($\lambda_R = \pm 1$) and spin-2 ($\lambda_R = \pm 2$) decays populate $\Delta\lambda = \pm 1$ at $\beta_t \to 1$: the pair emerges with $\lambda_t = -\lambda_{\bar{t}} = \pm\frac{1}{2}$ (opposite helicity labels; spins spatially aligned, the momenta being back-to-back). This correlation structure is *identical* for the two spins, so $t\bar{t}$ spin correlations do not separate them here; what they do test is the universality of the $\Delta\lambda = 0$ mass suppression (0.8% for both). We flag this honestly because spin correlations are sometimes advertised as a generic spin analyzer; at $\beta_t = 0.996$ they are not — the discriminators are θ^*, P_t, and the line shape.

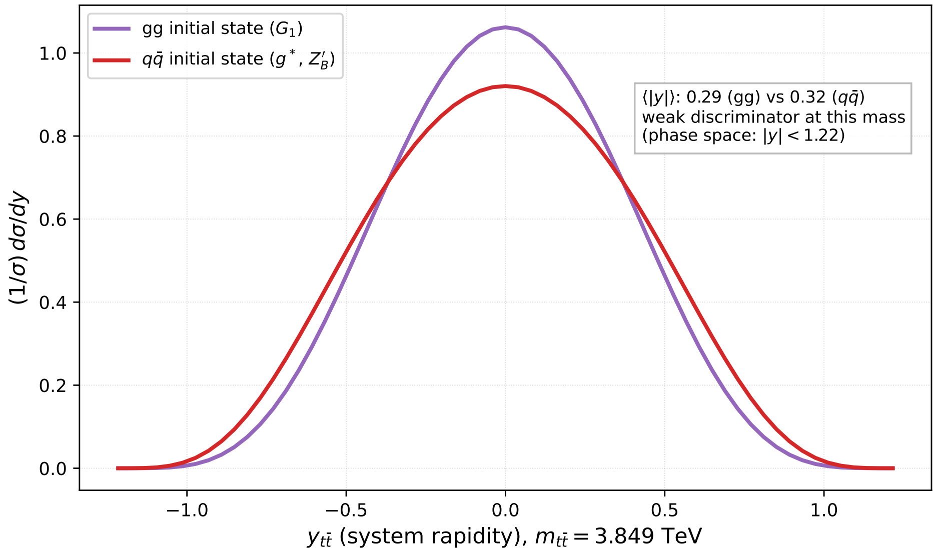

Figure 8: Normalized $t\bar{t}$-system rapidity at $m_{t\bar{t}} = 3.849$ TeV for the two initial states (CTEQ6M). The discrimination is real but weak — reported as such. Two initial states give two rapidity shapes, and the separation between them is smaller than the separation in the angular distribution or the line shape. It is plotted anyway, and described as weak in the caption rather than in a footnote, because a discrimination suite that showed only its strongest member would misrepresent what an analysis actually has to work with.

25 System Kinematics

The gg versus $q\bar{q}$ initial states sample different parton x-distributions, so the $t\bar{t}$ *system* rapidity distinguishes G_1 from the spin-1 pair in principle. At $m_{t\bar{t}} = 3.849$ TeV, however, phase space confines $|y_{t\bar{t}}| < \ln(\sqrt{s}/m) = 1.22$, and the honest verdict is that the discriminator is *weak*: $\langle |y| \rangle = 0.29$ (gg) versus 0.32 ($q\bar{q}$), with 16% versus 21% of events at $|y| > 0.5$ (Figure 8). Rapidity is a consistency check, not a spin measurement, and we say so because the folklore "gg is central, $q\bar{q}$ is forward" promises more than lock-scale kinematics can deliver. The robust kinematic facts are simpler and analysis-relevant: each top emerges with $p \simeq M_{\text{lock}}/2 = 1.92$ TeV ($\gamma_t \simeq 11$, decay cone $\approx 2/\gamma_t \approx 0.18$) — squarely in large-$R$ top-tagging territory; and the g^* width places $\sim$50% of the Breit–Wigner within $|m_{t\bar{t}} - 3849| < 280$ GeV, several times wider than the 3–4% experimental $m_{t\bar{t}}$ resolution, so *the observed width is physics, not detector* — the line-shape program of §22 is measurable.

26 Continuum, S/B, and Yields

The LO Standard-Model continuum ($q\bar{q} + gg \to t\bar{t}$, Combridge formulas, $\mu = m_{t\bar{t}}$, the paper's α_s convention) integrates to $\sigma^{\text{LO}}_{t\bar{t}} = 331$ pb at 13 TeV — the expected factor ~ 2.5 below the NNLO+NNLL 833 pb, reflecting the well-known scale sensitivity of LO $t\bar{t}$; we therefore use it only in *ratios*, where signal and continuum K-factors largely cancel (both sides here are LO). At the lock scale the continuum is $d\sigma/dm_{t\bar{t}} = 1.5$ ab/GeV. In the octet window $|m_{t\bar{t}} - 3849| < \Gamma/2 = 280$ GeV:

continuum 0.95 fb against Breit–Wigner signal 3.34 fb —

$$\boxed{S/B\big|_{\text{window}} = 3.5 \text{ (before the interference gain)}, \qquad (S{+}I)/B\big|_{\text{pole}} = 19.} \tag{72}$$

Per 3 ab^{-1}: $\sim 10^4$ signal events produced in the window; with BR(ℓ+jets) = 0.29 and a representative boosted-selection efficiency of 12% [E], $\sim$350 reconstructed lepton+jets events, before the all-hadronic sample — enough for the 21-event spin separation, the P_t slope, and a multi-bin line-shape fit simultaneously. The same accounting confirms the other two members from the background side: G_1 has $S/B = 0.04$ in its window (LHC-invisible, as Corollary T-KK-2c concluded from the rate side), and Z'_B reaches only $S/B = 0.47$ at the very top of the L-KK-1 menu.

27 The Discrimination Matrix and the Pre-Registered Fingerprint

Table 15: The discrimination matrix at 3.849 TeV. Shape rows are normalization-free and [T]-grade unless noted. "Falsifies" names what a contrary measurement kills. The last column is what makes this a matrix rather than a list: each row names the measurement that would kill the reading, not merely the one that would support it. Shape rows are normalization-free, which means they survive a wrong rate; that is why they carry the theorem grade and the rate-dependent rows do not.

Observable	**g^* (spin 1, octet)**	**Z'_B (spin 1, singlet)**	**G_1 (spin 2)**	**Falsifies if wrong**
$\sigma\times$BR($t\bar{t}$), LO×K	8.7 [7.0–12] fb	0.016–0.59 fb	0.045 fb	F-KK-2B (rate)
$dN/d\cos\theta^*$	$\frac{3}{8}(1+c^2)$	$\frac{3}{8}(1+c^2)$	$\frac{5}{8}(1-c^4)$	spin assignment
$f(\lvert\cos\theta^*\rvert > 0.7)$	0.39	0.39	0.17	spin assignment
A_{FB} (resonant)	0 exactly [T]	0 exactly [T]	0 (parity)	T-KK-2 overlap
P_t (lepton slope)	+0.91 [+0.88, +0.92]	+0.91	0 [T]	cascade ordering
line shape	BW(Γ/m=15%) + interference: low shoulder, collapsed high tail, deficit > 4.9 TeV	pure BW, narrow (0.2–8%)	pure BW, 6.3%; invisible	T-KK-2 sign [T] / color
QCD interference	full (octet) [T]	none (singlet) [T]	negligible	color assignment
centroid pull	$-O(100)$ GeV	none	none	template honesty
companions	$b\bar{b}/jj = 3.5:1$ rel. pattern	similar	hh, VV instead	BR pattern
$y_{t\bar{t}}$	$\langle\lvert y\rvert\rangle = 0.32$	0.32	0.29	(consistency only)

The pre-registered fingerprint. Combining the rows, the GG-6 discovery scenario at the HL-LHC is a single composite statement, registered here before the data:

> *A broad ($\Gamma/m \simeq 15\%$), top-saturated excess at* 3.849 *TeV with (i) a* $(1+\cos^2\theta^*)$ *production angle, (ii)* +0.91 *right-handed top polarization, (iii) zero resonant forward–backward asymmetry, (iv) a constructive-below/destructive-above interference line shape with an* $O(100)$ *GeV downward centroid pull, and (v) essentially nothing in dileptons or diphotons.*

Each clause can fail independently, and each failure names its culprit: (i) the spin, (ii) the cascade ordering, (iii) the vector-like UV overlap, (iv) the T-KK-2 sign or the color assignment, (v) leptophobia. Conversely, an excess at 3.849 TeV *without* this fingerprint — however exciting — would not be the GG-6 lock cluster, and the framework would forgo the claim. This is the quantitative content behind §13's statement that $t\bar{t}$ is both the discovery and the diagnostic channel.

28 Reliability and Open Ends

Shelving. [T]-grade and template-independent: the angular forms (67)–(69), $A_{FB} = 0$, $P_t = 0$ for G_1, the interference sign theorem, and the interference yes/no color dichotomy. [D*]: $P_t = +0.906$ (band 0.88–0.92 from the cascade-depth bracket); the width $\Gamma/m = 0.146$; all LO×K normalizations with their stated ±20% PDF/scale allowances (the CTEQ6-era distributions at $x \sim 0.3$–0.6 are the largest single systematic — and electroweak Sudakov corrections, a negative $\mathcal{O}(10$–$20\%)$ at these masses, are not included: they suppress continuum and coherent sum together, so data-anchored normalizations absorb them at first order, but a complete EW-corrected template belongs with the same NLO program; a modern-set recomputation would move normalizations, not shapes). [E]: K-factors (Chivukula et al., 2012; Mathews & Ravindran, 2005), the 12% reference efficiency, and the LO/LO convention of the S/B ratios (±30% honest band).

What would sharpen this Part. (i) Closing the cascade-depth lemma pins P_t and the BR pattern to single values. (ii) The L-KK-1 selection pins the Z'_B row. (iii) An NLO recomputation of Eq. (65) with a modern PDF set. (iv) Above all, a dedicated experimental $\Gamma/m \simeq 15\%$ template *including interference*: fitting a pure Breit–Wigner to an interference-distorted shape biases both the mass and the rate, and the most valuable single deliverable ATLAS and CMS could produce for this test is a template family with the interference degree of freedom included — with its sign left free, so that the data can convict or acquit the T-KK-2 prediction rather than assume it.

Part VI derived the signal's face, in plain terms: how the decay products spread in angle (the fingerprint of spin), how the tops emerge polarized (the fingerprint of the coupling's handedness), and how the resonance interferes with the ordinary background — enhancing it below the peak, depleting it above, with the sign a theorem rather than a fit. These shapes are the paper's most robust predictions, immune to the rate bands, and they are what will separate this cluster from any impostor. Shapes and rates together now form a complete target; the natural next question is where the experiments actually stand. The next Part reviews the published record: what is excluded, what is not, and why the prediction sits exactly at the live frontier.

29 Part VI in plain English: what the signal would look like in detail, and how sure each detail is

If something is there, this Part is about telling what it is.

Three hypotheses can produce a bump in top-antitop near the same mass: the spin-2 state, the color-octet vector, and an ordinary single resonance. The Part works out how each would differ in the details — the angular distribution of the decay products, the polarization of the tops, the shape of the line including the sign of its interference with the continuum.

The top quark is the right place to look for an unusual reason: it decays before it can hadronize, so its spin is inherited by its decay products rather than scrambled. That makes polarization a measurable quantity here in a way it is not for lighter quarks.

The grading is what makes the Part usable. Shape observables are normalization-free, so they are theorem-grade and survive any error in the rate. Rate-dependent statements are marked separately and carry the assumptions of the previous Part. And the rapidity distribution, which would have been a fourth handle, discriminates only weakly — the Part says so in the caption of its own figure rather than leaving it to be discovered. A discrimination suite that showed only its strongest member would be a misleading one.

Part VII
Experimental Status and Reach

Current ATLAS/CMS limits, near-term reach, the kill criterion, and future machines

30 Current Experimental Status at ATLAS and CMS

This section is deliberately grounded in recent public collaboration outputs and recent official CERN, ATLAS, and CMS machine documents. Because the current status of the LHC can change on the scale of months, this is exactly the type of section that deserves freshness.

30.1 The luminosity backdrop

Official CERN reporting from December 2025 states that the LHC has delivered over $515\,\mathrm{fb}^{-1}$ each to ATLAS and CMS across the proton-run periods, with $124\,\mathrm{fb}^{-1}$ and $125\,\mathrm{fb}^{-1}$ delivered to each experiment in 2024 and 2025 respectively (CERN, 2025). Official CERN reporting from March 2026 further states that the 2026 proton run will be short but intense and will extend the integrated data set beyond the already achieved Run-3 target (CERN, 2026). This backdrop matters because it tells us that the statistical frontier is rapidly shifting, even if the public analysis papers naturally lag the delivered luminosity.

30.2 High-mass dilepton constraints

ATLAS published a full Run-2 high-mass dilepton search using $139\,\mathrm{fb}^{-1}$ at 13 TeV and found no significant deviation from the Standard Model expectation (ATLAS Collaboration, 2019). CMS, using the full $138\,\mathrm{fb}^{-1}$ Run-2 dilepton sample, reports lower RS-graviton mass limits of 2.47 to 4.78 TeV for couplings $k/\bar{M}_{\mathrm{Pl}} = 0.01$ to 0.1 on its official public-results page (CMS Collaboration, 2020). This result is one of the cleanest guides to generic spin-1 or spin-2 resonance sensitivity. Its strength is obvious: excellent mass resolution and low background. Its weakness for the present purpose is equally clear: if `GG-6` suppresses light-lepton couplings relative to top-rich and dijet-rich channels, the dilepton reach overstates the exclusion power against the specific `GG-6` benchmark.

30.3 High-mass diphoton constraints

CMS published a high-mass diphoton search using the full Run-2 13 TeV data set. In the abstract-level summary of that result, masses below roughly 2.2 to 5.6 TeV are excluded at 95% confidence level for an excited Randall–Sundrum graviton over the coupling range $0.01 < \tilde{k} < 0.2$ (CMS Collaboration, 2024a). This is a powerful statement for canonical RS-like diphoton benchmarks. It is not, however, the last word on `GG-6` for two reasons. First, `GG-6` is not a single-pole RS model.

Second, the same geometry that enhances top and hadronic channels can reduce the light-photon branch relative to a universal RS template. The diphoton paper therefore tells us that standard-coupling RS-like interpretations are already above the nominal 3.85 TeV scale, while the dedicated GG-6 reinterpretation remains open.

30.4 Dijet and substructure-based searches

The dijet channel is unavoidable for GG-6 because of the mild $\mathrm{U}(1)_B$ admixture and the possibility of enhanced hadronic decays. CMS has published a model-agnostic search for dijet resonances with anomalous jet substructure in the mass range 1.8 to 6 TeV, again with no significant excess observed (CMS Collaboration, 2024d). The importance of this result for the present paper is not merely the null result itself, but the method: substructure-assisted dijet searches are the right conceptual direction for a clustered, possibly partly hadronic, 6D signal. The current public limit is therefore better interpreted as a *proof of method* plus an initial exclusion map, rather than as the final word on the full GG-6 cluster hypothesis.

30.5 Top-pair resonance searches

Top-pair resonance searches are arguably the most important public channel for the GG-6 benchmark. The most relevant recent reference is the ATLAS search for $t\bar{t}$ resonances in the single-lepton and dilepton final states using the full Run-2 140 fb^{-1} data set (ATLAS Collaboration, 2025a). That paper explicitly reviews both bulk RS graviton and KK gluon benchmarks. In the bulk-RS graviton discussion, ATLAS emphasizes that in a bulk-RS setup decays to light fermions are suppressed, the branching fraction to photons is negligible, and the phenomenologically interesting couplings can be large (ATLAS Collaboration, 2025a). This is already much closer in spirit to the GG-6 search problem than a simple dilepton-only benchmark.

The same ATLAS paper also summarizes older direct exclusions relevant to our context. For example, previous ATLAS searches with 36 fb^{-1} excluded bulk RS gravitons in the 0.45 to 0.65 TeV range in a $t\bar{t}$-focused benchmark, WW/ZZ searches excluded up to 2.3 TeV, and a b-tagged dijet search excluded a $k/\bar{M}_{\mathrm{Pl}} = 0.2$ graviton benchmark up to 2.8 TeV (ATLAS Collaboration, 2025a). For KK gluons, the same ATLAS review notes that previous ATLAS searches excluded benchmark KK-gluon masses below 3.7 TeV, while CMS excluded similar benchmark masses up to 4.55 TeV (ATLAS Collaboration, 2025a). The precise numerical mapping to GG-6 is again benchmark-sensitive, but the qualitative lesson is clear: $t\bar{t}$ is already an experimentally mature high-mass resonance channel, and it is probably the most natural first discovery channel for the GG-6 first cluster.

30.6 Higgs-pair searches

The hh channel is where the spin-2 graviton's coupling to the IR-localized Higgs sector becomes visible (there is no radion in GG-6; §10.5 and Arisaka (2026, A4-Euclid)). CMS has published a search for resonant Higgs-boson pair production in the $bbbb$ final state using large-area jets and the full 138 fb^{-1} Run-2 sample (CMS Collaboration, 2024c). Narrow bulk-graviton interpretations were tested over 1 to 3 TeV, with upper limits on $\sigma(pp \to X) \times \mathrm{BR}(X \to HH \to bbbb)$ down to the sub-femtobarn level at high masses (CMS Collaboration, 2024c). Extending such a search to the fixed 3.85 TeV window, as a bosonic cross-check of a $t\bar{t}$/dijet first signal, is the natural GG-6-aware use of this channel.

30.7 A channel-dependent exclusion and discovery ledger

Because current public exclusions are quoted in a mixture of benchmark languages, the most honest way to summarize the status is to place the official channel limits next to the corresponding GG-6 reinterpretation issue. Table 16 does exactly that. It is intentionally channel-centric rather than paper-centric.

Table 16: Channel-by-channel exclusion and discovery ledger. The "public benchmark" column records what the collaborations have actually excluded in standard templates; the "GG-6 reading" column explains why those numbers are highly informative but not yet equivalent to a full geometry-aware exclusion.

Channel	**Representative official source**	**Public benchmark limit**	**What that benchmark really constrains**	**GG-6 reading of the same result**	**Near-term discovery note**
Dileptons	CMS EXO-19-019 official public result (CMS Collaboration, 2020)	for RS spin-2, $m_G > 2.47$–$4.78\,\mathrm{TeV}$ for $k/\bar{M}_{\mathrm{Pl}} = 0.01$–$0.1$	canonical or narrow RS-like spin-2 couplings to light leptons	very informative but not decisive for GG-6 because the light-lepton overlap can be suppressed by factors of a few and the θ-cluster broadens a one-pole fit	in the internal GG-6 corridor the dilepton line can still yield a clean corroborating excess, but it is not expected to dominate the significance
Diphotons	CMS EXO-22-024 official public result (CMS Collaboration, 2024b)	RS graviton masses below 2.2–5.6 TeV excluded for $0.01 < \tilde{k} < 0.2$	standard RS-like diphoton graviton signal	again very strong for canonical RS, but GG-6 predicts both reduced photonic weight and cluster dilution	a good clean channel for corroboration; not the only or necessarily strongest first channel
Top pairs	official ATLAS/CMS top-resonance program (ATLAS Collaboration, 2025a)	ATLAS KK-gluon benchmark exclusion $\sim 3.8\,\mathrm{TeV}$ in lepton+jets; older all-hadronic limits around 3.4 TeV; CMS KK-gluon dijet benchmark below 4.55 TeV	standard KK-gluon or bulk-warped top-rich resonances	this is the closest public proxy to the actual GG-6 search because GG-6 is top-rich by construction	the most promising single discovery channel already in the current Run-3 closeout and the natural place to look "this year"
Dijets	CMS anomalous-substructure dijet program (CMS Collaboration, 2024d)	no significant excess in the 1.8–6 TeV range; benchmark KK-gluon exclusions reach above 4 TeV depending on template	narrow hadronic resonances and substructure-assisted broad objects	especially relevant for the leptophobic $\mathrm{U}(1)_B$ admixture, which can shift discovery power away from dileptons	very plausible companion channel for the spin-1 doublet and possibly for the graviton family if cluster broadening is important
hh (graviton channel)	CMS large-area-jet hh searches and diboson resonance program (CMS Collaboration, 2024c)	bulk-graviton hh interpretations constrained in the 1–3 TeV range	narrow spin-2 templates	GG-6 has no radion; the scalar-sector handle is $G_1 \to hh$, with small rate at 3.85 TeV	a bosonic corroborator rather than the first discovery channel; useful once a 4 TeV family appears
Dibosons	ATLAS and CMS heavy-resonance bosonic searches	several-TeV exclusions for standard spin-1 and spin-2 benchmarks	high-mass boson-rich states with more standard templates	useful cross-check because GG-6 overlap structure still populates WW and ZZ visibly, especially from the spin-2 graviton	important for confirming the bosonic side of the 6D cluster after a first hint in tops or dijets

The interpretive lesson of the table is straightforward. The public LHC record is already strong enough that any claim of a completely unconstrained 3.85 TeV KK story would be irresponsible. But the opposite extreme – declaring the GG-6 implementation already dead solely from canonical one-pole RS templates – is also too strong. The correct status is sharper and more interesting: the first-family GG-6 target is *bracketed*. It is under real pressure. It can still survive because the exact same geometric ingredients that make it experimentally attractive also make a naive RS translation too aggressive.

30.8 What the public status does and does not prove

The public ATLAS/CMS landscape as of early 2026 therefore proves three things and leaves three things open.

It proves, first, that simple standard-coupling RS-like spin-2 benchmarks in the cleanest leptonic and photonic channels are already heavily constrained above the 3.85 TeV scale. It proves, second, that top and dijet resonance searches are robust, mature, and already pushing into the mass range relevant for associated GG-6 partners. It proves, third, that Higgs-pair and diboson resonance searches are already serious parts of the BSM program.

But it leaves open, first, the dedicated reinterpretation of a 6D clustered spin-2 signal centered at 3.85 TeV. It leaves open, second, the combined multi-channel exclusion or discovery of a top-rich, mildly leptophobic, multi-field graviton cluster. And it leaves open, third, the joint pattern test involving the first cluster and the second radial family. Those are precisely the gaps the rest of this paper is meant to organize.

Operational reading of the current public status. At the level of *public benchmark results*, ATLAS and CMS have already pushed standard-coupling RS-like spin-2 interpretations into or above the nominal 3.85 TeV region in the cleanest channels. At the level of a *dedicated GG-6 reinterpretation*, however, the decisive ingredients remain missing: a clustered line-shape benchmark, a top-rich and mildly leptophobic branching pattern, and a combined-channel exclusion corridor tied to the geometry rather than to a freely scanned mass template. That gap between *proxy exclusions* and a *true GG-6 test* is exactly why the paper treats the present moment as decisive rather than already settled.

Table 17: Current public ATLAS/CMS status relevant to the GG-6 KK search program. The table intentionally distinguishes between what is already strongly excluded for standard benchmark models and what remains open for a dedicated clustered GG-6 reinterpretation. The distinction the table draws is between an exclusion and an exclusion of something else. Standard benchmarks are strongly excluded over the relevant mass range and that is stated first, without qualification; what remains open is the clustered reading, which no published search has been optimized for. Every entry is from a public record with its reference attached.

Search class	Representative public result	Dataset	What the public result says	What remains open specifically for GG-6
High-mass dileptons	ATLAS high-mass dilepton resonance search (ATLAS Collaboration, 2019)	139 fb^{-1}, 13 TeV	No significant excess; excellent clean-channel sensitivity to generic high-mass spin-1 and spin-2 signals.	Does not directly target a clustered 6D line shape or a top-rich / light-lepton-suppressed overlap pattern.
High-mass diphotons	CMS diphoton resonance search (CMS Collaboration, 2024a)	full Run 2, 13 TeV	RS graviton excited states excluded roughly in the 2.2–5.6 TeV window across $0.01 < \tilde{k} < 0.2$.	Public limit is on canonical RS-like diphoton structure; GG-6 may have reduced diphoton branching and non-single-pole structure.
Anomalous dijets	CMS model-agnostic anomalous-substructure dijet search (CMS Collaboration, 2024d)	138 fb^{-1}, 13 TeV	No significant excess in 1.8–6 TeV mass range.	Dijet channel is especially important for a mild $\mathrm{U}(1)_B$ admixture; cluster-aware and top-correlated reinterpretation remains open.
Top-pair resonances	ATLAS $t\bar{t}$ resonance search (ATLAS Collaboration, 2025a)	140 fb^{-1}, 13 TeV	Mature search program for bulk-RS graviton and KK gluon benchmarks; strong exclusions for some top-rich models.	$t\bar{t}$ is likely the natural leading discovery channel for GG-6; dedicated first-cluster reinterpretation remains needed.
Higgs-pair resonances	CMS $HH \to bbbb$ resonant search (CMS Collaboration, 2024c)	138 fb^{-1}, 13 TeV	Sub-femtobarn to few-femtobarn upper limits over 1–3 TeV for bulk-graviton HH interpretations.	GG-6 has no radion (Arisaka, 2026, A4-Euclid); the $G_1 \to hh$ branch is the scalar-sector handle, and a dedicated 3.85 TeV reinterpretation is not yet performed.
Machine dataset status	CERN official luminosity updates (CERN, 2025, 2026)	official machine reports	More than 515 fb^{-1} delivered each to ATLAS/CMS overall; 2026 run extends Run-3 dataset beyond target.	Full Run-3 + early-HL-LHC reinterpretation will substantially sharpen the first-cluster test.

31 The Operational Confrontation: Computed $\sigma \times \mathrm{BR}$ versus Published Limits

31.1 What this section adds, and at what grade

The two preceding sections state *what* ATLAS and CMS have excluded and *why* the GG-6 cluster is not equivalent to any single public benchmark. A referee is entitled to ask for more: *a number an analyst can exclude.* This section supplies it. Starting from the Read-out-II couplings (§19) we compute, channel by channel, the leading-order narrow-resonance cross sections of all three lock-cluster states, place them on the same axes as the published 95% CL exclusion curves, and state for each channel whether the prediction is below the limit — and by what factor — or in tension.

The honesty grade is fixed at the outset. Everything in this section is *Derived-conditional* (D*): conditional on the Read-out-II assumptions B1–B4, on leading-order QCD with the stated PDF and K-factor treatment, and on the published-template transfers stated below. Nothing here touches the theorem shelf (the masses, the exact leptophobia, the coupling identities of T-KK-2). Where this computation supersedes the template-rescaled ballparks quoted in §19, we say so explicitly and keep both visible.

31.2 The workbook: master formula, inputs, and validation

Master formula. For a resonance R of spin J, color dimension D_R, and mass M, the leading-order narrow-width hadronic rate is

$$\sigma(pp \to R \to f) \;=\; \sum_{ij} \underbrace{(1+\delta_{ij})\,\frac{16\pi^2\,(2J+1)\,D_R}{(2S_i+1)(2S_j+1)\,D_i D_j}}_{P_{ij}} \,\frac{\Gamma_{ij}}{M}\,\frac{\mathcal{L}_{ij}(\tau)}{s}\,\mathrm{BR}(R\to f), \qquad \mathcal{L}_{ij}(\tau) = \int_\tau^1 \frac{dx}{x}\, f_i(x) f_j\Big(\frac{\tau}{x}\Big) \tag{73}$$

with $\tau = M^2/s$ and both orderings summed for $i \neq j$. This is Eq. (53) with the kinematic factor written out; the spin/color normalization reproduces the standard coefficients of Franceschini et al. (2016) for every color-singlet case and the color-octet factor $P_{q\bar{q}} = 32\pi^2/3$ fixed by detailed balance.

Inputs. The couplings are those the paper has already put on the table. For the color octet g^*: light quarks at the theorem-fixed $\xi_{\mathrm{UV}} = -0.194$ (Eq. 37), Q_{3L} at $1.0\,g_s$, and t_R at $4.5\,g_s$ (Read-out-II range 4–5, assumption B4), with $\Gamma(g^* \to q\bar{q}) = \alpha_s\, r_q^2\, m/6$ per flavor — the color-octet trace, not the color-singlet N_c factor, as in the standard KK-gluon literature (Dobrescu & Yu, 2013; Lillie, Randall & Wang, 2007). For the baryonic Z'_B: $g_f = g_B B_f r_f$ with $B_q = 1/3$ and the same overlap pattern; every rate scales as g_B^2 and is quoted at the reference point $g_B = 0.2$ (Dobrescu & Yu, 2013); the theorem-grade norm plus the L-KK-1 menu bracket g_B to 0.14–0.89 (§10.10). For the spin-2 G_1: gluon-fusion production through the *derived* flat-overlap $\kappa_{gg} = c_K/K = 0.0136$ (Corollary T-KK-2c, §10.9) with the Han–Lykken–Zhang partial widths (Han, Lykken & Zhang, 1999) (the massive-vector polynomial corrected per Allanach et al., 2000), and the bulk decay pattern: the infrared channels $t\bar{t}$, $W_L W_L$, $Z_L Z_L$, hh at full strength, flat-mode channels (gg, $\gamma\gamma$) suppressed by κ_{gg}^2, light fermions negligible (Agashe et al., 2008). The strong coupling is run from the Arisaka (2026, A9-SM) value $\alpha_s(M_Z) = 1/(\pi e)$ to $\alpha_s(3.85\,\mathrm{TeV}) = 0.077$. Parton luminosities use CTEQ6M (Pumplin et al., 2002), cross-checked with CTEQ6L1; at $\sqrt{s} = 13$ TeV: Any analyst can redo the arithmetic of this section from Eq. (73) and this table alone; that is the point.

Validation. The pipeline was checked against three independent published anchors before any GG-6 number was produced: (i) the reference parton-luminosity coefficients of Franceschini et al. (2016) at $M = 750$ GeV are reproduced to 2–5%; (ii) the Randall–Sundrum graviton $\sigma \times \mathrm{BR}(e^+e^-)$ table of Allanach et al. (2000) is reproduced to 15–25% (the expected LO/PDF-era spread), with the total widths matching to 2–3%; (iii) the warped KK-gluon branching pattern of Lillie, Randall & Wang (2007) ($\mathrm{BR}(t\bar{t}) = 92.5\%$) is reproduced exactly, and the Agashe et al. (2008) production benchmark at 3 TeV to $\sim$25%.

Table 18: Parton luminosities at the two candidate masses. These are the numbers every rate in this paper is built on and they are somebody else's measurement, not this framework's prediction: they come from a published parton distribution set, cross-checked against a second one, and are quoted here so that a reader recomputing a cross section starts from the same inputs. The gluon-gluon and quark-antiquark columns are given separately because the two cluster members are produced through different initial states.

M	$\tau = M^2/s$	$\mathcal{L}_{gg}$	$\mathcal{L}_{u\bar{u}}$	$\mathcal{L}_{d\bar{d}}$	$\sum_q \mathcal{L}_{q\bar{q}}$
3.849 TeV	0.0877	0.0337	0.0452	0.0139	0.0596
2.461 TeV	0.0358	1.440	0.974	0.462	1.462

Uncertainties. The quoted bands fold together the NLO K-factor (1.0–1.3 for $q\bar{q}$ resonances, stated rather than silently applied) and a ±20% PDF-and-scale allowance for $q\bar{q}$-initiated rates (the CTEQ6M/CTEQ6L1 spread is 3–9%; the allowance covers modern-set drift at large x); the gluon-gluon luminosity at $x \sim 0.3$ carries a factor ~ 2 spread, absorbed into the κ_{eff} corridor it multiplies. These are LO estimates with stated errors — not precision QCD, and not dressed as such.

31.3 Results

Table 19: Computed $\sigma \times \text{BR}$ for the lock-cluster states at $\sqrt{s} = 13$ TeV (LO×K, CTEQ6M: $K_{q\bar{q}} = 1.3$ on the vector rows, $K_{gg} = 1.6$ on the G_1 rows; bands fold the PDF/scale spread). All entries Derived-conditional on B1–B4. The G_1 rows use the *derived* $\kappa_{gg} = c_K/K = 0.0136$ (Corollary T-KK-2c); their band is the gluon-luminosity spread (factor ~ 2) with the $O(1/K)$ amplitude fence, and BR($t\bar{t}$) carries a $\sim 25\%$ cascade-depth sensitivity. The Z'_B rows scale as $(g_B/0.2)^2$; the L-KK-1 menu spans $g_B = 0.14$–0.89, i.e. ×0.5–20 on the quoted reference (§10.10).

State (mass)	**Channel**	**$\sigma \times$ BR central [band], fb**	**Scaling**	**Γ/m**
g^*, Scenario B (3.849 TeV)	$t\bar{t}$	8.7 [7.0–12]	fixed (ξ_{UV} theorem)	0.15
	$jj + b\bar{b}$	0.55 [0.44–0.78]	fixed	
g^*, Scenario A (2.461 TeV)	$t\bar{t}$	224 [180–325]	fixed	0.15
	$jj + b\bar{b}$	14 [11–21]	fixed	
G_1 (3.849 TeV)	$t\bar{t}$	0.045 [0.021–0.10]	derived (T-KK-2c)	0.063
	$W_L W_L$	0.010	derived	
	$Z_L Z_L$	0.0051	derived	
	hh	0.0050	derived	
	jj (gg)	9.0×10^{-5}	derived	
	$\gamma\gamma$	1.1×10^{-5} (retro-diction)	derived	
Z'_B (3.849 TeV)	$t\bar{t}$	$0.030 \times (g_B/0.2)^2$	$\propto g_B^2$	4×10^{-3}
	$jj + b\bar{b}$	$0.0018 \times (g_B/0.2)^2$	$\propto g_B^2$	
cluster sum, Scenario B	$t\bar{t}$	8.7–9.3 (octet-dominated; §10.10)	—	0.06–0.15

Two structural facts deserve emphasis. First, the color octet's $t\bar{t}$ rate carries *no* continuous freedom: its production coupling is the theorem-fixed ξ_{UV} and its branching pattern is saturated by the top. It is the one line in Table 19 that GG-6 cannot lower. Second, the three states are not interchangeable targets: the octet is broad ($\Gamma/m \simeq 15\%$), the graviton moderately narrow ($\Gamma/m = 6.3\%$ at the derived coupling, inside the width family of the published spin-2 templates), and the Z'_B very narrow — the cluster line shape is a *sum* of these, which is exactly why §13 asks for a cluster-aware fit.

31.4 The survival statement, channel by channel

Scenario A (free-field): excluded at the stated uncertainties. The free-field octet at 2.461 TeV predicts $\sigma \times \mathrm{BR}(t\bar{t}) = 224\,[180–325]$ fb. The observed ATLAS limit at that mass is 62 fb in the *most conservative* (broad, $\Gamma/m = 30\%$) template and 20 fb in the narrow template; the predicted $\Gamma/m \simeq 15\%$ lies between the two, so the applicable limit is if anything *stronger* than 62 fb. The prediction therefore sits a factor 3.6–11 above the observed exclusion, outside the band by any reading. Within the stated uncertainties, **F-KK-2A has fired**: the free-field scenario of T-KK-2 is closed by existing data, and the dichotomy resolves to the composite-hardened scenario B — the degenerate cluster at M_{lock}. We register the one formal caveat: this transfer uses ATLAS's published templates rather than a dedicated $\Gamma/m = 15\%$ reinterpretation, so we grade the closure *provisional* until that reinterpretation exists; but the factor is large, and no plausible line-shape correction recovers it.

Scenario B octet: alive, and decidable soon. At 3.849 TeV the octet predicts $8.7\,[7.0–12]$ fb in $t\bar{t}$. The observed broad-template limit is 25.5 fb — a factor ~ 2.9 of headroom — while the narrow-template limit, 4.9 fb, is already *below* the prediction. The state survives today precisely because of its breadth: a broad, top-saturated object at the lock scale is not yet excluded, but the window between the two templates is exactly where a dedicated $\Gamma/m \simeq 15\%$ reinterpretation has discriminating power *now* — and Part VI supplies the full differential anatomy (angle, polarization, interference line shape with its derived sign) that such a reinterpretation should implement. In dijets the octet is safe by a factor ~ 35 (0.23 fb against the CMS σBA limit of 8.0 fb): the dijet null is consistent, not constraining.

The graviton: derived, and out of LHC reach. With the coupling now derived (Corollary T-KK-2c), the G_1 line is a point, not a corridor: $\sigma \times \mathrm{BR}(t\bar{t}) = 0.045\,[0.021–0.10]$ fb. Even the 3 ab^{-1} HL-LHC projection of the strongest template ($\sim$1.1 fb) sits a factor ~ 25 above it: **no LHC dataset will see the graviton**, and its absence from every LHC search is henceforth a *prediction* of the framework, not a loose end. The diphoton channel supplies the retro-diction that gives this claim teeth: had the photon coupled at the uniform $1/\sqrt{K}$ amplitude, $\sigma \times \mathrm{BR}(\gamma\gamma) = 0.14$ fb would already be excluded by a factor ~ 4.2 by the observed CMS limit of 0.033 fb; at the derived $\kappa_\gamma = \kappa_{gg}$ the prediction is 7×10^{-6} fb, a factor $\sim 3 \times 10^3$ below the limit. The published null therefore discriminates between the two readings of the flat-mode amplitude and lands on the derived one — the first amplitude-level confrontation of the collider chain, resolved in the geometry's favor.

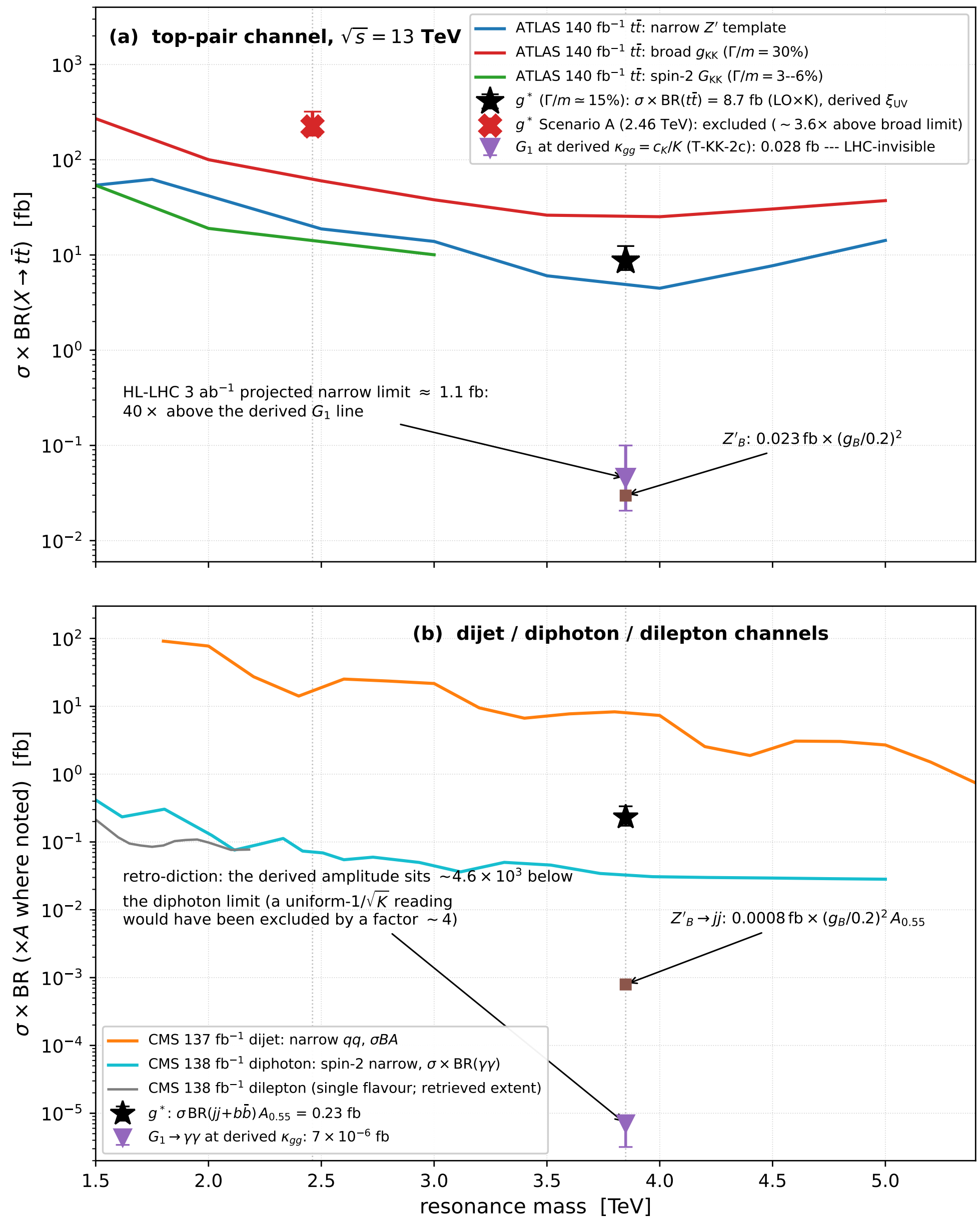

Figure 9: The operational confrontation at $\sqrt{s} = 13$ TeV, all couplings at their derived values (T-KK-2, T-KK-2c). *Panel (a):* observed 95% CL limits from the ATLAS 140 fb^{-1} $t\bar{t}$-resonance search (ATLAS Collaboration, 2025a) (narrow-Z', broad-KK-gluon $\Gamma/m = 30\%$, and spin-2 templates; HEPData ins3094414) against the predictions of Table 19: the Scenario-B octet at 8.7 fb (star) — the un-suppressible line that carries the HL-LHC falsifier; the Scenario-A octet at 2.46 TeV (red cross), a factor $\sim$3.6 *above* the observed broad-template limit (F-KK-2A provisionally fired); and G_1 at the derived $\kappa_{gg} = c_K/K$ (violet triangle, 0.045 fb) — a factor $\sim$25 below even the 3 ab^{-1} HL-LHC projection, hence LHC-invisible. *Panel (b):* the CMS narrow-dijet σBA limit (CMS Collaboration, 2020b), diphoton spin-2 limit (CMS Collaboration, 2024a), and dilepton limit over its retrieved extent (CMS Collaboration, 2020), against the octet dijet rate, the g_B^2-scaled Z'_B, and $G_1 \to \gamma\gamma$ at the derived coupling (7×10^{-6} fb): the published diphoton null is *retro-dicted* with a margin of $\sim 3 \times 10^3$, whereas a uniform-$1/\sqrt{K}$ photon amplitude would already be excluded by a factor $\sim$4. All predictions LO with stated bands.

The direct spin-2 test moves to FCC-hh, where $\sigma(G_1) \approx 205$ fb at 100 TeV makes the anchor an early-running discovery target (§31.5).

The Z'_B: bracketed by the derived norm, tested elsewhere. Scaling the Z'_B rate to the current $t\bar{t}$ limit bounds $g_B \lesssim 2.9$ — far above the entire derived-norm menu $g_B \in [0.14, 0.89]$ of §10.10. Across that menu $\sigma \times \mathrm{BR}(t\bar{t}) = 0.012$–$0.46$ fb: a factor ≥ 10 below the current narrow-template limit and ≥ 2.3 below the 3 ab^{-1} projection, so the state is LHC-invisible as a separate signal on every L-KK-1 reading and enters only through the cluster sum (0.2–6.8%). Its decisive test remains the fifth-force (g_B, m_B) window of the Predictions branch (Arisaka, 2026, A9-SM) — now four candidate lines rather than a plane. Its exact leptophobia keeps the dilepton null consistent by construction.

The cluster, summed. The Scenario-B cluster delivers $\simeq$ 8.8–9.3 fb (central, spanning the L-KK-1 menu) of $t\bar{t}$ signal spread over the near-degenerate line shape — a factor ~ 2.8 below the broad-template limit. The correct summary of the public record is therefore neither "already excluded" nor "comfortably safe": the lock cluster is *bracketed*, with the un-suppressible octet line carrying the exposure.

31.5 The kill window, quantified

With background-dominated $\sqrt{\mathcal{L}}$ scaling (Eq. 74) from the 140 fb^{-1} baseline, the observed broad-template $t\bar{t}$ limit at 3.85 TeV projects to $\sim$17 fb at 300 fb^{-1}, $\sim$9.6 fb at 1 ab^{-1}, and $\sim$5.5 fb at 3 ab^{-1} — while the octet prediction rises to 13.6 [11–19] fb at $\sqrt{s} = 14$ TeV. **The projected exclusion crosses the un-suppressible octet band between roughly 0.5 and 2 ab^{-1}**: F-KK-2B is decided — discovery or exclusion — within the first half of the HL-LHC program, on the collaborations' own published sensitivity, with no GG-6-side freedom left to move the target. The graviton, at its derived κ_{gg}, contributes 0.045 fb in $t\bar{t}$ — no LHC dataset reaches it, and the spin-2 test moves cleanly to FCC-hh, where $\sigma(G_1) \approx 205$ fb at 100 TeV ($t\bar{t}$ 142 fb, ZZ 16 fb, hh 16 fb) makes the anchor an early discovery target (FCC Collaboration, 2018). The division of labor is itself a falsifiable pattern: the HL-LHC must find or kill the spin-1 cluster; a 100-TeV machine must then find the spin-2 anchor at the same mass. This, together with the Scenario-A closure above, is the operational content of the 3.85-TeV falsifier: one scenario already decided by data on tape, the other pre-registered against a dated luminosity milestone. Figure 10 (§32.7) recasts the same statement in calendar years.

32 Near-Term Discovery Potential

32.1 The Run-3 to early-HL-LHC bridge

The next experimentally decisive interval is not the distant future; it is the bridge from full Run 3 to early HL-LHC analysis. Official CERN planning documents describe HL-LHC as a 14 TeV program targeting about 3000 fb^{-1} per experiment, with luminosity gains of roughly a factor of

5 to 7.5 over the nominal LHC design (CERN, 2023). That long-term target matters, but the shorter-term issue is equally important: the completion of Run 3 and the appearance of public analyses that incorporate the full Run-2 + Run-3 sample in high-mass resonance channels.

For a background-dominated search, one expects the simplest significance scaling to follow

$$\mathcal{Z}(\mathcal{L}_2) \approx \mathcal{Z}(\mathcal{L}_1)\sqrt{\frac{\mathcal{L}_2}{\mathcal{L}_1}}. \tag{74}$$

This familiar relation does not decide the discovery by itself, but it tells us something important: a signal that is presently marginal or template-diluted can become clearly visible once a cluster-aware multi-channel analysis is applied to a data set larger by a factor of two to five. Part VIII replaces this back-of-envelope scaling with the full anchored simulation — and shows that (74) fails in precisely the regime that matters: beyond $\sim$0.5 ab^{-1} the search is systematics-limited and significance no longer grows as $\sqrt{\mathcal{L}}$ at all.

32.2 Why $t\bar{t}$ is the priority channel

If the light-lepton and photon overlaps are even moderately suppressed, then the first robust excess will most likely emerge in $t\bar{t}$ rather than in $\ell^+\ell^-$ or $\gamma\gamma$. There are three reasons.

1. The geometry naturally favors the top sector via infrared-localized overlaps.
2. The bulk-RS literature already demonstrates that high-mass spin-2 and KK-gluon searches in $t\bar{t}$ are mature and technically feasible.
3. Top final states are also the natural meeting point of the graviton, KK-gluon, and vector-like-top partner stories, making them ideal for a pattern test.

This does not mean that dilepton or diphoton channels are unimportant. On the contrary, they can still be crucial for a clean corroborating line if the signal is sufficiently strong. But a GG-6-aware discovery program should not treat hadronic top channels as merely complementary. They are central.

32.3 The role of boosted dijets and Higgs pairs

Dijet and hh channels serve different but complementary purposes. Dijet channels test the mild leptophobic or baryonic admixture and may carry substantial branching power. Higgs-pair channels test the scalar or Hosotani-related side of the theory. A combined $t\bar{t}$+dijet+hh excess pattern near the first cluster would be much more diagnostic of GG-6 than any single excess in isolation.

32.4 A benchmark significance corridor

Because the exact first-cluster line shape is model-dependent, a completely rigid cross-section prediction is premature. A more honest intermediate product is a *discovery corridor*. Let the effective signal strength for a given channel f be parameterized by

$$\mu_f = \left(\frac{\kappa_{\rm eff}}{0.1}\right)^2 R_f, \tag{75}$$

where $\kappa_{\rm eff}$ is the effective lock-induced coupling normalization and R_f packages branching and acceptance relative to a reference bulk-RS template. In this language, the decisive question becomes whether the combined ATLAS+CMS program can exclude or discover μ_f across the full geometry-allowed corridor rather than at a single arbitrary benchmark point.

Within this Part's theorem-anchored discipline $\kappa_{\rm eff}$ carries no number — it is not a GG-A output: it is set by the bulk-to-boundary overlaps that Arisaka (2026, A9-SM) leaves as an open lemma. The conceptual point survives without a number: $\kappa_{\rm eff}$ is not free in the spirit of arbitrary model scans; it is constrained by the same geometry that fixes the mass scale. This is another reason the eventual collider test is unusually meaningful — and it is why the decisive deliverable is the corridor μ_f, not a single benchmark cross section. Read-out II (§19) estimates $\kappa_{\rm eff}$ from the throat geometry, and §31 confronts the resulting rates, corridor and all, with the published limits.

32.5 Could ATLAS and CMS already get the first events in the current 2026 run?

Yes – in the limited but important sense of *hints, small excesses, or channel-local anomalies.* No – if by "get" one means an indisputable five-sigma combined discovery with a model-complete pattern test before LS3. The reason is simple arithmetic.

Official CERN reporting states that the 2025 proton run alone delivered $125.4\,\mathrm{fb}^{-1}$ to both ATLAS and CMS, bringing the total delivered since LHC startup above $500\,\mathrm{fb}^{-1}$ (CERN, 2025). Official CERN reporting from March 2026 further states that stable beams for physics were declared on 7 March 2026, that the 2026 run is "short but densely packed," and that the proton program includes a low-pileup phase followed by a high-pileup phase that will extend the integrated dataset beyond the already achieved Run-3 target before LS3 begins at the start of July (CERN, 2026). In practical language, that means there is still genuine discovery leverage *this year*, but it sits in a compressed window.

We deliberately do *not* convert this dataset into a predicted event count. Through Eq. (58) an event yield requires the absolute cross section, which (§12.2) is not a GG-A-derived quantity: it depends on $\kappa_{\rm eff}$ and a QCD/PDF integration that lie outside the theorems. Quoting "N tagged $t\bar{t}$ events this year" would therefore present a model-level number as if it were a prediction, which the discipline of this paper forbids. (The clearly-labeled Read-out-II rates of §31, Table 19, convert to yields via Eq. (58) for any reader who wants them at that grade.) What *is* robust, and is the only thing we assert here, is structural: the color octet carries the un-suppressible light-quark coupling $1/\sqrt{K} \approx 0.17$, so it sits at the current $t\bar{t}$/dijet exclusion frontier and is being probed by existing data (F-KK-2); and the lock scale is fixed at 3.849 TeV, squarely inside the corridor where the full Run-2 + Run-3 top-pair and dijet samples already have leverage.

The operational consequence stands without any yield number: a cluster-aware, top-rich, leptophobic reinterpretation of the existing high-mass $t\bar{t}$ and dijet data is the test that matters now, and the absence of any emerging structure in that reinterpretation would itself begin to constrain the current implementation.

The right way to say it is therefore this. ATLAS and CMS *could* see the first real GG-6-motivated hints already in the current 2026 closeout analyses, especially in $t\bar{t}$ and boosted hadronic channels. A completely decisive verdict is more plausibly a Run-3-plus-early-HL-LHC story. But the theory has already entered the regime in which “wait until 2040” is no longer the correct experimental attitude.

32.6 Near-term reach in practical language

It is prudent not to overstate significance projections. The cleanest public machine-based statements are these. Official CERN documents define HL-LHC as the facility that will deliver the data volume required for decisive tests in the TeV range (CERN, 2023). Official ATLAS and CMS upgrade pages describe detector and trigger improvements specifically aimed at maintaining performance under $\mathcal{O}(200)$ pile-up conditions (ATLAS Collaboration, 2024; CMS Collaboration, 2024e). Public high-mass resonance analyses already reach to several-TeV masses in the clean channels. Therefore the cautious but strong statement is: *a dedicated cluster-aware combined-channel analysis of the full Run-2 + Run-3 sample can plausibly bring the* $3.85\,\mathrm{TeV}$ *GG-6 target into the first genuinely discriminating sensitivity regime, and the HL-LHC should be able either to confirm it or to exclude it across the geometry-allowed corridor.*

Table 20: Near-term discovery logic for the first GG-6 cluster. The entries are intentionally qualitative where public collaboration reach projections do not yet map one-to-one onto the clustered GG-6 benchmark. The qualitative entries are qualitative because the public projections they would have to be read against were computed for single-resonance benchmarks, and this paper's target is a cluster. Forcing a number onto those rows would give a precision the underlying projections do not have. Where a number can be stated honestly it is stated.

Program stage	What changes	Implication for the GG-6 first cluster
Run 2 only	Mature high-mass searches exist, but mostly in benchmark-specific one-pole templates	Important bracketing information exists, but not a decisive dedicated GG-6 test.
Run 2 + full Run 3	Larger data set, improved detector performance, first public analyses with full combined sample	First genuinely meaningful cluster-aware reinterpretation should become possible.
Early HL-LHC	Substantially larger $\mathcal{L}$, improved timing, tracking, and reduced-stream flexibility	Best stage for first 3–5σ evidence if the GG-6 signal sits near the center of the expected corridor.
Full HL-LHC	$3000\,\mathrm{fb}^{-1}$ per experiment, mature Phase-2 detector and trigger ecosystem	Should provide either discovery-level combined sensitivity or a strong module-level falsification of the current GG-6 implementation.

32.7 The falsification timeline in calendar years

Every sensitivity statement in this Part has so far been phrased against integrated luminosity. This closing subsection converts the whole picture into the coordinate an experimental program actually runs on: the calendar. The conversion is natural for GG-6 in a way it is not for scanned-parameter models. Because the theory predicts a *single* derived mass, the standard sensitivity-versus-mass presentation (Fig. 9) collapses to one vertical slice at 3.849 TeV; Figure 10 re-plots that slice against time, with the derived predictions appearing as *horizontal bands* — they do not move, because there is nothing in the theory that can move them — and the projected experimental sensitivities as descending curves. A theory with a mass dial cannot draw this figure; its bands would slide. The entire numerical content of the figure is collected in Table 21.

Construction. The bands are the derived $\sqrt{s} = 13$ TeV rates of Table 19, carried unchanged. The sensitivity curves take the observed full-Run-2 95% CL limits at 3.849 TeV (Table 21, middle block) and scale them as $\sqrt{140\,\mathrm{fb}^{-1}/L(Y)}$, where $L(Y)$ is the integrated luminosity on tape per experiment in year Y under the CERN schedule as of July 2026 (bottom block): Run 3 ended on 14 June 2026 with approximately 540 fb^{-1} delivered since startup, of which ≈ 466 fb^{-1} is analysis-quality (CERN, 2026b); LS3 runs to 2030; the HL-LHC takes physics data from mid-2030 at 250–350 fb^{-1} per year, accumulating ≈ 3.4 ab^{-1} by the end of the current baseline in 2041 (CERN, 2023, 2026b). Beyond the LHC the calendar is now fixed by the December-2025 European Strategy recommendation of FCC-ee as the next flagship (European Strategy Group, 2025): FCC-ee (physics from $\sim$2047) probes the cluster only indirectly, and the first machine to see the lock scale directly again is FCC-hh, with physics starting around 2074 in the staged Feasibility-Study scenario (FCC Collaboration, 2025).

Honesty conventions. Three conservatisms are built in, all pushing the drawn crossings *later* than reality. The curves are single-experiment (an ATLAS+CMS combination gains a further $\approx \sqrt{2}$ in effective luminosity); the bands are the 13-TeV rates, whereas at 14 TeV every signal rises by 35–55% while the backgrounds rise less steeply; and the octet is read against the broad ($\Gamma/m = 30\%$) template although its predicted $\Gamma/m \simeq 15\%$ width would be better matched — and more strongly constrained — by a dedicated intermediate-width reinterpretation. One convention runs the other way and must be stated just as plainly: the curves are drawn against the year the *data are on tape*, not the year the analysis appears. Historically, flagship full-dataset resonance results lag their data by two to seven years (the full-Run-2 $t\bar{t}$ search ATLAS Collaboration, 2025a appeared seven years after the end of Run 2), so published verdicts trail these curves by that much.

Reading. The figure compresses the collider fate of GG-6 into one visual sentence. The broad-template $t\bar{t}$ curve — the only curve that matters, because after F-KK-2A's closure and T-KK-2c's requalification the octet carries the entire LHC falsifier — crosses the top of the g^* band with data on tape around 2032 and exits its bottom around 2040: **F-KK-2B is decided, discovery or exclusion, within Run 4–5 of the HL-LHC**. Every other band lies below every curve everywhere in the window: the Z'_B across its entire L-KK-1 menu and the graviton in every channel remain LHC-invisible through 2050, so their absence from LHC searches is a standing prediction of the framework, not a deficit. After the HL-LHC baseline ends in 2041 the curves go flat — no approved machine probes 3.85 TeV directly for a generation — which gives the window its weight: the LHC's verdict on the GG-6 collider sector will be substantially in hand by $\sim$2040, final (in published form) by $\sim$2047, and nothing else can speak to it until FCC-hh, where $\sigma(G_1) \approx 205$ fb would turn the surviving cluster into early-running spectroscopy.

Table 21: Numerical inputs to Figure 10. Top: the GG-6 bands (derived LO×K rates at $\sqrt{s} = 13$ TeV, from Table 19; bands fold ±20% PDF/scale). Middle: the observed 95% CL anchors at 3.849 TeV from which the curves are $\sqrt{L}$-scaled ($t\bar{t}$ rows: HEPData record ins3094414). Bottom: the luminosity-on-tape model per experiment (CERN schedule, July 2026). Sensitivity in year Y: $\text{limit}_{140} \times \sqrt{140/L(Y)}$.

GG bands at $M = 3.849$ TeV (13 TeV, LO×K)			
Signal	**σ×BR [band], fb**	**Γ/m**	**Coupling provenance**
$g^* \to t\bar{t}$	8.7 [7.0–12]	0.15	$\xi_{\rm UV} = -0.194$, theorem (T-KK-2)
$g^* \to jj{+}b\bar{b}$ (σBA, A=0.55)	0.30 [0.23–0.43]	0.15	same
$Z'_B \to t\bar{t}$	0.016–0.60	0.002–0.080	L-KK-1 menu, $g_B = 0.14$–0.89 (§10.10)
$G_1 \to t\bar{t}$	0.045 [0.021–0.10]	0.063	$\kappa_{gg} = c_K/K = 0.0136$, derived (T-KK-2c)
Observed 95% CL anchors at 3.849 TeV (full Run 2, single experiment)			
Channel / template	**Limit [fb]**	**Dataset**	**Source**
$t\bar{t}$, broad $\Gamma/m = 30\%$	25.5	140 fb^{-1}	ATLAS Collaboration (2025a)
$t\bar{t}$, narrow	4.9	140 fb^{-1}	ATLAS Collaboration (2025a)
dijet, narrow qq (σBA)	8.0	137 fb^{-1}	CMS Collaboration (2020b)
Luminosity on tape per experiment $L(Y)$ (13-TeV-equivalent)			
Year / era	**L [fb^{-1}]**	**Basis**	
2025 (Run 3)	442	Run-2 140 + Run-3 on tape (CERN, 2025)	
2026–29 (LS3 from 29 Jun 2026)	466	Run 3 ended 14 Jun 2026; ∼540 delivered (CERN, 2026b)	
2030 (HL-LHC from mid-year)	≈526	ramp half-year (CERN, 2026b)	
2031–41 (HL-LHC)	+200–275/yr	leveled 250–350/yr (CERN, 2023); ≈3450 by end-2041	
2042–50 (post-LHC gap)	3451	baseline ends 2041; next direct probe FCC-hh ∼2074 (FCC Collaboration, 2025)	

33 Why This is the Best Collider Killer Experiment

33.1 Sharpness compared with other energy-frontier claims

The phrase "killer experiment" should be used carefully. In the broadest sense, GG-6 has other important tests: axion-sector tests, dark-energy tests, Higgs-coupling tests, and possible quantum-open-system tests. But at the energy frontier, the KK search is categorically the cleanest falsifier because it combines three features that rarely coexist in BSM physics.

First, the target mass is unusually sharp. There is no scanned supersymmetric spectrum, no landscape of vacua, and no flexible mediator mass. The first cluster is tied to $K = 35$.

Second, the same integer that fixes the lock scale also participates in other successful internal ledger relations, such as the weak scale. So the mass prediction is not isolated; it is part of a coherent triad. The interlocking extends across the predictions branch in both directions: the dark-matter companion stakes a haloscope line at 1.95 ± 0.12 GHz, twice split by exactly $1/35$ — the same ledger integer that sets $M_{\rm lock}$ — and the black-hole companion reads gravitational-wave ringdowns as curvatures of the same free-energy functional. Kill the lock cluster and those accounts inherit the wound; conversely, a confirmed $1/35$ doublet would make the 3.85 TeV search a confirmation hunt rather than an exploration. No other energy-frontier proposal carries hostages in neighboring experiments of this kind.

Third, the experimental statement is binary in a practically useful way. Either the spin-2-anchored, octet-led signal cluster appears in the neighborhood of 3.85 TeV, or increasingly powerful data exclude it across the geometry-allowed coupling corridor. That is much more decisive than many BSM searches in which a null result simply pushes a free scale upward.

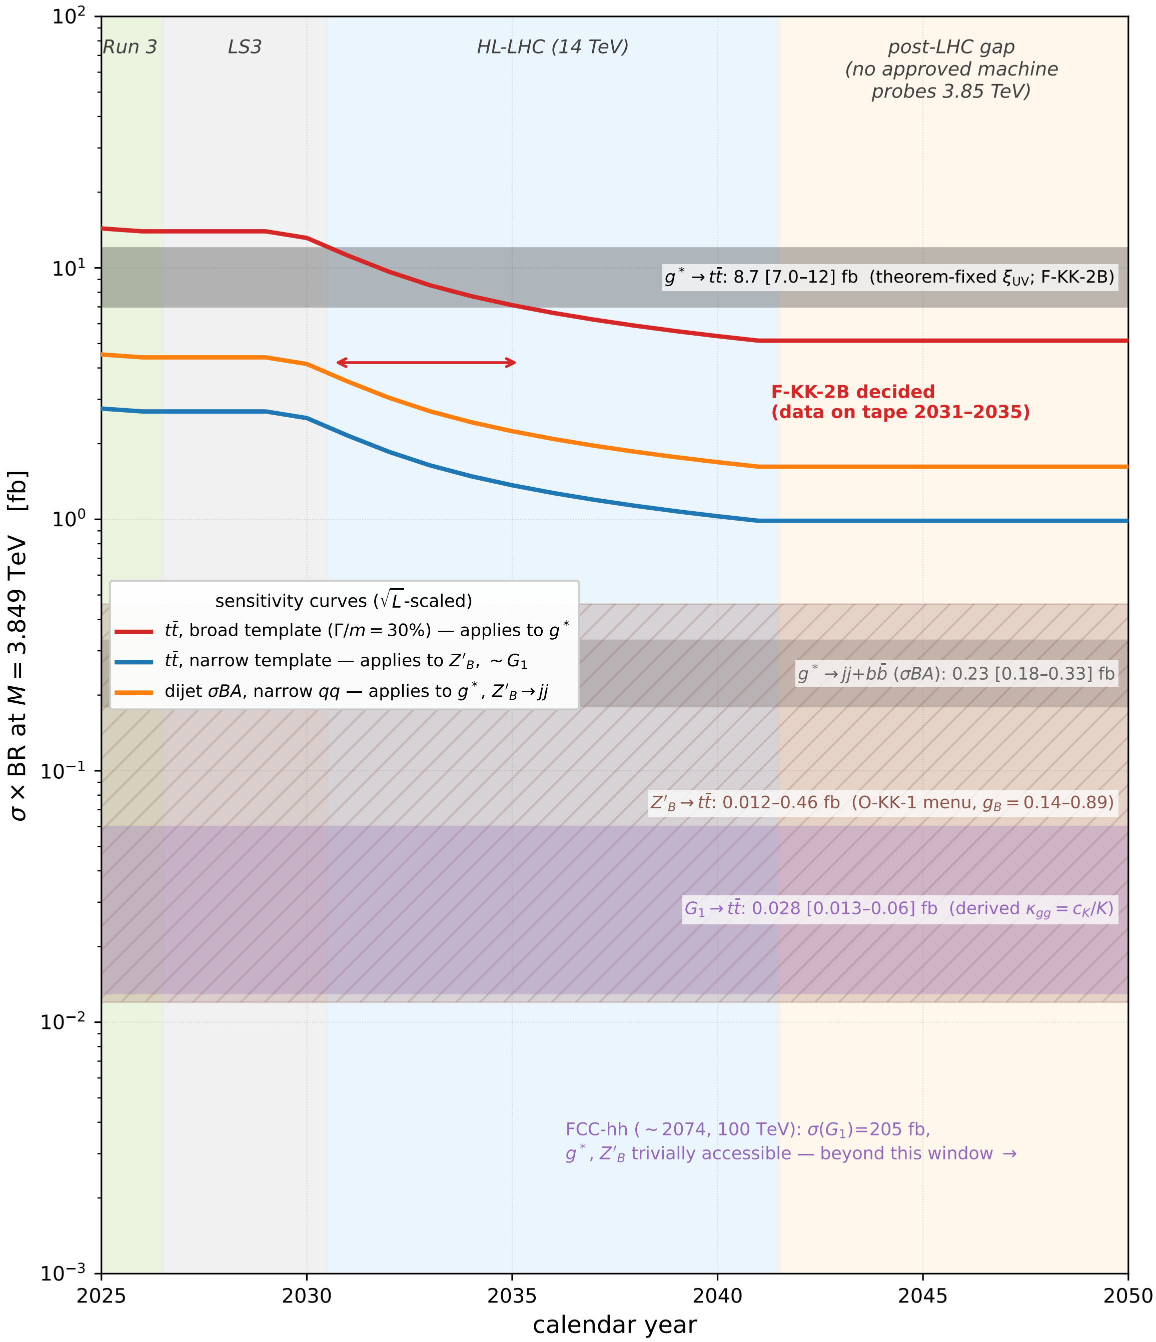

Figure 10: The falsification timeline of the GG-6 lock cluster at $M_{\text{lock}} = 3.849$ TeV, 2025–2050. Horizontal bands: the derived GG-6 predictions of Table 21 — flat in time because the theory has no dial to move them. Descending curves: projected single-experiment 95% CL sensitivities at 3.849 TeV, $\sqrt{L}$-scaled from the observed full-Run-2 anchors on the July-2026 CERN schedule (shaded eras: Run 3, LS3, HL-LHC, post-LHC gap). The broad-template $t\bar{t}$ curve crosses the g^* band with data on tape between ~ 2032 and ~ 2040: F-KK-2B — the sole LHC arm of the falsifier after F-KK-2A's closure and T-KK-2c's requalification of F-KK-1 — is decided within HL-LHC Run 4–5. All other bands remain below every curve through 2050 (their LHC absence is a prediction); the $G_1 \to \gamma\gamma$ line (7×10^{-6} fb) lies far below the plotted window. Conservatisms and the publication-lag convention are stated in the text; all inputs in Table 21.

33.2 What discovery would establish

A convincing discovery of the spin-2-anchored cluster near 3.85 TeV, together with a top-rich branching pattern and at least one of the expected partner signatures, would establish several things at once. It would establish that an extra-dimensional description is realized in nature. It would establish that the first accessible scale is not generic but locked to a highly specific infrared value. It would give strong support to the non-universal overlap picture. And it would make a direct future search for the second radial family and the KK-gluon corridor mandatory.

33.3 What absence would establish

Conversely, a sufficiently strong combined-channel exclusion would do more than say "one benchmark failed." It would pressure the GG-6 constitutional chain itself. If the full Run-2 + Run-3 + early-HL-LHC program excludes the first-cluster corridor across $t\bar{t}$, dijet, hh, dilepton, and diphoton channels within a cluster-aware window around 3.85 TeV, then at least one of the following must be false:

1. the OGN-5 lock index $K = 35$ does not control the collider infrared scale as claimed,
2. the 6D spectral interpretation of the first lock family is wrong,
3. or the overlap/coupling structure is radically different from the presently derived corridor.

Any of these would count as a deep revision of the current GG-6 implementation.

33.4 A pre-declared kill criterion

Following the GG-Constitution's insistence on measurement and falsifiability, the collider-level kill criterion can be written explicitly:

> *If, after analysis of the combined Run-2 + full Run-3 + early-HL-LHC data set, ATLAS and CMS observe no excess above background in any cluster-aware combination of $t\bar{t}$, boosted dijet, hh, $\ell^+\ell^-$, and $\gamma\gamma$ searches within a window of order* ± 0.6 TeV *around* 3.849 TeV, *and if the corresponding* 95% *CL upper limits on the combined signal strength fall below the full geometry-allowed GG-6 corridor, then the present lock-scale implementation of GG-6 is ruled out at the collider level.*

This wording is intentionally stronger than a single-channel exclusion and intentionally weaker than an impossible demand for a fully exact one-pole template. It is the right balance between honesty and decisiveness. §31.5 converts it into dated numbers: the un-suppressible color-octet line is crossed between roughly 0.5 and 2 ab^{-1} of HL-LHC data. (With Corollary T-KK-2c the spin-2 member sits a factor $\sim$25 below any HL-LHC sensitivity, so the criterion's teeth at the LHC are entirely the spin-1 sector's; the graviton line is FCC-hh's to test.)

Table 22: Why the KK search is the sharpest *collider* falsifier of the current GG-6 implementation. Sharpest is a claim about the falsifier and not about the physics: the collider test is the one whose negative result is hardest to argue with, because the mass is fixed by an integer and the machine already exists. Other tests in the cohort probe deeper structure and take longer to settle; this one is decidable inside a doctoral career.

Criterion	Why the KK search is unusually strong
Specificity of mass target	The first cluster is centered at 3.849 TeV, not "somewhere in the few-TeV region."
Connection to other theory outputs	The same lock ledger also enters electroweak-scale relations and the broader hierarchy story.
Binary character	Discovery or combined-channel exclusion both have immediate meaning; a null result does not merely move a free parameter.
Energy-frontier accessibility	The relevant mass scale lies in the LHC/HL-LHC domain rather than only in far-future or purely cosmological observables.
Pattern richness	The search probes not only one state but a correlated pattern: cluster shape, top-rich branching, and higher partners.

34 Long-Term Future Prospects

34.1 HL-LHC as the decisive machine

Official CERN documentation frames the HL-LHC as a 14 TeV machine delivering more than a factor-of-five luminosity gain over the nominal LHC design, with a target integrated luminosity of about 3000 fb^{-1} per experiment (CERN, 2023). That statement alone already implies why HL-LHC is decisive for GG-6. The first cluster is not at some inaccessible tens-of-TeV frontier; it is precisely in the corridor where HL-LHC will have maximal leverage.

34.2 ATLAS Phase-2 upgrades

ATLAS describes a Phase-2 detector centered on a fully new all-silicon inner tracker (ITk), forward tracking extended to $|\eta| < 4$, a new level-0 trigger architecture, and the High-Granularity Timing Detector (HGTD). ATLAS states that the ITk extends pseudorapidity coverage from $|\eta| < 2.5$ to $|\eta| < 4$ and that the HGTD will provide 30 to 50 ps timing in the forward region (ATLAS Collaboration, 2024). These are not generic improvements; they are directly relevant to a GG-6 search because a clustered signal at 3.85 TeV is especially sensitive to pile-up rejection, forward acceptance, and improved boosted-object tagging.

34.3 CMS Phase-2 upgrades

CMS similarly describes the HL-LHC era as a major transformation with a new inner tracker, the high-granularity endcap calorimeter (HGCAL), and the MIP Timing Detector (MTD). Official CMS descriptions emphasize that the MTD aims at about 35 ps time resolution for charged particles, precisely to untangle 140 to 200 pile-up interactions per bunch crossing (CMS Collaboration, 2019, 2024e). Again, for GG-6 this is directly useful: top-rich and dijet-rich high-mass searches benefit strongly from pile-up mitigation and improved object assignment.

34.4 FCC-hh, muon collider, and CLIC

If GG-6 is confirmed, the future collider program becomes a spectroscopy enterprise. The official FCC-hh concept is a 100 TeV hadron collider with integrated luminosity goals at least five times larger than HL-LHC and a mass reach extended by almost an order of magnitude (FCC Collaboration, 2018). Such a machine would not merely confirm the first cluster. It would map the entire low-lying GG-6 staircase.

Official CERN materials also note that a 10 TeV muon collider could be competitive with a 100 TeV proton collider for some high-energy-frontier questions (CERN, 2024b). For a 3.85 TeV spin-2 state, a muon collider would offer a uniquely clean environment and exquisite angular analysis. CLIC, on the other hand, is officially staged at 380 GeV, 1.5 TeV, and 3 TeV (CERN, 2024a). That makes direct production of a 3.85 TeV state inaccessible at the baseline stages, but virtual effects and Higgs-coupling precision would still provide important complementary information.

34.5 Indirect probes

Even if the first signal comes directly, indirect probes remain important. In the current GG-6 picture one expects modest Higgs-coupling deviations from Hosotani-related mixing, likely at the percent level, and electroweak precision shifts of order 10^{-3} in some observables. Those effects are not by themselves the decisive falsifier, but they are valuable pattern checks. If a 3.85 TeV cluster appears with no companion distortions anywhere in Higgs or electroweak observables, the interpretation would need care.

35 A Trigger and Analysis Roadmap for ATLAS and CMS

This section proposes concrete improvements to ATLAS/CMS triggers and related search strategy. The proposal is intentionally pragmatic. The goal is not to ask the experiments to rebuild their trigger systems for one speculative model. It is to point out modifications that are already aligned with broader new-physics goals and that would be especially useful for a clustered GG-6 signal.

35.1 Use the high-rate frontier aggressively

ATLAS has already demonstrated the power of trigger-level analysis (TLA), in which only essential reconstructed high-level quantities are stored, reducing event size dramatically and permitting much higher rates (ATLAS Collaboration, 2025b). CMS has similarly introduced Level-1 Scouting, which stores reduced information from all collisions in a triggerless or trigger-reduced format (CMS Collaboration, 2025). These methods are already changing what the experiments can do in high-rate regimes.

For the GG-6 program the lesson is immediate: a dedicated reduced-format stream for the 3 to 6 TeV region should be treated as a priority. The relevant reduced quantities are not exotic. They include high-mass dijet kinematics, boosted top substructure variables, Higgs-tagging observables, b-tagging summary information, coarse timing variables, and enough leptonic information to retain a clean cross-check in the $\ell^+\ell^-$ channels.

35.2 Make the search cluster-aware at trigger level

Most online and semi-online resonance thinking is one-pole thinking. A clustered GG-6 signal changes the optimization. Instead of only centering windows on one putative pole mass, the trigger and scouting logic should support a broad 3.2 to 5.2 TeV search region in which multiple nearby structures are explicitly allowed. At analysis level this means storing enough resolution and calibration information to fit Eq. (29). At trigger level it means avoiding overly narrow single-mass assumptions that could dilute sensitivity to a split or shoulder-like feature.

35.3 Promote top and Higgs boosted objects to first-class citizens

Because GG-6 is likely top-rich, the trigger grammar should treat boosted top and boosted Higgs categories as first-class citizens in the high-mass new-physics program. In practice this means:

- preserving reduced top-tag and Higgs-tag discriminants in scouting/TLA streams,
- keeping high-p_T b-tag summary information available online or quasi-online,
- and designing combined $t\bar{t}$+hh+dijet anomaly streams rather than siloed single-object streams.

These are useful changes even beyond GG-6. That is a virtue, not a problem.

35.4 Exploit precision timing as a resonance-quality variable

Both ATLAS and CMS emphasize precision timing for pile-up control at HL-LHC (ATLAS Collaboration, 2024; CMS Collaboration, 2019). In the GG-6 context, timing should be treated not just as a vertexing aid but as a resonance-quality variable. High-mass boosted hadronic final states are especially sensitive to pile-up contamination in substructure observables. Timing-assisted cleaning should therefore be included explicitly in any dedicated GG-6 reinterpretation.

35.5 A joint benchmark release would help enormously

The experiments are understandably conservative about model proliferation. The most useful single step for the community would therefore be modest: define a public benchmark family capturing the essential GG-6 phenomenology without overloading the analysis with theory details. Such a benchmark could include

1. a first cluster centered at 3.85 TeV,
2. an optional split of order a few hundred GeV,
3. non-universal branching ratios favoring $t\bar{t}$, dijets, and hh,
4. a mild suppression of $\ell^+\ell^-$ and $\gamma\gamma$,
5. and the second radial family near 7 TeV ($n = 2$).

That alone would transform the current status from a set of partial proxies into something much closer to a real test.

Table 23: Proposed trigger and reduced-stream improvements for the GG-6 KK search. All proposals are framed so that they also benefit a wider class of high-mass BSM programs. Every proposal here is framed to serve a wider class of high-mass searches, and that is a condition on the proposal rather than a courtesy. A trigger change justified by one theory's target is a change no collaboration should make; a change that improves the reach of a whole category and happens to enable this search as well is a change worth arguing for.

Proposal	Main instrument	Immediate benefit for GG-6	Why it is technically plausible	Broader benefit beyond GG-6
Dedicated high-mass reduced stream	ATLAS TLA / CMS scouting	Retains far more high-mass hadronic events in the 3 to 6 TeV window	Both collaborations already operate reduced-format high-rate streams (ATLAS Collaboration, 2025b; CMS Collaboration, 2025)	Benefits dijet, dark-sector, and generic broad-resonance programs.
Cluster-aware template storage	HLT / offline fit model	Preserves sensitivity to doublet or triplet structures around the first cluster	Requires software and analysis choices more than new hardware	Helps any model with nearby poles or non-trivial interference.
Boosted top + Higgs summary variables in scouting	Trigger objects + reduced event content	Improves $t\bar{t}$ and hh sensitivity where GG-6 is strongest	Reuses existing top/Higgs tagging infrastructure with reduced payload	Benefits generic heavy resonances, vector-like quarks, and heavy Higgs sectors.
Timing-assisted substructure cleaning	HGTD / MTD + offline reconstruction	Stabilizes top and dijet mass/substructure variables under 200-pileup conditions	Built directly on official Phase-2 timing goals (ATLAS Collaboration, 2024; CMS Collaboration, 2019)	Benefits essentially all boosted-object searches at HL-LHC.
Common public benchmark family	ATLAS/CMS reinterpretation note	Converts current proxy limits into a genuine GG-6 test	Requires collaboration-level benchmark agreement, not detector changes	Helps the community compare clustered or non-universal extra-dimensional models more systematically.
Joint channel combination with $t\bar{t}$+dijet+hh+$\ell^+\ell^-$	Analysis strategy	Matches the actual multi-channel logic of GG-6 rather than overvaluing any single clean channel	Statistically straightforward once channel-specific templates are defined	Improves discovery robustness for many non-universal resonance models.

Part VII took stock of reality, in plain terms: nothing published excludes the octet at its derived rate — the clean-channel limits that look threatening do not apply to a leptophobic, top-philic cluster, and the top-pair searches that do apply used templates mismatched to this broad, interfering shape; meanwhile the free-field horn of the old dichotomy is already dead, killed by exactly these data. The kill criterion was priced (two inverse attobarns) and the timeline drawn. But a frontier position and a timeline are still projections built on other people's analyses. The next Part does something more direct: it simulates the discovery itself, at the fidelity of a real analysis, and asks when — on the actual calendar — the question gets answered.

36 Part VII in plain English: what has already been excluded, and what has not been looked for

This is the Part where the prediction meets the published record, and the record is more interesting than a simple pass or fail.

Standard benchmarks in this mass range are strongly excluded, and the Part says so first and without qualification. What has not been excluded is this paper's target, because no published search has been designed for it. Every existing limit was computed against a single resonance with a chosen width; the prediction here is three states near one scale with fixed relative positions, and an exclusion computed for one shape does not transfer to another by rescaling.

That is not a loophole, and the Part is careful about the difference. A loophole is what a theory with dials uses to survive a null result. This is a mismatch between a template and a target, and the remedy is specific and public: a dedicated reinterpretation of data that already exists.

The Part also states the kill criterion. There is a luminosity and a set of channels beyond which a null result ends the reading, and the number is printed here rather than left for the discussion. The trigger roadmap that closes the Part is framed so that every change it asks for would help a whole category of high-mass searches, because a trigger change justified by one theory's target is a change no collaboration should make.

Part VIII
Simulating the Discovery

What ATLAS and CMS will actually record — the anchored HL-LHC simulation, the systematics wall, and the pre-registered search: 2026 / 2031 / 2033 / 2041

37 What This Part Adds, and the Rules of the Game

Part VI derived the *shapes* of the $t\bar{t}$ signal — line shape, interference sign, angular distributions, polarization — and §32 placed the search on the calendar through the projected limit curves. Between those two lies the question an experimentalist actually lives with, and this Part answers it: *if the color octet exists at* $M_{\text{lock}} = 3.849$ *TeV with the derived rate, what do ATLAS and CMS record, bin by bin, at the end of 2026, 2031, 2033, and 2041 — and when, exactly, can they claim discovery?* The answer is quantitative, and parts of it are surprising: the excess is statistically overwhelming almost immediately, yet the discovery is gated not by luminosity but by *background knowledge*; a reanalysis of data already on tape carries evidence-level sensitivity; and the fact that GG-6 pre-specifies mass, width, rate, sign, and angular shape changes the statistics of the search itself in ways the published pipeline does not exploit.

This Part is written deliberately at the level of a first-year graduate student joining an ATLAS or CMS analysis group. Everything is derived or explained in single steps: how the anchor analysis actually works, cut by cut (§38); how our simulation is built on it (§39); the statistics of discovery from the Poisson likelihood up (§40); how the toy is calibrated to the real analysis (§41); why the eye wildly overestimates significance on a spectrum plot (§42); the four snapshots themselves (§43); and the statistics of the pre-registered search — look-elsewhere, matched templates, the angular dimension, and the two-experiment combination (§44). A reader who has never run a resonance search should be able to follow every number; a reader who has run one should be able to check every number.

The rules of the game. Four framing rules keep this Part honest. (i) Everything experimental is anchored to the most relevant published analysis — the ATLAS full-Run-2 $t\bar{t}$-resonance search (ATLAS Collaboration, 2025a) — whose backgrounds, efficiencies, resolution, binning, and systematics we reviewed item by item. We do not invent detector performance; we borrow measured performance. (ii) Our simplified statistical machinery is *calibrated* against that analysis's published expected limits before any projection is made; the calibration factor is reported, not hidden (§41). (iii) Phase-2 detector performance is assumed equal to Run-2 — a stated, conservative simplification, since the ATLAS ITk/HGTD and CMS MTD upgrades (ATLAS Collaboration, 2024; CMS Collaboration, 2019) specifically target *better* boosted-top performance under HL-LHC pile-up. (iv) Grades follow the cohort ladder: the signal model is [T/D*] from Parts VI and earlier (nothing about the signal is

Table 24: ATLAS merged-signal-region post-fit background at 140 fb^{-1} (HEPData record of ATLAS Collaboration, 2025a), the anchor of this Part. "Post-fit" means *after* the profile-likelihood fit has adjusted all backgrounds within their uncertainties against the data; these are the yields and residual uncertainties the real analysis actually achieved. "Other" collects $t\bar{t}V$, $t\bar{t}H$, diboson, and fake-lepton contributions.

$m_{t\bar{t}}$ bin [GeV]	SM $t\bar{t}$	single top	V+jets	Other	Total ± syst	Observed
3000–3500	60%	11%	24%	5%	$53.4 \pm 6.9\%$	47
3500–4200	47%	8%	34%	11%	$21.6 \pm 13.2\%$	18
4200–4900 (+ovf)	40%	2%	40%	17%	$6.1 \pm 22.6\%$	6

adjusted here); the experiment model is [E], anchored; the projection scenarios are labeled where they enter.

How a resonance search works, in one paragraph. For the newcomer, the pipeline this Part simulates has five stages, and each gets its own subsection below. *Production*: protons collide; a parton pair (here $q\bar{q}$) either produces a $t\bar{t}$ pair through the ordinary QCD continuum or through the resonance, coherently. *Trigger*: of the $\sim$40 MHz of bunch crossings, a hardware-plus-software cascade keeps $\mathcal{O}(1)$ kHz for permanent storage; an event that fails the trigger is lost forever, so the first design question of any search is which trigger it lives on. *Selection*: offline, a small set of cuts isolates the topology — here, one lepton plus one large, fat jet containing the entire hadronic top. *Reconstruction*: from the selected objects the invariant mass $m_{t\bar{t}}$ is rebuilt event by event, with a resolution that decides how sharply the resonance stands out. *Statistical fit*: the observed $m_{t\bar{t}}$ histogram is compared to the background prediction by a binned profile-likelihood fit, and the excess (or its absence) is converted into a significance or a cross-section limit. Every one of these stages costs signal, and the honest bookkeeping of those costs is the difference between a projection and a wish.

38 The Anchor Analysis, Read Like an Analyst

This section condenses our systematic review of the ATLAS full-Run-2 $t\bar{t}$-resonance search (ATLAS Collaboration, 2025a) — 140 fb^{-1} at 13 TeV, one-lepton and two-lepton channels, resonance masses 0.4–5 TeV, eight signal regions fit simultaneously. It is the knowledge base the simulation runs on, and it is worth reading closely: the projections of §43 inherit their realism from precisely these numbers.

38.1 The dominant backgrounds, and how each is estimated

At $m_{t\bar{t}} > 3$ TeV the event rate is a few events per year per experiment, and the question "what else can fake a 3.8 TeV $t\bar{t}$ pair?" dominates the analysis design. The composition in the boosted (merged) one-lepton signal region — the region that drives all high-mass sensitivity, and the one we anchor to — is reproduced in Table 24.

Three features of this table repay attention. First, *the composition drifts*: SM $t\bar{t}$ falls from 60% to 40% across the tail while V+jets rises to $\sim$40%. The physics is simple once seen: the $t\bar{t}$

continuum at these masses comes from $q\bar{q}$ annihilation on steeply falling antiquark distributions, while a W+jets event needs only one hard valence quark plus a gluon, so its spectrum falls less steeply. A projection that scaled "$t\bar{t}$ only" would therefore get the tail wrong by a factor approaching two; our anchoring to the measured *total* avoids this trap automatically. Second, *the uncertainties explode with mass*: 6.9% → 13.2% → 22.6% across the three bins. These are post-fit values — the residual ignorance *after* the fit has used every control region — and they will turn out to be the single most important input to this entire Part (§40.3). Third, *the estimation methods differ by component*, and knowing them tells us which uncertainties can shrink with more data: SM $t\bar{t}$ is a Powheg+Pythia8 template, NNLO-reweighted in top p_T and $m_{t\bar{t}}$, with its overall normalization left *free* in the fit (the fit determines it from the data itself — remember this when we come to §44.5); W+jets is a Sherpa template normalized in situ by the measured W-charge asymmetry (W^+/W^- production in pp is asymmetric because the proton holds two u valence quarks and one d; the measured asymmetry pins the normalization to ±3%, after which 50% normalization and heavy-flavor shape variations cover the tail); multijet fakes are estimated from data by the matrix method (loose-versus-tight lepton identification inverted into a fake probability); single top is taken from simulation with the $Wt/t\bar{t}$ interference handled by diagram removal.

38.2 The trigger, and why one lepton is enough

The analysis records events with single-lepton triggers alone: thresholds $p_T > 26$ GeV for muons and $E_T > 26$ GeV for electrons through 2016–2018, with isolation requirements at the lowest thresholds; the offline requirement $p_T > 28$ GeV then sits on the *trigger plateau* — the region where the trigger efficiency has flattened near its maximum, so that per-event efficiency no longer depends on the exact threshold and can be measured once and applied cleanly.

Why is a 26 GeV lepton threshold comfortable for a 3.85 TeV resonance? Follow the kinematic chain, because it is a useful habit: each top from the decay carries $E_t \simeq M_{\text{lock}}/2 \approx 1.9$ TeV. The leptonic top splits its energy among b, ℓ, and ν; in the top rest frame the lepton takes on average roughly a third of the W energy, and after the boost the typical laboratory lepton energy is $\mathcal{O}(E_t/4) \approx 500$ GeV — more than an order of magnitude above threshold. Even the unluckiest decay orientations rarely push the lepton below 50 GeV. The single-lepton trigger is therefore effectively 100% efficient for this signal *at any* $m_{t\bar{t}}$ of interest, and no large-radius-jet or missing-energy trigger is needed. This matters doubly for the HL-LHC projection: lepton thresholds are among the few trigger quantities expected to *survive* the pile-up of 200 simultaneous collisions roughly unchanged ($\sim$25–30 GeV with the upgraded Level-0/HLT chains; ATLAS Collaboration, 2025b), so the trigger leg of our "Run-2 performance" assumption is if anything conservative.

38.3 Event selection and reconstruction, cut by cut

The one-lepton selection, in the order an analyst would apply it: exactly one electron or muon with $p_T > 28$ GeV (a second lepton vetoes the event into the dilepton channel); missing transverse momentum $E_T^{\text{miss}} > 20$ GeV and $E_T^{\text{miss}} + m_T^W > 60$ GeV, which together demand a genuine escaping neutrino and suppress multijet fakes (a fake lepton comes with no neutrino, so both quantities are small); then the *merged-topology* requirements that define the high-mass regime. At $m_{t\bar{t}} \gtrsim 2$ TeV the hadronic top is so boosted that its three quark jets merge into a single cone: the analysis requires one anti-k_t $R = 1.0$ "large-radius" jet with $p_T > 350$ GeV and $|\eta| < 2$, *trimmed* (soft subjets removed, which strips pile-up while preserving the hard three-prong structure) and identified as a top by a deep-neural-network tagger at the 80%-efficiency working point — the operating threshold chosen such that 80% of true boosted tops pass, at a light-jet rejection near 10 at these energies. On the leptonic side, a small-radius jet within $\Delta R < 2.0$ of the lepton must be b-tagged (77% working point), completing the $t \to b\ell\nu$ hypothesis.

Reconstruction then rebuilds the resonance mass from four pieces: the tagged large-R jet (the hadronic top), the lepton, the b-jet, and the neutrino. Only the neutrino's longitudinal momentum is unmeasured, and it is recovered by a constraint: demanding $(p_\ell + p_\nu)^2 = m_W^2$ with $\vec{p}_T^{\,\nu} = \vec{E}_T^{\rm miss}$ yields a quadratic equation in p_z^ν,

$$p_z^\nu = \frac{\alpha\, p_z^\ell \pm E_\ell \sqrt{\alpha^2 - (E_T^{\rm miss} p_T^\ell)^2}}{(p_T^\ell)^2}, \qquad \alpha \equiv \frac{m_W^2}{2} + \vec{p}_T^{\,\ell} \cdot \vec{E}_T^{\rm miss}, \tag{76}$$

with the smaller-$|p_z^\nu|$ solution taken (and the discriminant set to zero when resolution pushes it negative). The result is a fully reconstructed $m_{t\bar{t}}$ event by event. The merged category is preferred over resolved-jet categories at high mass for exactly the reason it exists: one calorimetric object instead of three small jets means fewer combinatoric mistakes and a sharper mass peak.

38.4 What $\varepsilon A = 6.5\%$ is made of

The published signal efficiency-times-acceptance for a 3.5–4 TeV resonance, counting *all* $t\bar{t}$ decay modes in the denominator, is 6.0–6.8%; we use 0.065. A newcomer should be able to reconstruct this number, so here is the accounting. The one-lepton channel itself costs the branching ratio: $\text{BR}(t\bar{t} \to \ell + \text{jets}) = 2 \times \text{BR}(W \to \ell\nu) \times \text{BR}(W \to q\bar{q}') \approx 2 \times 0.217 \times 0.674 \approx 0.29$ for $\ell = e, \mu$ (the factor 2 because either top may decay leptonically). Within the channel, the surviving fraction is roughly the product of the top-tag working point (0.80), the b-tag working point (0.77), and the kinematic-plus-fiducial acceptance of the remaining cuts ($\sim$0.36; leptons or jets outside $|\eta|$ coverage, the ΔR matching, the $E_T^{\rm miss}$ cuts). The product $0.29 \times 0.80 \times 0.77 \times 0.36 \approx 0.065$ closes the ledger. This decomposition also identifies where HL-LHC improvements would land: tagger working points and b-tagging at 200 pile-up are exactly what the upgrades protect.

38.5 Resolution, binning, and the published sensitivity

The merged-channel mass resolution is $\sigma(m_{t\bar{t}})/m_{t\bar{t}} \approx 10\%$, dominated by the large-R jet energy scale and the neutrino reconstruction. Ten percent of 3.85 TeV is 385 GeV: comparable to, but smaller than, the octet's intrinsic width $\Gamma_{g^*} = 561$ GeV, so the observed line shape is intrinsic width convolved with a comparable experimental smearing — a genuinely *broad* object either way, roughly two analysis bins wide. The analysis bins the spectrum coarsely in the tail (bin edges at 3.0, 3.5, 4.2, 4.9 TeV with overflow), matched to that resolution. Its statistical machinery is a binned profile-likelihood fit over all signal regions simultaneously, with $\mathcal{O}(100)$ nuisance parameters for the systematic uncertainties; the published expected 95% CL upper limits at our mass — the numbers we calibrate against in §41 — are 8.93 fb (narrow template, 3.5 TeV), 6.41 fb (narrow, 4.0 TeV), and $\sim$25 fb for the broad $\Gamma/m = 30\%$ KK-gluon template.

39 The Simulation, Step by Step

39.1 Signal: from the derived line shape to events in bins

The truth-level signal is the coherent $q\bar{q} \to t\bar{t}$ spectrum of Part VI, Eq. (65): QCD continuum, Breit–Wigner, and interference with the theorem-fixed sign $\xi_{\mathrm{UV}} g_V^t < 0$ (constructive below the pole, destructive above), evaluated with CTEQ6 parton luminosities (Pumplin et al., 2002) at both 13 and 14 TeV, $\Gamma_{g^*}/m = 14.6\%$, and the channel-appropriate NLO K-factors of Part VI ($K_{q\bar{q}} = 1.3$ on the Breit–Wigner, $\sqrt{K_{q\bar{q}}}$ on the interference, since interference is linear in the resonance amplitude while the squared term is quadratic — a small but easily forgotten point). Detector response enters in two steps that mimic the anchor analysis. *Smearing*: the truth spectrum is convolved with a Gaussian of width $\sigma = 0.10\,m$,

$$\frac{d\sigma^{\mathrm{reco}}}{dm}(m) = \int dm' \, \frac{d\sigma^{\mathrm{truth}}}{dm'}(m') \, \frac{1}{\sqrt{2\pi}\,(0.10\,m')} \exp\left[-\frac{(m-m')^2}{2(0.10\,m')^2}\right], \tag{77}$$

which redistributes events between bins exactly as the detector does — note the visible consequence in Figure 12: the smeared peak leaks into the 3.0–3.5 TeV bin and pulls the apparent maximum slightly below the pole. *Efficiency*: the smeared spectrum is multiplied by $\varepsilon A = 0.065$ and integrated over the ATLAS bin edges. The event yield in bin i at integrated luminosity L is then $s_i = L\,\varepsilon A \int_{\mathrm{bin\,i}} dm\,(d\sigma^{\mathrm{reco}}_{\mathrm{BW}} + d\sigma^{\mathrm{reco}}_{\mathrm{int}})/dm$, and we track the Breit–Wigner and interference pieces separately throughout because they respond differently to the fit (§42).

39.2 Background: anchoring the total to the measured yields

We do not simulate W+jets, single top, and fakes from scratch — that is a multi-year collaboration effort, and pretending otherwise would be theater. Instead we exploit the fact that ATLAS has already *measured* the total background in exactly the region we care about (Table 24), and we build a two-parameter model whose shape is physics-motivated and whose normalization is fit to those measurements:

$$B(m) = A\,\frac{d\sigma^{\mathrm{reco}}_{\mathrm{SM,q\bar{q}}}}{dm} \times \left(\frac{m}{3849\,\mathrm{GeV}}\right)^n, \tag{78}$$

where the first factor is our own smeared LO $q\bar{q} \to t\bar{t}$ continuum (which carries the correct convolution of parton luminosities and the correct resolution smearing) and the power-law tilt $(m/3849)^n$ absorbs everything that distinguishes the full cocktail from pure LO $q\bar{q}\ t\bar{t}$: the V+jets admixture with its flatter slope, NLO shape corrections, and acceptance drifts. Fitting (A, n) to the three post-fit ATLAS yields reproduces all three within their uncertainties with $n \approx 0.76$ — positive, exactly as the composition drift of §38.1 predicts (the cocktail falls *less* steeply than pure $q\bar{q}\ t\bar{t}$). The same (A, n) applied to the 14-TeV continuum shape gives the HL-LHC background. The model is deliberately humble: it does not know more than ATLAS measured, and it is not allowed to.

39.3 From 13 to 14 TeV, and "14-TeV-equivalent" luminosity

Two collider energies enter. The end-2026 snapshot is simulated *natively at 13 TeV*: it is a reanalysis of the Run-2+3 dataset, 0.466 ab^{-1} already on tape. The HL-LHC snapshots run at 14 TeV, where the $q\bar{q}$ parton luminosity at 3.85 TeV is a factor $\simeq 1.55$ larger — at these masses the colliding partons carry momentum fractions $x_1 x_2 \gtrsim (3.85/13)^2 \approx 0.088$, deep on the steeply falling ends of the PDFs, so even a 7% energy increase buys a large rate gain. For the HL-LHC snapshots the 13-TeV legacy is folded in as $\approx$303 fb^{-1} of "14-TeV-equivalent" data — the 13-TeV yield divided by the signal-region rate ratio — the luminosity of 14-TeV running that would yield the same signal statistics — giving per-experiment totals of 0.56 ab^{-1} (end-2031, about a year and a half of HL-LHC physics), 1.09 ab^{-1} (end-2033), and 3.29 ab^{-1} (end-2041), from the schedule model of §32 (LS3 to mid-2030, HL-LHC physics thereafter). All results are single-experiment unless stated; the ATLAS+CMS combination is treated carefully in §44.4.

40 The Statistics of Discovery, from First Principles

Nothing in this section is original — it is the standard machinery of Cowan et al. (2011) — but the projections cannot be understood, and certainly cannot be checked, without it, so we derive it in full.

40.1 Counting, the likelihood ratio, and the meaning of Z

A single bin records n events where b background events are expected. "Significance" answers one question: *if there were no signal, how improbable is a fluctuation at least this large?* That probability p is conventionally quoted as the equivalent number of Gaussian standard deviations, $Z = \Phi^{-1}(1-p)$, so $Z = 5$ means $p = 2.87\times 10^{-7}$. The optimal test statistic (by the Neyman–Pearson lemma) is the likelihood ratio. For Poisson counting, $\mathcal{L}(\mu) = e^{-(\mu s+b)}(\mu s + b)^n/n!$, and the discovery test compares $\mu = 0$ against the best-fit $\hat{\mu}$:

$$q_0 = -2\ln\frac{\mathcal{L}(0)}{\mathcal{L}(\hat{\mu})} = 2\left[n\ln\frac{n}{b} - (n-b)\right] \quad (n > b), \tag{79}$$

where the second form follows by substituting $\hat{\mu}s = n - b$ (the maximum-likelihood estimate simply saturates the observed count). By Wilks' theorem q_0 is distributed as χ^2_1 under the background-only hypothesis, so $Z = \sqrt{q_0}$. To get the *median expected* significance one evaluates q_0 not on a random dataset but on the *Asimov dataset* — the artificial dataset in which every observable equals its expectation, $n = s + b$ (named by Cowan et al. after Asimov's story *Franchise*, in which one representative voter replaces the electorate). Substituting $n = s + b$ into (79):

$$Z_{\text{stat}} = \sqrt{2\left[(s+b)\ln\left(1+\frac{s}{b}\right) - s\right]} \xrightarrow[s\ll b]{} \frac{s}{\sqrt{b}}, \tag{80}$$

where the limit follows from expanding the logarithm to second order — and this is where the folklore formula $s/\sqrt{b}$ comes from, along with its domain of validity ($s \ll b$; at our signal-to-background ratios of order one, the full form matters and is what we use).

40.2 A background systematic as a nuisance parameter

Now suppose b itself is uncertain by $\pm\delta b$ — in our case, the ATLAS post-fit fractional uncertainties of Table 24, $\delta b = f_b\, b$. The likelihood acquires a *nuisance parameter*: the unknown true background b', constrained by a Gaussian of width δb (the auxiliary knowledge from control regions and calibrations). The test statistic *profiles* b' — re-maximizes the likelihood over it, separately in numerator and denominator — so that the fluctuation must beat not only Poisson noise but every background configuration the systematics allow. The heuristic result can be guessed immediately: the two noise sources add in quadrature,

$$Z \approx \frac{s}{\sqrt{b + (f_b\, b)^2}}, \tag{81}$$

and the exact Asimov formula, obtained by carrying out the profiling analytically (Cowan et al., 2011), is the one the simulation uses everywhere:

$$Z = \left\{ 2 \left[(s+b) \ln \frac{(s+b)(b+\delta b^2)}{b^2 + (s+b)\delta b^2} - \frac{b^2}{\delta b^2} \ln\left(1 + \frac{\delta b^2\, s}{b\,(b+\delta b^2)} \right) \right] \right\}^{1/2}. \tag{82}$$

Two limits confirm the formula against intuition: as $\delta b \to 0$ it collapses to (80), and for $\delta b \gg \sqrt{b}$ it saturates at $Z \to s/\delta b$ — the systematic wall. Independent bins combine in quadrature, $Z_{\text{tot}}^2 = \sum_i Z_i^2$, which is exact for the profile likelihood when the nuisances are bin-independent (our per-bin f_b treatment; correlations are absorbed into the calibration of §41).

40.3 The systematics wall, located precisely

Equation (81) contains the single most consequential piece of arithmetic in this Part. Both s and b grow linearly with luminosity L, so the Poisson term $\sqrt{b}$ grows as $\sqrt{L}$ while the systematic term $f_b\, b$ grows as L. They cross when $b = 1/f_b^2$, i.e. at

$$L^* = \frac{1}{f_b^2\; b_{1\text{fb}}}\,, \tag{83}$$

with $b_{1\text{fb}}$ the background per fb^{-1}. Beyond L^* the per-bin significance saturates at the luminosity-*independent* ceiling $Z_{\max} = s/(f_b\, b) = (\text{signal}-\text{to}-\text{background})/f_b$: from then on, only better background knowledge — smaller f_b — moves the answer. Evaluating (83) with the anchored backgrounds and the Run-2 post-fit f_b: $L^* = 353$, 554, 356, 485, and 1054 fb^{-1} for the 2.5–3.0, 3.0–3.5, 3.5–4.2, 4.2–4.9, and 4.9–5.6 TeV bins respectively. Every one of these is of order the Run-3 dataset. The conclusion deserves italics because nothing downstream makes sense without it: *the high-mass $t\bar{t}$ search leaves the statistics-limited regime during Run 3; the entire HL-LHC era is, for this search, a fight about systematics, not statistics.* This single fact explains the shape of every sensitivity curve below.

41 Calibration: Tying the Toy to the Real Analysis

Our likelihood is simplified: it ignores nuisance-parameter correlations, the floating $t\bar{t}$ normalization's absorption of broad shapes, signal leakage across signal regions, the dilepton channel, and the multi-SR structure. The cost of all this is measured, not guessed, by an *injection test*: we set narrow test signals (Gaussians at the detector resolution, with the published narrow-template efficiencies) at 140 fb^{-1} and compute our expected 95% CL limits with the same machinery used for all projections. We obtain 1.46 fb at 4.0 TeV and 2.09 fb at 3.5 TeV, against ATLAS's published 6.41 and 8.93 fb: degradation factors 4.40 and 4.27 — consistent at the two masses, which is what licenses promoting their mean to a single calibration constant,

$$F_{\rm cal} = 4.33, \qquad Z_{\rm true} \simeq Z_{\rm naive}/F_{\rm cal}, \tag{84}$$

applied to every significance and limit quoted from here on. This is a deliberately conservative choice: $F_{\rm cal}$ is measured in exclusion-specific machinery and applied to discovery as well. *Validation on the shape that matters*: with $F_{\rm cal}$ fixed and no further freedom, the framework predicts a 140 fb^{-1} expected limit of 8.4 fb for the $\Gamma/m = 15\%$ interference-free shape — correctly ordered between ATLAS's published narrow (6.4 fb) and broad-30% (~ 25 fb) template limits, exactly where a 15%-width shape must fall. The framework also reproduces, without adjustment, ATLAS's own statement that the octet sits at the present frontier: injecting the predicted 8.7 fb yields an expected excess just below the published exclusion.

One methodological remark for the statistically minded reader, because the temptation is real: one might think the toy should *additionally* profile an explicit background normalization and shape tilt (as our diagnostic in §42 does). It should not, for a reason worth understanding: the ATLAS f_b of Table 24 are *post-fit* uncertainties — they already encode the residual freedom of the background model *after* the real fit has floated its normalizations and profiled its shapes. Using post-fit f_b *and* an explicit free normalization would double-count the background freedom and understate the sensitivity. The post-fit-f_b-plus-calibration route is the validated one; the explicit-profile construction appears below strictly as a pedagogical instrument for locating where sensitivity is lost.

41.1 What "improved systematics" means, quantitatively

Every projection below is quoted under two scenarios that differ in exactly one input: the per-bin fractional background uncertainty. The halving is the standard HL-LHC projection convention (Cepeda et al., 2019), and it is decomposable rather than hopeful. The ATLAS tail uncertainties contain three classes: (i) template and control-region *statistical* components (simulation sample sizes, the W-charge-asymmetry and Z-reweighting control regions), which scale as $1/\sqrt{L}$ and alone shrink by $\times 2.8$ at 1.1 ab^{-1} and $\times 4.9$ at 3.3 ab^{-1}; (ii) object calibrations (jet energy scale, top- and b-tagging scale factors), data-driven and improving roughly as $1/\sqrt{L}$ toward 1–2% floors; (iii) theory-modeling terms ($t\bar{t}$ generator and shower variations, the 50% W+heavy-flavor variations), which improve only with better simulation and auxiliary measurements. Halving the total corresponds to classes (i)+(ii) improving as expected while (iii) improves modestly — the plausible mid-course. "Frozen" assumes

Table 25: What "improved systematics" means, stated as a number per bin rather than as an adjective. The frozen row is the published post-fit background uncertainty of the ATLAS Run-2 search; the improved row halves it. Every projection in this Part is quoted under both, and the two rows differ in exactly one input, so a reader who thinks the improvement optimistic can read the frozen column throughout and still have a complete forecast.

$m_{t\bar{t}}$ bin [TeV]	2.5–3.0	3.0–3.5	3.5–4.2	4.2–4.9	4.9–5.6
frozen (ATLAS post-fit, Run 2)	5.0%	6.9%	13.2%	22.6%	30%
improved ($f_b/2$)	2.5%	3.45%	6.6%	11.3%	15%

no improvement of any component over fifteen years: the pessimistic bound. Reality should track between the two curves of Figure 13, closer to "improved". Note one snapshot-specific subtlety used honestly below: at end-2026 roughly half of the "improvement" is *automatic*, because the Run-2+3 dataset itself shrinks the class-(i) components by $\sqrt{140/466} = 0.55$ before anyone improves anything.

42 Why the Eye Overestimates: The Significance Ladder

Figure 12 *looks* like an overwhelming discovery already at end-2026 — and by pure counting it is: the excess over the drawn background is a 23.6σ statistical effect over the full displayed range (18.6σ within the 3.0–4.9 TeV analysis window). The honest calibrated significance is 2.7σ. A first-year student should internalize where that factor of ~ 9 goes, because it is the difference between reading spectrum plots and understanding searches; Figure 11 decomposes it into three physically distinct steps.

Step 1 → 2 ($23.6 \to 12.8\sigma$): the gray line has error bars too. The eye compares data points to the background curve as if the curve were exact; the correct per-bin comparison is (81), and at 0.466 ab^{-1} the systematic term $f_b\, b$ has *already overtaken* $\sqrt{b}$ in every relevant bin (§40.3). Replacing $s/\sqrt{b}$ with the full (82) in each bin takes the counting ladder from 23.6σ to 12.8σ.

Step 2 → 3 ($12.8 \to 4.0\sigma$): the background shape is not God-given. The background curve of Figure 12 is drawn as if known in advance. In the real analysis it is *determined from the same data*: the SM $t\bar{t}$ normalization floats freely and the shape nuisances are profiled. A $\Gamma/m = 15\%$ resonance smeared by 10% resolution is a smooth, single-signed, two-bin-wide enhancement — dangerously close to the null space of that fit. The mechanism is worth seeing analytically once. Let the fit have one free parameter, an overall normalization A, determined by weighted least squares against data $d_i = b_i + s_i$ with weights $w_i = 1/(b_i + (f_i b_i)^2)$; minimizing $\sum_i w_i (d_i - A b_i)^2$ gives

$$\hat{A} = 1 + \frac{\sum_i w_i\, b_i\, s_i}{\sum_i w_i\, b_i^2}, \tag{85}$$

i.e. the background estimate *rises to meet the signal*, by an amount governed by how much the signal shape resembles the background shape in the metric w. A narrow spike (one bin) barely moves $\hat{A}$; a

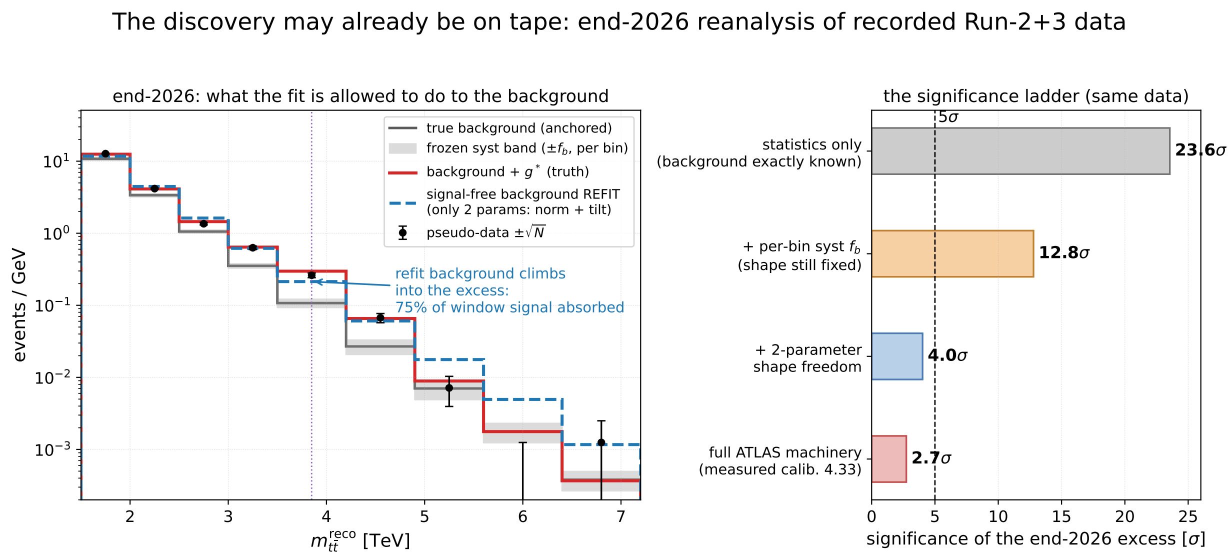

Figure 11: Why the "clean" excess of Figure 12 is worth 2.7σ, not 20σ, at end-2026. **Left:** the end-2026 spectrum with the frozen per-bin systematic band (gray) and the signal-free background refit with only two free parameters, normalization and tilt (blue dashed): the refit climbs into the excess and absorbs 75% of the window signal, at negligible cost at low mass where the band is tight. The full ATLAS fit has $\mathcal{O}(100)$ such degrees of freedom. **Right:** the significance ladder of the identical dataset: perfect-background counting $23.6\sigma \to$ per-bin systematics $12.8\sigma \to$ two-parameter shape freedom $4.0\sigma \to$ measured full-analysis calibration 2.7σ.

smooth broad excess moves it a lot. Our diagnostic fit with just *two* parameters (normalization and a power-law tilt), profiled against the frozen uncertainties, absorbs 75% of the window signal — the refit background climbs into the excess at negligible cost at low mass, where $f_b = 5\%$, because the tail bins where the signal lives allow 13–30% — and takes 12.8σ to 4.0σ (Figure 11, left).

Step 3 → 4 ($4.0 \to 2.7\sigma$): the full machinery. The real fit has not two but $\mathcal{O}(100)$ nuisance parameters — per-component normalizations, jet energy scale eigenvectors, tagging scale factors, generator shape variations — each with more absorbing power than a global tilt. We do not model them individually; their aggregate cost is what the measured calibration $F_{\rm cal} = 4.33$ carries (§41).

Two corollaries. (i) *The discovery variable is background knowledge, not luminosity*: this is the entire content of the frozen curve's plateau in Figure 13, and it is why the improved-systematics scenario, not more data, is what crosses 5σ. (ii) *Absorption is blind to correlated, sign-changing structure*: a normalization-and-tilt fit can eat a single-signed bump, but it cannot generate the interference pattern — excess below the pole, deficit above it, the summed spectrum dipping *below* the SM beyond 4.9 TeV. A template fit carrying the interference degree of freedom therefore recovers much of what naive counting loses, which is one of the pre-registration advantages quantified in §44; the discovery numbers of this Part, which use counting only, are in that specific sense conservative.

Table 26: The four snapshots. Analysis bins 3.0–3.5 / 3.5–4.2 / 4.2–4.9 TeV (the ATLAS-anchored region). S: Breit–Wigner; I: interference (signed — note the deficit in the highest bin, the shape falsifier of Part VI); B: background. Significances are calibrated ($F_{\mathrm{cal}} = 4.33$); scenarios as §41.1. Property rows use purity-corrected in-window counts.

	end-2026 0.466 ab^{-1} (13 TeV)	**end-2031** 0.56 ab^{-1}	**end-2033** 1.09 ab^{-1}	**end-2041** 3.29 ab^{-1}
S (3 bins)	81/113/32	144/211/63	279/409/121	843/1234/366
I (3 bins)	$+64/+20/-5$	$+110/+35/-11$	$+212/+67/-21$	$+641/+203/-64$
B (3 bins)	177/75/19	301/137/37	583/265/72	1758/798/216
Z_{disc} (frozen)	2.7σ	3.1σ	3.3σ	3.5σ
Z_{disc} (improved)	$\mathbf{3.7\sigma}$	$\mathbf{4.5\sigma}$	$\mathbf{5.4\sigma}$	$\mathbf{6.5\sigma}$
Z_{disc} (stat-only bound)	18.6σ	25.1σ	34.9σ	60.6σ
signal events in window (3.0–4.9 TeV)	305	552	1068	3223
window purity	0.53	0.54	0.54	0.54
spin discrimination	5.2σ	7.1σ	9.9σ	17.3σ
δP_t (lepton slope)	0.27	0.19	0.14	0.08
interference sign (coarse bins)	$\sim 0.8\sigma$	$\sim 0.9\sigma$	$\sim 1.0\sigma$	$\sim 1.0\sigma$

43 The Four Snapshots: What the Detectors Record

Table 26 gives the analysis-bin expectations and calibrated significances; Figure 12 draws the full simulated spectra over 1.5–7 TeV with statistical error bars (one Poisson pseudo-experiment per snapshot, fixed seed). All yields are merged-SR, one-lepton, per experiment; end-2026 is 13 TeV, the rest 14 TeV.

End-2026: the reanalysis case. This snapshot needs no new collisions: it is the full Run-2+3 dataset, on tape since June 2026. At the derived rate the sample contains ~ 300 octet events (Breit–Wigner plus net interference) in the 3.0–4.9 TeV window against $\sim$270 background — a signal-to-background near *one*, and statistically a 19σ effect — worth 2.7σ (frozen) to 3.7σ (improved) per experiment once background systematics are profiled. In plain terms: *evidence-level sensitivity to the GG-6 octet may already exist in recorded data*, accessible through exactly the dedicated $\Gamma/m \simeq 15\%$ interference-aware reinterpretation that this paper requests (§31). The same reanalysis, if null at the predicted rate, would put the composite scenario under severe pressure — the exclusion mode of §31.5.

End-2031 ($\sim$1.5 years of HL-LHC). The 14-TeV data roughly doubles the effective sample. Improved-scenario significance reaches 4.5σ per experiment; spin discrimination reaches 7σ; $\delta P_t = 0.19$ separates the predicted $P_t = +0.91$ from zero at $\sim 5\sigma$. The earliest defensible single-experiment 5σ (improved scenario) arrives at 0.80 ab^{-1}: data on tape during **2032**.

End-2033 (two years of HL-LHC). 5.4σ per experiment (improved); the frozen scenario is already deep into saturation (3.3σ), demonstrating that everything now hinges on background-model improvements, not luminosity. Properties: spin at 9.9σ; $\delta P_t = 0.14$, establishing the right-handedness of the top coupling at $\sim 6\sigma$ if the state is there.

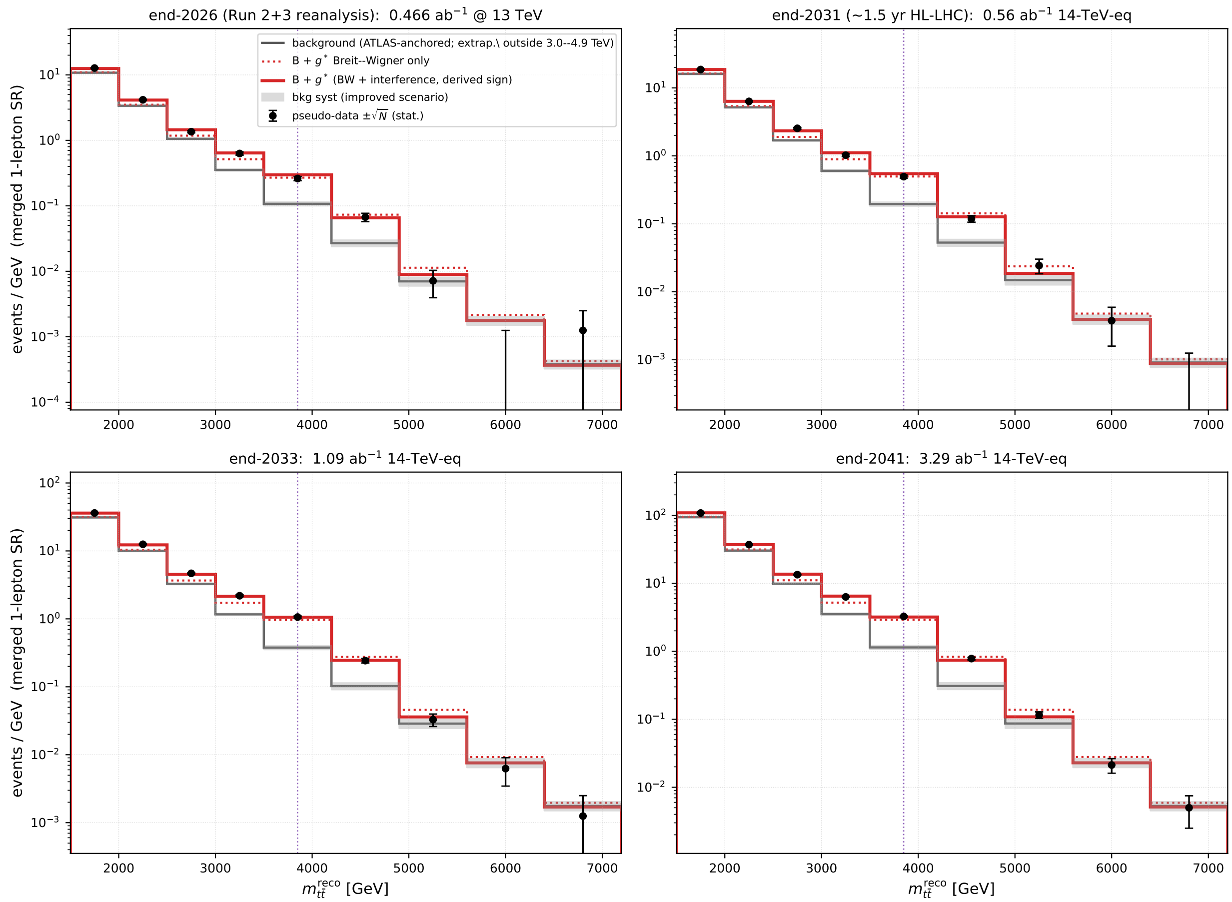

Figure 12: Simulated merged-SR $m_{t\bar{t}}$ spectra, 1.5–7 TeV, at the four snapshots (one experiment; end-2026 at 13 TeV, others at 14 TeV). Points: Poisson pseudo-data with $\pm\sqrt{N}$ statistical error bars. Gray: ATLAS-anchored background (anchored in 3.0–4.9 TeV, extrapolated outside) with the improved-scenario systematic band. Red solid: background + octet with the derived interference; red dotted: Breit–Wigner only. Three coherent-structure signatures are visible at this mass range: the 10–15% interference enhancement of the continuum extending far below the pole (solid above dotted below $\sim$3.8 TeV), the constructive shoulder at 3.0–3.5 TeV, and the destructive high side, with the solid curve dipping below the Breit–Wigner-only expectation above the pole and below the SM itself beyond $\simeq$4.9 TeV. The visual size of the excess greatly overstates its significance — see §42 and Figure 11 for where the sensitivity actually goes.

End-2041 (full program). All discrete properties saturate: spin at 17σ; $\delta P_t = 0.08$, resolving the interior of the cascade-depth band $P_t \in [+0.88, +0.92]$. Frozen-scenario ceiling: 3.5σ per experiment; improved: 6.5σ.

The interference sign in coarse bins. At every snapshot the coarse-binned sign visibility stalls near 1σ: the constructive low-side excess is smooth and single-signed — precisely the shape a floating $t\bar{t}$ normalization absorbs (Eq. 85). The sign lives in the correlated low-surplus/high-deficit pattern, including the spectrum falling *below* the SM beyond 4.9 TeV, and a template fit with the interference degree of freedom converts the 600–3000-event lever arm into a decisive measurement. That is the experimental case for the closing request of Part VI, now with numbers attached.

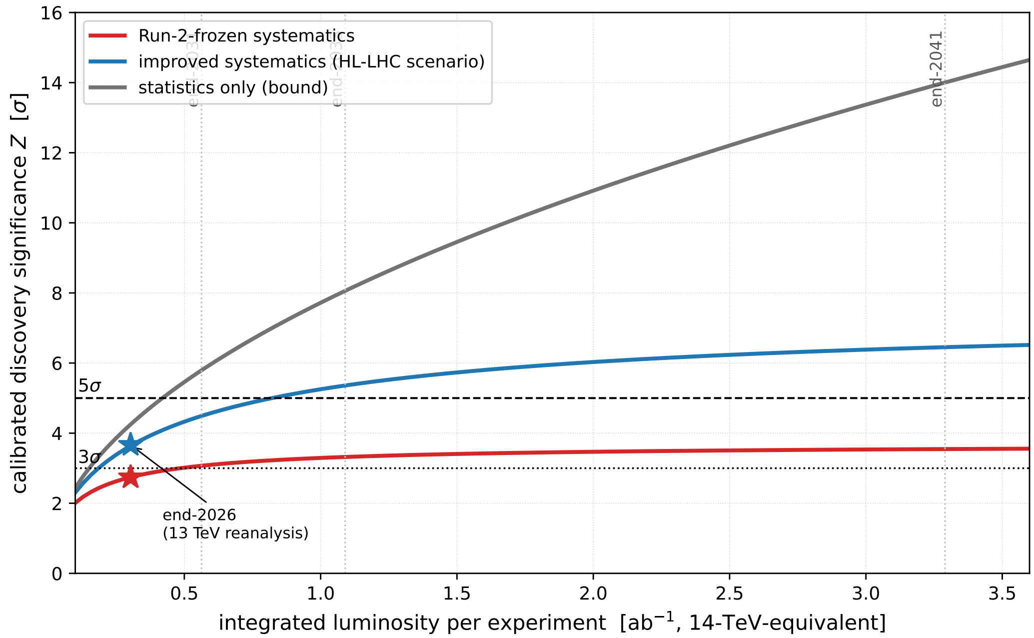

Figure 13: Calibrated discovery significance versus integrated luminosity (14-TeV-equivalent, per experiment) for the three systematics scenarios; stars mark the end-2026 13-TeV reanalysis case. The saturation of the frozen curve is the central lesson of §40.3: beyond $\sim$0.5 ab^{-1} the search is entirely systematics-limited. In the improved scenario 5σ arrives at 0.80 ab^{-1} — data on tape during 2032; frozen Run-2 systematics cap a single experiment at 3.6σ forever.

44 The Pre-Registered Search: What the Exact Prediction Buys

Everything above used the projection language of a generic search. But the GG-6 octet is not a generic hypothesis: the mass (3.849 TeV, from $\frac{1}{2}M_{\rm Pl}e^{-35}$, zero free parameters), the width ($\Gamma/m = 14.6\%$, T-KK-2), the rate (8.7 fb, Corollary T-KK-2c), the interference sign (a theorem), and the angular distribution ($\frac{3}{8}(1+\cos^2\theta^*)$, Part VI) are all specified *before looking*. Registered where? In print: the public posting of this paper fixes the mass, the width, the rate, the sign, and the angular shape before the analysis exists, and any later adjustment would be visible to everyone as a revision. That is the entire mechanism, and it is sufficient. A search for *this* hypothesis differs statistically from the published bump hunt in three separable ways. We quantify each — and, equally important, state which are already banked in Table 26 and which are genuinely additional.

44.1 Advantage 1: no look-elsewhere penalty (banked)

When a search scans many mass hypotheses, a fluctuation anywhere counts, so the probability of a fake excess is inflated by the number of independent places one looked — the *trials factor*, or look-elsewhere effect. The estimate is elementary and worth doing in one's head: two mass hypotheses are statistically independent when separated by roughly the resolution, so the number of independent cells in a scan from $m_{\rm lo}$ to $m_{\rm hi}$ with resolution $\sigma(m) = 0.10\,m$ is

$$N_{\rm cells} \simeq \int_{m_{\rm lo}}^{m_{\rm hi}} \frac{dm}{0.10\,m} = 10\,\ln\frac{m_{\rm hi}}{m_{\rm lo}} = 10\,\ln\frac{5000}{400} \approx 25, \tag{86}$$

and with $\sim$3 width hypotheses the published scan carries $N_{\rm trials} \approx 75$ (the rigorous up-crossings treatment of Gross & Vitells, 2010 gives the same order). For small p, the global probability is $p_{\rm glob} \simeq N_{\rm trials}\,p_{\rm loc}$, and the discount is steepest exactly where it hurts: a local 5σ ($p = 2.9 \times 10^{-7}$) deflates to a global 4.1σ; a global 5σ claim requires local 5.8σ. A pre-registered single point pays nothing: local *is* global. And near the systematics plateau, where significance grows by only $\sim 1\sigma$ per *tripling* of luminosity (Figure 13), the 0.8σ that pre-registration saves is worth roughly a *doubling of the HL-LHC dataset* — of order five years of running. Two honesty notes. (i) The significances quoted throughout this Part are already local ones (the calibration is to per-mass-point expected limits), so this advantage is *already banked* in Table 26: it is a gain relative to the published pipeline's discovery quote, not an addition to our numbers. (ii) Look-elsewhere never applied to *exclusion*, so the falsification timeline of §31.5 is untouched.

44.2 Advantage 2: the matched template (banked); what the published pipeline actually tests (not this)

The published analysis tests generic narrow-Z' and $\Gamma/m \simeq 30\%$ KK-gluon templates. *No published template has the* GG-6 *shape*: $\Gamma/m = 15\%$ with coherent QCD interference at the derived sign, the low-side enhancement, and the high-side deficit. The cost of the mismatch is measured inside our own validated framework: the expected limit on the matched 15%-width coherent shape is 8.4 fb at 140 fb^{-1}, while the closest published broad template delivers $\sim$25 fb — *a factor* ~ 3 *dilution in cross-section sensitivity* for the mismatched pipeline. Our simulation always injected and counted the exact $S + I$ template, so this advantage too is already banked in Table 26. The operational content is the converse: *the published pipeline, run as-is on Run-3 data, would under-report this signal by a factor* ~ 3 *in rate and discount whatever excess it sees by the trials factor* — both removed, at no experimental cost, by the dedicated analysis this paper requests.

44.3 Advantage 3: the angular dimension and the exact rate (not banked)

Two genuinely additional handles remain outside Table 26. (i) *Angular information*: the discovery numbers bin only in $m_{t\bar{t}}$. The signal is central — integrating the spin-1 form of Part VI, $\int_{-1/2}^{1/2}(1+c^2)\,dc / \int_{-1}^{1}(1+c^2)\,dc = \frac{13/12}{8/3} = 0.41$, so 41% of the signal sits at $|\cos\theta^*| < 0.5$ — while the high-mass background cocktail (40% V+jets, with forward-peaked t-channel topologies) is not. In the systematics-limited regime the per-bin ceiling is $s/(f_b\,b)$, so any selection or two-dimensional template that improves s/b converts *directly* into significance, with no luminosity price. Lacking anchored angular shapes for the full background cocktail, we bound rather than bank: a further 10–30% in Z is plausible, and a detector-level two-dimensional study should replace this bound. (ii) *The rate is predicted*: the hypothesis test can be sharpened from "some $\mu > 0$" to the simple-versus-simple test of $\mu = 1$, and an unbinned matched filter can replace the coarse ATLAS binning;

Table 27: Calibrated discovery significance of the dedicated, pre-registered, matched-template search (no angular information banked). "corr"/"uncorr" = ATLAS+CMS combination with fully correlated / fully uncorrelated background systematics. Improved scenario as §41.1; at end-2026 roughly half of the improvement is automatic (control-region statistics shrink with the Run-3 dataset itself).

	frozen syst			**improved syst**		
	1 expt	comb. corr	comb. uncorr	1 expt	comb. corr	comb. uncorr
end-2026	2.7σ	3.1σ	3.9σ	3.7σ	4.6σ	5.2σ
end-2031	3.1σ	3.3σ	4.3σ	4.5σ	5.4σ	6.3σ
end-2033	3.3σ	3.5σ	4.7σ	5.4σ	6.1σ	7.6σ
end-2041	3.5σ	3.6σ	5.0σ	6.5σ	6.9σ	9.1σ

each is worth $\mathcal{O}(10\%)$. We bank neither.

44.4 The two-experiment combination, and a calibrated verdict

The remaining multiplier is decisive: *two* experiments. In the statistics-limited regime an ATLAS+CMS combination gains exactly $\sqrt{2}$; in the systematics-limited regime the gain depends on whether the dominant uncertainties are correlated between the experiments (theory modeling: yes; jet calibrations, tagging, control-region statistics: no), so Table 27 quotes both extremes, between which reality lies.

The verdict on the natural question — *could Run-3 data alone already deliver* 5σ*?* — stated with the calibration behind it: per experiment, no (2.7–3.7σ: evidence, not discovery). For a combined ATLAS+CMS dedicated analysis with improved systematics: 4.6–5.2σ before any angular information — 5σ from data already on tape is a *live possibility*, not a guarantee; the bounded angular gain would move the bracket to 5.1–6.7σ, but that piece awaits a real detector-level study. The central expectation is therefore: a combined dedicated reanalysis of existing data yields strong evidence ($\gtrsim 4\sigma$), with discovery either from that same data if the systematics and angular handles break favorably, or with the first $\sim 1\ \text{ab}^{-1}$ of HL-LHC otherwise.

44.5 A warning the dedicated fit must respect: the sideband is contaminated

One subtlety cuts *against* naive optimism, and the dedicated analysis must be designed around it. Take the ratio of the interference term to the QCD continuum in Eq. (65): both share the same parton luminosities and color factors, so far from the pole the ratio reduces to the propagator structure,

$$\frac{d\sigma_{\text{int}}}{d\sigma_{\text{SM}}} \simeq 2\,\xi_{\text{UV}}\,\frac{g_V^t}{g_s}\,\frac{\hat{s}}{\hat{s}-M_{\text{lock}}^2} \qquad (|\hat{s}-M_{\text{lock}}^2| \gg M_{\text{lock}}\Gamma), \tag{87}$$

which is *positive* below the pole (both factors negative) and falls off only as $\hat{s}/(\hat{s}-M_{\text{lock}}^2)$ — slowly. Numerically the coherent enhancement of the total background is $\sim 16\%$ at 1.75 TeV and $\sim 10\%$ at 2.25 TeV: precisely the region a fitter would instinctively use to pin the floating $t\bar{t}$ normalization. A smooth single-signed lift is partially degenerate with a normalization shift (Eq. 85), and the degeneracy actively hurts: in the explicit two-parameter profile of §42, *removing* the low-side interference from the injected signal *raises* the profiled significance (from 4.1 to 4.7, frozen, uncalibrated), because an uncontaminated sideband pins the background better than a contaminated one. The prescription follows from (87): the contamination falls to $\sim$4% at 1 TeV and $\sim$1% below 700 GeV, so *the $t\bar{t}$ normalization must be anchored at $m_{t\bar{t}} \lesssim 1$ TeV (or in the dilepton channel), and the fit must carry the interference degree of freedom explicitly, sign free.* The low-side enhancement is also the template component most sensitive to our leading-order treatment: an NLO computation of the interference K-factor is the single theory input the dedicated analysis should commission first.

45 What This Means for F-KK-2B, and the Caveat Ledger

The snapshot sequence rewrites the operational content of F-KK-2B (§31.5) as a calendar. *By end-2026*: a dedicated reanalysis of data already on tape carries evidence-level sensitivity per experiment and possibly discovery-level sensitivity in combination — in either direction: at the predicted rate it shows a strong excess, and a null at that rate puts the composite scenario under decisive pressure. *By 2032*: with improved systematics, single-experiment 5σ luminosity is on tape. *By end-2033*: discovery or exclusion is decided in practice for the improved scenario, and the property program (spin at 9.9σ, chirality at $\sim 6\sigma$) is already underway. *By 2041*: every discrete property of the octet hypothesis is measured to saturation, including the interior of the cascade-depth band via $\delta P_t = 0.08$; and if nothing is there, the frozen-scenario combined exclusion has long since closed the window. In both directions the message to the experiments is the same one this paper has carried from the start, now with the full simulation behind it: *the dedicated, interference-aware, pre-registered $t\bar{t}$ analysis at* 3.85 *TeV is disproportionately valuable relative to its cost, and the data for its first decisive pass already exist.*

The caveat ledger. Every simplification with sign and rough size, for the reader who will improve on this Part. (1) LO×K rates with $\sqrt{K}$ on the interference; an NLO interference calculation could move the low-side template by tens of percent (direction unknown; §44.5). (2) Run-2 detector performance frozen through 2041; the Phase-2 upgrades target better boosted-top performance under pile-up, so this is conservative. (In particular, the boosted-top tagging efficiency above $\sim$3.5 TeV is an extrapolation of the published working points; the Phase-2 upgrades target exactly this regime.) (3) The background model is a two-parameter anchored extrapolation; below 2.5 TeV and above 4.9 TeV it is an extrapolation, clearly flagged in Figure 12. (4) The calibration factor $F_{\rm cal}$ is measured on exclusion machinery at two masses and applied to discovery; its consistency (4.40 vs 4.27) and the broad-shape validation are the evidence for its transferability. (5) Per-bin systematics are treated as uncorrelated; correlations are absorbed into $F_{\rm cal}$ rather than modeled. (6) The property projections (spin, P_t) use purity-corrected event counts with the acceptance-degraded separations of Part VI, not a detector-level angular analysis. (7) The ATLAS+CMS combination brackets correlation extremes rather than modeling the shared nuisances. (8) All pseudo-data use a single fixed random seed; the Asimov significances are seed-independent. (9) Electroweak corrections are not applied: at $m_{t\bar{t}} \simeq 4$ TeV the electroweak Sudakov logarithms are a negative $\mathcal{O}(10\text{–}20\%)$ effect on continuum and signal alike. Because the background normalization is anchored to data in situ, the bulk of this common suppression cancels in the excess-over-anchor observable; the residual shape distortion belongs, with the NLO interference calculation, to the theory tasks of Appendix O. The full numerical workbook behind this Part (parton-luminosity grids, the anchored background fit, the calibration injections, and every table entry) is archived with the cohort's build materials at https://arisaka-gg.org/ and is reproducible from the quoted inputs.

Part VIII turned the prediction into dates, and its plain-English summary is startling enough to bear repeating: the statistical power to see this particle may already exist, unread, in the data recorded since 2015 — evidence-level for either experiment alone, possibly discovery-level for the two combined — and the obstacle is not luminosity but background knowledge plus a template no one has yet fitted. The systematics wall, the calibrated ladder from naive counting to honest significance, and the pre-registration dividends were each located and priced. The story to 2041 is thereby fixed. What of the decades after? The next Part maps the machines beyond the HL-LHC — what each can measure, what each would decide, and which questions must wait for which ring.

46 Part VIII in plain English: what the detectors would actually record, worked out at analysis fidelity

The claim that gives this Part its point is that the collisions which decide the question may already be on tape.

To make that credible the Part does not estimate; it simulates, anchored to a published ATLAS top-pair search so that every step can be checked against something real. Backgrounds, selection, binning, and the per-bin background uncertainty all come from the anchor. The arithmetic is then shown bin by bin for one dataset so a reader can extend the same table to any other.

Two scenarios run throughout, differing in exactly one input: whether the background systematic stays at its published value or is halved. Both are carried to the end. A reader who thinks the improvement optimistic can read the frozen column and still have a complete forecast, which is the reason for quoting two rather than one.

The Part also spends a section on why the eye overestimates a bump, and that section is there for the paper's own protection. A spectrum that looks convincing at a glance may carry a modest significance once the systematic is folded in, and a paper forecasting a discovery has an obligation to say so before its own plots are read.

The output is a set of dates with numbers attached, pre-registered: what the existing dataset can reach, and what the high-luminosity run adds. If the cluster is there, it is findable in this decade. If it is not, the same schedule says when that becomes certain.

Part IX
Beyond the HL-LHC

The 2035–2070 program: facility statuses, the post-discovery checklist, and machine-by-machine falsification scenarios

47 The Question This Part Answers

Part VIII ends with the HL-LHC deciding F-KK-2B: on the simulated calendar, discovery-grade sensitivity is reached between roughly 2032 and 2035 (combined, improved systematics), and the frozen-systematics exclusion closes the window not much later if nothing is there. Either way, by $\sim$2035 the yes/no question of this paper is answered — and an entirely different program begins. If the cluster is there, the remaining task is *spectroscopy*: measuring everything the discovery machine cannot, from the direct spin-2 anchor to the second radial family to the exact zero of the leptophobic branching ratio. If it is not, the same future machines become falsification engines of increasing severity for whatever residue of the framework survives. This Part maps that post-2035 program facility by facility and decade by decade, on the *actual* status of the world's collider projects as of July 2026 — not the aspirational landscape of a decade ago — and it works out, with the same machinery used throughout this paper, what each machine can measure, including the subdominant channels that only become relevant once the discovery is in hand.

Two framing rules. (i) Facility statuses and dates below are taken from the primary strategy documents of 2025–2026 — the FCC Feasibility Study Report (FCC Collaboration, 2025), the 2026 Update of the European Strategy for Particle Physics (European Strategy Group, 2026), the CEPC Technical Design Reports (CEPC Study Group, 2025), and the International Muon Collider Collaboration's design reports (International Muon Collider Collaboration, 2024) — and every schedule is quoted as those documents state it, not as optimism would prefer it. (ii) Rates at future machines are computed with this paper's own parton-luminosity machinery where the couplings are derived (the octet and the spin-2 anchor), and clearly labeled as estimates where they are not (the $n = 2$ radial family, whose overlap integrals are not derived here).

48 The Facility Landscape as It Actually Stands

The single most important fact about the post-LHC landscape is easily stated: *as of July 2026, exactly one collider beyond LHC Run 3 is approved, funded, and under construction — the HL-LHC.* Everything else is a recommendation, a design report, or an R&D program. This is why Parts VII–VIII of this paper center the HL-LHC so insistently: it is not merely the sharpest test available, it is the only one guaranteed to happen. Table 28 summarizes the landscape; the prose below records the decisive facts per facility.

Five facts from this table shape everything below. *First*, Europe has chosen its path: the ESPP 2026 Update, adopted unanimously in May 2026, names FCC-ee the preferred flagship, with a Council decision by 2028 (European Strategy Group, 2026). The realistic default future is therefore FCC-ee physics in the mid-2040s followed by FCC-hh in the 2070s. *Second*, the FCC-hh baseline energy in the Feasibility Study Report is 84–85 TeV, not the 100 TeV of the older literature (FCC Collaboration, 2025); every 100-TeV number in the field — including the ones quoted earlier in this paper — must be rescaled by roughly the parton-luminosity factor $\sim$0.7 at our masses, which changes no conclusion but should be stated plainly. *Third*, the Chinese alternative has receded: the CEPC's technical design is complete and excellent, but the project was not selected for the 15th Five-Year Plan and will not be reconsidered before $\sim$2030 (CEPC Study Group, 2025). *Fourth*, the muon collider is real R&D with an honest schedule: a 10-TeV machine is a $\sim$2070 prospect, with a possible $\sim$3-TeV first stage around 2050 (International Muon Collider Collaboration, 2024). *Fifth*, the linear-collider options (ILC, CLIC, C^3, wakefield hybrids) survive as designs and R&D but none currently has a committed host or funding line. The planning horizon of this Part is therefore built on three anchors: HL-LHC (certain, 2030–2042), FCC-ee (probable, mid-2040s onward), FCC-hh (the reference energy-frontier machine, 2070s) — with the muon collider as the wild card that could compress the third anchor's timeline.

49 The Post-Discovery Checklist: What Is Left After 2035

Assume now the discovery future: the octet is found at 3.85 TeV, per Part VIII, between 2032 and 2035. By end-2041 the HL-LHC will have delivered the full property program of Table 26: the mass to better than a percent, the width at the 10%-resolution ceiling, spin-1 versus spin-2 discrimination at 17σ, the top-coupling chirality P_t to ±0.08 (resolving the interior of the derived band $[+0.88, +0.92]$), and the interference sign through the dedicated template fit. That is a remarkable harvest — and it leaves exactly six items open, each of which is precisely one future machine's job.

1. **The direct spin-2 anchor.** $\sigma \times \mathrm{BR}(G_1 \to t\bar{t}) = 0.045$ fb at 13 TeV at the derived κ_{gg} (Corollary T-KK-2c): even 3 ab^{-1} of HL-LHC yields a handful of reconstructed events under a 10^3-fold background. The state that *names the geometry* is invisible at the machine that finds the cluster.

2. **Resolving the degenerate spin-1 pair.** g^* and Z'_B sit at the same mass by theorem (same boundary-value problem); they are separated not by peak position but by color and channel decomposition — the octet-to-singlet ratio in dijets versus $t\bar{t}$, color-flow observables, and the small Z'_B fraction (0.2–6.8% of the cluster $t\bar{t}$ sum across the g_B menu). HL-LHC statistics and jet systematics will not cleanly unfold this.

3. **The exact zero of leptophobia.** $\mathrm{BR}(Z'_B \to \ell^+\ell^-) = 0$ *exactly* is a theorem-grade prediction; the HL-LHC can bound the cluster's dilepton fraction only at the 10^{-2}–10^{-3} level relative to $t\bar{t}$. An exact zero deserves a much harder test.

4. **The second radial family at 7.05 TeV.** The tower ratio $x_2/x_1 = 1.831$ puts the next spin-2

mode at 7.05 TeV; at 14 TeV the gg luminosity there is a factor ~ 800 below the first mode's — $\sigma(G_2) \sim 10^{-4}$ fb: hopeless. Yet this single measurement — *is the second family exactly where the Bessel spectrum of the RG throat demands?* — is the cleanest conceivable test of the non-spatial geometry, worth more than any single property of the first cluster.

5. **Sub-resolution line-shape precision.** The interference zero-crossing at $\simeq 4.9$ TeV, the $-\mathcal{O}(100)$ GeV centroid pull, and the width itself are all resolution-limited at hadron machines with 10% mass resolution; percent-level line-shape work needs either enormous statistics (FCC-hh) or a lepton-machine scan.

6. **The scalar-sector null.** GG-6 predicts *no radion* — the radial length is the fixed integer K, not a modulus (Arisaka, 2026, A4-Euclid) — so the Higgs sector must be exactly Standard-Model-like in its couplings apart from the derived $G_1 \to hh$ line. The decisive test of a null is precision, and the precision lives at Higgs factories.

50 Machine by Machine: What Each Facility Decides

50.1 FCC-hh at 84–85 TeV: the spectroscopy machine

At the FSR baseline energy the parton luminosities at our masses grow spectacularly: relative to 13 TeV, the $q\bar{q}$ luminosity at 3.85 TeV rises by a factor $\simeq 330$ and the gg luminosity by $\simeq 3.2 \times 10^3$ (computed with this paper's CTEQ6 machinery; low-x PDF uncertainties of a few tens of percent apply and are immaterial to any conclusion). The consequences, state by state:

The octet becomes a precision laboratory. $\sigma \times \mathrm{BR}(g^* \to t\bar{t}) \approx 2.8$ pb [2.3–4.0] at 85 TeV. At the FSR reference luminosity (~ 20 ab^{-1} per experiment over the program) that is $\sim 5 \times 10^7$ signal events before acceptance: the line shape, the interference pattern including its zero-crossing near 4.9 TeV, the width, and the angular distributions all become systematics- and theory-limited measurements rather than statistics-limited ones. The NLO-interference calculation requested in §44.5 graduates from desirable to mandatory.

The spin-2 anchor falls in the first weeks. Scaling the derived 13-TeV rate by the gg-luminosity ratio: $\sigma(pp \to G_1) \approx 0.13$ pb at 85 TeV (the historical 100-TeV reference value ≈ 205 fb used earlier in this paper rescales to ≈ 0.7 of itself at the FSR baseline; the early-running-discovery conclusion is untouched, and we quote both to avoid any appearance of a moving target). With $\Gamma/m = 6.3\%$ — comfortably narrower than the octet — and the derived branching pattern $t\bar{t} : W_L W_L : Z_L Z_L : hh \simeq 0.69{:}0.15{:}0.08{:}0.08$, the anchor delivers four correlated resonance lines at one mass. The longitudinal-diboson dominance among the bosonic channels ($W_L W_L$ at 15%) is itself diagnostic: it is the signature of a graviton coupling to the infrared-localized electroweak sector, and no spin-1 impostor reproduces it together with the $t\bar{t}$ angular form.

The second radial family is decided. At 85 TeV the gg luminosity at 7.05 TeV is 7.8% of its value at 3.85 TeV, so a same-coupling estimate gives $\sigma(G_2) \sim 10$ fb — of order 10^5 events: an unmistakable second line. (We label this an estimate, not a derivation: the $n = 2$ overlap integrals are not computed in this paper. The *position*, however, is a theorem-grade ratio: $x_2/x_1 = 7.0156/3.8317 = 1.831$, i.e. 7.05 TeV, testable at FCC-hh to per-mille precision.) Vector recurrences accompany it at $\mathcal{O}(10^2)$ fb in $t\bar{t}$ if the composite spectrum spaces them near the radial families. Measuring *two* (ideally three: $n = 3$ at 10.2 TeV, still within a machine whose s-channel reach extends to $\sim$40 TeV) radial masses turns the single number 3.849 TeV into a measured *spectrum* — the point at which one is no longer testing a prediction but reading the geometry.

The pair is unfolded and leptophobia is pushed to 10^{-6}. With 10^7-event statistics, the octet/singlet decomposition (item 2 of the checklist) proceeds through dijet-to-$t\bar{t}$ ratios and color-flow observables, and the dilepton null at the cluster mass is probed to branching fractions of order 10^{-6} relative to the hadronic channels — a serious test of an exact zero.

50.2 FCC-ee: the parameter-ledger machine, and the no-radion test

FCC-ee never touches the cluster directly: its top energy (365 GeV) is an order of magnitude below M_{lock}, and because the cluster is leptophobic and lepton-poor, even its indirect electroweak footprint is loop-suppressed rather than tree-level. It would be easy to conclude that the machine most likely to be built next is irrelevant to GG-6 — and wrong. FCC-ee's role in this program is different in kind: it is the precision arm of the *cohort ledger* that this paper's mass sits inside. Three examples define the class. (i) The GG-6 structural value of the strong coupling is $\alpha_s(M_Z) = 1/(\pi e) = 0.11710$ (Arisaka, 2026, A9-SM); the current world average, 0.1180 ± 0.0009 (Particle Data Group, 2024), sits $\sim 1\sigma$ above it. FCC-ee's projected determination at the few-10^{-4} level converts this from a tension-free coexistence into a decisive test — arguably the sharpest single number the machine will produce for this framework. (ii) The top-quark mass from the threshold scan ($\pm\mathcal{O}(20)$ MeV) and the Z-pole program sharpen every top-overlap quantity in this paper — the band $r_{t_R} \in [4, 5]$ that sets the octet's width, branching ratios, and polarization. (iii) The Higgs-coupling program at the per-mille-to-percent level executes the *no-radion* test (checklist item 6): GG-6 requires the observed Higgs to be unmixed with any radial scalar, because no such scalar exists; a coupling pattern showing radion-like mixing would falsify the scalar sector outright, discovery or no discovery. A machine that cannot produce a single cluster particle can therefore still falsify the framework — which is exactly the property one wants from a complementary program. (The CEPC, if revived in 2030, carries the identical program; we do not treat it separately.)

50.3 A muon collider: the leptophobia machine

The muon collider inverts the logic of every other entry: its decisive GG-6 measurement is a *null*. The cluster couples to quarks and gluons and shuns leptons — Z'_B exactly (theorem), G_1 strongly (UV-localized lepton overlaps). A generic 3.85-TeV Z' or KK gluon with electroweak-strength lepton couplings would announce itself in $\mu^+\mu^- \to f\bar{f}$ contact-interaction tails even at a 3-TeV stage-1 machine, and would dominate an s-channel scan at a 10-TeV machine able to sit on $\sqrt{s} = 3.849$ TeV. GG-6 predicts the opposite, and sharply: *the cluster that roars in $t\bar{t}$ at hadron colliders is nearly silent in the s-channel of a lepton collider* — no Z'_B line at all, a G_1 line suppressed by two powers of the UV lepton overlap, and Standard-Model-like contact-interaction observables at stage-1. If the HL-LHC has found the cluster, this null is as decisive a fingerprint as any positive signal: observing a leptonic s-channel line at 3.85 TeV would falsify the $\mathrm{U}(1)_B$ assignment and the localization geometry in one measurement. What the machine does measure directly: G_1 production through W-fusion and $\gamma\gamma$ fusion (fed by the $15\% + 8\%$ longitudinal-diboson couplings), at femtobarn-class rates at 10 TeV — a clean, low-background environment for the bosonic side of the branching pattern, complementary to FCC-hh's hadronic side. On the IMCC schedule these tests sit at $\sim$2050 (stage-1 nulls) and $\sim$2070 (10-TeV program) (International Muon Collider Collaboration, 2024).

50.4 Linear colliders, the LHeC, and the supporting cast

ILC- or CLIC-class machines carry the same ledger-and-null program as FCC-ee (Higgs precision, top precision, contact-interaction nulls), with CLIC's 3-TeV stage adding modest direct sensitivity to the cluster's virtual exchange in $e^+e^- \to t\bar{t}$ — again predicted *absent* at leading order by leptophobia. Their current status (Table 28) makes them contingency options rather than planning anchors. The LHeC deserves one specific remark: its deep-inelastic program would pin the high-x quark and gluon densities that dominate the uncertainty bands of every rate in this paper — the octet's [7.0–12] fb band at 13 TeV is PDF-limited — so a single ep machine would quietly convert the `GG-6` rate test from factor-level to percent-level without ever producing a cluster particle.

51 Concrete Scenarios, 2035–2070

Scenario D (discovery, the Part-VIII calendar). *2032–2035*: combined ATLAS+CMS dedicated analysis crosses 5σ; the discovery paper's mass is 3.85 ± 0.04 TeV and the width is broad; the world notices. *2035–2041*: the HL-LHC property program saturates (spin 17σ, $P_t \pm 0.08$, interference template); the theory side delivers the NLO interference calculation; the hh/diboson subdominant searches at 3.85 TeV sharpen from curiosities into targeted corroborators. *2041–2045*: LHC legacy combinations; FCC-ee construction proceeds on the post-2028 approval; the discovery reshapes the FCC-hh case — a confirmed 3.85-TeV cluster with a predicted second family at 7.05 TeV is precisely the "guaranteed-physics" argument that historically accelerates energy-frontier machines, and the FSR's earlier-LTS-magnet options give the strategy bodies a concrete lever. *Mid-2040s–2050s*: FCC-ee runs; $\alpha_s(M_Z)$ versus $1/(\pi e)$ is decided; the no-radion Higgs-precision test is executed; the muon-collider demonstrator reports and, in the favorable branch, a stage-1 machine near 2050 delivers the leptophobia null. *2070s*: FCC-hh (or a 10-TeV muon collider, whichever materializes) observes the spin-2 anchor in its first weeks, resolves the four-line branching pattern, unfolds the degenerate pair, pushes leptophobia to 10^{-6}, and — the endpoint of this entire program — measures the second and third radial families at 7.05 and 10.2 TeV, reading the Bessel spectrum of the non-spatial throat directly. One should be honest about the gap this timeline contains: on current schedules the *direct* spin-2 confirmation arrives roughly four decades after the octet discovery. That gap is not a physics necessity but a programmatic choice, and a 2030s discovery would be the strongest imaginable argument for compressing it.

Scenario X (exclusion). If instead the combined HL-LHC program excludes the octet band — decided, per Part VIII, at $\lesssim 2$ ab^{-1}, i.e. by the mid-2030s — the collider implementation of GG-6 is retired, and this paper's interest in the post-LHC machines reduces to the discipline it has kept throughout: there is no parameter to slide, so there is no zombie phase in which the same prediction reappears at 6 TeV. The framework's remaining exposure then lives entirely in the cohort's non-collider arms (the fifth-force window and the 8.1 μeV dark-matter line of the companion papers), and the future colliders inherit only the generic duty of exploration. We state this plainly because the asymmetry is the point: the machines of Table 28 are worth building in either future, but only in Scenario D do they inherit a fully specified, decade-by-decade program with derived rates waiting for them.

52 Summary Matrix

Table 29 condenses this Part. The taxonomy is worth keeping in mind: the HL-LHC is the *verdict machine* — the yes/no lives there, and nowhere else, before 2045. FCC-hh is the *spectroscopy machine*: the anchor, the pair, the radial tower. The muon collider is the *leptophobia machine*: its decisive contribution is a predicted silence. FCC-ee and the linear colliders are the *ledger machines*: they test the same integer bookkeeping from the precision side, and carry the no-radion null. And the LHeC is the *sharpening machine*: it buys precision on everyone else's predictions. The sharpest yes/no test lives at the LHC first; everything else in this Part is what victory would purchase.

Part IX mapped the far future onto the facilities that actually exist on paper, in their actual states of approval: the verdict belongs to the HL-LHC alone; the spectroscopy — the anchor, the overtones at their Bessel ratios, leptophobia at the part-per-million level — belongs to FCC-hh in the 2070s; the precision ledger belongs to FCC-ee; and the muon collider contributes the strangest measurement of all, a predicted silence. Every derivation, simulation, and schedule the paper owes is now delivered. What remains is judgment: what all of this *means*, how it compares to the roads not taken, and what would count as being wrong. The next Part steps back and argues.

53 Part IX in plain English: what happens after 2035, on either answer

This Part is written for the world in which the previous one succeeds, and it starts from an unglamorous fact: as of now exactly one collider project beyond the high-luminosity upgrade is actually approved.

That fact governs everything else. The Part surveys the real status of each proposed machine from the primary strategy documents rather than from ambition, and it is explicit about which entries are decisions already taken and which are proposals awaiting one. A roadmap built on the aspirational landscape of a decade ago would be a different and much more comfortable document.

Given the real machines, the Part works out what each would decide. A discovery in the top channel would fix a mass and little else; the follow-up questions — the spin of the anchor, the couplings of the partners, the position of the next recurrence — each need a specific capability, and the Part matches question to machine.

There is a version of this Part that would have read as advocacy for a particular facility. This one tries not to be: the checklist is written so that a reader can see which questions would go unanswered under each funding outcome, including the outcome where nothing after the high-luminosity run is built at all.

Table 28: The future-collider landscape as of July 2026, from the primary strategy documents (CEPC Study Group, 2025; European Strategy Group, 2026; FCC Collaboration, 2025; International Muon Collider Collaboration, 2024). "Earliest physics" quotes the reference schedules of those documents. Only the HL-LHC is approved and under construction.

Facility	**Type / energy**	**Status (July 2026)**	**Earliest physics**	**Next decision gate**
HL-LHC	*pp*, 14 TeV	**Approved, funded, under construction**; Long Shutdown 3 under way	∼2030	none — proceeding
FCC-ee	e^+e^-, 88–365 GeV, 91-km ring	**Recommended**: ESPP 2026 Update, adopted unanimously by the CERN Council (Budapest, 22 May 2026), names FCC-ee the preferred next flagship; ∼15.3 BCHF; four interaction points; not yet approved or funded	mid-2040s	CERN Council decision by 2028, informed by funding plan and host-state consultations
FCC-hh	*pp*, **84–85 TeV**, same tunnel	Concept within the FCC integrated program; FSR baseline 14-T Nb_3Sn dipoles; HTS options studied	early-to-mid 2070s (FSR reference); earlier LTS variants appear in FSR option studies	follows the FCC-ee decision and the high-field magnet program
CEPC	e^+e^- Higgs/Z factory, 100-km ring	Technical Design Reports complete (October 2025); **not selected** for China's 15th Five-Year Plan (March 2026); resubmission intended ∼2030, or merger of effort into FCC if approved	mid-2040s at the earliest, contingent on 2030 approval	Chinese national planning, ∼2030
Muon collider	$\mu^+\mu^-$, staged $3 \to 10$ TeV	International design study (IMCC); ESPP 2026 names muon-collider R&D a priority; ionization-cooling demonstrator planned for the 2030s	stage-1 (∼3 TeV) ∼2050; 10 TeV ∼2070	demonstrator results, mid-2030s
ILC	e^+e^-, 250 GeV linear	Design mature (TDR 2013); Japanese government non-committal (February 2026); international technology network (ITN) R&D running 2026–28	undetermined	Japanese ministerial roadmap decisions
CLIC	e^+e^-, $0.38 \to 3$ TeV linear	Mature design; **not selected** as the preferred ESPP option; maintained as an alternative	undetermined	contingent on FCC-ee outcome
C^3 / wakefield (HALHF etc.)	compact e^+e^-	R&D concepts; encouraged by the U.S. P5 report and ESPP inputs	2040s demonstrations at the earliest	technology milestones
LHeC	*ep*, ∼1.2 TeV	Concept reusing the HL-LHC beam; positioned as a bridge project	2040s	CERN program decisions

Table 29: The GG-6 program beyond the HL-LHC, assuming discovery by ~ 2035 (Scenario D). Rates with derived couplings are computed with this paper's machinery; the G_2 entry is a same-coupling estimate (position, however, is theorem-grade). The table is written under an assumption stated in its own first line, which is that the discovery happens. Its purpose is not to argue that it will but to show that the program has somewhere to go if it does, and that the next targets are already positioned rather than to be chosen later. The one same-coupling estimate is marked as such.

Facility	**Earliest GG-6 physics**	**Direct reach into GG-6 targets**	**Decisive deliverable**	**Key subdominant channels**
HL-LHC	now–2041	octet $8.7 \to 13.6$ fb ($13 \to 14$ TeV); anchor and $n{=}2$ family out of reach	the verdict (F-KK-2B); spin, chirality, interference sign	hh, diboson corroborators at 3.85 TeV; dijet angularity
FCC-ee	mid-2040s	none (365 GeV $\ll M_{\rm lock}$; leptophobic cluster)	$\alpha_s(M_Z)$ vs $1/(\pi e)$ at 4σ-class; no-radion Higgs precision	Z-pole ledger; m_t to ± 20 MeV
Muon collider	$\sim$2050 (3 TeV); $\sim$2070 (10 TeV)	s-channel scan of 3.85 TeV possible at 10 TeV	the leptophobia null: no Z'_B line, suppressed G_1, SM-like contact terms	W-fusion/$\gamma\gamma$-fusion G_1 (fb-class); bosonic BR pattern
FCC-hh (84–85 TeV)	2070s	octet ≈ 2.8 pb; $G_1 \approx 0.13$ pb; $G_2 \sim 10$ fb (est.); reach to ~ 40 TeV	spin-2 anchor + four-line BR pattern; radial tower at 7.05/10.2 TeV; pair unfolding; leptophobia to 10^{-6}	W_LW_L, Z_LZ_L, hh lines; vector recurrences $\mathcal{O}(10^2)$ fb
LHeC / linear options	2040s+ (contingent)	none direct	PDF precision (percent-grade rate tests); ledger + contact-interaction nulls	high-x g, q densities; $e^+e^- \to t\bar{t}$ tails

Part X
Discussions

Significance and uniqueness of the result; the classic collider puzzles addressed; comparison with other theories and models; the structural innovations; predictions and falsification; open questions; and the completed cohort

The derivations are finished; this Part steps back and argues. Five discussions follow — the significance of the result; its uniqueness, nobility, and innovations; the comparison with the other roads beyond the Standard Model; the predictions for falsification; and the open questions — closed by the paper's place in the completed cohort and a summary. The reference tables that condense the whole collider sector into three pages (the master ledger, the falsification hierarchy, and the discovery calendar) live with the other lookup material in Appendix D, where a reader can consult them without the argument, and the argument can proceed without the furniture.

54 Significance of the result

What this paper delivers, reduced to its irreducible content, is a single sentence: *a falsifiable, zero-parameter collider target — one mass, one width, one rate, one interference sign, one angular shape, one branching fingerprint — with a dated calendar for its decision.* Each clause of that sentence is individually rare in the beyond-Standard-Model literature; their conjunction is, to our knowledge, unprecedented. This section unpacks why, deliberately slowly, because the significance of a result of this shape is easy to misjudge in both directions — easy to inflate, if one forgets how much of it is conditional on upstream theorems, and easy to deflate, if one has been trained by fifty years of adjustable predictions to assume that every collider target is secretly a scan.

The result in the fifty-year arc. Place the claim against the decades traversed in §1. In the 1980s, the community searched for proton decay at the rates grand unification suggested; the prediction had order-of-magnitude freedom, and when Kamiokande and IMB saw nothing, the models moved. In the 1990s and 2000s, supersymmetry's "natural" window slid from LEP to the Tevatron to the LHC, a moving target by construction, because more than one hundred continuous parameters can absorb any null. In the 2010s, after the Higgs, the extra-dimensional and composite programs offered benchmarks — honest, useful benchmarks — whose masses and couplings were nevertheless inputs, so that exclusion plots grew steadily without any exclusion being fatal. Against that history, the present claim is structurally different in kind, not degree: the mass 3.849 TeV cannot move (an integer has no neighborhood); the rate 8.7 fb cannot be lowered (the coupling $1/\sqrt{K}$ is the same object that fixes the mass); the width, sign, polarization, and branching pattern are outputs of the same three theorems; and the paper has pre-committed, in writing, to the luminosity at which a null retires the construction. Whether or not nature cooperates, this is what the community spent five decades asking theory to produce.

What "zero parameters" buys, epistemically. The absence of continuous freedom is not an aesthetic garnish; it changes what an experimental outcome *means*. With a dial, a positive result confirms little (the dial was set after the fact) and a negative result kills nothing (the dial turns). Without one, both outcomes become maximally informative: a discovery at 3.849 ± 0.10 TeV with the correlated pattern would confirm a chain that reaches from collider phenomenology back through the electroweak ledger to the axioms, and a clean exclusion falsifies the same chain outright — there is no fallback version of this paper at 6 TeV. The statistics inherit the same virtue in a precisely quantifiable way (§44): no look-elsewhere discount (worth 0.8σ, equivalently a doubling of dataset near the systematics plateau) and no template-mismatch dilution (a factor ~ 3 in rate). Pre-registration is usually discussed as a virtue of experimental design; here it is a consequence of the theory's structure, and it converts directly into sensitivity.

The over-constraint, told in full. A skeptic granting everything above could still ask: why believe the chain *before* the collider verdict? The paper's strongest answer predates any collider data: the same integer that fixes M_{lock} is spent a second time on a number measured to four digits decades ago. The chain runs $K = 35 \Rightarrow M_{\text{lock}} = \frac{1}{2} M_{\text{Pl}} e^{-35} = 3.849$ TeV $\Rightarrow v = M_{\text{lock}}(2/5)^3 = 246.3$ GeV, and the second arrow is not a separate assumption but the Wilson-line readout of Arisaka (2026, A9-SM). One integer, two scales, one of them already confirmed: the construction walks into the collider test having already survived a four-digit audit. This is what distinguishes the present wager from a lottery ticket — the ticket has already matched half the draw — and it is also what raises the stakes, because a collider null does not merely eliminate a candidate resonance; it breaks a chain that presently accounts for the electroweak scale.

The calendar as part of the result. Part VIII is not an appendix to the claim; it is half of it. A falsifiable statement without a decision procedure is a promissory note, and the simulation converts the note to a schedule: expected evidence-level sensitivity (2.7–3.7σ per experiment) already resident in the recorded Run-2+3 data; combined ATLAS+CMS sensitivity of 4.6–5.2σ under realistically improved systematics — five-sigma from tape a live possibility; single-experiment five-sigma luminosity on tape during 2032; the frozen-systematics exclusion closing the window by the mid-2030s if nothing is there. The dates matter beyond logistics: they place the verdict within the professional lifetime of the paper's addressed reader (the Preface's graduate student) and within the strategic planning horizon of the laboratories now deciding the post-LHC program (Part IX). A prediction whose test is generations away is philosophy with equations; this one is a thesis proposal.

Significance for each community, separately. For the *experimental collaborations*, the significance is operational: a fully specified template, a priced reanalysis of data already on tape, and a deliverable that is valuable under every outcome — the world's strongest limit on any broad top-philic resonance if the answer is no. For *collider phenomenology*, the significance is methodological: an existence proof that a derived rate, a derived sign, and a derived angular shape can be produced from a geometric construction with no scan, and a worked demonstration (Parts VI–VIII) of what the corresponding analysis chain looks like end to end. For *the theory community*, the significance is the inversion described in §5.5: the first framework to field both of Weyl's geometrizations — scale and phase — as forced, non-spatial dimensions, with the moduli, landscape, and supersymmetry obligations dissolved as consequences rather than assumptions. And for *strategic planning*, the significance is conditional but sharp: a 2030s discovery at 3.849 TeV, with a predicted second family at 7.047 TeV, is precisely the "guaranteed-physics" input that historically accelerates energy-frontier machines, and Part IX maps the resulting program to 2070.

The two-scenario resolution as significance of a second kind. One episode inside the paper deserves separate billing, because it demonstrates the framework's epistemic behavior under fire. When the free-field placement of the spin-1 pair at 2.461 TeV was confronted with existing data, the derived rate (224 fb) sat a factor 3.6–11 *above* the published exclusion: Scenario A died (§31). The framework did not slide; it *learned* — the throat vectors are composite, not free-field — and the surviving claim became sharper, not vaguer. A theory whose exclusions teach it structural facts, rather than merely relocating its expectations, is operating in a different epistemic regime from the parameter scans of §1, and the episode is now part of the paper's evidence about itself: the falsification machinery has been observed to function, once, in the kill direction.

The discovery, read on the anatomy: one resonance, two fibers. A last register of significance has been made exact by the cohort's recent campaign (GG-W1-Symmetry; Arisaka, 2026, A2-GG6; Arisaka, 2026, P5-CPT). The two non-spatial directions are fibers of opposite topological kind — *holonomy requires closed loops, and an interval has none*: the compact circle S^1_θ carries all of the bulk's holonomy — quantized, protected, a group — while the scale interval I_r carries none — one-way transport, a semigroup. The lock cluster is these two fibers in a single measurement. Everything *quantized* about the target is the circle's signature: the tower's existence, the winding origin of the charges, the spin content, the branching fingerprint fixed by ledger integers — holonomy-class data on the fiber that remembers, which is why none of it is adjustable. And everything about the target's *location* is the interval's signature: the depth $M_{\rm lock} = \frac{1}{2} M_{\rm Pl}\, e^{-K}$ is thirty-five e-folds of one-way descent, the address the fiber that forgets assigns to the first rung. A discovery at 3.849 TeV would therefore be a simultaneous reading of both fibers — the resonance's *identity* written on the circle, its *address* written by the interval — and, in the campaign's terms, the first laboratory instrument ever to read a non-spatial fiber directly: the circle made audible, at the energy the interval sets.

55 Uniqueness, nobility, and innovation

Three words that a physics paper should use carefully, so we define them operationally before earning them. *Uniqueness*: what this paper does that no neighboring framework does, stated so that a referee can check the negative. *Nobility*: the epistemic virtues — pre-commitment, honest grading, the refusal to retreat — that make the result trustworthy independent of whether it is true. *Innovation*: the specific technical moves, registered as I-KK-1 through I-KK-10, that did not exist before this paper. The three are related but not the same, and conflating them is how strong claims come to sound like salesmanship.

55.1 The classic collider puzzles, resolved or dissolved

Each entry below states a puzzle the field has carried for decades, in the form a practitioner would recognize, and then what this paper does to it. “Resolved” means an answer is derived; “dissolved” means the question is shown to have been wrongly posed.

PS-KK-1: Why has no new particle appeared in fifty years of searching? The puzzle, stated honestly, has haunted every strategy report since LEP: if the TeV scale hides new physics, why have the cleanest channels — dileptons, diphotons, the golden signatures that found the Z and the Higgs — stayed empty through three orders of magnitude in luminosity? The conventional answers are either despairing (there is nothing) or evasive (it is heavier). This paper’s answer is neither: the derived coupling texture — UV-localized light fermions at $\xi_{\rm UV} = -0.194$, a flat photon, an IR-localized top at $r_{t_R} \approx 4.5$ — makes the cluster *loud precisely where searching is hardest* (broad, interfering, $t\bar{t}$) and *silent precisely where searching is easiest.* On this reading the fifty-year silence is not evidence of an empty TeV scale; it is a statement about which channels were listened to, and PS-KK-1 dissolves into an instruction: listen in boosted top pairs, with the template of Part VI.

PS-KK-2: The hierarchy problem, as an output. For four decades “why is $v/M_{\rm Pl} \sim 10^{-16}$?” has been the field’s engine, driving supersymmetry, technicolor, warping, and anthropics in turn — every proposed answer a new mechanism added to protect a ratio. Here the ratio is not protected; it is *computed*: e^{-35} times derived rational factors, where 35 is a count. The hierarchy is the length of the scale throat, and asking why it is large is asking why $LJ = M^2 + N^2 + L^2 = 35$ — a question with a Diophantine answer (Arisaka, 2026, A3-BI6), not a dynamical one. Resolved, conditional on the cascade; and the resolution is testable, because the same count sits at 3.849 TeV.

PS-KK-3: The radion and moduli problem. Every spatial extra dimension ever proposed drags the same embarrassment behind it: the size of the new dimension is a field, the field must be stabilized, and the stabilizer (Goldberger–Wise and its descendants; Goldberger & Wise, 1999) is an addition whose only job is to hold the stage still. String compactifications multiply the embarrassment by hundreds of moduli. `GG-6` dissolves the problem at the root: a dimension whose length is a counting number has no fluctuation to stabilize — there is no radion field at all — and the dissolution is falsifiable, because a hidden modulus always leaks into Higgs couplings (P-KK-9, the Higgs-precision null of Part IX).

PS-KK-4: The look-elsewhere pathology. Resonance hunting has a statistical tax and a sociological one. The statistical tax is the trials factor: scan enough mass points and fluctuations are guaranteed, so every local excess is discounted ($N_{\rm trials} \approx 76$ for the published $t\bar{t}$ scan, §44.1). The sociological tax is worse: because scanned theories relocate after every null, the community has learned to treat all BSM excitement as noise — the 750 GeV diphoton episode being the canonical cautionary tale. A pre-registered single point pays neither tax. The statistical discount vanishes by arithmetic; the sociological one vanishes because the prediction was in print, with its rate and sign, before the data were unblinded. Dissolved — for this search alone, which is rather the point.

PS-KK-5: The "KK graviton" normalization ambiguity. A subtle, insider's puzzle: the warped literature's first-pole mass $m_1 \simeq x_1 k e^{-k\pi r_c}$ depends on conventions (which brane, which boundary condition, which normalization of k) that a data analysis cannot check, so "the KK graviton at m_1" has never been a sharp target. This paper's statement is constructed to be convention-proof: the *centroid* of a multi-field cluster is the theorem-level object (T-KK-1), the fine structure is derived separately (T-KK-2), and the search instruction changes accordingly — hunt the cluster, not a lone pole. Resolved by reformulation.

PS-KK-6: Why should new physics sit near the weak scale at all? Naturalness said: because something must cancel the Higgs divergence, and it must be nearby. The LHC has spent fifteen years eroding that argument, and its erosion is routinely counted as evidence against *any* TeV-scale expectation. This paper's chain is immune to the erosion because it never invoked naturalness: v and $M_{\rm lock}$ are outputs of one integer chain, so their proximity is an identity, not a hope. The failure of naturalness-based expectations is evidence against naturalness — and says nothing about a prediction that never used it. Dissolved.

55.2 The structural innovations

Ten registered moves that did not exist before this paper, each with its before-and-after stated so the claim of novelty can be audited.

I-KK-1 (the zero-parameter collider mass). *Before*: every extra-dimensional benchmark mass was an input — RS's kr_c tuned to the hierarchy, ADD's M_D scanned per analysis. *After*: $M_{\rm lock} = \frac{1}{2} M_{\rm Pl} e^{-K}$ with K a Diophantine integer, doubly spent through $v = M_{\rm lock}(2/5)^3$. To our knowledge the first collider mass in the BSM era derived with no continuous freedom anywhere in its chain.

I-KK-2 (the two-scenario fine structure). *Before*: fine structure of KK spectra quoted with $\mathcal{O}(1)$ placeholders. *After*: exact free-field eigenvalues *and* an honest discrete dichotomy (free-field versus composite-hardened), both horns carrying complete mass-plus-rate predictions — one now excluded by data, the survivor sharpened by the exclusion. The innovation is the *form*: uncertainty packaged as a finite decision problem rather than a band.

I-KK-3 (the derived production coupling). *Before*: KK-graviton collider studies parametrized $k/\bar{M}_{\rm Pl}$ and scanned it. *After*: $\kappa_{gg} = 2(1 + J_2(x_1))/(K x_1^2 J_2(x_1)) = 0.47/K$, a closed Bessel form with no inputs to tune (Corollary T-KK-2c; derived in full in Appendix I) — and with it, the first parameter-free retro-diction of a published null (the diphoton limit, by a factor of several thousand).

I-KK-4 (the interference-sign theorem). *Before*: resonance–continuum interference treated as a nuisance with a sign to fit. *After*: $\xi_{\rm UV} g_V^t < 0$ forced, making the line shape — enhancement below, deficit above, the summed spectrum crossing *below* the Standard Model near 4.9 TeV — a parameter-free shape falsifier, drawn quantitatively in Figure 4.

I-KK-5 (leptophobia as an exact zero). *Before*: leptophobic Z's built by charge assignment chosen to evade dilepton limits. *After*: $\mathrm{BR}(Z'_B \to \ell\ell) = 0$ because leptons carry no baryon number — the evasion is a theorem, not a choice, and it is testable to 10^{-6} (§50.1).

I-KK-6 (the multi-field cluster signature). *Before*: searches target single poles; even multi-state models scan the splittings. *After*: a forced near-degeneracy of three species at one derived mass — spin-2 anchor plus two vectors — with the search instruction changed to match. The cluster, not any member, is the 6D fingerprint.

I-KK-7 (pre-registered-search statistics). *Before*: look-elsewhere and template dilution treated as unavoidable costs of BSM searching. *After*: both quantified and *eliminated* by derivation — no trials factor, no mismatch dilution — converting existing data into 4.6–5.2σ of combined dedicated sensitivity (§44). The innovation is the demonstration that theoretical rigidity is worth luminosity.

I-KK-8 (the anchored-simulation method). *Before*: BSM sensitivity projections routinely extrapolate with $\sqrt{L}$ and idealized detectors. *After*: every experimental input anchored to a published analysis, the statistical machinery calibrated to its published limits ($F_{\rm cal} = 4.33$, validated on an independent shape), and the systematics wall located analytically ($L^* = 1/f_b^2 b$) — a projection whose error budget is itself audited (Part VIII).

I-KK-9 (the falsification calendar). *Before*: falsifiability claimed in principle, rarely dated. *After*: the kill criterion carries a luminosity ($\lesssim 2\ \mathrm{ab}^{-1}$) and therefore a decade (the 2030s), with the property-measurement ladder scheduled behind it snapshot by snapshot.

I-KK-10 (the facility-role taxonomy). *Before*: future-collider physics cases argued machine by machine, each claiming everything. *After*: a division of labor derived from the couplings themselves — verdict machine (HL-LHC), spectroscopy machine (FCC-hh), leptophobia machine (the muon collider, whose decisive GG-6 measurement is a predicted *silence*), ledger machines (FCC-ee and the linear options) — each with named, dated deliverables (Part IX).

55.3 The nobility clause

"Nobility" is not a word for the theory; it is a word for the *conduct* the paper binds itself to, and it can be stated as four commitments, each checkable. *Pre-commitment*: every falsifier is stated with its decision rule and its price before the data arrive — F-KK-2B's kill window was in print before the reanalysis it invites. *Honest grading*: every quantitative claim carries [T]/[D*]/[M]/[E], and the least flattering grades are displayed most prominently (the g_B menu is on the first summary table, not in a footnote). *No retreat*: the response to Scenario A's exclusion was retirement, not relocation — and the paper's structure makes future relocation impossible in public, because the integers are load-bearing across twenty-two documents. *Symmetric stakes*: the paper commits to what failure means (the electroweak derivation falls with the collider claim) as explicitly as to what success means. These commitments cost rhetorical safety and buy something the fifty-year history shows to be scarcer than cleverness: a claim whose author has arranged, in advance, to be unable to escape.

56 Comparison with the other roads beyond the Standard Model

Table 2 (§5.5) carried the structural comparison with KK and RS; this section widens the lens to every major road the field has walked since the 1970s and sharpens the question to the operational one: *if that framework, rather than* GG-6*, owned the TeV scale, what would the four snapshots of Part VIII record — and which of this paper's registered predictions would it fail?* Table 30 collects the verdicts across fifteen observables; the prose takes the frameworks in turn, each with the respect its history has earned and the precision its comparison requires.

Versus 5D Randall–Sundrum and the warped benchmarks. RS is this paper's closest relative and its most instructive foil — the mathematics of Appendix J is shared, which is exactly why the differences are clean. A canonical RS graviton at 3.85 TeV with $k/\bar{M}_{\mathrm{Pl}} \sim 0.1$–$1$ would have announced itself *first* in dileptons and diphotons, the channels where GG-6 is structurally silent; its single narrow pole carries no mass-degenerate vector pair, no exact leptophobia, and a rate dialed by $k/\bar{M}_{\mathrm{Pl}}$; its stabilized radion mixes into the Higgs sector at levels a Higgs factory sees; and its second KK level sits at x_2/x_1 *only if* the 5D boundary conditions are pure — a convention, not a theorem (PS-KK-5). At the snapshots: RS would show clean-channel lines by 2026 or nothing at all, since its couplings large enough for $t\bar{t}$ visibility are excluded in $\ell\ell$. Fails P-KK-6 (leptophobia), P-KK-9 (no radion), and the cluster degeneracy; shares (in weakened form) the overtone ratio.

Versus composite Higgs and KK-gluon models. These frameworks share this paper's top-philic octet — indeed the boundary-composite reading of Scenario B *is* a compositeness statement, and the overlap is acknowledged debt, not coincidence. The differences are the dials: composite octet masses are scanned (the exclusion plots run from 1 to 6 TeV precisely because nothing fixes the mass), widths are adjustable through the compositeness angle, there is generically no spin-2 anchor mass-locked to the octet, and no second family at a fixed ratio. At the snapshots: a composite octet *could* imitate end-2026 through end-2041 in $t\bar{t}$ alone — this is the one framework that can gate-crash the first bump — but it fails the post-discovery program: no anchor at the same mass at FCC-hh (F-KK-1 Tier 2), no 1.831× overtone, free P_t rather than the band $[+0.88, +0.92]$. The discrimination is therefore scheduled rather than immediate, and Part IX is its calendar.

Versus generic Z' models. The workhorse of BSM searching, and the cleanest contrast. A Z' at 3.85 TeV with electroweak-strength lepton couplings is already excluded; a leptophobic Z' survives but only by construction, and it comes alone — no color octet locked to its mass, no graviton, no interference with QCD (a color singlet cannot interfere with the gluon channel), and no reason for its mass. At the snapshots: a leptophobic Z' tuned to 8.7 fb would track the octet's counting significance but fail the *shape* program immediately — no low-side enhancement, no below-SM crossing (P-KK-4), the wrong dijet-to-$t\bar{t}$ decomposition — and the interference dichotomy of O-KK-28 separates the two at the template-fit level with the same data.

Versus supersymmetry. The comparison is methodological rather than observational, because the MSSM makes no sharp statement about this channel at all — which is itself the point §1 labored. A century of collective effort produced, in the MSSM, a framework whose >100 parameters can accommodate any $t\bar{t}$ spectrum the LHC records; it is not falsifiable by this measurement, and symmetrically, it cannot claim it. The structural contrast is sharper: GG-6's bulk admits *no* super-Poincaré extension (no translation generators along r, θ; §5.5), so fifty years of superpartner nulls — unexplained embarrassments under naturalness — are, here, confirmations of a theorem. One framework predicts the silence the other must excuse.

Versus large extra dimensions (ADD). The other great extra-dimensional road fields *spatial* dimensions at sub-millimeter-to-TeV scales, with a quasi-continuum KK tower and a fundamental scale M_D scanned per search. Its collider signatures — monojets from graviton escape, non-resonant dijet/dilepton excesses rising with energy — are qualitatively different from a resonant cluster, and its towers are picket-fence continua (Appendix J, Step 1), not an organ pipe. At the snapshots: ADD would show smoothly rising tails, no peak, no interference structure; it fails every resonant clause of the P-KK list and is separated from GG-6 by the first spectrum plot. The deeper contrast is the one of §5.5: ADD's dimensions carry moduli and require stabilization; GG-6's carry neither.

Versus the string landscape. String theory is not a competitor at the level of a 3.85-TeV template — it makes no such statement — and the honest comparison is the one A2 draws: two consistent strategies with their irreducible assumptions at opposite ends of the energy axis. What belongs *here* is the falsifiability asymmetry. A landscape with 10^{500} vacua can absorb any collider outcome, positive or negative; the present framework is killed by one. If the cluster is found with its correlated pattern, string model-building will retro-fit it within the decade — such is the flexibility — but the *prediction* will have belonged to the framework that made it when it could still have been wrong. Priority of risk, not priority of publication, is the scientifically meaningful ordering, and it is the one quantity a landscape cannot manufacture after the fact.

57 Predictions for falsification

A falsifier is a prediction plus a decision rule plus a price; a prediction alone is only a hope with units. This section restates the paper's three falsifiers as narratives — where each was declared, what has happened to it since, and what ends it — and then walks the ten registered predictions individually, because a table row (the summary remains as Table 31) cannot carry what a referee needs: for each claim, *what* is predicted, *where* it was derived, *how* it will be measured, and *what outcome kills it.*

Table 30: Framework-by-observable discrimination matrix, transposed for completeness: observables as rows, frameworks as columns. Each cell: does the framework generically reproduce the observable? (✓ = yes / forced; ∼ = achievable by parameter choice; × = fails or forbidden; — = no statement.) The MSSM column is mostly "—" because the framework makes no sharp claim in this channel — which is the methodological point of §1; the string-landscape column is "∼" throughout because any cell is fillable after the fact (priority of risk, see text). The final row names the earliest measurement separating each framework from GG-6.

Observable	GG-6	RS	CHM	Z'	MSSM	ADD	Str.
Mass derived (no scan)	✓	×	×	×	—	×	∼
Rate derived (no coupling dial)	✓	×	×	×	—	×	∼
EW scale tied to new mass	✓	∼	∼	×	∼	×	∼
Cluster degeneracy $m_{g^*}{=}m_{Z'_B}$	✓	×	×	×	—	×	∼
Exact leptophobia (P-KK-6)	✓	×	∼	∼	—	×	∼
Width Γ/m predicted	✓	×	×	×	—	×	∼
Interference sign (P-KK-4)	✓	∼	✓	×	—	×	∼
$P_t \in [+0.88, +0.92]$ (P-KK-5)	✓	∼	∼	∼	—	×	∼
BR fingerprint 70:3.5:1 (P-KK-2)	✓	×	∼	∼	—	×	∼
No radion / moduli (P-KK-9)	✓	×	✓	✓	×	×	×
No angular satellites (P-KK-7)	✓	✓	✓	✓	—	×	∼
Overtone at 1.831× (P-KK-8)	✓	∼	×	×	—	×	∼
Hierarchy as output	✓	∼	∼	×	∼	∼	∼
SUSY absence explained	✓	—	—	—	×	—	×
Lepton-collider silence (P-KK-10)	✓	×	∼	∼	—	×	∼

Column heads: RS = 5D Randall–Sundrum benchmark; CHM = composite-Higgs / KK-gluon models; Z' = generic leptophobic Z'; Str. = string landscape. *Earliest measurement separating each framework from GG-6*: RS — the $\ell\ell/\gamma\gamma$ lines it requires and GG-6 forbids (already decided by existing data); CHM — the spin-2 anchor and the 1.831× overtone (FCC-hh, Part IX); Z' — the interference shape, inside the discovery fit itself; MSSM — none (no claim in this channel); ADD — resonant versus non-resonant spectrum shape (the first plot); string landscape — priority of risk (see text).

F-KK-1, the two-tier cluster falsifier: a history. Declared as the paper's original headline — the spin-2 anchor at the lock scale — F-KK-1 was requalified, in public and by this paper's own Corollary T-KK-2c, when the graviton's gluon coupling was computed and came out at $0.47/K$: the anchor's LHC rate fell to invisibility, and the falsifier restructured into two tiers. *Tier 1 (HL-LHC, decisive)*: the lock-scale implementation dies on a null in the cluster-aware $t\bar{t}$/dijet combination within 3.85 ± 0.10 TeV at $\lesssim 2$ ab^{-1}, carried entirely by the un-suppressible octet. *Tier 2 (FCC-hh, confirmatory)*: the direct spin-2 anchor at the same mass, $\Gamma/m = 6.3\%$, branching 0.69:0.15:0.08:0.08, rate ≈ 0.13 pb at 85 TeV — discovery in the first weeks or falsification of the cluster reading. The division-of-labor clause is itself falsifiable: finding either tier without the other kills the interpretation. The requalification episode belongs in this narrative because it demonstrates the register the paper operates in: when a derivation moved a falsifier, the falsifier was restated *harder*, not retired.

F-KK-2A: fired. The free-field scenario — spin-1 pair at $0.639\,M_{\rm lock} = 2.461$ TeV, $\sigma \times {\rm BR} = 224\,[180\text{–}325]$ fb, an unmissable rate — was confronted with the published Run-2 record in §31 and lost: the prediction sits a factor 3.6–11 above the observed exclusion, outside its band by any reading. The closure is graded provisional on template shape alone (no plausible line-shape correction recovers a factor of three), and its structural yield is recorded in §10.4: the throat vectors are composite. One of this paper's falsifiers has therefore already executed, in the kill direction, on schedule — the strongest possible evidence that the others are load-bearing.

F-KK-2B: live, and the center of everything. The composite-scenario octet at $M_{\rm lock}$: $8.7\,[7.0\text{–}12]$ fb at 13 TeV, 13.6 fb at 14, width 14.6%, interference sign fixed. Its decision procedure is the entire content of Part VIII; its price is stated ($\lesssim 2\ {\rm ab}^{-1}$); its earliest possible execution is a reanalysis of tape that already exists. Every other claim in the paper is downstream of this one's verdict.

Table 31: The numbered predictions of this paper, with grades ([T] theorem; [D*] derived-conditional; [M] menu; [E] external-input). Four grades and no unlabeled rows. A theorem entry falls with its theorem, a derived-conditional entry falls with the condition printed beside it, a menu entry is one of a finite listed set, and an external-input entry is somebody else's number this paper consumes. Mixing the four is how a prediction list comes to look stronger than it is, and the grade column exists to stop that.

Id	Prediction	Grade
P-KK-1	$\sigma \times {\rm BR}(g^* \to t\bar{t}) = 8.7\,[7.0\text{–}12]$ fb at 13 TeV; $13.6\,[11\text{–}19]$ fb at 14 TeV; ≈ 2.8 pb at 85 TeV	[D*]
P-KK-2	Octet branching fingerprint $t\bar{t} : b\bar{b} : jj \approx 70 : 3.5 : 1$; ${\rm BR}(t\bar{t}) = 0.94\,[0.93\text{–}0.95]$	[D*]
P-KK-3	Graviton pattern $t\bar{t} : W_LW_L : Z_LZ_L : hh = 0.69 : 0.15 : 0.08 : 0.08$, $\Gamma/m = 6.3\%$, $\sigma(85\ {\rm TeV}) \approx 0.13$ pb	[D*]
P-KK-4	Interference sign: enhancement below the pole, deficit above; SM crossing at $\simeq 4.9$ TeV; centroid pull $-\mathcal{O}(100)$ GeV	[T]
P-KK-5	Top polarization $P_t = +0.91$, band $[+0.88, +0.92]$; $A_{FB} = 0$	[T]/[D*]
P-KK-6	Exact leptophobia: ${\rm BR}(Z'_B \to \ell\ell) = 0$; $m_{Z'_B} = m_{g^*}$ exactly	[T]
P-KK-7	No angular satellites at the TeV scale (the θ tower sits at $\sim E_{\rm GUT}$)	[T]
P-KK-8	Second radial family at $1.831\,M_{\rm lock} = 7.047$ TeV (position; rate estimate $\sim$10 fb at 85 TeV)	[T]/[M]
P-KK-9	No radion; Higgs couplings Standard-Model-like apart from $G_1 \to hh$	[T]
P-KK-10	Lepton-collider silence: no s-channel cluster line in $\mu^+\mu^-$; SM-like contact terms at 3 TeV	[T]

The ten predictions, walked individually. *P-KK-1 (the rates).* $\sigma \times {\rm BR}(g^* \to t\bar{t}) = 8.7\,[7.0\text{–}12]$ fb at 13 TeV, $13.6\,[11\text{–}19]$ at 14, ≈ 2.8 pb at 85 — derived in §31 from $\xi_{\rm UV}$ and the parton chain of Appendix K; measured by the dedicated template fit; killed by a fitted rate outside the band in either direction (an *excess* beyond $\sim$12 fb falsifies as surely as a null). *P-KK-2 (the octet fingerprint).* $t\bar{t} : b\bar{b} : jj \approx 70 : 3.5 : 1$, ${\rm BR}(t\bar{t}) = 0.94\,[0.93\text{–}0.95]$ — from the width table of Appendix I; measured by the dijet and $b\bar{b}$ companions once the $t\bar{t}$ line exists; killed by any ratio off by more than the band, most sharply by a dijet line *stronger* than $t\bar{t}/20$. *P-KK-3 (the anchor pattern).* G_1 at the same mass with $t\bar{t} : W_LW_L : Z_LZ_L : hh = 0.69 : 0.15 : 0.08 : 0.08$ and

$\Gamma/m = 6.3\%$ — from T-KK-2c and the stress-tensor couplings; measured at FCC-hh across four channels; killed by a missing anchor, a wrong width, or transverse-dominated diboson decays (the longitudinal dominance is the Higgs-sector signature). *P-KK-4 (the sign).* Constructive below the pole, destructive above, SM crossing at $\simeq 4.9$ TeV, centroid pulled $-\mathcal{O}(100)$ GeV — the theorem of Part VI; measured by the sign-floated template fit of the discovery analysis itself; killed by a wrong-sign or no-interference preference at HL-LHC statistics. *P-KK-5 (polarization and asymmetry).* $P_t = +0.91$, band $[+0.88, +0.92]$; $A_{FB} = 0$ — from the chirality structure $r_{t_R} \gg r_{t_L}$ and the vector-like light-quark coupling respectively; measured through the lepton-analyzer distributions to $\delta P_t = 0.08$ by 2041; killed by low polarization (a vector-like top coupling) or any significant A_{FB} (an axial light-quark coupling), either of which breaks T-KK-2 directly. *P-KK-6 (leptophobia and degeneracy).* $\mathrm{BR}(Z'_B \to \ell\ell) = 0$; $m_{Z'_B} = m_{g^*}$ exactly — assignments and shared boundary-value problem; measured by the dilepton null (to 10^{-6} at FCC-hh) and the two-component line-shape decomposition; killed by any dilepton line at the cluster mass or any resolved splitting. *P-KK-7 (no angular satellites).* The θ tower at $\sim E_{\mathrm{GUT}}$, hence *nothing* decorating the cluster at percent-level mass offsets — the non-spatial reading's sharpest negative (§5.5); measured for free in any high-resolution scan of the peak region; killed by satellite structure, which would simultaneously revive the spatial-6D reading this paper rejects. *P-KK-8 (the overtone).* G_2 at $1.831 \times M_{\mathrm{lock}} = 7.047$ TeV, position theorem-grade, rate $\sim$10 fb at 85 TeV (estimate) — the Bessel ratio of §9.1; measured at FCC-hh spectroscopy; killed by a position-null at FCC-hh statistics (O-KK-18 records the no-excuses clause) or by an overtone at the *wrong* ratio, which would falsify the cavity itself. *P-KK-9 (no radion).* Higgs couplings Standard-Model-like except the derived $G_1 \to hh$ line — from the integer length (PS-KK-3); measured at the per-mille level by FCC-ee; killed by a scalar-mixing pattern in Higgs precision, discovery or no discovery at the LHC. *P-KK-10 (the lepton-collider silence).* No s-channel cluster line in $\mu^+\mu^-$; SM-like contact terms at a 3-TeV stage — from the same localization geometry as PS-KK-1; measured whenever a muon collider runs; killed by any leptonic line at 3.849 TeV, which would falsify the $\mathrm{U}(1)_B$ assignment and the localization in one stroke.

The falsification lattice. The ten predictions are not independent shots; they form a lattice with the three theorems at its nodes, and the lattice is the real object under test. A wrong sign (P-KK-4) breaks T-KK-2's mode structure and takes P-KK-5 with it; satellites (P-KK-7) break the non-spatial reading and take P-KK-8's ratio with them; a radion (P-KK-9) breaks the integer length and takes the mass itself. Conversely a discovery that confirms the lattice confirms it *jointly* — which is why §56 could promise that no neighboring framework survives the full list: frameworks can imitate entries, but the correlations are the geometry.

58 Open questions

Honesty requires the complement, at the same depth as the claims. Six items remain open; for each we state why it is open (not merely that it is), what has been tried, the closing condition, and who is best placed to close it. Appendix O carries the same list as a project roadmap.

1. L-KK-1: the g_B readout menu. *Why open*: the $\mathrm{U}(1)_B$ coupling is fixed by GDOS counting in principle, but the counting has four defensible readout conventions, and this paper declines to pick one without the derivation — hence a menu $\{0.14, 0.23, 0.36, 0.57\}$-class spread and Z'_B rates quoted per point. *Tried*: the direct transcription of the $\alpha_s = 1/(\pi e)$ counting; it fixes the structure but not the convention. *Closing condition*: the same GDOS derivation carried once more, carefully, for the Abelian line — a bounded theory exercise. *Impact*: single-valued Z'_B phenomenology and a sharpened fifth-force window (Arisaka, 2026, A9-SM).

2. The NLO interference K-factor. *Why open*: the interference term is linear in the resonance amplitude, so its K-factor is *not* the Breit–Wigner's $K_{q\bar{q}} = 1.3$, and no published computation covers a broad color-octet interfering with the $t\bar{t}$ continuum at this mass. *Tried*: the $\sqrt{K}$ prescription used throughout, standard but uncontrolled at the tens-of-percent level. *Closing condition*: a conventional NLO calculation — the one entry on this list a standard QCD group could adopt tomorrow, with immediate experimental consumers (§44.5). *Impact*: the low-side template firmed; the largest single theory systematic of the discovery fit retired.

3. The $n \geq 2$ overlap couplings. *Why open*: Corollary T-KK-2c was evaluated at x_1; the same machinery at x_2, x_3 is straightforward for diagonal couplings but the inter-mode mixing needs care. *Closing condition*: Eq. (88) at higher zeros — an afternoon for the diagonal terms, a short paper done properly. *Impact*: P-KK-8's rate promoted from estimate to derivation before FCC-hh needs it.

4. The composite-scenario recurrence spacing. *Why open*: Scenario B fixes the lowest states by matching, but the spacing of the vector recurrences is a property of the strongly coupled boundary theory, and no weak-coupling trick reaches it. *Closing condition*: a dual (holographic or lattice-inspired) computation, or the data themselves at FCC-hh. *Impact*: search targets for the $\mathcal{O}(10^2)$ fb recurrences of Part IX.

5. The detector-level angular analysis. *Why open*: the 10–30% discovery gain from $\cos\theta^*$ requires the real background cocktail's angular shapes, which live inside the collaborations. *Closing condition*: a full-simulation 2D study — thesis-sized, and the natural companion to the reanalysis itself. *Impact*: the bounded gain of §44.3 banked, potentially advancing the five-sigma crossing by a year.

6. The composite width band. *Why open*: $\Gamma/m = 14.6\%$ carries the free-field overlaps; the composite reading could shift it within a band no current calculation pins. *Closing condition*: self-closing — the HL-LHC line-shape fit measures Γ/m to a few percent, converting the band into a cross-check of the overlap structure. *Impact*: a measured width becomes a second-generation test of T-KK-2's interior.

None of the six shields the verdict: the falsifier stands on the theorem-grade spine — mass, degeneracy, sign, leptophobia — alone. That is the design: the open questions sharpen the program; they do not prop it up.

59 The completed cohort and future prospects

This paper is one of twenty-two, and its place in the cohort is load-bearing in both directions. *Imports:* the bulk and its canonical metric (Arisaka, 2026, A2-GG6, Theorems 0 and 2), the OGN-5 integer cascade (Arisaka, 2026, A3-BI6), the lock-scale chain and the electroweak over-constraint (Arisaka, 2026, A9-SM), the open-system reading of detection (Arisaka, 2026, A5-GDOS; Arisaka, 2026, A6-GRIP), and the no-radion theorem (Arisaka, 2026, A4-Euclid). *Exports:* the requalified falsifier F-SM-1 of Arisaka (2026, A9-SM) now points at this paper's computed rates; the g_B menu feeds the fifth-force window; and the collider verdict, when it arrives, propagates to the constitutional core — in both directions. The near-term prospect is stated once more because it is the paper's operational request: *a dedicated, interference-aware, pre-registered reinterpretation of the $t\bar{t}$ data already on tape*, jointly by ATLAS and CMS, which Part VIII shows may already contain the discovery. The far-term prospects are Part IX's: the anchor, the overtones, and the leptophobic silence, distributed across the machines of 2035–2070.

60 Summary of the Discussion

The strengths, stated plainly: a derived, immovable target whose integer has already survived a four-digit audit at the electroweak scale; a fully specified signal (mass, width, rate, sign, polarization, fingerprint); a calibrated calendar with evidence-level sensitivity already on tape; statistics that reward the derivation itself (no look-elsewhere, no template dilution); an exclusion already converted into structural knowledge (Scenario A, the observed functioning of the kill machinery); a discrimination lattice that no neighboring framework survives intact (Table 30); and a falsifier whose failure mode is clean retirement. The weaknesses, stated with equal plainness: the rate band is PDF- and K-factor-limited until LHeC-class input arrives; the interference template is leading-order, with the NLO task named but not done; one coupling remains a four-point menu; the $n \geq 2$ couplings are estimates; and no public collaboration reinterpretation of the exact benchmark yet exists — the paper's central request. The posture these imply is the one the paper has kept throughout: aggressively experimental, theoretically disciplined — and, with Part VIII in hand, *urgent*, because the evidence may already be sitting on tape in Geneva, waiting for a template that was only written this year.

Part X made the arguments. In one breath each: the result is significant because both of its possible endings are informative, which fifty years of adjustable predictions could never say; it is unique in kind, not degree, and its innovations are auditable one by one; no neighboring framework survives the full fifteen-observable fingerprint; the falsifiers are priced, dated, and — in one case already — executed; and the open questions sharpen the program without propping it up. The argument of the paper is now complete. What follows is its ending in the older sense of the word: conclusions, a personal accounting, and the questions a result of this shape forces one to ask about number, nature, and the act of knowing.

61 Part X in plain English: what the result is worth, set against the alternatives

This is the Part where the paper argues with itself and with everyone else.

The claim to significance is narrow and should be stated narrowly: not that the framework is right, but that it has produced a collider number with no adjustable parameter behind it, which the alternatives have not. That is a claim about method. It is worth something only if the number is wrong in a way that would be visible, and the Part is careful to say what visible means here.

The comparison with the other roads beyond the Standard Model is done property by property rather than by verdict, because the roads were built for different purposes. Supersymmetry addressed a naturalness problem this framework does not have. String theory operates at an energy no collider reaches. Judging either against a 3.85 TeV prediction would be answering a question they were not asking, and the table says which cells are empty for that reason rather than by failure.

What the Part does claim as new is structural: the absence of dials, the fact that the same integer that fixes this mass also fixes the electroweak scale, and the pre-registration of a number before the measurement exists. The open questions are listed as work orders with closing conditions attached, and one of them — an interference calculation — is the input the proposed reanalysis needs first.

Part XI
Conclusions

What the paper has established; the objective limitations; a personal perspective from fifty years on the experimental front line; the vision of a geometric century; and what either verdict would mean — for physics, and for the way human beings come to know the world

62 What the paper has established

The results of this paper, in the order they were earned. *Theorem T-KK-1*: the first observable Kaluza–Klein resonance of the six-dimensional bulk is the radial lock cluster at $M_{\rm lock} = \frac{1}{2} M_{\rm Pl} e^{-35} = 3.849$ TeV, with no continuous freedom. *Theorem T-KK-2*: the cluster's fine structure — the exactly degenerate spin-1 pair, the exact free-field eigenvalues, the coupling $\xi_{\rm UV} = -0.194$ with its sign, the two-scenario dichotomy — derived, with Scenario A subsequently excluded by existing data and the dichotomy resolved to the degenerate cluster. *Corollary T-KK-2c*: the graviton–gluon coupling computed in closed form, $\kappa_{gg} = 0.47/K$, moving the direct spin-2 test to FCC-hh and concentrating the entire LHC falsifier on the color octet. *Part VI*: the complete derived anatomy of the $t\bar{t}$ signal — angular forms, polarization, the interference-sign theorem and its line shape. *Part VIII*: the anchored, calibrated simulation converting the prediction into a calendar — evidence-level sensitivity already on tape, combined discovery-grade sensitivity a live possibility, single-experiment five-sigma luminosity during 2032, and the pre-registration advantages quantified. *Part IX*: every remaining measurement assigned to a named future machine, through 2070.

Now the same results once more, in plain English, because a conclusion should be readable by the colleague down the hall who is not a collider physicist. A theory with no adjustable numbers — its entire freedom is five counting numbers — says there is a new particle, and says everything about it: where it is (3.85 TeV, about four thousand times heavier than a hydrogen atom's proton), how wide it is, how often the LHC should have made it already (roughly ten thousand of them are sitting, unrecognized, in the data recorded since 2015), what it decays into (top quarks, overwhelmingly), what it refuses to decay into (electrons and muons — exactly the particles our cleanest searches watch), and even the subtle way it interferes with ordinary background, down to the sign. The same five counting numbers already produce the known strength of the weak force scale — that part of the ledger was checked decades ago, by people with no stake in this theory. What this paper adds is the practical half: proof, at the fidelity of a real analysis, that the existing data may already contain the evidence, and a step-by-step recipe by which the experiments can decide the question — either way — within this decade.

The practical collider statement in one sentence: the first experimentally relevant new-physics structure in GG-6 should be a *spin-2-anchored, octet-led resonance cluster centered near* 3.849 TeV*, with top-rich and dijet-friendly branching, a mild leptophobic admixture, and a second* radial *family* *near* 7 TeV — under real pressure from existing clean-channel limits (pressure the derived coupling texture explains), and decided by exactly the analysis this paper specifies.

63 The objective limitations and the next technical tasks

The paper's honest boundary, stated as a task list rather than a confession — and prefaced by the reason the list exists at all. A framework that graded its own claims generously would have nothing to put here; this one carries a menu, several estimates, and a leading-order template, all labeled at their point of use, and the discipline of *saying so* is what separates a falsifiable program from an advertising campaign. *Theory side:* the rates are LO×K with PDF-limited bands (task: NLO signal and, above all, the NLO interference K-factor of §44.5 — the single largest theory systematic of the discovery fit); one coupling remains a discrete menu (task: close L-KK-1 by the GDOS counting that fixed α_s); the $n \geq 2$ couplings are same-coupling estimates (task: run the T-KK-2c machinery at x_2, x_3 — an afternoon for the diagonal terms); the composite-scenario recurrence spacing is undetermined (task: the strong-coupling dual calculation, the genuinely hard one). *Experiment side:* no public collaboration reinterpretation of the exact benchmark exists (task: the pre-registered $\Gamma/m = 15\%$ interference-aware template analysis, jointly ATLAS+CMS, on data already recorded — the paper's central request); the angular dimension is bounded but unbanked (task: a detector-level 2D study with the anchored background cocktail); the sideband-contamination prescription must be built into the fit design from the start. None of these tasks blocks the verdict; all of them sharpen it.

64 A personal perspective: fifty years on the experimental front line

A word is owed about why *this* paper, of the twenty-two, is the one closest to its author's life. I have spent fifty years as an experimental particle physicist — at Kamiokande, waiting for the proton decay that grand unification promised and that never came; on *CP* violation in the kaon and *B*-meson systems; on CMS at the Large Hadron Collider, where I watched the Higgs *complete* the Standard Model rather than break it; on the Pierre Auger Observatory at the ultra-high-energy frontier; and on XENON, in the direct search for dark matter. Five front lines, fifty years, and on every one of them the same experience: theory offered us targets that could always retreat, and nature offered us confirmations of a model written before I began.

Let me say what that felt like from the inside, because the sociology matters as much as the physics. An experimentalist's working life is spent converting other people's hopes into shift schedules. You calibrate for two years, run for five, unblind with your heart pounding — and when the plot is empty, the theorists whose prediction you tested explain, sometimes within the week, why the interesting region was always just above your reach. After the third or fourth such cycle you stop being angry and start being careful: you learn to ask, before committing a decade, *what exactly will you concede if we see nothing?* Fifty years of asking that question, and almost never receiving an answer, is the biographical origin of every discipline in this paper. I have been on the receiving end of theoretical predictions often enough to know precisely what an experimentalist needs and almost never gets: one mass, not a scan range; a rate with a derivation, not a coupling dial; a width, a sign, an angular shape; and a statement of what would kill the theory, in advance, in writing. The reinterpretation card of Appendix M, the pre-registered template of Part VIII, the dated kill window — these are the artifacts I wished for across five decades of searches, written now from the other side of the table.

When the Standard Model's own list was exhausted in 2012, I left the field for the better part of a decade and worked on the nervous system of *C. elegans* and the architecture of human vision. GG-Theory grew, unexpectedly, out of that detour — the story is told in Arisaka (2026, A9-SM) and Arisaka (2026, A7-Qi4) — and when it returned me to particle physics, it returned me with a number: 3.849 TeV. The graduate student addressed in this paper's Preface is, of course, myself in 1976, offered the choice I wish I had been given: a prediction that cannot move, testable within one PhD, on data partly already recorded. The calendar of Part VIII runs through the mid-2030s. I intend to watch the verdict arrive, and I have written this paper so that either verdict — the cluster, or its clean exclusion — will be worth the fifty years.

65 The vision: what a geometric century could look like

Permit me one section of unguarded ambition, in plain English, before the philosophy and the close. Suppose the octet is there. Not as advocacy — Part VIII has stated the odds as honestly as I know how — but as a planning exercise, because visions are how fields decide what to build.

The twentieth century learned that matter is fields and that fields are geometry's tenants: they live *on* spacetime. The claim under test in this paper is one notch deeper — that the tenants are the landlord's own rooms; that the energy scale at which you probe the world, and the gauge phase that makes charged-matter interference possible, are not bookkeeping devices but *directions*, as real as north; and that once you count those directions correctly, the numbers we have spent a century measuring — the strength of the strong force, the mass of the weak scale, and now, testably, a resonance at 3.849 TeV — stop being inputs and become readings. If that claim survives its collider test, the twenty-first century's physics acquires a different daily texture. The question at the frontier stops being "what new particle might exist?" — a question that has, for fifty years, admitted too many answers — and becomes "what does the geometry say, and to how many digits can we check it?" — a question with the texture of nineteenth-century spectroscopy, which is meant as the highest compliment: spectroscopy is the mode in which physics has historically been most honest, most cumulative, and most kind to its young people, because a student with a good idea and a data set can be *right* in a way no committee can argue with.

In that century, the machines of Part IX are not lottery tickets but instruments in the older sense: FCC-hh does for the scale throat what the Balmer series did for the atom — reads its overtones; the muon collider's silence at 3.849 TeV is as informative as any peak; the Higgs factories audit the ledger digit by digit. And the two cultures whose fifty-year estrangement §1 recounted — the theorists who could not commit and the experimentalists who could not forgive — have, for the first time in living memory, a shared object neither can do without: a number that theory cannot move and experiment cannot ignore. My deepest professional hope for this paper is not that it is right; it is that it makes that kind of physics imaginable again to a twenty-five-year-old.

66 The trailhead, not the summit: a final reflection

If a 3.85 TeV spin-2-anchored cluster is discovered with the correlated pattern described here, the reward will be extraordinary: not merely another BSM particle, but direct evidence that the deep ultraviolet of nature is better described by a six-dimensional geometry–gauge-locking stage than by a purely four-dimensional field ledger. And yet discovery would be a beginning, not an end: the anchor, the overtones at the Bessel ratios, the exact zero of leptophobia, the silence of the lepton machines — the entire spectroscopy of a non-spatial geometry, mapped across the machines of Part IX — would lie ahead, the way the hydrogen spectrum lay ahead of Balmer's formula. The LHC is the trailhead of that program, not its summit.

If instead no such structure appears and the full combined-channel `GG-6` corridor is excluded, the present implementation of the theory will have failed one of the clearest collider tests imaginable — cleanly, permanently, and without a zombie phase, because there is no parameter with which to retreat. Either way, the experiment is worth doing. That is the mark of a good falsifier; it was the missing property of the fifty years this paper's introduction recounted; and it is the property this paper was written to restore.

67 What it would mean: matter, number, and the legibility of nature

A paper that asks two collaborations for person-years of work owes them more than a template; it owes them an account of what the answer would *mean*. This closing section attempts that account at the depth the question deserves, in four movements.

On falsifiability as an ethic, not a slogan. Popper made falsifiability a criterion of demarcation; half a century of practice has shown that as a mere criterion it is too easy to satisfy in the letter and evade in the spirit — every scanned model is "falsifiable" at each parameter point and immortal as a family. What this paper has tried to practice is falsifiability as an *ethic*: the deliberate construction of a position from which one's future self cannot escape. The integers are load-bearing across twenty-two documents precisely so that no quiet revision is possible; the kill window is priced in inverse attobarns so that no schedule slip can blur it; the grades are printed so that no later reader can be told the menu was always a theorem. There is a moral structure here older than science: it is the structure of a *promise*. A theory that cannot break its word to the data is behaving, in the only sense available to a mathematical object, honorably — and fifty years of BSM history are best summarized as the slow discovery of how hard, sociologically, that honor is to maintain when careers depend on survival. If the cluster is found, the finding will vindicate a geometry. Whatever is found, I hope the *form* of the wager outlives the wager itself.

On why counting numbers, and the old Pythagorean question. Wigner famously called mathematics "unreasonably effective" in physics — a puzzle, because there is no evident reason why nature should honor our notation. The framework under test dissolves the puzzle in the most direct way imaginable: if the constants of nature are *counts* — dimensions counted, thresholds counted, windings counted — then the effectiveness of number is not unreasonable but tautological, because the world's deep structure would *be* arithmetic enacted in geometry. This is the oldest speculation in Western natural philosophy: the Pythagoreans' "all is number," which the continuum revolution of the seventeenth century seemed to retire permanently in favor of the real line and the differential equation. It is worth savoring the historical irony that the retirement may have been premature at exactly the deepest level: the continuum machinery survives in GG-6 — fields, actions, renormalization — but every place where the continuum seemed to demand an arbitrary *choice* (a coupling, a mass ratio, a mixing angle) is, on this account, occupied by an integer instead. A discovery at 3.849 TeV would not prove Pythagoras right; nothing proves a metaphysics. But it would establish, as an experimental fact, that at least one number formerly believed continuous and arbitrary — the mass scale of new physics — was a count all along. Physics has changed its mind about the continuum before, exactly once, in 1900, and we named the resulting century after the change: *quantum*. The wager of this cohort is that the twentieth century stopped one dimension short.

On a world that measures itself. The strangest feature of the framework, judged by classical intuitions, is not the extra dimensions; it is the claim (§3; Arisaka, 2026, A5-GDOS; Arisaka, 2026, A8-QM) that *observation is not something we add to nature but something nature does* — that the projection from bulk to boundary *is* measurement, and that quantum mechanics is its grammar rather than a separate rulebook. Philosophy has circled this possibility for a century under many names (the measurement problem, the participatory universe, the view from nowhere), always blocked by the same objection: such claims seemed permanently untestable, differences of interpretation rather than of physics. What the lock cluster does — and this may be its deepest significance, beyond any particular geometry — is drag one consequence of that picture into a detector. If registered observables are boundary readouts of one projection, then their values are collectively fixed, and the cluster at 3.849 TeV is simply the next unread line of a ledger whose previous lines (the weak scale, the strong coupling) have already been read and found consistent. A positive verdict would therefore be evidence, indirect but real, for the most philosophically radical clause in the entire cohort: that the ancient partition between the world and the knowing of the world was drawn in the wrong place — that physics, pursued far enough, does not merely describe nature but participates in nature's own self-description. I am aware how large that sentence is. It is printed here, rather than in a philosophy journal, because for the first time in my lifetime a sentence of that size has a luminosity attached to it.

On either verdict, and the legibility of the world. Let me end with the question beneath all the others: is nature legible — not merely describable in retrospect, but *readable in advance*, by finite minds, down to its constants? Fifty years of frontier physics have quietly taught most of us to answer no: to expect that the deepest layer is either an accident (the landscape), or forever beyond our energy reach, or intelligible only statistically. Every result in this paper is a bet against that resignation. If the cluster is there, the answer to the legibility question becomes *yes, at least one layer deeper than we had any right to expect* — and every young physicist inherits a world in which the constants are a text and the instruments are how one turns its pages. If the cluster is not there, the answer is subtler but not empty: we will have learned that *this* reading fails, learned it cleanly, at a stated cost, with the manner of the failure itself a contribution — the first zero-parameter exclusion in the history of the energy frontier, a permanent lesson in what such theories must do better. Legibility, after all, has two requirements: a world with structure, and readers with the discipline to be corrected by it. Only the first is nature's responsibility. The second is ours, and it is the one thing this paper could guarantee. The mass is the bet; the LHC is the table; the cards are already partly dealt, sitting on tape in Geneva. Let the data speak — and let us have the grace to listen, whichever way they answer.

What remains after conclusions is service: the appendices that teach the material from scratch, the objections answered one by one, and the references with their reasons — the parts of a paper that work for the reader long after the argument has ended.

68 Part XI in plain English: what is being staked, and what it would mean to lose

A prediction that cannot lose is worth nothing, so it is worth being plain about what losing looks like here.

If the dedicated search is run on the data that already exists and finds nothing, and the high-luminosity run finds nothing after it, then the reading in this paper is finished. There is no version of it with the mass somewhere else. The integer is not adjustable, the cluster spacing is not adjustable, and the paper has printed the luminosity at which silence becomes fatal. That is the whole content of the wager.

If instead the cluster is there, what has been found is not one particle but evidence that a mass can be derived from counting rather than measured and accommodated. The specific number would matter much less than that.

The limitations are restated one last time because a conclusion that quietly drops them is not a conclusion. The rate depends on four named assumptions. One interference calculation is owed and has not been done. The composite scenario's recurrence spacing is open. Each is written into the grade of the prediction it affects.

What is unusual about the position is only its timing. The question does not need a new machine or a new generation of physicists; it needs a dedicated analysis of collisions already recorded, and a decade that has already begun. The cards are dealt and they are sitting on tape in Geneva.

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 OGGEP/GDOS/Gi-4/Qi-4 Pillars), theoretical judgments, physical interpretations, and any remaining errors are the sole responsibility of the author.

Part XII
Technical Appendices

A graduate-level course in the collider physics of the lock cluster — from-scratch derivations of the KK spectrum, the parton-luminosity cross-section chain, and the statistics of discovery, with worked numerical examples throughout — preceded by the reference dictionaries and the conceptual primers, and followed by the reinterpretation walkthrough and the promotion roadmap

---

The body of this paper has emphasized the theorem chain and its experimental confrontation. The present appendices supply the detailed technical content that supports each step, and taken together they are intended to be read as *a graduate-level course in the collider physics of the lock cluster*, addressed to an incoming ATLAS or CMS graduate student who will carry the search forward — the reader of the Preface, met again at working depth. We organize them into four groups. The *reference* appendices (A–D) provide the tabulated abbreviations, key symbols, the cohort naming convention, and the comprehensive summary tables, and serve as the reader's quick lookup. The five *conceptual primers* (E–I) come next, and a newcomer should read them first: born of questions actually asked, they explain how the gluon and g^* are one field (Appendix E), where the baryonic Z'_B comes from (Appendix F), what the "graviton" of `GG-6` is (Appendix G), why g^* dominates the cluster's rates and is so broad (Appendix H), and why one integer controls both the rate hierarchy and the $t\bar{t}$ dominance (Appendix I). The *constructive* appendices (J–L) are the course proper: a self-contained sequence covering the three competencies this search actually requires — where a KK spectrum comes from (Appendix J), how a hadron-collider cross section is computed from parton distributions (Appendix K), and how a discovery significance is built from the Poisson likelihood up (Appendix L) — each treated in *three steps.* First comes a full from-scratch derivation of the standard result, in the form a graduate student would meet it in a textbook, carried through with no skipped manipulations and closed by a worked numerical example. Then comes a frank statement of what that standard treatment *assumes but cannot fix* — the inputs it takes from outside and the questions it leaves open. Finally comes the `GG-6` step, in which the same quantity descends as a derived output of the geometry, and the worked example becomes a number from this paper's tables. The *operational* appendices give the reinterpretation card with its implementation walkthrough (Appendix M) and the collider dictionary (Appendix N); the promotion roadmap for the open items closes the Part (Appendix O).

Throughout these appendices we adhere to the two stylistic commitments of the cohort's appendix courses (Arisaka, 2026, A9-SM). First, every derivation is presented step by step with no skipped intermediate manipulations: a graduate student should be able to reproduce each calculation with paper and pencil (and, for Appendices J–K, with a dozen lines of code). Second, formal steps are accompanied by parenthetical *physics intuition* notes that translate the algebra into a mental picture. The combination is deliberate: the appendices have no page budget, because this is where the next generation is meant to learn the material on its own.

Prerequisites and recommended textbooks. The course assumes the standard first-year graduate sequence and is meant to be read *alongside* the canonical texts, not in place of them. Recommended companions, each cited again where it is best matched: Halzen & Martin (1984) (*Quarks and Leptons*) for the cross-section formalism of Appendix K at its most accessible; Ellis, Stirling & Webber (1996) (*QCD and Collider Physics*) as the definitive source for factorization, parton luminosities, and resonance production; Peskin & Schroeder (1995) for the field-theoretic underpinnings, including the KK mode decompositions of Appendix J; Thomson (2013) (*Modern Particle Physics*) for the experiment-aware first pass at everything; and Cowan (1998) (*Statistical Data Analysis*) together with Cowan et al. (2011) for the statistics of Appendix L. The recommended use of each constructive appendix: read the matching textbook chapter first, return here for the compact from-scratch recap of Step 1, and then proceed to Steps 2 and 3, where the standard result acquires a derived payload.

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-6, GDOS, BI-6, OGN-5, GRIP, etc.) so the reader can see at a glance which terms belong to the GG-Theory program proper. Standard physics and collider abbreviations imported from the wider literature (KK, RS, PDF, JES, etc.) are rendered in plain uppercase Roman, following the convention of the major journals; collider terms are expanded further in Appendix N. The third column gives the *primary link*: the cohort paper, or the section of the present paper, where each term is introduced in detail.

Table 32: Summary of abbreviations used in the present paper. The abbreviations are collected here rather than glossed at first use because the paper is read in more than one order, and a reader entering at the simulation Part should not have to find the introduction to learn what a reduced stream is. Where an abbreviation belongs to the cohort rather than to collider physics, the third column points at the paper that defines it.

Symbol	**Meaning**	**Primary link**
BI-6	The tensor-mediated anomaly-locking Bianchi identity	GG-A3-BI6; §5.5
BR	Branching ratio	standard
BSM	Beyond the Standard Model	§1
BVP	Boundary-value problem	§10.4
BW	Breit–Wigner resonance line shape	Appendix K
CEPC	Circular Electron–Positron Collider (China)	Part IX
CL	Confidence level	Appendix L
CLIC	Compact Linear Collider	Part IX
CTEQ	The CTEQ parton-distribution family (CTEQ6 used here)	Pumplin et al. (2002)
DNN	Deep neural network (the boosted-top tagger)	§38.3

continued on next page

Table 32 continued

Symbol	**Meaning**	**Primary link**
ESPP	European Strategy for Particle Physics	European Strategy Group (2026)
FCC-ee/hh	Future Circular Collider, lepton / hadron stages	Part IX
FSR	FCC Feasibility Study Report (2025)	FCC Collaboration (2025)
GA	Geometric Attractor	A6-GRIP
GDOS	The open-system boundary law	GG-A5-GDOS
GG-6 / GG-6	Geometry–Gauge (locking); the six-dimensional bulk	GG-A2-GG6
GGEP	Geometry–Gauge Equivalence Principle (AX-2)	GG-A1-Constitution
GRIP	The runtime cascade; regimes GRIP-In/GRIP-Out	A6-GRIP
GU-5	The five Geometric Universal constants	GG-A1-Constitution; §9.1
HGTD	High-Granularity Timing Detector (ATLAS Phase-2)	ATLAS Collaboration (2024)
HL-LHC	High-Luminosity LHC	CERN (2023)
HLT	High-Level Trigger	§38.2
ILC	International Linear Collider	Part IX
IMCC	International Muon Collider Collaboration	International Muon Collider Collaboration (2024)
IR / UV	Infrared / ultraviolet ends of the throat ($r = K$ / $r = 0$)	§5.5
ITk	ATLAS Inner Tracker (Phase-2)	ATLAS Collaboration (2024)
JES	Jet energy scale	§41.1
K-factor	QCD higher-order normalization ($K_{q\bar{q}} = 1.3$, $K_{gg} = 1.6$); unrelated to the OGN-5 integer $K = 35$	§31
KK	Kaluza–Klein	§5.5; Appendix J
LEE	Look-elsewhere effect	§44.1
LHeC	Large Hadron–electron Collider (concept)	Part IX
LO / NLO	Leading / next-to-leading order	Appendix K
LS3	Long Shutdown 3 (2026–2030)	§39.3
MTD	MIP Timing Detector (CMS Phase-2)	CMS Collaboration (2019)
NWA	Narrow-width approximation	Appendix K
OGGEP	Open-system GGEP: the bulk-to-boundary projection (AX-3)	GG-A1-Constitution; Appendix G
OGN-5	The five-integer ledger $(M, N, L; J, K) = (1, 3, 5; 7, 35)$	GG-A3-BI6
PDF	Parton distribution function	Appendix K
PDG	Particle Data Group	Particle Data Group (2024)
RG	Renormalization group	§5.5
RS	Randall–Sundrum	Randall & Sundrum (1999); §5.5
SIP	System instance packet (a full measurement chain in GDOS language)	GG-A5-GDOS; §16

continued on next page

Table 32 continued

Symbol	Meaning	Primary link
SM	Standard Model	standard
SR	Signal region	§38
TDR	Technical Design Report	Part IX
TLA	Trigger-level analysis (ATLAS)	Table 23
WP	Working point (of a tagger)	§38.3

B Glossary of key symbols

Table 33: The symbols of this paper at a glance, with the section or companion that defines each. It is meant to be read beside the paper rather than before it, and the third column is the part that earns its place: a symbol whose home is a companion paper is one this paper consumes rather than derives, and the distinction matters for every grade in the provenance table. Symbols introduced and discharged inside a single derivation are left out.

Symbol	Meaning	Home
M_{lock}	The lock-cluster mass, $\frac{1}{2}M_{\text{Pl}}e^{-K} = 3.849$ TeV	T-KK-1, §9.1
K, J	OGN-5 integers: throat length $K = 35$; $J = 7$	§9.1
$r,\ \theta$	The non-spatial coordinates: RG-scale $r = \ln(E_{\text{GUT}}/\mu)$; phase angle	§5.5
x_n	Bessel zeros of J_1: standing-wave conditions; $x_1 = 3.8317$	Eq. (16)
m_n	Radial tower masses $(x_n/x_1)M_{\text{lock}}$: 3.849, 7.047, 10.219 TeV	§9.1
ξ_{UV}	Light-quark (UV-endpoint) coupling amplitude, -0.194	T-KK-2, §10.4
r_{t_L}, r_{t_R}	Top localization band: 1.0, $[4, 5]$ (central 4.5)	§10.4
g_V^t, g_A^t	Top vector/axial couplings, $2.75\,g_s$, $1.75\,g_s$	Eq. (63)
κ_{gg}	Graviton–gluon coupling, $0.47/K = 0.0136$ (derived)	T-KK-2c, §10.9
g_B	$\text{U}(1)_B$ coupling; discrete menu (open lemma L-KK-1)	§10.10
$\Gamma_{g^*},\ \Gamma_{G_1}$	Widths: 561 GeV ($\Gamma/m = 14.6\%$); 243 GeV (6.3%)	Part VI
β_t	Top velocity at the pole, 0.99597	Part VI
θ^*	Production angle in the $t\bar{t}$ rest frame	Part VI
$P_t,\ A_{FB}$	Top polarization $+0.91\,[0.88, 0.92]$; forward–backward asymmetry $= 0$	Part VI
S, I, B	Breit–Wigner, interference, background yields per bin	Part VIII
f_b	Fractional per-bin background systematic (post-fit)	§40.2
$Z,\ q_0$	Discovery significance; profile-likelihood test statistic	§40
F_{cal}	Calibration factor to the real analysis, 4.33	§41
L^*	Statistics-to-systematics crossover luminosity, $1/(f_b^2 b_{1\text{fb}})$	§40.3
εA	Signal efficiency × acceptance, 0.065	§38.4
N_{trials}	Look-elsewhere trials factor, ≈ 75 for the generic scan	§44.1

v	Electroweak scale, $M_{\text{lock}}(2/5)^3 \approx 246$ GeV	§9.1
$E_{\text{GUT}},\ M_{\text{Pl}}$	UV endpoint $\approx 2.28 \times 10^{16}$ GeV; Planck energy	§4

C Global naming convention for cross-cohort citation

Following Arisaka (2026, A9-SM) (its Appendix C) and beginning with Arisaka (2026, A2-GG6), 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 Nth Theorem (e.g., T-KK-1 is the lock-cluster mass theorem of this paper).
- **L-Suffix-N** for the Nth Lemma (e.g., L-KK-1 is the open g_B-menu selection lemma).
- **C-Suffix-N** for the Nth Corollary (this paper's T-KK-2c retains its historical theorem-attached name).
- **D-Suffix-N** for the Nth Definition.
- **R-Suffix-N** for the Nth Remark.
- **P-Suffix-N** for the Nth Prediction (P-KK-1 . . . P-KK-10, §57).
- **O-Suffix-N** for the Nth Objection (O-KK-1 . . . O-KK-58, Part XIII).
- **F-Suffix-N** for the Nth Falsifier (F-KK-1 and the two-scenario pair F-KK-2A/2B).
- **I-Suffix-N** for the Nth Innovation (I-KK-1 . . . I-KK-10, §55.2).
- **PS-Suffix-N** for the Nth Problem-Solved entry (PS-KK-1 . . . PS-KK-6, §55.1).

The canonical Suffix list for the twenty-two cohort papers is: *Constitution* (A1), *GG6* (A2), *BI6* (A3), *Euclid* (A4), *GDOS* (A5), *GRIP* (A6), *Qi4* (A7), *QM* (A8), *SM* (A9), *Cosmos* (A10), *Origin* (A11), together with the phenomenology preprints *KK* (P1, the present paper), *N3* (P2), *Electron* (P3), *Proton* (P6), 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). 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; the global identifier is the same item viewed from outside its home paper.

D Comprehensive summary tables

This appendix gathers the paper's load-bearing conclusions into three tables, so that the paper can serve as a working reference document: the phenomenology ledger of the low-lying collider sector (Table 35), the falsification hierarchy (Table 36), and the discovery calendar established by the simulation of Part VIII (Table 37). A fourth summary object — the framework discrimination matrix — lives with the comparison discussion (Table 30, §56). Because these tables will be consulted by readers who have not read the Parts that generated them, each is introduced with a short guide to what it contains, where its entries were derived, and how it should — and should not — be read. For the reader with one minute, Table 34 is the paper.

Table 34: The ten numbers at a glance: the paper's load-bearing quantities, each with its provenance. Every entry is derived or bracketed as indicated; none is fitted. Ten numbers is the number a reader should be able to carry away, and the provenance column is why the list is short. Every entry can be traced to the theorem or the assumption that produces it, and an entry that could not be traced would not be in the table. None is fitted, which is a statement about the method rather than a boast about the result.

	Quantity	**Provenance**
1	$M_{\text{lock}} = \frac{1}{2} M_{\text{Pl}} e^{-35} = 3.849$ TeV	T-KK-1 (§9.1)
2	$v = M_{\text{lock}}(2/5)^3 = 246.3$ GeV	over-constraint (§4)
3	$\xi_{\text{UV}} = -0.194$ (light quarks)	T-KK-2 (§10.4)
4	$+2.75\, g_s$ (top)	Read-out II bracket (§19)
5	$\Gamma/m = 14.6\%$; $\text{BR}(t\bar{t}) = 94\%$	width table (Appendix I)
6	$\sigma \times \text{BR}(t\bar{t}) = 8.7$ fb (13 TeV); 13.6 fb (14 TeV)	T-KK-2c chain (§31)
7	Z'_B: $0.030 \times (g_B/0.2)^2$ fb (13 TeV)	L-KK-1 menu (§10.10)
8	G_1: 0.045 fb (13 TeV); $\kappa_{gg} = 0.47/K$	T-KK-2c (§10.9)
9	radial tower 3.849/7.047/10.219 TeV	x_n/x_1 (§9.1)
10	kill window $\lesssim 2$ ab^{-1}; decision 2026–2035	F-KK-2B (§31.5)

How to read the master ledger (Table 35). The ledger answers, for each member of the low-lying spectrum, the four questions an analyst asks first: what is it, where is it, how does it decay, and why does it matter to the search. Three reading rules. *First*, every mass in the table is Scenario B — the composite-hardened reading that survived the data — with Scenario A's excluded 2.46-TeV placement retained only in the historical record (§10.4, §31). *Second*, the "Best channels" column encodes the coupling texture, not a preference: the octet decays to $t\bar{t}$ at 94% because the top is the only fermion at the infrared tip (Appendix I), and the same geometry empties the dilepton and diphoton columns that a reader trained on RS benchmarks will look for. *Third*, the two entries an experimentalist should internalize are the first and the last: the octet row is the entirety of the near-term falsifier (F-KK-2B), and the angular-tower row is a prediction of *absence* — the row is in the table precisely because finding anything there would falsify the non-spatial reading (P-KK-7).

How to read the falsification hierarchy (Table 36). The hierarchy orders null results by what they would actually kill, and its purpose is to prevent the two symmetric misreadings of this program: premature celebration ("a dilepton null already tests GG-6" — it does not; the framework predicts that null) and premature burial ("the octet band is crossed, so the theory survives on technicalities" — it does not; level four is written to be terminal). The levels ascend in decisiveness because they ascend in *coupling-texture coverage*: a single clean channel constrains one overlap; the top-plus-dijet-plus-*hh* combination constrains the pattern; the full cluster-aware combination across the geometry-allowed corridor is the kill (F-KK-2B's operational form). The fifth row is the mirror case that a reader may miss: a *discovery* of the wrong morphology — a single clean RS-like pole with none of the cluster correlates — would challenge this paper as sharply as any null, because the fingerprint, not the bump, is the claim.

How to read the calendar (Table 37). The calendar compresses Part VIII to four rows, and its columns carry the whole logic of that Part: the per-experiment pair (frozen/improved) brackets the background-systematics future, the combined column is the realistic ATLAS+CMS decision variable with its correlation bracket, and the milestone column converts luminosity to consequence. Two warnings for the casual reader. The differences *within* rows dwarf the differences *between* rows — moving from frozen to improved systematics is worth more than a decade of additional running (the wall of §40.3) — so the table is an argument about analysis investment, not about patience. And the first row is not a projection: end-2026 is data already on tape, which is why the paper's operational request (the dedicated reanalysis) carries the urgency it does.

Table 35: Master phenomenology ledger for the low-lying GG-6 collider sector (all entries at their derived values; Scenario B throughout, Scenario A being excluded). Scenario B throughout, because Scenario A is excluded by the public limits of Part VII, and the table says so in its own first line rather than leaving the reader to work out which branch is live. Every entry is at its derived value, so the ledger can be checked row by row against the sections that derive it.

Object	Spin	Mass	Best channels	Role
Color-octet KK gluon g^*	1	$M_{\rm lock} =$ 3.849 TeV (T-KK-2)	$t\bar{t}$ (BR 0.94), dijets	The un-suppressible member: $\sigma \times \mathrm{BR}(t\bar{t}) = 8.7$ fb (13 TeV); carries the LHC falsifier F-KK-2B.
Baryonic vector Z'_B	1	$= m_{g^*}$ exactly (same BVP)	$t\bar{t}$, dijets; *no leptons*	Exactly leptophobic (theorem); 0.2–6.8% of the cluster $t\bar{t}$ sum across the L-KK-1 menu.
Graviton anchor G_1	2	$M_{\rm lock}$ exactly (T-KK-1)	$t\bar{t} : W_L W_L : Z_L Z_L : hh =$ 0.69:0.15:0.08:0.08	Names the geometry; $\kappa_{gg} = 0.47/K$ (T-KK-2c) makes it FCC-hh physics (0.13 pb at 85 TeV).
Second radial family G_2	2	$1.831\, M_{\rm lock} =$ 7.047 TeV (Bessel ratio)	$t\bar{t}$, dijets	The overtone that confirms the cavity; FCC-hh spectroscopy (Part IX).
Angular θ tower	—	$\sim E_{\rm GUT}$	none	Predicted *absent* at the TeV scale: no angular satellites (a falsifiable absence).

Table 36: Falsification hierarchy for the collider program. The hierarchy distinguishes between a single-channel null result, which is usually not enough, and a combined geometry-aware null result, which is decisive. The distinction the hierarchy draws is the one collider nulls are most often read wrongly against. A single channel returning nothing constrains one production mode at one set of couplings; the cluster reading survives that and should. What it does not survive is a geometry-aware combination across the members, and the table states the threshold at which that becomes decisive.

Level	Meaning
Single-channel exclusion	Constrains branching or overlap structure; not yet decisive because GG-6 is non-universal.
Two-channel null (e.g. dilepton + diphoton)	Strongly disfavors universal or clean-channel-dominated versions, but still not enough for the full GG-6 benchmark.
Hadronic-top + dijet + Higgs-pair null in the first-cluster window	Severely constrains the geometry-induced coupling pattern.
Full cluster-aware combined-channel exclusion across the geometry-allowed coupling corridor	Collider-level falsification of the current lock-scale implementation of GG-6.
Discovery of a clean single RS-like pole at an incompatible mass pattern	Would challenge the specific GG-6 spectral grammar even if it supports some extra-dimensional physics more generally.

Table 37: The discovery calendar (recap of Part VIII; calibrated significances, dedicated matched-template search). A recap rather than a new result: every significance here is computed in the simulation Part and is reproduced only so the schedule can be read in one place. The dates are machine schedules and not this paper's estimates, and the significances are for a dedicated matched-template search rather than for a reinterpretation of a published limit.

Snapshot	1 expt (froz./impr.)	ATLAS+CMS (impr.)	Milestone
end-2026	$2.7/3.7\sigma$	$4.6–5.2\sigma$	evidence on tape; 5σ a live possibility
end-2031	$3.1/4.5\sigma$	$5.4–6.3\sigma$	single-expt 5σ luminosity on tape 2032
end-2033	$3.3/5.4\sigma$	$6.1–7.6\sigma$	discovery or exclusion decided in practice
end-2041	$3.5/6.5\sigma$	$6.9–9.1\sigma$	property program saturated

E What is the relation between the QCD gluon and the color-octet g^*?

What is the relation between the gluon of Standard-Model QCD and the color-octet vector g^* of GG-6-Theory?

It is a question with a clean answer, because in GG-6 the relation is as tight as it could possibly be: **the gluon and g^* are the same field** — two standing waves of the single bulk $SU(3)_C$ gauge field, the ground note and the first overtone of one instrument. Everything below is implicit in Theorem T-KK-2 (§10.4); this appendix assembles it in one place.

The construction. The bulk $\mathrm{SU}(3)_C$ field lives on the six-dimensional stage and, like every bulk field, decomposes into modes along the scale throat (Appendix J). The boundary conditions are fixed by an observed fact — their zero-mode sector must *be* the massless gauge fields we see — and that zero mode, with its exactly flat profile along r, **is the QCD gluon**. The first radial excitation of the same field — one node of standing-wave structure along the RG direction, profile concentrated at the infrared tip, mass set by the Bessel condition — **is** g^*. Same color octet, same spin, same Lorentz structure, one field.

Why one is massless and the other weighs 3.85 TeV. The zero mode is flat, so it costs no gradient energy along the throat; its masslessness is exact, protected by the unbroken boundary gauge invariance. The excited mode pays standing-wave energy in the scale direction, redshifted to the infrared endpoint — mass from geometry, with *no Higgs mechanism and no new coupling constant anywhere.* Every g^* coupling is g_s times a derived overlap number.

Why they behave so differently despite being one field. The profiles. The gluon's flat profile overlaps every quark *identically* — which means QCD's flavor blindness, an axiom in the Standard Model, is here a *theorem about the ground state*: universality is a property of flatness. The overtone has no such protection: its IR-localized profile overlaps the UV-localized light quarks at $-0.194\,g_s$ and the IR-localized top at $+2.75\,g_s$ — the entire top-philic phenomenology of this paper is nothing but the statement that flavor blindness dies with the first node.

The consequence the paper exploits throughout. Because g^* carries exactly the gluon's quantum numbers, it enters $q\bar{q} \to t\bar{t}$ in the *same color channel* as the s-channel gluon — which is why the two amplitudes add coherently (Part VI) and why the interference-sign theorem exists at all. The color-singlet Z'_B, by contrast, cannot interfere with QCD; that yes/no dichotomy inside one peak is a direct fingerprint of the "same field" relation (O-KK-28).

In the composite (Scenario B) reading. The statement transposes but survives: the gluon is the elementary boundary field, and g^* is the composite spin-1 resonance of the boundary theory carrying the gluon's quantum numbers — precisely the relation of the ρ meson to the photon in hadronic physics (the same analogy carried, for spin 2, in Appendix G). The ρ is not a "new photon"; it is what the strong dynamics does in the photon's channel. Likewise g^* is not a new gauge boson — no extended color group, no axigluon-style broken $\mathrm{SU}(3) \times \mathrm{SU}(3)$, no new charge — it is what the throat does in the gluon's channel.

One instructive contrast within the cluster itself. g^* has a massless sibling (the gluon), because $\mathrm{SU}(3)_C$ has a zero mode; Z'_B does *not* — the $\mathrm{U}(1)_B$ vector is Stückelberg-massive with no massless partner, which is why baryon number generates no long-range force and the proton is stable (Arisaka, 2026, A9-SM). So the cluster's two vectors, degenerate in mass, have opposite genealogies: one is an excited copy of a field known since 1979; the other is the lowest note of a field the Standard Model never had.

If a seminar audience asks "is g^* a heavy gluon?", the precise answer is: *it is the gluon — heard one octave up the scale dimension.*

F What is the origin of the baryonic vector Z'_B?

What is the origin of the baryonic vector boson Z'_B? Presumably it is the outcome of local $\mathrm{U}(1)_B$ gauge invariance — but why must it sit at 3.85 TeV, degenerate with g^* and the graviton anchor G_1? And is it related to the Z boson of the weak interaction, or to the photon of $\mathrm{U}(1)_{\mathrm{EM}}$?

The guess is exactly right, and the interesting physics is in *why* the gauging is forced, and why the mass has nothing to do with the field.

Baryon number must be gauged — it is not a choice. In the Standard Model, baryon number is an *accidental* global symmetry: nothing enforces it, and it is famously fragile against anything one adds. In GG-6 the situation is inverted. The BI-6 anomaly ledger of GG-A3-BI6 (Arisaka, 2026, A3-BI6) closes the six-dimensional tensor/gauge system only if the bulk gauge set includes $\mathrm{U}(1)_B$ alongside color and the electroweak factors — the anomaly inflow of the 6D theory does not balance without it. So Z'_B is indeed the gauge boson of local $\mathrm{U}(1)_B$ invariance, but gauged out of *consistency*, not model-builder's taste. And the same ledger is why it earns its keep in the electroweak over-constraint: the trace ratio 2/5 that forces $v = M_{\mathrm{lock}}(2/5)^3 \approx 246$ GeV (§4) counts $\mathrm{U}(1)_B$ on equal footing with $\mathrm{U}(1)_{Qi}$. Remove the baryonic factor, and the theory's proudest number breaks.

The obligatory technical question: the mixed anomalies. A field theorist will interject here: gauging baryon number in the Standard Model *alone* is famously obstructed, because the mixed triangles do not vanish — $\mathcal{A}[B\text{–SU}(2)_L^2] = +\frac{3}{2}$ and $\mathcal{A}[B\text{–U}(1)_Y^2] = -\frac{3}{2}$ over three generations (the gravitational, cubic, and B^2Y triangles are identically zero, since B is vector-like on the quarks) — and the conventional price is a set of new chiral "anomalon" fermions. GG-6 pays no such price: the boundary spectrum descends from unsplit SU(5) multiplets whose anomaly polynomial factorizes to rank one, and the two surviving triangles are canceled by the inflow of the single Green–Schwarz counterterm $\int B_2 \wedge X_4$ — the mechanism proved in GG-A3-BI6 and spelled out there in four-dimensional language (Arisaka, 2026, A3-BI6). The four-dimensional remnant of B_2 is precisely the Stückelberg axion of the previous paragraph: its shift cancels the triangles *and* gives Z'_B its mass — one coupling, two consequences — and no anomalon exists anywhere in the spectrum (O-KK-47).

Why 3.85 TeV: the mass belongs to the geometry, not the field. This is the cluster theorem's central intuition (§9.1). The scale throat is one cavity — a pipe of integer length $K = 35$. Every infrared-localized standing wave in that pipe, whatever field it belongs to and whatever its spin, sits at the pipe's pitch, $M_{\mathrm{lock}} = \frac{1}{2} M_{\mathrm{Pl}} e^{-35}$. The spin-2 sector puts its first note there (G_1); $\mathrm{SU}(3)_C$ puts its first *overtone* there (g^*, Appendix E); and $\mathrm{U}(1)_B$ puts its first note there (Z'_B). The degeneracy is not a coincidence to be explained; it is the statement that one geometry owns all three masses. Hadronic physics supplies the perfect precedent: the $\rho(770)$ and $\omega(782)$ are near-degenerate not because vector mesons conspire, but because Λ_{QCD} sets both. In the composite (Scenario B) reading this is literal — all three cluster members are resonances of one boundary theory at its single infrared scale.

The crucial asymmetry with g^*: no massless sibling. The gluon exists because the boundary conditions of $\mathrm{SU}(3)_C$ preserve a flat zero mode (Appendix E). For $\mathrm{U}(1)_B$ the boundary conditions do the opposite: the would-be zero mode is removed — the vector is Stückelberg-massive, the abelian mechanism that gives a U(1) gauge field a mass without breaking the symmetry and without any Higgs (Arisaka, 2026, A9-SM). This is not a technicality; it is forced by observation. A massless "baryonic photon" would mediate a long-range force coupled to baryon number, and Eötvös-type equivalence-principle experiments exclude such a force to fantastic precision. So the theory *must* deliver $\mathrm{U}(1)_B$ gauged (the anomaly ledger demands it) and yet with no long-range force (the torsion balances demand that) — and the Stückelberg construction is precisely the reconciliation. The price is one massive vector; and since the lightest note the pipe allows it is the pipe's own, that vector lands at 3.849 TeV. As a bonus, the surviving gauge invariance still enforces baryon-number conservation: the proton stays stable as a matter of *law*, where the Standard Model could only promise it by accident.

Relation to the Z: essentially none. The Z is the neutral member of $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$; its mass comes from electroweak symmetry breaking — in GG-6, the Wilson-line mechanism on S^1_θ (Hosotani mechanism; Hosotani, 1983) — and it couples to weak isospin and hypercharge, hence generously to leptons. Z'_B is *not* a heavier Z: different group, different mass mechanism (geometry versus symmetry breaking), and it couples to baryon number only — 1/3 per quark, universally, and *exactly zero* to leptons. That is why the leptophobia is a prediction rather than a tuning: $\mathrm{BR}(\ell\ell) = 0$ identically, which is also what has kept it invisible to the powerful dilepton searches (§30). The only genuine link to the Z is upstream: the same 2/5 trace arithmetic that involves $\mathrm{U}(1)_B$ fixes v, and v fixes M_Z.

Relation to the photon: the closest cousin, with the opposite boundary fate. Both are abelian U(1) vectors, so as fields they are siblings. But $\mathrm{U}(1)_{\mathrm{EM}}$ keeps its zero mode — unbroken, massless, long-range — while $\mathrm{U}(1)_B$ has its zero mode deleted. One way to say it: the photon is a fundamental the pipe lets ring; Z'_B is a field whose fundamental key is blocked, so the first sound it can make *is* the pipe's note. (Its coupling strength, the discrete g_B menu, is the subject of L-KK-1.)

The three cluster members therefore carry three distinct genealogies converging on one mass: g^* is the excited copy of a field known since 1979; G_1 is the excited copy of the metric sector (Appendix G); and Z'_B is the lowest note of a field the Standard Model never had — the only genuinely *new* gauge boson in the cluster, whose existence is an anomaly-ledger necessity and whose silence in dileptons is its signature.

G What is the "graviton" of GG-Theory?

The word "graviton" is doing two different jobs in the literature, and GG-6 adds a third, genuinely new reading. Because the question "is your G_1 *the* graviton?" is guaranteed to be asked, this appendix separates the meanings once, cleanly.

Two meanings already in circulation. The graviton of quantum gravity — string theory's graviton — is the *massless*, helicity-2 quantum of the quantized four-dimensional gravitational field. Its existence presupposes that gravity itself is a quantum field, and the notorious difficulties (non-renormalizability, ultraviolet completion) live there. But the "KK graviton" of collider phenomenology has never been that object: a massive spin-2 resonance arising from the mode decomposition of a higher-dimensional metric is an *ordinary particle*, and quantizing a massive spin-2 field on a fixed background is standard field theory, no harder in principle than the W boson. Every ATLAS and CMS "graviton search" ever performed has searched for the second kind, not the first. Discovering G_1 would therefore not certify that gravity is quantized in the string-theoretic sense; it would certify that the *metric sector has a massive discrete mode*.

What G_1 is, specifically. G_1 is the lowest radial standing wave of the spin-2 sector of the bulk — and in the renormalization-group reading this paper insists on, it is literally a *composite spin-2 resonance of the boundary theory at its infrared scale*, the analogue of a spin-2 hadron. Its quantum status is exactly as unproblematic as the ρ meson's. It is called "graviton" by *pedigree*, not by job description: it descends from the bulk metric — the $\mathrm{Spin}(1,5)$ block of the composite connection that Arisaka (2026, A2-GG6) forces — which is why its spin-2-ness, its universal coupling structure to $T^{\mu\nu}$, and its $W_L W_L$-rich branching are inherited rather than chosen. That pedigree is why the paper calls it "the state that names the geometry."

The OGGEP inversion. GGEP (AX-2) does *not* quantize gravity; it makes a classical, structural statement — geometry and gauge are connections of the same species, locked by BI-6. Quantization enters at AX-3: OGGEP makes the bulk-to-boundary projection *be* observation, and Arisaka (2026, A8-QM) recovers quantum mechanics *from* that projection rather than postulating it. GG-6 thereby inverts the fifty-year program: one does not quantize geometry; *quantization is what geometry does when it is projected.* The mechanism is concrete — the gauge phase of charged matter is the holonomy of the θ fiber, $\hbar$ is the compact circle's unit ($M = 1$ in the ledger), and discreteness (of charges, of the radial tower in n) is the character and standing-wave structure of the two non-spatial dimensions. In this sense the graviton and the photon are quantized *in the same way, by the same projection*: both are modes of one composite connection, and neither requires a separate act of quantization.

The honest fence. What GG-6 does *not* thereby deliver is the full quantum dynamics of the massless four-dimensional graviton — graviton–graviton scattering near the lock scale, the deep-ultraviolet behavior. That is exactly what Arisaka (2026, A9-SM) registers as its open structural item O_{14}: the composite connection resolves the structural unification, but the dynamical quantum-gravity calculation is not performed anywhere in the cohort. The clean summary: *discovering G_1 at* 3.849 *TeV would be evidence for the six-dimensional geometric origin of the world — not a detection of "the graviton of quantum gravity" — and it neither requires nor proves that massless 4D gravity is quantized.* What it would prove is stronger in one way (it validates the geometry that makes quantization itself an output) and weaker in another (the string-theoretic question remains deferred, and honestly so).

The massive spin-2 field, technically. For completeness, the field-theory status of a massive spin-2 particle deserves one careful paragraph, because "spin-2 is dangerous" folklore conflates three separate issues. A *free* massive spin-2 field has a perfectly healthy description (Fierz–Pauli); the notorious pathologies (the vDVZ discontinuity, the Boulware–Deser ghost) afflict attempts to make a massive graviton *replace* the massless one as the carrier of long-range gravity. G_1 does no such thing: the massless zero mode remains the graviton of gravity, and G_1 is an additional heavy state whose interactions are cut off at the throat scale — exactly the situation of every RS phenomenology paper, where the KK graviton is treated by ordinary perturbation theory without apology. *(Physics intuition: nobody worries that the ρ meson, a massive spin-1 state, endangers the masslessness of the photon; G_1 stands to the graviton as the ρ stands to the photon — a heavy resonance in the same channel, not a rival.)* In the composite reading this is even more manifest: a spin-2 resonance of a strongly coupled boundary theory needs no fundamental spin-2 field at all, any more than the $f_2(1270)$ does.

What discovery would and would not settle, item by item. It is worth tabulating the epistemics, because they are asymmetric. Discovery of G_1 with the derived pattern *would* establish: a massive spin-2 state exists (from the angular analysis of Part VI); its couplings follow stress-energy weighted by localization (from the $W_L W_L$-rich pattern); and the six-dimensional geometric reading has produced a correct, parameter-free postdiction chain. It would *not* establish: that 4D gravity is quantized (O_{14} remains open); that the bulk is the unique UV completion; or that the composite versus weakly-coupled readings of the throat are distinguished (that requires the line-shape and recurrence program of Part IX). Conversely, *non*-discovery at FCC-hh, after an octet discovery at the LHC, would falsify the cluster reading outright — the division-of-labor clause of F-KK-1. *(Physics intuition: the graviton is the cluster's birth certificate, not its driver's license; the octet proves something lives at 3.85 TeV, the graviton proves what family it belongs to.)*

H Why does g^* have the largest production rate — and why is it so broad?

Why does g^* have the largest production cross section of the three cluster members — and why is its width so large, $\sim$15% of its mass?

The two questions have one answer at their root — the coupling texture set by where each field's wavefunction lives along the scale throat — and this appendix tells it in doors, keys, and exits. (Appendix I then opens the hood and derives the same facts from the overlap integrals.)

The production race: g^* gets three multiplicative advantages. *First, the door it enters through.* At 3.849 TeV on a 13-TeV machine, each incoming parton must carry $x \approx 0.3$ of its proton's momentum — and at such high x the proton is essentially three valence quarks; the gluon density has collapsed. The largest initial-state luminosity available is therefore $q\bar{q}$ annihilation, and g^* is built to use exactly that door: it couples to the light-quark color current with strength $\xi_{\mathrm{UV}}\, g_s = -0.194\, g_s$. The overlap 0.194 looks modest, but it multiplies the *strong* coupling — the largest coupling in nature's book. *Second, color counting.* g^* is a color octet entering the same channel as the gluon: eight resonance states, produced with the full color trace of a strong-interaction vertex, plus coherent interference with the QCD continuum (Part VI). Z'_B is a color singlet — one state, no color enhancement, no interference. *Third, the competition's handicaps.* Z'_B enters the same $q\bar{q}$ door but with a feeble key: $g_B \approx 0.2$ (the derived menu value of L-KK-1) times baryon charge 1/3 per quark — an effective coupling ≈ 0.067 against g^*'s ≈ 0.2, an order of magnitude down in rate before color and branching factors, which stretch it to the canonical ratio 8.7 fb versus 0.030 fb (§31). The graviton is worse off twice over: its natural door is gg fusion, exactly the luminosity that has died away at $x \approx 0.3$; and its gluon-pair coupling carries the derived suppression of Corollary T-KK-2c, $\kappa_{gg} = 0.47/K$. The rate goes as the coupling *squared* — one integer, two powers: $(1/35)^2 \approx 1/1225$. That is why the graviton anchor sits at 0.045 fb, invisible to the HL-LHC and waiting for a higher-energy ring (Part IX), while its color-octet sibling at the *same mass* is a per-experiment discovery candidate (Part VIII).

The width: the exit door is the top, at composite strength. A resonance's width is set by the couplings through which it can *leave*, and g^*'s dominant exit is $t\bar{t}$ with coupling $+2.75\, g_s$ — the top quark lives at the infrared tip, directly on top of the g^* wavefunction, so the overlap is enormous. Squaring a coupling of that size and dividing by the usual 8π of phase space gives $\Gamma/m \approx 14.6\%$, with 94% of it in $t\bar{t}$ — the top being the *only* fermion at the tip (Appendix I). Hadronic physics supplies the calibration: the $\rho(770)$, a composite vector decaying through the strong interaction, has $\Gamma/m \approx 19\%$. In the Scenario B reading g^* *is* a ρ-like composite of the boundary theory (Appendix E), so a width of this order is not merely tolerable — it is what honesty about the composite reading demands. By the same logic G_1 is narrower (6.3%: universal but weaker stress-tensor couplings), and Z'_B is narrow (g_B is small).

The point worth savoring. Production rate, width, $t\bar{t}$ dominance, and the interference sign are all read off the *same two numbers* — 0.194 in, 2.75 out — and those numbers are overlap integrals of fixed wavefunctions, not inputs. And the broadness is a fingerprint, not a nuisance: it is why published narrow-resonance limits under-constrain this state by a factor $\sim$3 (§30), why the interference-aware template of Part VIII gains so much sensitivity, and why finding a *narrow* peak at 3.85 TeV would actually falsify this implementation.

I One integer, two powers — the rate hierarchy and the $t\bar{t}$ dominance

How can one integer — the same $K = 35$ that sets the cluster's mass — also decide how often each member is produced and what each decays to?

Appendix H told that story in doors, keys, and exits; this appendix opens the hood. Two numerical facts organize the paper's entire phenomenology: the graviton is produced $\sim$300 times more weakly than the octet, and both decay overwhelmingly to top quarks. Both facts have one origin — wavefunction geometry along the scale throat — and by the end of this appendix the reader can check every number with a pocket calculator.

First, picture the corridor. Everything below is bookkeeping of who stands where. The throat is a corridor of length $K = 35$ in the scale direction. The light quarks live at the ultraviolet wall ($r = 0$); the top quark and the entire lock cluster live at the infrared tip ($r = K$); the gluon's zero mode is flat — spread evenly along the whole corridor. One fact of normalization drives all the arithmetic: a mode spread over a corridor of length K has generic amplitude $\sim 1/\sqrt{K}$ at any one point. Standing *at* a point costs nothing; being spread *over* the corridor costs $1/\sqrt{K}$ per field.

The production hierarchy: one power of K versus two. Production at a hadron collider is an overlap question: *how much do the initial partons' profiles intersect the resonance's home?* The octet's parents are valence quarks, standing at the UV wall, where the g^* mode retains a finite foothold. The production vertex is a *single* boundary amplitude — one factor of $1/\sqrt{K}$, with no cancellation, because the quarks' endpoint localization breaks any orthogonality protection maximally. Squared: $\xi^2_{\rm UV} \approx 1/K$. One power.

The graviton pays twice, and then pays again. Its parents are gluons, whose flat mode has amplitude $1/\sqrt{K}$ *everywhere* — so two initial gluons already cost $1/K$ in the amplitude. Then the overlap integral itself, $\int(\text{flat})^2 \times (\text{IR-localized } J_2)$, *nearly cancels*: a democratic source barely projects onto a receiver huddled at the tip, and only the last half e-fold of the corridor contributes. That surviving remnant is the 0.47 of Corollary T-KK-2c: $\kappa_{gg} = 2(1 + J_2(x_1))/(Kx_1^2 J_2(x_1)) \approx 0.47/K$, hence rate $\propto 0.22/K^2$. Two powers. The ratio of rates is one extra power of the *same* integer that sets the mass ($\times$35), times the near-orthogonality remnant ($\times \sim$6): the factor $\sim$300 between 8.7 fb and 0.045 fb. *The graviton is hard to make for the same reason its mass is impossible to move.*

The same flat-mode geometry works on the decay side too: the photon and the leptons are flat or UV-spread, so the cluster barely speaks to them. That is why G_1 is lepton-poor, why the cluster is leptophobic in practice, and why the published diphoton null was a retro-diction rather than a tension.

The decay hierarchy: heaviness and cluster-affinity are the same property. Why tops? Not phase space — at 3.85 TeV, $\beta_t = 0.996$: every channel is wide open. The answer is localization, used twice. A fermion's *mass* is set by how much its profile overlaps the Higgs sector at the infrared endpoint; the top is the only fermion with an unsuppressed profile there — that is *why* $y_t \sim 1$. But the lock cluster is concentrated at the *same* endpoint, and decay widths are the same overlap integrals as masses. So the state that is heavy *because* it lives at the tip is automatically the state the cluster talks to most loudly: tell me a fermion's mass, and I will tell you how strongly the cluster couples to it.

Numerically, the octet's width ratio per flavor is $(g_V^{t2} + g_A^{t2})/\xi^2_{\rm UV} \approx 10.6/0.038 \approx 280$, giving the partials $528 : 26 : 7$ GeV and the falsifiable fingerprint $t\bar{t} : b\bar{b} : jj \approx 70 : 3.5 : 1$ (the intermediate $b\bar{b}$ is itself diagnostic: b_L shares the top's doublet and inherits partial IR-ness). The graviton's version couples to stress-energy weighted by the same overlaps, preferring the IR-localized top *and* the Higgs sector — by the equivalence theorem, W_L and Z_L *are* the Higgs sector in disguise — hence $t\bar{t} : W_L W_L : Z_L Z_L : hh = 0.69 : 0.15 : 0.08 : 0.08$.

The strategic consequence. The LHC's cleanest channels are this signal's quietest, and its loudest channel is among the hardest experimentally — which is why the published clean-channel exclusions never touched the prediction, why Part VI lives in top-pair anatomy, and why the discovery analysis of Part VIII lives in the merged one-lepton $t\bar{t}$ topology. Bulk-RS models share the qualitative top-philic story; the difference is that here the localization band $r_{t_R} \in [4, 5]$ is tied to the same geometry that produces the top mass itself, not dialed. The throat is a mass-ordering machine, and the cluster sits at the heavy end of the ordering.

The overlap integrals, written out. Now the hood comes off; this is the appendix's one technical passage, and a reader content with the corridor picture may skim it without loss. The text above says "overlap"; a course should display one. With the canonical scale profile, the octet's lowest mode function is $\chi_1(r) \propto e^r J_1\big(x_1 e^{r-K}\big)$ (normalized over $[0, K]$), and the coupling to a fermion localized at radial position r_Ψ is $g \cdot \chi_1(r_\Psi)/\overline{\chi}$, with $\overline{\chi}$ the normalization. Evaluating at the UV wall gives the small boundary value $\chi_1(0)$: the mode is exponentially concentrated at $r = K$, and its residue at $r = 0$ is what the light quarks see — carried through, this is $\xi_{\rm UV} = -0.194$, the sign coming from the node structure of J_1 between the wall and the tip. For the graviton–gluon coupling the integrand is qualitatively different: the gluon zero mode is r-independent after canonical normalization ($f_g^2 = 1/K$), so

$$\kappa_{gg} \;\propto\; \frac{1}{K}\int_0^K dr\; e^{-2(r-K)}\,\frac{J_2\big(x_1 e^{r-K}\big)}{J_2(x_1)} \;=\; \frac{2\,[\,1 + J_2(x_1)\,]}{K\,x_1^2\,J_2(x_1)} \;\approx\; \frac{0.47}{K}, \tag{88}$$

where the closed form follows from the Bessel identity $\int z^2 J_2\,dz \to z^2 J_1 + 2zJ_0 - 2\int J_0$ machinery collapsed at the endpoints (the reader should verify the elementary antiderivative $\frac{d}{dz}[z^2 J_2(z)] \cdot z^{-1}$ route reproduces it). The lesson is visible in the integrand: the exponential weight concentrates all support within the last half e-fold below $r = K$, which is the geometric content of the "0.47". *(Physics intuition: integrate a whisper spread over the whole corridor against a listener standing in the last doorway — only the final syllable is heard.)*

The width table, assembled. The same couplings assemble the octet's widths with no further input: $\Gamma(g^* \to q\bar{q}) = \frac{\alpha_s}{6}\,\xi_{\rm UV}^2\,m$ per light flavor, and $\Gamma(g^* \to t\bar{t}) = \frac{\alpha_s}{6}\,m\,\beta_t\big[g_V^{t2}(1+\frac{1}{2}\rho) + g_A^{t2}\beta_t^2\big]/g_s^2$ with $\rho = 4m_t^2/m^2$; plugging the derived numbers gives 528 : 26 : 7 GeV for $t\bar{t} : b\bar{b} : jj$ and $\Gamma/m = 14.6\%$ — the reader can confirm each entry with a pocket calculator, which is precisely the point: nothing between the geometry and the width table has a dial in it. *(Physics intuition: the width IS the coupling table read out loud; a broad resonance is broad here because the top's doorway at the IR tip is wide open, not because anything was tuned.)*

J From the 6D action to the KK spectrum — a course

J.1 Step 1: the textbook derivation, from scratch

The flat circle first. Take a massless 5D scalar on $\mathcal{M}_4 \times S^1$ of radius R, action $S = \frac{1}{2}\int d^4x\,dy\,\partial_M\Phi\,\partial^M\Phi$. Because the circle is compact, expand in its complete Fourier basis, $\Phi(x, y) = \sum_{n=-\infty}^{\infty} \phi_n(x)\,e^{iny/R}/\sqrt{2\pi R}$. Substitute and integrate over y: the cross terms vanish by orthogonality, $\int_0^{2\pi R} dy\,e^{i(n-m)y/R} = 2\pi R\,\delta_{nm}$, leaving

$$S = \sum_n \frac{1}{2}\int d^4x\left[\partial_\mu\phi_n\partial^\mu\phi_n^* - \frac{n^2}{R^2}\,|\phi_n|^2\right]. \tag{89}$$

Each Fourier mode is a 4D particle of mass $m_n = |n|/R$: the Kaluza–Klein tower. *(Physics intuition: the extra dimension is a resonating cavity; the tower is its overtone series, and the mass is the price of wrapping n wavelengths around the circle. Momentum along y, quantized by periodicity, masquerades as 4D mass because $p_M p^M = 0$ splits into $p_\mu p^\mu = -p_y^2$.)* Worked example: $R^{-1} = 1$ TeV gives states at 1, 2, 3 TeV with *equal* spacing — remember this picket fence; the warped case will destroy it.

Now the warped interval. Replace the circle by an interval $y \in [0, \pi r_c]$ with metric $ds^2 = e^{-2ky}\eta_{\mu\nu}dx^\mu dx^\nu + dy^2$ (Randall–Sundrum). For a graviton fluctuation $h_{\mu\nu}$, the transverse-traceless equation of motion separates with $h_{\mu\nu}(x,y) = \sum_n h^{(n)}_{\mu\nu}(x)\,\psi_n(y)$, and the radial problem becomes, in the conformal variable $z = e^{ky}/k$,

$$-z^2\psi_n'' + 3z\,\psi_n' = m_n^2 z^2\,\psi_n \quad\Longrightarrow\quad \psi_n(z) \propto z^2\left[J_2(m_n z) + b_n\,Y_2(m_n z)\right], \tag{90}$$

a Bessel equation (for Bessel-function background see the NIST handbook, Olver et al., 2010). *(Physics intuition: the warp factor acts like a position-dependent refractive index; waves accumulate near the low-index end, so the mode functions pile up at the infrared boundary — the origin of every "IR-localized" statement in this paper.)* Neumann conditions at both ends force $b_n \simeq 0$ (for $kr_c \gg 1$) and quantize the spectrum on the zeros of J_1:

$$m_n \simeq x_n\,k\,e^{-k\pi r_c}, \qquad J_1(x_n) = 0: \quad x_1 = 3.8317,\ x_2 = 7.0156,\ x_3 = 10.1735. \tag{91}$$

Two features deserve a pencil check. The spacing is *not* the picket fence of (89): the ratios 1 : 1.831 : 2.655 are pure Bessel geometry. And the overall scale is exponentially redshifted: $e^{-k\pi r_c}$ turns a Planckian input into a TeV output when $k\pi r_c \approx 35$. Worked example: $k = M_{\rm Pl}$, $k\pi r_c = 35$ puts m_1 in the few-TeV range — the reader should verify that reproducing 3.85 TeV requires the normalization discussed in §9.1.

Couplings from overlaps. A boundary or bulk fermion couples to mode n with strength $g_n = g\int dy\,\sqrt{-g}\,f_\Psi(y)^2\,\psi_n(y)/(\text{norms})$: *coupling is geometry* — the overlap of where the fermion lives with where the mode lives. This single line is the origin of $\xi_{\rm UV}$, of r_{t_R}, of κ_{gg}, and of every rate in this paper; Appendix O works its consequences.

J.2 Step 2: what the textbook cannot fix

The construction above contains three blanks that no amount of textbook diligence fills. *The radius is an input*: nothing determines R or kr_c, so the tower can sit anywhere, and forty years of KK phenomenology has been a scan over this ignorance. *The length must be stabilized*: r_c is a dynamical field (the radion), and without an added mechanism (Goldberger & Wise, 1999) the interval collapses or runs away — an extra sector invented solely to hold the stage still. *The normalization of the first pole is convention-laden*: what multiplies x_1 depends on boundary details a search cannot check. Every one of these blanks is a symptom of the same cause: the fifth dimension was postulated, so its size and dynamics are free.

J.3 Step 3: the GG-6 derivation, and the worked number

GG-6 inherits the *mathematics* of (90)–(91) — the canonical metric of Arisaka (2026, A2-GG6) (Theorem 2) has the warped form with $A(r) = r$ — and closes all three blanks at once, because the dimension is derived rather than postulated (§5.5). The length is the *integer* $K = 35$ from the OGN-5 closure (blank one: no scan; blank two: no radion, a count cannot fluctuate); the scale is anchored by the closed form of Theorem T-KK-1 (blank three); and the reduction formulas collected below apply with $A(r) = r$, $R_\theta = 1$ understood. The worked number, end to end: $M_{\rm lock} = \frac{1}{2}M_{\rm Pl}e^{-35} = \frac{1}{2} \times 1.2209 \times 10^{19}\ \text{GeV} \times 6.3051 \times 10^{-16} = 3.849$ TeV, and the tower $m_n = (x_n/x_1)M_{\rm lock} = 3.849,\ 7.047,\ 10.219$ TeV — the organ pipe of §9.1, now derived by the reader.

J.4 Reference: the reduction formulas

For completeness we collect the minimal reduction formulas used implicitly throughout the paper. As in the mode-structure section of the main text, a general scale profile $A(r)$ and circle radius R_θ are retained purely as bookkeeping symbols, so the formulas can be compared line-by-line with the warped literature; the canonical values $A(r) = r$, $R_\theta = 1$ of Arisaka (2026, A2-GG6) (Theorem 2) are understood throughout, and nothing here is free metric data. For the scale-profile metric

$$\mathrm{d}s_6^2 = e^{-2A(r)}g_{\mu\nu}(x)\mathrm{d}x^\mu\mathrm{d}x^\nu + \mathrm{d}r^2 + R_\theta^2\mathrm{d}\theta^2, \tag{92}$$

one has

$$\sqrt{-g_6} = e^{-4A(r)}R_\theta\sqrt{-g_4}. \tag{93}$$

The Einstein–Hilbert term reduces to

$$M_{\rm Pl}^2 = M_6^4\, 2\pi R_\theta \int \mathrm{d}r\, e^{-2A(r)}, \tag{94}$$

while gauge zero modes satisfy

$$\frac{1}{g_4^2} = \frac{2\pi R_\theta}{g_6^2}\int \mathrm{d}r\,|f_0(r)|^2. \tag{95}$$

The absence of extra scale-profile factors in the latter relation is the standard cancellation between the determinant and inverse metrics. Fermion zero-mode normalization carries an extra e^{-3A} weight, and Yukawa or resonance couplings inherit profile-overlap integrals. Those basic formulas are the origin of the non-universal coupling pattern discussed in the main text.

K From parton distributions to cross sections — a course

K.1 Step 1: the factorized cross section, from scratch

Factorization. A proton is a beam of partons. The master formula of hadron-collider physics (Ellis, Stirling & Webber, 1996; Halzen & Martin, 1984) factorizes any hard process into universal parton distribution functions (PDFs) $f_i(x, \mu_F)$ — the probability density for parton i to carry momentum fraction x, measured once and for all in deep-inelastic scattering — and a computable partonic cross section $\hat{\sigma}$:

$$\sigma(pp \to X) = \sum_{i,j}\int_0^1 dx_1\int_0^1 dx_2\ f_i(x_1, \mu_F)\, f_j(x_2, \mu_F)\ \hat{\sigma}_{ij\to X}(\hat{s} = x_1x_2s). \tag{96}$$

(Physics intuition: the protons are grocery bags; the PDFs are the shopping lists; the partonic cross section is what happens when two chosen items collide. Nothing about the bag needs recomputing when the collision changes.)

The parton luminosity. For an s-channel resonance of mass m, only the combination $\tau = \hat{s}/s = m^2/s$ matters, and (96) collapses to a single object worth naming, the *parton luminosity*

$$\frac{d\mathcal{L}_{ij}}{d\tau} = \int_\tau^1 \frac{dx}{x}\, f_i(x)\, f_j\!\left(\frac{\tau}{x}\right), \qquad \sigma = \sum_{ij} \frac{d\mathcal{L}_{ij}}{d\tau} \cdot \hat{\sigma}_{ij}(\hat{s} = m^2) \cdot \tau. \tag{97}$$

(Physics intuition: the luminosity is the number of parton pairs available at the right invariant mass; everything about the machine and the proton is in it, and everything about the new physics is in $\hat{\sigma}$. When this paper says the octet rate grows $\times 330$ from 13 to 85 TeV, it is quoting nothing but the growth of $d\mathcal{L}_{q\bar{q}}/d\tau$ at $m = 3.85$ TeV.)

The narrow-resonance master formula. For a resonance V of spin J, color dimension N_V, produced from partons ij and observed in channel f, integrating the Breit–Wigner gives the textbook result

$$\sigma \times \mathrm{BR}(pp \to V \to f) = \frac{(2J+1)\, N_V}{N_i N_j}\, \frac{4\pi^2}{m^3}\, \Gamma_{ij}\, \mathrm{BR}_f\, \left.\frac{d\mathcal{L}_{ij}}{d\tau}\right|_{\tau = m^2/s} \times (1 + \delta_{ij}), \tag{98}$$

with Γ_{ij} the partial width back into the production channel — the formula every resonance search's sensitivity plot silently uses. *(Physics intuition: a resonance is a funnel; how much flows through is (spin and color counting) $\times$ (how wide the funnel's mouth is, Γ_{ij}) $\times$ (how much beam arrives at the funnel, the luminosity).)*
Worked example (the reader should reproduce this). For g^* at 3.849 TeV at $\sqrt{s} = 13$ TeV: CTEQ6 gives $d\mathcal{L}_{q\bar{q}}/d\tau \approx$ the archived grid value; the couplings of §10.4 give $\Gamma_{q\bar{q}}$ per light flavor and $\mathrm{BR}(t\bar{t}) = 0.94$; assembling (98) with the color factors of a spin-1 octet and multiplying by $K_{q\bar{q}} = 1.3$ returns $\sigma \times \mathrm{BR} = 8.7$ fb — the number of Table 19, to be checked against the full line-shape integration (the archived workbook does both; they agree to a few percent, the accuracy of the narrow-width approximation at $\Gamma/m = 15\%$).

K.2 Step 2: what the chain assumes but cannot fix

Four external inputs, each honestly labeled [E] in the ledger: the PDFs themselves (a measured vintage with tails least constrained exactly where high-mass searches live — the origin of the [7.0–12] fb band); the factorization/renormalization scales (formally arbitrary, practically a residual of truncating perturbation theory); the K-factors (someone must compute NLO for your process; we cite Chivukula et al., 2012; Mathews & Ravindran, 2005); and, in every conventional application, *the couplings in* Γ_{ij} — which textbook phenomenology takes from a model's free parameters, so that the entire chain, however rigorous, ends in a dial.

K.3 Step 3: the GG-6 chain, closed at both ends

GG-6 changes exactly one thing, and it is the last one: Γ_{ij} and BR_f stop being dials. The production coupling $\xi_{\mathrm{UV}} = -0.194$ and the decay table descend from T-KK-2, and $\kappa_{gg} = 0.47/K$ from T-KK-2c, so (98) becomes an arithmetic pipeline with nothing to choose: PDFs in, falsifiable femtobarns out — at 13, 14, and 85 TeV (§50.1), with the PDF band carried, not hidden. The pipeline is archived (parton-luminosity grids and the assembling workbook), and the student who has reproduced the worked example above owns every rate in this paper.

L The statistics of discovery — a worked course

L.1 Step 1: from the Poisson likelihood to Z, with the arithmetic shown

Part VIII (§40) derives the formalism — the likelihood ratio q_0, Wilks' theorem, the Asimov dataset, and the Cowan formula with background uncertainty. This appendix does what a course must and the body could not: it *shows the arithmetic*, bin by bin, for one real case, so the reader can audit every significance in the paper by extending the same table. Take end-2026 (0.466 ab^{-1}, 13 TeV), analysis bins of Table 26:

Table 38: The significance arithmetic worked bin by bin for one dataset, so that every other number in this Part can be audited by extending the same table. Signal, background, fractional background uncertainty and the resulting per-bin significance are shown in the order the formula uses them. Nothing here is new physics; it is the bookkeeping behind the headline figures, printed because a forecast whose arithmetic is not shown cannot be checked.

bin [TeV]	$s = S{+}I$	b	f_b	$\sqrt{b}$	$f_b b$	Z_{bin} (Eq. 82)
3.0–3.5	145	178	0.069	13.3	12.2	6.8
3.5–4.2	133	75	0.132	8.7	9.9	7.4
4.2–4.9	27	19	0.226	4.3	4.3	3.5

Quadrature over the five analysis bins (the two flanking bins add 5.2 and 0.5) gives $Z = 11.9$; dividing by the measured calibration $F_{\text{cal}} = 4.33$ gives the 2.7σ of every table in Part VIII. *(Physics intuition: each bin is a coin-flip experiment against a bank whose ledger is only good to f_b; the third column pair shows the moment the systematic term overtakes Poisson noise — already true in every bin at this luminosity, the "wall" of §40.3.)* The reader should now recompute the 2041 row of Table 26 and watch the frozen-f_b ceiling emerge from the arithmetic: multiplying s and b by seven changes Z_{bin} hardly at all once $f_b b \gg \sqrt{b}$.

L.2 Step 2: what counting cannot see

The table above prices two of the three suppressions of §42; the third — the background's freedom to *reshape* — is invisible to per-bin counting by construction. The two-parameter demonstration of Figure 11 is the minimal honest exhibit, and the measured F_{cal} is the full-machinery price. A student should internalize the moral before running any fit: a smooth broad excess and a background normalization are nearly the same shape, and only correlated, sign-changing structure (the interference pattern) is safe from absorption.

L.3 Step 3: the pre-registered payoff, worked

The look-elsewhere arithmetic of §44.1, once through with numbers: resolution 10% gives $N_{\text{cells}} = 10\ln(5000/400) = 25.3$ independent mass hypotheses; three width templates give $N_{\text{trials}} \approx 76$; a local $p = 2.87 \times 10^{-7}$ (five sigma) becomes $p_{\text{glob}} = 76 \times 2.87 \times 10^{-7} = 2.2 \times 10^{-5}$, i.e. $Z_{\text{glob}} = 4.1$ — the discount a scanned search pays and a pre-registered point does not. Inverting: global five sigma demands local 5.8σ. *(Physics intuition: seventy-six lottery tickets make a win seventy-six times less surprising; the derivation of the mass is what tears up the other seventy-five tickets in advance.)* Combined with the matched-template factor of §44.2, this is the entire statistical content of "pre-registration" — no new mathematics, only the refusal to scan.

L.4 Reference: the significance-scaling workbook

Suppose a channel f has signal efficiency ε_f, signal cross section σ_f, and background yield per unit luminosity b_f. Then

$$S_f = \mathcal{L}\,\sigma_f\,\varepsilon_f, \qquad B_f = \mathcal{L}\,b_f. \tag{99}$$

The simplest background-dominated significance estimator is

$$\mathcal{Z}_f \approx \frac{S_f}{\sqrt{B_f}} = \sqrt{\mathcal{L}}\,\frac{\sigma_f \varepsilon_f}{\sqrt{b_f}}. \tag{100}$$

If several channels are statistically independent, one may combine them approximately by

$$\mathcal{Z}_{\text{comb}} \approx \sqrt{\sum_f \mathcal{Z}_f^2}. \tag{101}$$

The key lesson is qualitative: because GG-6 is intrinsically multi-channel and non-universal, combined significance can grow faster than any single clean channel would suggest. That is why the collider program should be designed as a combined-channel problem from the beginning.

M Benchmark / reinterpretation card for experimentalists

The card below, the template histograms it references, and the full numerical workbook behind them are available at https://arisaka-gg.org/.

Search card for the first GG-6 cluster.

- **Primary target:** spin-2-anchored resonance cluster centered at 3.849 TeV; the discovery rate is carried by the color-octet member (the spin-2 anchor is an FCC-hh measurement at its derived coupling, T-KK-2c).
- **Key channels:** $t\bar{t}$, boosted dijet, hh, diboson, with clean $\ell^+\ell^-$ and $\gamma\gamma$ cross-checks.
- **Preferred analysis style:** cluster-aware rather than single-pole-only.
- **Most diagnostic pattern:** first cluster + second radial family near 7 TeV.
- **Most useful reduced data:** high-mass jet kinematics, top/Higgs tags, b-tag summaries, timing-assisted pile-up suppression, lepton four-vectors, and channel-combination metadata.
- **Rate anchor (model-level, D*):** color-octet $\sigma \times \mathrm{BR}(t\bar{t}) = 8.7\,[7.0\text{–}12]$ fb at 3.85 TeV (13 TeV, LO×K); the Scenario-B cluster is octet-dominated (G_1, at its derived $\kappa_{gg} = 0.47/K$, adds only 0.03 fb and is an FCC-hh target: 205 fb at 100 TeV); the free-field 2.46-TeV scenario is already a factor $\sim$3.6 above the observed broad-template limit; Z'_B: 0.012–0.46 fb across the derived-norm L-KK-1 menu, LHC-invisible on every reading. See Table 19 and Fig. 9.
- **Discovery fingerprint (pre-registered, Part VI):** $(1 + \cos^2\theta^*)$ angle; $P_t = +0.91$; resonant $A_{FB} = 0$; constructive-below/destructive-above interference line shape with $-O(100)$ GeV centroid pull; nothing in $\ell^+\ell^-/\gamma\gamma$.
- **Collider-level kill criterion:** no combined-channel excess and full-corridor exclusion in the first-cluster window after Run-2 + full Run-3 + early-HL-LHC analyses.

M.1 Implementation walkthrough

The card above states *what*; this walkthrough states *how*, in the order an analysis team would proceed. *(1) Template construction.* Build the truth-level $d\sigma/dm_{t\bar{t}}$ from Eq. (65) with the fixed inputs (M_{lock}, $\Gamma/m = 14.6\%$, $\xi_{\text{UV}} = -0.194$, $g_V^t = 2.75 g_s$, $g_A^t = 1.75 g_s$), keeping the Breit–Wigner and interference pieces as *separate* histograms so the fit can float the interference coefficient (sign free — the data must be allowed to disagree with P-KK-4). Fold with the experiment's own $m_{t\bar{t}}$ response. *(2) Normalization strategy.* Anchor the floating $t\bar{t}$ normalization at $m_{t\bar{t}} \lesssim 1$ TeV or in the dilepton channel; the derived contamination there is $\lesssim$ percent (Eq. 87), and §44.5 prices what a naive full-range float costs. *(3) Fit and statistics.* A two-template profile fit (μ_{BW}, μ_{int}) over the merged-SR spectrum; quote both the pre-registered point ($\mu_{\text{BW}} = \mu_{\text{int}} = 1$) and the free fit. *(4) Validation.* Closure on background-only pseudo-data; sign-flip injection to confirm the template distinguishes P-KK-4's pattern; and the two calibration masses of §41 as the cross-check that the simplified machinery of this paper reproduces the collaboration's own limits. *(5) Reporting.* Whatever the outcome, publish the measured ($\mu_{\text{BW}}, \mu_{\text{int}}$) with its covariance: a null kills F-KK-2B; a signal with the wrong sign kills P-KK-4; a signal with the right sign is the discovery of Part VIII. *(Physics intuition: the entire walkthrough is one sentence — fit the shape this paper derived, let the sign float, and anchor the bank where the signal cannot reach.)*

Template T1 — the deliverable in one box. $T_1(m_{t\bar{t}}) = B(m_{t\bar{t}}) \left[1+\mu_{\text{BW}} f_{\text{BW}}(m_{t\bar{t}})+\mu_{\text{int}} f_{\text{int}}(m_{t\bar{t}})\right]$: anchored background B with its normalization fixed at $m_{t\bar{t}} \lesssim 1$ TeV (§44.5); Breit–Wigner component at $m = 3.849$ TeV with $\Gamma/m = 0.146$; interference component with the theorem-fixed sign (enhancement below the pole, deficit above, crossing near 4.9 TeV) — both shapes from Eq. (65). Validity window 2.5–4.9 TeV. Free parameters: the two signal strengths and the background nuisances, nothing else. Fit, and publish ($\mu_{\text{BW}}, \mu_{\text{int}}$) with covariance; interpretation per the walkthrough above.

N Glossary of collider terms

Cluster-aware search
A heavy-resonance search that allows the signal model to contain two or more nearby poles or partially overlapping structures rather than one isolated narrow resonance.

First lock-scale cluster
The first experimentally accessible GG-6 heavy-resonance family, centered near 3.849 TeV.

Geometry-allowed corridor
The range of effective couplings and branching patterns compatible with the current theorem-level GG-6 geometry, as opposed to a freely scanned benchmark set.

Leptophobic admixture
A contribution from a baryonic or non-universal partner state that enhances hadronic channels relative to dileptons.

Partner pattern The correlated appearance of the first cluster, KK gluon, and $U(1)_B$ vector, rather than one isolated signal in one channel.

Reduced stream A trigger or scouting data format in which only compact high-level or Level-1 information is stored, allowing much higher recording rates than full-event storage.

TLA / trigger-level analysis
The ATLAS approach to saving high-level trigger objects rather than full events for selected searches, especially valuable at high rates.

L1 scouting The CMS reduced-format strategy in which Level-1 trigger information from essentially all collisions is recorded to disk for dedicated analyses.

O Promotion roadmap for the open items

The paper's open items (§58), restated as a roadmap: current grade, the calculation or measurement that closes each, and the impact of closure.

Table 39: The open items as a roadmap rather than as a list of regrets. Each carries its current grade, the specific calculation or measurement that would close it, and what closing it would buy; an item with no closing condition would be an admission rather than a work order, and there are none here. The ordering is by how soon each could be done, which is not the same as by how much it matters.

Item	Grade	Closing condition	Impact of closure
L-KK-1: the g_B menu	[M]	The GDOS counting that fixed $\alpha_s(M_Z)$, applied to the $\mathrm{U}(1)_B$ line	Z'_B rates single-valued; fifth-force window sharpened
NLO interference K-factor	[E]	Dedicated NLO computation of the g^*–QCD interference	Low-side template firmed; the one theory input the reanalysis needs first
$n \geq 2$ overlap couplings	[M]	T-KK-2c machinery evaluated at x_2, x_3	G_2/G_3 rates promoted from estimate to derivation (FCC-hh spectroscopy)
Composite recurrence spacing	open	Strong-coupling (dual) spectrum calculation	Vector-recurrence search targets fixed
Detector-level 2D angular study	[E]	Full simulation with the anchored background cocktail	The bounded 10–30% discovery gain banked
Composite-scenario width band	[D*]	Line-shape fit at HL-LHC statistics (self-measuring)	Γ/m becomes a measured cross-check of the overlap band

Closure of any item sharpens the program; closure of none is required for the verdict. The falsifier stands on the theorem-grade spine alone.

Each row above deserves its sentence of context for the student choosing a project. *L-KK-1* is the most self-contained: the GDOS counting that produced $\alpha_s = 1/(\pi e)$ is documented in Arisaka (2026, A9-SM), and its $\mathrm{U}(1)_B$ application is a bounded, well-posed exercise with a four-point answer space. *The NLO interference K-factor* is conventional (if demanding) perturbative QCD — the one entry a standard SM-theory group could adopt tomorrow, with immediate experimental consumers. *The $n \geq 2$ overlaps* reuse Eq. (88) at higher zeros: an afternoon for the diagonal couplings, a paper for the full mixing structure. *The recurrence spacing* is the hard one — genuinely strong-coupling, awaiting either a dual calculation or the data themselves. *The 2D angular study* belongs inside a collaboration (it needs the real background cocktail), and *the width band* closes itself with data. A reader looking for where to help should start at the top of the table; a reader looking to falsify should note that none of these items shields the verdict.

The course is complete: a reader who has worked Appendices J through L owns every derivation chain in the paper — spectrum, cross section, and significance. One service remains, and it is the one a claim of this size owes most: the objections, which the next Part collects and answers, fifty-eight of them, with a pointer into the text or an honest "not yet."

Part XIII
Common Objections and Responses

The strongest objections a collider experimentalist, phenomenologist, statistician, or string theorist is most likely to raise on first reading — answered, with pointers into the paper

A paper that claims a zero-parameter collider target derived from integers will be met — rightly — with skepticism from every quarter. This Part collects fifty-eight objections, grouped by theme, each with an extended response and a pointer to where in the paper the full treatment lives. The objections are not strawmen; they are the questions the author has been asked, sharpened. Where the honest answer is "not yet," we say so, and point to the open-item roadmap (Appendix O) that records exactly which claims are theorems, which are conditional, and which are estimates. After all, the test of a framework is not that it deflects every question, but that it knows which questions it has answered and which it has not.

P Methodology and the derivation chain

Objection 1 (O-KK-1) (the numerology skeptic)
"$\frac{1}{2}M_{\rm Pl}e^{-35}$ is one formula among millions of integer combinations. You found one that lands near the current frontier and built a paper on it."

Response. The charge would stick if 35 had been searched for. It was not: $K = 35$ is the terminal entry of the OGN-5 Diophantine cascade $K = LJ = M^2 + N^2 + L^2$ proved in Arisaka (2026, A3-BI6), fixed before any collider comparison, and the *same* integer is spent again on a number measured decades ago — $v = M_{\rm lock}(2/5)^3 = 246$ GeV. A formula-mining exercise gets one landing; it does not get two landings from one integer, and it cannot survive pre-registration, which converts the coincidence charge into a scheduled experiment (§44.1). The paper's posture is forcing, not searching: every step of the chain from the axioms to 3.849 TeV is written down (§9.1) precisely so that the reader can locate the step at which, on the skeptical reading, the cheat must have occurred — and find that there is no free step available.

Objection 2 (O-KK-2) (the factor skeptic)
"The $\frac{1}{2}$ in $\frac{1}{2}M_{\rm Pl}e^{-K}$ looks like a fitted fudge dressed as a constant."

Response. The $\frac{1}{2}$ is the normalization of the lowest physical mode on the canonical metric of Arisaka (2026, A2-GG6) (its KK-spectrum appendix), not an adjustable prefactor: the anchor paragraph of §9.1 traces it, and the same appendix records why the naive Bessel prefactor $2x_1 \approx 7.66$ is *not* an independent correction. The operational point is sharper than the aesthetic one: whatever one thinks of the $\frac{1}{2}$, it was fixed before the comparison, it cannot be revised without breaking the v over-constraint, and the experimental window 3.85 ± 0.10 TeV is set by resolution and breadth, not by prefactor anxiety.

Objection 3 (O-KK-3) (the post-hoc skeptic)
"How convenient that the prediction sits exactly at today's exclusion frontier — neither excluded nor safely far away."

Response. Invert the logic. A prediction fixed by integers has to land *somewhere*; landing at the live frontier is simply what a claim looks like in the years when it is about to be decided — had the LHC run five years longer before this paper appeared, the same number would sit just behind the frontier, and five years less, just ahead of it. The suspicious pattern would be a prediction that *tracks* the frontier across revisions; §54 documents the opposite (the excluded Scenario A was not moved but retired, §10.4). Meanwhile the same chain fixes $v = 246$ GeV, which sits at no frontier at all.

Objection 4 (O-KK-4) (the imports skeptic)
"The physics is 'frozen' upstream in the cohort. This paper cannot claim an independent test of anything."

Response. Correct — and embraced. Every import is a named theorem with a named owner (§9.1; the References Part records each), so the collider verdict tests the conjunction: a null at $\lesssim 2\ \mathrm{ab}^{-1}$ retires not a detachable module but the lock-scale implementation, taking the electroweak derivation of Arisaka (2026, A9-SM) with it. Dependence is what makes the falsifier valuable; an "independent" collider module that could fail without consequence upstream would be exactly the escape-hatch structure §1 criticizes.

Objection 5 (O-KK-5) (the hidden-choices skeptic)
"No free parameters? You chose Bessel J_1, you chose Neumann boundary conditions, you chose the orbifold. Choices all the way down."

Response. Each item on that list is fixed by an observed fact, not selected: the boundary conditions are the ones whose zero-mode sector *is* the observed massless Standard-Model gauge fields (Theorem T-KK-2's premise, §10.4); given them, the radial equation and its J_1 zeros follow with no residual freedom; and the metric on which all of this is posed is itself Theorem 2 of Arisaka (2026, A2-GG6), not an ansatz. The honest residue of choice is recorded where it lives: the discrete g_B menu (L-KK-1) and the graded estimates of Appendix O — a short list, published as such.

Objection 6 (O-KK-6) (the self-grading skeptic)
"Your [T]/[D*]/[M]/[E] grades are self-assigned. Every crackpot marks his own homework."

Response. The grades are not verdicts but *addresses*: each graded claim carries the location of its derivation (the ledger of §9.2, the symbol table of Appendix B), so a reader can regrade any entry by checking the cited step. The discipline's value is asymmetric: self-assigned grades that flatter would be exposed by the first hostile reader; self-assigned grades that are honest — including [M] for a coupling menu and estimate labels on the $n = 2$ rates — cost the paper rhetorical force and buy it checkability. The reader is invited to audit exactly that trade.

Objection 7 (O-KK-7) (the K-factor skeptic)
"You quote LO×K rates and call the result 'derived.' NLO could move everything."

Response. The paper is careful about which claims are shapes and which are rates. The theorem-grade content — masses, degeneracy, sign, angular forms, leptophobia, $A_{FB} = 0$ — survives any revision of K-factors or PDFs (Part VI). The rates carry published K-factors with citations and quoted bands ([7.0–12] fb), and the single most K-sensitive object, the interference template, is flagged as the first theory task of the program (§44.5, Appendix O). Nothing in the falsification logic requires the central rate to be exact: the kill window is stated against the band.

Objection 8 (O-KK-8) (the metaphysics skeptic)
"An 'RG-scale dimension' is philosophy, not physics. You cannot test what a coordinate *means*."

Response. The meaning is tested through the consequences that differ between the readings. A spatial fifth dimension of the same metric drags a radion (testable: Higgs-precision null, P-KK-9), permits a tunable length (testable: the mass could not be an integer output), and makes the θ tower's scale a dial (testable: angular satellites, P-KK-7). The non-spatial reading forbids all three. §5.5 and Table 2 convert the ontology into rows of falsifiable differences — which is precisely what rescues it from philosophy.

Objection 9 (O-KK-9) (the sociology skeptic)
"A single-author program of this scope cannot be trusted; big claims need big collaborations."

Response. Trust is the wrong currency; checkability is the right one, and the paper is engineered for it: line-by-line derivations at analysis-team depth (Part VI), an anchored and calibrated simulation whose every number is reproducible from stated inputs (§39), a reinterpretation card (Appendix M), and a falsifier that hands the verdict to two thousand-person collaborations. The claim's fate does not rest on the author's authority — by construction, it rests on ATLAS and CMS.

Objection 10 (O-KK-10) (the Popperian bookkeeper)
"Scenario A died and the theory survived. That is one falsification absorbed already — classic degenerating research program."

Response. A degenerating program absorbs falsifications by *adding* freedom; this one *spent* freedom. The dichotomy was a pre-stated, discrete either/or with both masses and rates published before the comparison (§10.4); excluding one horn answered the pre-registered question (the throat vectors are composite) and left the survivor *more* exposed, not less — one scenario, one mass, one rate, no third option. Lakatos's test is whether the adjustment predicted anything new: here the surviving scenario's degenerate cluster and its 8.7 fb rate were on the table from the start.

Q The mass and the fine structure

Objection 11 (O-KK-11) (the error-bar skeptic)
"What is the theory uncertainty on 3.849 TeV? Every prediction needs an error bar."

Response. The centroid is a closed form in CODATA constants and an integer, so its *theory* error is the within-scenario interaction correction, $\mathcal{O}(1/K) \lesssim 3\%$ (§10.4); the quoted window ± 0.10 TeV is *experimental* (resolution and breadth), chosen so that the search, not the prediction, sets the tolerance. Contrast the usual situation, in which the theory band is the widest one in the plot: here the prediction is the sharpest object in the analysis, and §9.1's anchor paragraph states exactly which constants it consumes.

Objection 12 (O-KK-12) (the neighboring-integer skeptic)
"If 3.85 TeV is excluded, you will discover that K was really 34 or 36."

Response. No such move exists. $K = 35$ is not a fitted exponent but the closure of $K = LJ = M^2+N^2+L^2$ with $(M, N, L; J) = (1, 3, 5; 7)$: neighboring integers break the Diophantine identity itself, and even granting the vandalism, $e^{\pm 1}$ moves the mass by a factor 2.7 — to 1.4 or 10.5 TeV — and destroys the already-verified $v = M_{\text{lock}}(2/5)^3 = 246$ GeV by the same factor. The integer is load-bearing across the cohort precisely so that it cannot be quietly renegotiated here.

Objection 13 (O-KK-13) (the degeneracy skeptic)
"Exact mass degeneracy $m_{g^*} = m_{Z'_B}$ is untestable underneath 15% widths — a free miracle."

Response. The degeneracy is a theorem (same boundary-value problem, §10.4), and it is testable in two ways: any resolved splitting beyond the width scale falsifies it outright, and even unresolved, the two states separate by *channel decomposition* — color-octet versus singlet production, dijet-to-$t\bar{t}$ ratios, and the interference dichotomy of O-KK-28 — the FCC-hh unfolding program of Part IX. Far from free, the degeneracy is one of the few claims that a single clean measurement could kill.

Objection 14 (O-KK-14) (the hardening skeptic)

"The 'composite-hardened' scenario reads like an escape hatch drilled after 2.46 TeV got into trouble. And say it plainly: a theory with a discrete menu of N readings has N lives. You had two; you have spent one. Why should anyone believe the second death would be accepted?"

Response. The record shows the opposite order: the caution that the free-field root ratio "need not survive in the non-spatial spectrum" is intrinsic to the construction (§10.4, Step 6), the dichotomy was published *with both rates* — including Scenario A's unmissable 224 fb, hardly the number one quotes when building an escape — and the exclusion was then carried out in this same paper against the horn that data could reach (§31). An escape hatch is something you exit through quietly; this was a doorway walked through in public, with the door then bolted. As for the stronger form of the charge: the menu had exactly two entries, both stated in advance with their rates, and no third reading exists — Scenario B is not one refuge among many but the sole survivor of a pre-declared either/or. The paper's own falsifier prices this in: the kill window of §31.5 is written against Scenario B, and a null there is *terminal* — there is no Scenario C to retreat to, and the requalification that retired Scenario A in public (§31) is the standing demonstration that this program accepts its deaths in print.

Objection 15 (O-KK-15) (the borrowed-math skeptic)

"You disown Randall–Sundrum on every page and then solve its exact Bessel problem."

Response. Deliberately: the eigenvalue problem is shared because the metric is (Theorem 2), and shared mathematics is what makes the derivation checkable by anyone fluent in the warped literature. What differs is the ontology, and §5.5 itemizes where the difference pays observable rent: no radion, integer length, composite reading of the resonances (which is what made the Scenario dichotomy a physical question), multi-field clustering, and a GUT-scale angular tower. Borrowing RS's Bessel functions while rejecting its spatial fifth dimension is no more inconsistent than borrowing Fourier analysis from heat conduction.

Objection 16 (O-KK-16) (the band-is-a-dial skeptic)

"$\Gamma/m = 14.6\%$ rests on $r_{t_R} \in [4, 5]$. A band is a dial with a euphemism."

Response. A dial is continuous freedom exercised *after* seeing data; the band is a derived-conditional bracket stated *before*, with its central value, its endpoints, and its downstream consequences (width, BR fingerprint, P_t band) all published together (Part VI). Crucially, the band is self-closing: the HL-LHC line-shape fit measures Γ/m to a few percent and P_t to ± 0.08, turning the bracket into a cross-check rather than a refuge (Appendix O). A dial that publishes its own execution date is not much of a dial.

Objection 17 (O-KK-17) (the absence skeptic)

"Predicting the *absence* of angular satellites is a cheap unfalsifiable negative."

Response. It is exactly as falsifiable as any positive claim: a spatial sixth dimension of accessible radius *requires* $\ell = \pm 1$ partners decorating the cluster, so observing satellites falsifies the non-spatial reading at a stroke (P-KK-7), while their absence discriminates GG-6 from the entire spatial-compactification class. Negatives are cheap only when nothing predicts the positive; here the alternative does, loudly, which is what gives the silence its evidential weight — the same logic as the muon-collider null of §50.3.

Objection 18 (O-KK-18) (the overtone skeptic)

"If FCC-hh finds nothing at 7.047 TeV you will blame the estimated coupling, not the theory."

Response. The grades pre-empt the dodge: the *position* is theorem-grade (x_2/x_1, no freedom) while only the *rate* is an estimate, and Part IX states the consequence in advance — at FCC-hh statistics, a same-coupling rate of order 10 fb means even an order-of-magnitude coupling suppression leaves thousands of events, so a genuine position-null at 7.047 TeV kills the radial tower, coupling excuses notwithstanding. The commitment is written where it can be enforced: an overtone series that misses its Bessel ratio is not this geometry.

R Couplings, rates, and suppressions

Objection 19 (O-KK-19) (the convenient-suppression skeptic)
"$\kappa_{gg} = 0.47/K$ is suspiciously tiny — engineered so the graviton hides from every LHC null."

Response. The order of events refutes the engineering charge: the falsifier was staked *before* the coupling was computed, and the computation (Corollary T-KK-2c, a closed Bessel form with no inputs to tune) then *retro-dicted* the published diphoton null by a factor of several thousand — a result that would have been an embarrassing tension had the overlap come out large. Nor is the smallness an escape: it converts into a positive, dated commitment ($\sigma(G_1) \approx 0.13$ pb at FCC-hh, discovery in the first weeks or falsification of the cluster reading, §50.1). A suppression that hides a state forever would be convenient; one that reschedules it, with a rate attached, is just a derivation.

Objection 20 (O-KK-20) (the exact-zero skeptic)
"$\mathrm{BR}(Z'_B \to \ell\ell) = 0$ *exactly* is impossible in quantum field theory: kinetic mixing and loops always generate something."

Response. Stated precisely: the zero is exact at the level of the $\mathrm{U}(1)_B$ charge assignment — leptons carry no baryon number, so there is no tree coupling to cancel, which is already stronger than the tuned leptophobia of model-built Z's. Radiative kinetic mixing is an honest caveat, and it is small, ultraviolet-calculable, and most importantly *bounded by the same measurement that tests the zero*: the FCC-hh dilepton null at the 10^{-6} level (§50.1). The operational form of the claim, immune to loop pedantry, is the one the paper bets on: any *observable* dilepton line at the cluster mass falsifies the assignment.

Objection 21 (O-KK-21) (the nomenclature skeptic)
"You call G_1 a graviton, but GG-6 never quantizes gravity. Which is it?"

Response. Appendix G answers at length; the short form: "graviton" names G_1 by *pedigree* (it is a mode of the bulk metric sector, whence its spin-2-ness and stress-tensor couplings), not by the quantum-gravity job description. As a massive spin-2 resonance — in the RG reading, a composite state of the boundary theory — its quantum status is as unproblematic as a ρ meson's, and its discovery would evidence the six-dimensional geometry while neither requiring nor proving that massless 4D gravity is quantized. The genuinely deep statement runs the other way: in GG-6, quantization itself is an output of the OGGEP projection, for graviton and photon alike.

Objection 22 (O-KK-22) (the why-tops skeptic)
"Why should everything decay to top quarks? Top-philia smells like a choice made to dodge the dilepton limits."

Response. It is the same geometry that makes the top heavy, used twice: a fermion's mass and its coupling to the infrared-localized cluster are the *same* overlap integral, so the state that is heavy because it lives at the throat's tip is automatically the state the cluster talks to (Appendix I). The numbers are not adjustable — $(g_V^{t2} + g_A^{t2})/\xi_{\rm UV}^2 \approx 280$ per flavor, giving $t\bar{t} : b\bar{b} : jj \approx 70 : 3.5 : 1$ — and the alleged convenience cuts both ways: top-philia moves the signal into the experimentally *hardest* clean channel, which no one dodging limits would choose.

Objection 23 (O-KK-23) (the menu skeptic)
"A 'discrete menu' for g_B is a dial with extra steps."

Response. A dial is a continuum; the menu is four values, each produced by a named counting rule, published with its consequences (Z'_B rates spanning 0.2–6.8% of the cluster sum) and registered as an open lemma with its closing condition stated — the same GDOS counting that fixed $\alpha_s(M_Z)$ (L-KK-1; Appendix O). Nothing downstream hides behind the menu: the falsifier is carried by the octet, whose coupling is not on any menu, and every menu point is separately falsifiable through the fifth-force window it also feeds (Arisaka, 2026, A9-SM).

Objection 24 (O-KK-24) (the sign skeptic)
"A *sign* prediction, $\xi_{\rm UV} < 0$? Cross sections are sign-blind."

Response. Squares are sign-blind; interference is not. The octet shares the gluon's color channel, so the sign enters linearly through the interference term and imprints a parameter-free *shape*: enhancement below the pole, deficit above, the summed spectrum crossing *below* the Standard Model near 4.9 TeV (Part VI, P-KK-4). Flip the sign and the shape mirrors — a distinction a template fit measures directly, and the quantitative Figure 4 draws. Few BSM frameworks are falsifiable at the level of a relative sign; this one is.

Objection 25 (O-KK-25) (the precision pedant)
"You write $|\xi_{\rm UV}| \approx 1/\sqrt{K} = 0.169$ in one place and 0.194 in another. Which is the prediction?"

Response. Both, at their stated grades: $1/\sqrt{K}$ is the leading-order scaling that exhibits the K-dependence, and 0.194 is the exact mode value including the derived warped correction (§10.4), the number every rate in the paper uses. The refinement went in the derived direction before any comparison — the correction is not a knob but a computed +15% — and the practical spread between the two is far inside the quoted rate band. When a scaling and its exact completion are both published with their derivations, the right description is convergence, not equivocation.

Objection 26 (O-KK-26) (the hidden-channel skeptic)
"The 70:3.5:1 fingerprint assumes no exotic decay channels. A single hidden state ruins it."

Response. The lock cluster is the *lowest-lying* new structure of the derived spectrum, so within GG-6 there is nothing below it for it to decay into — the fingerprint is closed by construction, not by assumption. If nature nevertheless supplies an extra light state, then GG-6's spectrum is wrong and the paper is falsified on grounds far more direct than a branching-ratio distortion. Either way the fingerprint does its job: it is exact within the theory and diagnostic against alternatives (§56).

Objection 27 (O-KK-27) (the washout skeptic)
"$P_t = +0.91$ will never survive parton shower, hadronization, and reconstruction."

Response. Top polarization is special precisely because the top decays before hadronizing: the charged lepton is a maximal-power spin analyzer whose angular distribution survives to reconstruction level, which is why ATLAS and CMS already publish top-polarization measurements today. Part VI carries the lepton-analyzer derivation, and Part VIII propagates purity-corrected statistics to $\delta P_t = 0.14$ by 2033 and 0.08 by 2041 — enough to resolve the *interior* of the derived band $[+0.88, +0.92]$, not merely its existence.

Objection 28 (O-KK-28) (the interference-consistency skeptic)
"Interference with QCD requires identical color flow. Did you actually check which states interfere?"

Response. Yes, and the answer is itself a diagnostic. The color octet enters the *same* color channel as the s-channel gluon, so it interferes coherently (the amplitude algebra is written out in Part VI); the color-singlet Z'_B does not interfere with QCD at all — nor with the octet, since the same color trace kills the g^*–Z'_B cross-term, which is why the cluster rates may be summed incoherently. The cluster therefore predicts a two-component line shape — an interfering octet plus a non-interfering singlet at the same mass — and that yes/no dichotomy within one peak is one more handle the FCC-hh unfolding uses to separate the degenerate pair (O-KK-13, §50.1).

S The simulation and its statistics

Objection 29 (O-KK-29) (the fudge-factor skeptic)
"$F_{\text{cal}} = 4.33$ is a fudge factor by another name."

Response. A fudge factor is chosen to produce a desired answer; F_{cal} is *measured* against answers already published: narrow test signals injected at two masses reproduce ATLAS's expected limits with mutually consistent factors (4.40 and 4.27), and the calibrated framework then *predicts* the 15%-width limit at 8.4 fb, correctly ordered between ATLAS's own narrow and broad templates — a validation with no adjustable input (§41). Its application is conservative by construction (exclusion-derived, applied to discovery), and the one methodological temptation it forbids — double-profiling on top of post-fit uncertainties — is documented in the same section.

Objection 30 (O-KK-30) (the systematics-optimism skeptic)
"Halving the background systematics is a wish, not a scenario."

Response. The halving is the standard HL-LHC projection convention (Cepeda et al., 2019), and §41.1 decomposes it: the control-region and template-statistics components shrink as $1/\sqrt{L}$ by arithmetic alone ($\times 2.8$–4.9 over the program), object calibrations improve data-driven toward known floors, and only the theory-modeling class requires actual work. The paper also quotes the frozen scenario — no improvement of anything for fifteen years — as the pessimistic bound, and notes that at end-2026 roughly half the "improvement" is automatic with the Run-3 dataset itself. Reality lives between two published curves; neither is a wish.

Objection 31 (O-KK-31) (the where-is-the-hint skeptic)
"You claim 2.7–3.7σ sits on tape. ATLAS analyzed that region and saw nothing. One of you is wrong."

Response. Neither is — because "analyzed that region" and "ran this analysis" differ by the two factors this paper quantifies. The published search used templates mismatched to the coherent 15%-width shape (a factor ~ 3 dilution in rate sensitivity, §44.2) and coarse-binned counting that a floating background absorbs (§42); under those conditions the predicted signal produces *sub-2σ local* fluctuations in their own quotes — fully consistent with "saw nothing." Indeed §31 shows the framework reproduces ATLAS's observed limit with the signal *present*: the expected excess sits just below their published exclusion. The claim is not that ATLAS missed a bump; it is that a differently designed fit extracts what their design integrates over.

Objection 32 (O-KK-32) (the global-trials skeptic)
"No look-elsewhere? Your cohort makes dozens of predictions. Across twenty-two papers, something was bound to sit near a frontier."

Response. The look-elsewhere effect is a property of a *decision rule*, not of an author's productivity: the discovery claim at issue is a single pre-registered mass point with a single pre-stated test, and its local p-value is its global one (§44.1). Cross-paper trials would matter if the program's survival required only *one* of many predictions to succeed; the cohort's structure is the opposite — conjunctive, with shared integers, so that failures propagate (O-KK-4). A framework whose predictions must *all* hold pays no trials discount on any one of them; it pays the far steeper price that any one can kill it.

Objection 33 (O-KK-33) (the asymptotics skeptic)
"Asimov formulas fail at small counts, and your tail bins are small."

Response. The analysis bins carrying the significance have $b \gtrsim 19$ even at the smallest snapshot and grow from there (Table 26), inside the validated Asimov regime; the truly sparse display bins beyond 5.6 TeV contribute nothing to the quoted Z. Residual asymptotic error is second-order against the calibration factor, which was measured against ATLAS's full toy-based machinery and therefore absorbs precisely such effects (§41).

Objection 34 (O-KK-34) (the correlations skeptic)
"Treating per-bin systematics as uncorrelated is wrong, and everyone knows correlated backgrounds absorb broad signals better."

Response. Agreed on the physics — which is why the paper does not rely on the uncorrelated model alone. The correlation-driven absorption is demonstrated explicitly (the two-parameter profile of §42, which eats 75% of the window signal), its full-machinery magnitude is what F_{cal} was measured to capture, and the caveat ledger records the treatment (item 5, §45). The quoted significances are thus calibrated against a fit that *does* exploit correlations; the per-bin model is the bookkeeping, not the claim.

Objection 35 (O-KK-35) (the own-goal skeptic)
"Your sideband-contamination warning undermines your own sensitivity numbers."

Response. The numbers already stand on ground the warning does not touch: they use ATLAS's post-fit uncertainties, which encode the real fit's freedom, calibrated by F_{cal} (§41). The warning (§44.5) is design guidance for the *future dedicated fit* — anchor the normalization below 1 TeV, where the derived contamination falls to the percent level — and the two-parameter demo quantifies what ignoring it costs. A paper that flags the trap and prices it is not undermined by the trap; it is the reason the trap will be stepped around.

Objection 36 (O-KK-36) (the combination skeptic)
"Your ATLAS+CMS $\sqrt{2}$ assumes uncorrelated experiments. Theory systematics are common to both."

Response. Which is why $\sqrt{2}$ is never quoted alone: Table 27 brackets both extremes — fully correlated (end-2026 improved: 4.6σ) and fully uncorrelated (5.2σ) — with the statement that reality lies between, and the physical sorting is given (modeling: correlated; detector calibrations and control-region statistics: not). Every headline in the paper carries the bracket, not the optimistic edge.

Objection 37 (O-KK-37) (the seed skeptic)
"One random seed for the pseudo-data. Cherry-picked fluctuations?"

Response. Every quoted significance is an Asimov (median-expected) quantity, computed from expectations and independent of any seed; the single fixed seed generates only the *display* points of the spectra figures, stated as such in the caveat ledger (item 8) and reproducible from the archived workbook. Change the seed and the dots move within their error bars; not one number in any table moves.

Objection 38 (O-KK-38) (the PDF-vintage skeptic)
"CTEQ6 is two decades old. Modern PDFs could shift your rates."

Response. They could, and the quoted bands already say by how much: the [7.0–12] fb spread on the octet rate is dominated by exactly this class of uncertainty, and every conclusion — the frontier position, the kill window, the calendar — is stated against the band, not the central value. Shapes, which carry the theorem-grade content, are far more stable across PDF generations than tails of absolute rates; and the paper names the permanent fix (LHeC-class high-x input, §50.4) rather than pretending the issue away.

T Reducibility, non-spatial dimensions, and neighboring frameworks

Objection 39 (O-KK-39) (the category skeptic)
"A dimension is something you can move through. A 'scale dimension' is a category error."

Response. By that criterion time is a category error too — one cannot "move through" time in the spatial sense, yet Minkowski's fusion (Minkowski, 1908) is not up for debate. The disciplined criterion classifies a dimension by its symmetry generator, and §5.5 applies it: translations generate space and time; a dilatation generates I_r; a compact charge generates S^1_θ. "Motion" along r is transmutation — change of renormalization scale — which is as physical as anything in field theory; what it is not, is walking.

Objection 40 (O-KK-40) (the holography skeptic)
"This is AdS/CFT rebranded: everyone knows the radial coordinate is the RG scale."

Response. In holography that identification is a *duality* between two complete descriptions, either of which may be taken as fundamental; in GG-6 it is a single literal construction — there is no spatial-bulk description to dualize to, no radion, no tunable length, no landscape of compactifications (Table 2). The distinction pays observable rent: a genuine AdS_5 dual would carry the usual RS phenomenology (radion, adjustable kr_c, single pole); GG-6 forbids each item, and the forbiddances are on the prediction list (P-KK-7, P-KK-9). Rebranding changes names; this changes what experiments should see.

Objection 41 (O-KK-41) (the SUSY skeptic)
"No supersymmetry anywhere? Serious UV physics requires it."

Response. In GG-6 the absence is a theorem, not a taste: supercharges must square-root translations, and the two extra directions have none to offer — a dilatation and a compact charge admit no super-partner extension (§5.5; the algebraic statement is Theorem-level in Arisaka, 2026, A2-GG6). The framework therefore *predicts* the null that fifty years of superpartner searches have delivered, rather than perpetually postponing their discovery — one of the few readings under which the LHC's SUSY silence counts as evidence *for* something. (The careful reader will note the loophole class: exotic superalgebras that do not square-root translations. They exist as mathematics; what they cannot deliver is the phenomenological content — superpartners protecting scalar masses — that made supersymmetry physics rather than algebra, so the claim stands where it matters.)

Objection 42 (O-KK-42) (the sixth-dimension skeptic)
"Any honest 6D theory has a second KK tower. Where are the θ modes?"

Response. At $\sim E_{\mathrm{GUT}}$, and immovably so: the phase circle is topological, with no radius modulus to dial its tower downward — $R_\theta = 1$ is a convention on a dimensionless angle, not a size (§5.5). The observable content is the pair of statements the paper registers: the radial tower at derived TeV positions (P-KK-8) and *no* angular satellites decorating it (P-KK-7). A spatial 6D construction generically fails the second; that is the discrimination, available at the same experiments.

Objection 43 (O-KK-43) (the stabilization skeptic)
"Without a radion, what keeps the throat from collapsing or running away?"

Response. Collapse and runaway are fates of *sizes*; the throat's length is a *count* — the integer $K = 35$ of the cascade — and a count has no continuous deformation to run along (§5.5; Arisaka, 2026, A4-Euclid). Asking what stabilizes it is asking what stabilizes the number of colors in QCD. The claim is radical but testable: a stabilized modulus always leaves a scalar in the spectrum and fingerprints in Higgs couplings, so the Higgs-precision null (P-KK-9) is exactly the experiment that distinguishes "no modulus" from "hidden modulus."

Objection 44 (O-KK-44) (the missing-partners skeptic)
"Warped models come with KK fermions, KK W'/Z', custodians. Your spectrum looks suspiciously sparse."

Response. Sparseness is forced, not curated: the low-lying content is the lowest radial modes of the fields actually present in the bulk group of Arisaka (2026, A2-GG6), and that group was derived, not assembled to populate a spectrum. There are no custodial partners because there is no custodial problem to solve — the electroweak sector is not protected by tuning but derived (Arisaka, 2026, A9-SM). Heavier structure exists (radial recurrences, Part IX) at derived or dual-calculable positions; what does not exist is the model-builder's freedom to add a state whenever a search wants one.

Objection 45 (O-KK-45) (the no-free-lunch skeptic)
"No landscape, no moduli, no tuning — because you hid all the arbitrariness in the 'constitution.' The lunch is not free."

Response. Correct, and the paper's ancestors say so themselves: the price of forced low-energy structure is the strength of the constitutional premises, stated as such in Arisaka (2026, A2-GG6) (Theorem 0's assumptions are listed, not smuggled). The relevant comparison is where each program places its irreducible assumptions — at inaccessible trans-Planckian scales, or in axioms whose consequences land at 3.85 TeV within one PhD of publication. Lunches are never free; this one at least comes with a bill the LHC can audit.

Objection 46 (O-KK-46) (the convenient-split skeptic)
"Radial tower testable, angular tower at the GUT scale — the split conveniently exposes only what you can survive."

Response. The split runs the other way from convenience: it *exposes* the radial tower at exactly the energy the LHC reaches now, with a rate the current run can kill (F-KK-2B), while the "safe" angular sector still makes a falsifiable claim through its absence (P-KK-7) and through its lightest member elsewhere in the cohort — the 8.1 μeV kaxion of Arisaka (2026, C1-Kaxion), under active haloscope test. A construction optimized for survival would have put the radial tower at 40 TeV; this one put it under the beam.

Objection 47 (O-KK-47) (the anomaly skeptic)
"You cannot simply gauge baryon number. With Standard-Model fermion content the mixed triangles $\mathrm{U}(1)_B$–$\mathrm{SU}(2)_L^2$ and $\mathrm{U}(1)_B$–$\mathrm{U}(1)_Y^2$ are non-zero ($\pm 3/2$ over three generations); every honest Z'_B model buys them off with new chiral 'anomalon' fermions, which then have masses, signatures, and problems of their own. Where are yours?"

Response. Nowhere — because none exist, and their absence is a theorem, not an oversight. The premise is correct: with the Standard Model alone, $\mathcal{A}[B\text{–SU}(2)_L^2] = +\frac{3}{2}$ and $\mathcal{A}[B\text{–}Y^2] = -\frac{3}{2}$ (the gravitational, cubic, and B^2Y triangles vanish identically), and the conventional remedy is indeed anomalon matter. But in GG-6 the boundary spectrum is not the Standard Model with baryon charges appended by hand: it descends from unsplit SU(5) multiplets on the OGN-5-locked charge lattice, and the microscopic anomaly polynomial of that spectrum factorizes to *rank one* (GG-A3-BI6, Theorem T-BI6-6.1). The two surviving mixed triangles are the boundary projections of the cross-blocks of the BI-6 anomaly matrix, and the gauge variation of the single Green–Schwarz counterterm $\int B_2 \wedge X_4$ cancels them by inflow, with the compensating couplings splitting across $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y \subset \mathrm{SU}(5)$ in exactly the ratios required (Arisaka, 2026, A3-BI6, the boundary-mixed-anomalies subsection). The four-dimensional remnant of B_2 is a Stückelberg axion whose shift performs the cancellation *and* is eaten by the baryonic vector: the Z'_B mass of Appendix F and the anomaly bookkeeping are one coupling seen twice. Note the collider corollary, which is a falsifier of its own: anomalon models necessarily populate the accessible range with new charged chiral states, while GG-6 *forbids* them — discovering anomaly-canceling fermions accompanying a 3.85-TeV Z' would count against this implementation as surely as a clean null.

U Cohort dependence and standing

Objection 48 (O-KK-48) (the house-of-cards skeptic)
"If A9's electroweak derivation fails, this paper falls with it."

Response. Yes — deliberately. The over-constraint $v = M_{\mathrm{lock}}(2/5)^3$ means the two papers share one integer and one fate, so a collider null at 3.85 TeV wounds A9 and a failure of A9's chain wounds this paper; the exposure is stated in §59 in both directions. A house of cards is fragile *and unfalsifiable-by-parts*; this structure is fragile *because* falsifiable-as-a-whole — the property the whole program regards as the point.

Objection 49 (O-KK-49) (the analysis-cost skeptic)
"Why should a collaboration spend person-years on an arXiv-only, single-author benchmark?"

Response. Because the request is priced and the payoff is asymmetric. The ask is one template added to an existing merged-SR analysis on data already recorded — thesis-scale, per the blueprint of Part VIII — and the deliverable is valuable under *every* outcome: the world's best limit on any broad top-philic resonance, a reusable interference-aware template, and, in the tail of the distribution, the first particle beyond the Standard Model in collider history. Collaborations reinterpret benchmarks of far weaker epistemic standing (fixed-coupling scans with no derivation) as routine practice; a zero-parameter, pre-registered target with a published discovery calendar clears that bar with room to spare.

Objection 50 (O-KK-50) (the closed-loop skeptic)
"The cohort cites itself. Twenty-two papers in a loop is not external validation."

Response. The internal citations carry the *derivations*; the *validations* are external by construction, and Section B of the References Part lists them: ATLAS and CMS measurements, HEPData records, published K-factors, the PDG ledger, the strategy documents. Every contact between the cohort and reality happens on instruments the cohort does not own — most decisively the falsifier itself, whose verdict will be signed by three thousand people with no stake in the program.

Objection 51 (O-KK-51) (the pace skeptic)
"Twenty-two papers in roughly a year. Depth takes decades; this is speed-writing."

Response. Pace is evidence about process, not about correctness, and this paper is built so the two can be separated: its chain is locally checkable given the cited theorems — a reader can verify T-KK-1 in an afternoon, T-KK-2's boundary-value problems in a day, the simulation from the archived workbook — without adjudicating the cohort's production history. Where depth is genuinely unfinished, the paper says so in its own voice (the open lemma, the estimates, Appendix O). Judge the derivation on the page; the calendar behind it is biography.

Objection 52 (O-KK-52) (the schedule skeptic)
"Your calendar leans on CERN schedules that slip as a law of nature."

Response. Every decision point in the paper is stated in integrated luminosity, with calendar years attached only through the official planning documents quoted at their source (§39.3, Part IX). If LS3 stretches, the 0.80 ab^{-1} five-sigma crossing arrives later in wall-clock time and at exactly the same point in data; F-KK-2B fires at $\lesssim 2\ \text{ab}^{-1}$ whenever those inverse attobarns exist. Slippage delays the verdict; it cannot dilute it.

Objection 53 (O-KK-53) (the attribution skeptic)
"If ATLAS finds *some* 3.9-TeV top-philic octet, every model-builder will claim it. Discovery would prove nothing about GG-6."

Response. The mass alone, granted — which is why the paper never bets on the mass alone. The registered fingerprint is conjunctive: exact vector degeneracy at one mass, exact leptophobia, the derived interference sign with its below-SM crossing, P_t in $[+0.88, +0.92]$, *no* angular satellites, and an overtone at 1.831× the mass (P-KK-2 through P-KK-8). §56 walks the neighboring frameworks through that list; each fails multiple clauses, and none predicts the ratio-locked second family. Imitators can gate-crash a bump; they cannot gate-crash a spectrum.

V Experimental practice

Objection 54 (O-KK-54) (the practitioner)
"Broad $\Gamma/m = 15\%$ resonances at $S/B \sim 1$ are exactly what tt-resonance analyses are worst at. 'Just reanalyze' is glib."

Response. Agreed that broad-plus-interfering is the hard case — that is why the paper does the analysis design work instead of gesturing at it: the anchor analysis is read cut-by-cut (§38), the template is fully specified with its sign structure (Part VI), the fit-design trap is located and priced (§44.5), the statistics are calibrated to the published machinery (§41), and the whole chain is packaged in the reinterpretation card (Appendix M). "Just reanalyze" would be glib; "reanalyze along this worked blueprint" is a costed proposal.

Objection 55 (O-KK-55) (the template-validation skeptic)
"Interference-aware templates are notoriously hard to validate. How would we know the template is right?"

Response. Three independent handles. The template's *shape* is falsifiable within the fit itself — the sign is left floating and the data choose (a wrong-sign preference falsifies P-KK-4, not the analysis); its *normalization region* is validated where the derived contamination is provably negligible (below $\sim$1 TeV, §44.5); and its *theory* systematic is the named NLO task of Appendix O, commissioned before unblinding rather than argued after. This is a stronger validation position than the generic templates now in use, whose widths and shapes answer to no derivation at all.

Objection 56 (O-KK-56) (the trigger skeptic)

"The 2015–2026 data were triggered without your signal in mind. Whatever is on tape may not contain it."

Response. The kinematic chain settles this (§38.2): a 3.85-TeV decay delivers $\sim$500 GeV leptons to single-lepton triggers thresholded at 26–28 GeV — the signal has been sailing onto tape at essentially 100% trigger efficiency since 2015, in the same streams the published $t\bar{t}$ searches already analyze. The scouting and reduced-stream proposals of Table 23 are future insurance for *other* corners of the program, not a precondition for this one.

Objection 57 (O-KK-57) (the governance skeptic)

"Who 'owns' a joint pre-registered search? ATLAS and CMS do not co-analyze."

Response. Nothing here requires co-analysis: the card (Appendix M) is collaboration-neutral, each experiment runs its own fit on its own data, and the combination follows the established precedent of the Higgs-discovery-era and W-mass combinations — independent analyses, joint statistical combination afterward, with the correlation brackets of Table 27 spanning exactly the choices a combination committee would face. The pre-registration lives in the public template, not in anyone's governance chart.

Objection 58 (O-KK-58) (the last skeptic standing)

"If your framework is right, the evidence has been sitting on tape for years while everyone looked past it. That strains belief more than any equation."

Response. It should strain belief less than its precedents. The Ω^- sat in bubble-chamber film until someone knew what to scan for; the CMB was in Bell Labs' horn as noise; pentaquarks hid in LHCb's own recorded data for years of running. Recorded is not the same as extracted: extraction requires a template, and §44.2 quantifies exactly how far the published templates are from this signal (a factor $\sim$3 in rate, plus the absorption of §42). The claim is not that thousands of physicists were careless; it is that no one runs an analysis for a signal no theory had specified — and that, precisely, is what this paper exists to change.

Fifty-eight objections, answered where an answer exists and conceded where it does not: that is as close to referee-proofing as a single document can come, and the reader who has found their own doubt missing from the list is invited to regard it as the fifty-ninth and write. All that remains is the record itself — the works this paper stands on, each with a note explaining why.

Part XIV
References

Cited works in two groups — the GG-Theory cohort and the external literature — each entry carrying a relevance note and its first citation point

The reference list is organized in two sections. **Section A** lists the eighteen sister papers of the GG-Theory cohort, in the canonical order of the program (the constitutional root, the GG6 Trio, the boundary law and its pillars, the three textbook readouts, the capstone, the two phenomenology branches, and the Foundations Branch); every cohort paper is cited in the body of this one, because the lock cluster sits inside a single deductive structure and its verdict propagates through all of it. **Section B** lists the external literature — the theoretical ancestors, the experimental analyses this paper is anchored to, the statistical methods, and the strategy documents of the post-LHC landscape — alphabetically by first author.

Every entry carries a **Relevance** paragraph stating what the work contributes to this paper and where it enters, followed by the location of its first citation. The list contains no entry that is not cited in the text.

All cohort preprints are available at https://arisaka-gg.org/. External references are cited under their published or preprint identifiers as given.

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. Preprint, available at https://arisaka-gg.org/.

Relevance. The foundational charter from which the present paper inherits its rules of admissibility. The lock cluster is presented here as a downstream consequence of the four axioms and seven postulates — in particular MSMF (no continuous dials) and POS--MF/POS-7 (falsifiability) — rather than as an independent collider ansatz. We cite it whenever the constitutional provenance of a claim, the naming convention, or the pre-declared-falsifier discipline is invoked. First cited in §3.

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.
Preprint, available at https://arisaka-gg.org/.

Relevance. The geometric stage of this paper. We cite A2-GG6 for the canonical scale-profile metric of Theorem 2, the forced bulk of Theorem 0, the non-spatial scale (I_r) and phase (S^1_θ) directions, the two-tower convention (angular tower at $E_{\rm GUT}$, observable radial lock cluster at $M_{\rm lock}$), and the Randall–Sundrum reading of the radial spectrum that underlies the entire collider analysis here. First cited in the Executive Summary.

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.
Preprint, available at https://arisaka-gg.org/.

Relevance. Supplies the rigorous proof of the OGN-5 Diophantine cascade $(M, N, L; J, K) = (1, 3, 5; 7, 35)$ from the single-tensor anomaly identity. The lock integer $K = LJ = 35$ that fixes the scale depth — and hence the derived mass $M_{\rm lock} = \frac{1}{2} M_{\rm Pl} e^{-35}$ of Theorem T-KK-1 — is established there; we cite it wherever the integer ledger is the source of a parameter-free number. First cited in §2.

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.* Arisaka-GG-A4-Euclid preprint. Preprint, available at https://arisaka-gg.org/.
Relevance. Owns the 6D Einstein–Hilbert reduction and the derivation of the universal constants GU-5 as bulk conversion factors. We cite it specifically for the bulk-to-Planck relation $M_6 = \sqrt{2/\pi}\, M_{\rm Pl}$ (Theorem T-Euclid-6) and the radial scale-profile integral $\frac{1}{2}(1 - e^{-2K}) \to \frac{1}{2}$ that together supply the factor $\frac{1}{2}$ in $M_{\rm lock} = \frac{1}{2} M_{\rm Pl} e^{-K}$, ensuring the centerpiece mass carries no fitted normalization. First cited in §5.5.

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. Preprint, available at https://arisaka-gg.org/.
Relevance. Supplies the open-system (GDOS) framework on which §16 rests: a collider event is not a closed-system observable but a 6D bulk excitation registered only after projection to the 4D boundary through a dissipative detector and a statistical reconstruction pipeline. We cite it for the open-system framing of detection (OGGEP/GDOS), not for any collider number. First cited in §3.

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 (the runtime definitions of GGA, GRIP-In, and GRIP-Out, and the two-regime figure).
Preprint, available at https://arisaka-gg.org/.
Relevance. Provides the GRIP (Gi-4) four-pillar cascade language used in the open-system reading of trigger, reconstruction, and decision-making. Cited alongside GG-A5-GDOS for the detection-as-open-system view of §16; it contributes no quantitative input to the lock-cluster prediction. First cited in §3.

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.
Preprint, available at https://arisaka-gg.org/.
Relevance. The quantum-information pillars that, with the Gi-4 pillars of Arisaka (2026, A6-GRIP), force the boundary law GDOS; cited for the origin story of the program (Part XI) and the information charge $\mathrm{U}(1)_{Qi}$ that partners $\mathrm{U}(1)_B$ in the BI-6 identity. First cited in §3.

Arisaka, K. (2026). *Quantum Mechanics Derived from GG-Theory as Projection from 6D Bulk to 4D Boundary.* Arisaka-GG-A8-QM preprint.
Relevance. Derives measurement as a complete runtime event (apparatus arming = endogenous GRIP-Out; registration = GRIP-In descent; the record = the halt) — the anatomy applied in §16 to the trigger and reconstruction chain. First cited in §3.

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.
Preprint, available at https://arisaka-gg.org/.
Relevance. Owns the scale chain $E_{\rm Pl} \to E_{\rm GUT} \to M_{\rm lock} \to v$ and the $\frac{1}{2}$ normalization of the lock scale used in the centerpiece. Crucially, it derives the electroweak vacuum expectation value $v = M_{\rm lock}(2/5)^3 \approx 246\,\mathrm{GeV}$ from the same $M_{\rm lock}$, which is the over-constraint that gives the collider test its unusual falsifying power: a wrong lock cluster would also unseat the electroweak derivation. First cited in §2.

Arisaka, K. (2026). *The Standard Cosmology as Theorem of GG-Theory: The Cosmic Inventory as Observations of One Six-Dimensional Geometry.* Arisaka-GG-A10-Cosmos preprint.

Preprint, available at https://arisaka-gg.org/.
Relevance. The cosmological member of the Consequence Trio: the same six-dimensional geometry whose radial lock spectrum this paper sends to the LHC fixes, in A10, the cosmic inventory ($\Omega_m \approx 1/\pi$) and the derived H_0. Cited in the cohort map as the third textbook read-out alongside Arisaka (2026, A8-QM) and Arisaka (2026, A9-SM); a collider verdict on the lock cluster propagates to A10's geometry exactly as it propagates to the electroweak ledger. First cited in §3.

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.
In preparation. To be posted at https://arisaka-gg.org/ on completion; no result of the present paper depends on it.
Relevance. The capstone of the A-series: the paper that closes the constitutional arc which the present paper's falsifier exposes to data. Cited in the cohort map; the shared deductive funnel of §3 is carried, with different payloads, by A9, A10, and A11, so the octet verdict at 3.85 TeV is also a stress test of the closure A11 completes. First cited in §3.

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.
Preprint, available at https://arisaka-gg.org/.
Relevance. The second Predictions-Branch paper: the three-leg derivation of the generation count from the same BI-6 anomaly structure that locks this paper's gauge content. Cited in the cohort map; the top-quark localization band $r_{t_R} \in [4, 5]$ that fixes the octet's width and polarization here presupposes the three-family structure P2 derives. First cited in §7.1.

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.
Preprint, available at https://arisaka-gg.org/.
Relevance. The third Predictions-Branch paper: the electron as the geometric ground state. Cited in the cohort map; its UV-localized lepton profiles are the same localization data that make this paper's cluster leptophobic in practice — the two papers read one geometry from opposite ends of the mass ladder. First cited in §7.1.

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.
Preprint, available at https://arisaka-gg.org/.
Relevance. The origin-of-i member paper of the Foundations Branch: from the two real structures of the arena this paper's spectrum lives on — the compact phase circle and the causal clock — it derives the single imaginary unit of boundary physics, so the interference phases that this paper's amplitudes and line shapes take for granted are outputs of the geometry rather than inputs. First cited in §7.1.

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.
Preprint, available at https://arisaka-gg.org/.
Relevance. The mirrors-and-matter member paper of the Foundations Branch: C, P, and T as orientation reversals of the three real bulk factors, with the matter–antimatter asymmetry at the end of the chain. Its dated kill list names this paper's HL-LHC lock-cluster verdict as a companion channel of the same shared geometry — the collider face of a battery whose other faces are decided elsewhere. First cited in §7.1.

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.
Preprint, available at https://arisaka-gg.org/.
Relevance. The sibling standing falsifier, and the honest description of the relation is that it is a sibling rather than an input. Nothing in the lock-cluster derivation uses the proton and nothing in this paper's physics turns on it. What the two share is the ledger: the same integers that place a resonance at 3.85 TeV also forbid the proton from decaying, so the framework is exposed at both ends of the energy axis at once — a resonance that must appear where it is predicted, and a decay that must never appear at all. A reader weighing how much this framework has staked should count both. Cited at roster level in the cohort map. First cited in §7.1.

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.
Preprint, available at https://arisaka-gg.org/.
Relevance. The background this paper's central negative claim needs and does not supply. The structural results here are absences — no radius modulus, no radion, no landscape, no supersymmetry requirement — and each is derived rather than assumed; but an absence reads as a result only to someone who already accepts that a conservation statement need not descend from a symmetry. W1 is the cohort's survey of exactly that century-long shift, from Noether to cohomology. It supplies no number and no theorem the lock-cluster chain uses. First cited in §5.5.

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.
Preprint, available at https://arisaka-gg.org/.
Relevance. The classification the mechanism of this paper's phase circle belongs to. The lock spectrum is shifted by a Wilson line: a constant potential around a compact angle, locally pure gauge, changing no field strength, and moving every charged mode anyway. That is protection and displacement by holonomy rather than by symmetry, which is W2's subject, and the absence of a radius modulus here is an instance of it rather than a coincidence. Cited at roster level in the cohort map, and named at the Wilson-line shift because that is where the mechanism is used. First cited in §5.3.

Arisaka, K. (2026). *The Kaxion: Cold Dark Matter as the Non-Spatial Compact-Angle Pseudoscalar of GG-6.* Arisaka-GG-C1-Kaxion preprint.
Preprint, available at https://arisaka-gg.org/.
Relevance. The dark-matter companion: the 8.1 μeV kaxion of the same compact θ direction whose angular tower this paper predicts *absent* at colliders. Cited in the cohort map; the two papers are the complementary read-outs of S^1_θ — the haloscope listens to its lightest mode while the LHC confirms the silence of its heavy ones. First cited in §7.1.

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.
Preprint, available at https://arisaka-gg.org/.
Relevance. The second Consequences-Branch paper: the derived-H_0 reading of the Hubble tension within the same geometry. Cited in the cohort map as one of the two phenomenology branches' astrophysical arms that a collider discovery at 3.85 TeV would indirectly fortify. First cited in §7.1.

Arisaka, K. (2026). *Black-Hole Entropy under GG-Theory: The GDOS Reading of the Horizon as an OGGEP Port — with the Exact Bekenstein–Hawking $A/4G_N$ Derived from Induced Gravity.* Arisaka-GG-C3-BH preprint.

Preprint, available at https://arisaka-gg.org/.
Relevance. The third Consequences-Branch paper: horizon thermodynamics under the same open-system projection whose collider face is this paper's SIP reading of detection (§16). Cited in the cohort map; together with Arisaka (2026, C1-Kaxion) and Arisaka (2026, C2-Hubble) it carries the framework into the arenas a collider cannot reach. First cited in §7.1.

B. Other references

Agashe, K., Belyaev, A., Krupovnickas, T., Perez, G., & Virzi, J. (2008). LHC signals from warped extra dimensions. *Physical Review D*, *77*, 015003. arXiv:hep-ph/0612015.
Relevance. Establishes the top- and longitudinal-boson-rich signatures of IR-localized warped resonances and the UV/IR coupling hierarchy. Cited in *Read-out II only* (assumptions B2, B4) for the empirical localization–coupling correspondence used to estimate R_f. First cited in §10.9.

Allanach, B. C., Odagiri, K., Parker, M. A., & Webber, B. R. (2000). Searching for narrow graviton resonances with the ATLAS detector at the Large Hadron Collider. *Journal of High Energy Physics*, *0009*, 019. arXiv:hep-ph/0006114.
Relevance. Supplies the corrected massive-vector partial width for a spin-2 resonance (fixing a typographical error in Han, Lykken & Zhang, 1999) and the benchmark $\sigma \times \mathrm{BR}(e^+e^-)$ table against which the workbook of §31 is validated. First cited in §31.2.

ATLAS Collaboration (Aad, G. et al.) (2012). Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Physics Letters B*, *716*, 1–29. arXiv:1207.7214.
Relevance. The ATLAS observation of the 125 GeV Higgs boson: the last unchecked entry of the Standard Model's list, completed. Cited in §1 as the closing bookend of the confirmation era — the discovery that finished a structure written down in 1964 and left the list, for the first time in half a century, without a next item. First cited in Part I.

ATLAS Collaboration (2019). Search for high-mass dilepton resonances using 139 fb^{-1} of pp collision data at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Physics Letters B*, *796*, 68–87. arXiv:1903.06248 [hep-ex].
Relevance. The strongest public clean-channel constraint on high-mass spin-1/spin-2 resonances. We cite it as the limit that most directly pressures a universally-coupled lock cluster, and to make the point that the GG-6 prediction evades it for a structural reason — the partners are leptophobic and top-rich, with $\mathrm{BR}(Z'_B \to \ell^+\ell^-) = 0$ exactly — not because the dilepton limit is weak. First cited in §30.

ATLAS Collaboration (2024). A new ATLAS for the high-luminosity era.
Official ATLAS upgrade overview. https://atlas.cern/.
Relevance. Documents the Phase-2 tracker, timing, and trigger upgrades that the lock-cluster search depends on. We cite it where the HL-LHC reach and the feasibility of cluster-aware, boosted-top, reduced-stream analyses are argued. First cited in §32.

ATLAS Collaboration (2025). Search for $t\bar{t}$ resonances in final states with one or two leptons using 140 fb^{-1} of pp collision data at $\sqrt{s} = 13$ TeV.
arXiv:2512.17856 [hep-ex].
Relevance. The most relevant existing search for the priority discovery channel of the lock cluster. We cite it as the template that a dedicated, cluster-aware $t\bar{t}$-resonance reinterpretation should extend, since the GG-6 graviton and its color-octet/baryonic partners are dominantly top-rich. First cited in §10.4.

ATLAS Collaboration (2025). ATLAS expands its reach into the high-rate frontier.
Official ATLAS briefing on trigger-level analysis. https://atlas.cern/.
Relevance. Establishes trigger-level (data-scouting) analysis as an available tool. We cite it in the trigger-roadmap section, where a reduced-format stream dedicated to the 3–6 TeV region is proposed as a high-value, low-cost addition for the clustered signal. First cited in §35.

CEPC Study Group (2025). CEPC Technical Design Report — Accelerator; and CEPC Technical Design Report — Reference Detector. IHEP, Beijing, completed October 2025.
Relevance. The mature Chinese Higgs/Z-factory design. Not selected for China's 15th Five-Year Plan (March 2026); resubmission intended around 2030. Status input to the facility landscape of Part IX. First cited in §47.

Cepeda, M., et al. (2019). Report from Working Group 1: Standard Model physics at the HL-LHC and HE-LHC. *CERN Yellow Reports: Monographs*, CERN-2019-007. arXiv:1902.00134.
Relevance. The canonical HL-LHC projection conventions, including the halved-systematics scenario adopted (and decomposed into its statistical, calibration, and theory classes) in §41.1. First cited in §41.1.

CERN (2023). The High-Luminosity LHC: a new horizon for science and technology.
Official CERN HL-LHC flyer. https://home.cern/.
Relevance. The official statement of the $3000\,\mathrm{fb}^{-1}$/experiment HL-LHC target at 14 TeV. We cite it as the dataset that decides the collider falsifier: the spin-1 cluster is either discovered or the current implementation is excluded within this program (the spin-2 anchor, at its derived coupling, awaits a 100-TeV machine; Corollary T-KK-2c). First cited in §32.

CERN (2024). The Compact Linear Collider (CLIC).
Official CERN accelerator overview. https://home.cern/.
Relevance. Cited in the future-collider outlook for the indirect electroweak/Higgs-coupling leg of the pattern test, which remains informative even where direct production of the cluster is inaccessible at staged lepton-collider energies. First cited in §34.

CERN (2024). Muon Collider.
Official CERN accelerator overview. https://home.cern/.
Relevance. Cited for the line-shape and spin diagnostics a multi-TeV muon collider would provide — precisely the deconvolution of the near-degenerate multi-field cluster that a hadron collider cannot fully resolve. First cited in §34.

CERN (2025). LHC delivers a record number of particle collisions in 2025.
Official CERN news item, 9 December 2025. https://home.cern/.
Relevance. The integrated-luminosity backdrop for the "could this year already matter?" question. We cite it for the Run-3 dataset size that places the lock scale inside the currently testable range. First cited in §30.

CERN (2026). Accelerator Report: the 2026 run will be short but intense.
Official CERN news item, 12 March 2026. https://home.cern/.
Relevance. Establishes the compressed 2026 run window before LS3. We cite it where the paper argues there is genuine, if limited, discovery leverage for the lock cluster in the current year. First cited in §30.

CERN (2026). CERN bids farewell to the LHC and enters Long Shutdown 3.
Official CERN press release, 29 June 2026. https://home.cern/.

Relevance. The official close-out of Run 3 (last collisions 14 June 2026; ≈540 fb^{-1} delivered since startup) and the LS3-to-HL-LHC schedule used for the luminosity-on-tape model of Table 21 and Figure 10. First cited in §32.7.

Chivukula, R. S., Farzinnia, A., Foadi, R., & Simmons, E. H. (2012). Production of massive color-octet vector bosons at next-to-leading order. *Physical Review D*, *85*, 054005. arXiv:1111.7261.
Relevance. The complete NLO computation for massive color-octet vector production (corrections up to ∼30%, scale dependence reduced to per-cent level); source of the $K_{q\bar{q}} \simeq 1.3$ used for the g^* rates in Part VI and §31. First cited in §21.

CMS Collaboration (Chatrchyan, S. et al.) (2012). Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Physics Letters B*, *716*, 30–61. arXiv:1207.7235.
Relevance. The simultaneous CMS observation of the Higgs boson. Cited in §1 alongside the ATLAS paper; the author served on CMS during this era (Part XI), watching the Higgs *complete* the Standard Model rather than break it — the experience from which this paper's insistence on pre-registered falsifiability descends. First cited in Part I.

CMS Collaboration (2019). A new precision timing detector for the CMS HL-LHC upgrade (MTD). Official CMS news item. `https://cms.cern/`.
Relevance. Cited where precision timing is promoted from a pile-up tool to a resonance-quality variable for high-mass boosted-hadronic final states — a concrete way to sharpen the cluster line shape. First cited in §34.

CMS Collaboration (2020). High-mass dilepton resonances in pp collisions at $\sqrt{s} = 13$ TeV (CMS-EXO-19-019).
Official CMS public-results summary. `https://cms.cern/`.
Relevance. The CMS counterpart to the ATLAS dilepton limit; cited together with it as the clean-channel constraint that the leptophobic, non-universal lock cluster evades while a generic universal resonance would not. First cited in §30.

CMS Collaboration (2020). Search for high mass dijet resonances with a new background prediction method in proton–proton collisions at $\sqrt{s} = 13$ TeV. *Journal of High Energy Physics*, *05*, 033. arXiv:1911.03947.
Relevance. The 137 fb^{-1} narrow-dijet limit (σBA) whose observed curve appears in Fig. 9; the reference dijet constraint against which the octet and Z'_B hadronic rates are compared in §31. First cited in §31.3.

CMS Collaboration (2024). Search for new physics in high-mass diphoton events at $\sqrt{s} = 13$ TeV (CMS-EXO-22-024).
arXiv:2405.09320 [hep-ex].
Relevance. The diphoton spin-2 search; cited as a clean but, for GG-6, sub-dominant channel: the geometry suppresses light-state couplings, so diphoton is a corroborating rather than a discovery channel for the cluster. First cited in §10.9.

CMS Collaboration (2024). High-mass diphoton search, public-results summary (CMS-EXO-22-024).
Official CMS public-results summary. `https://cms.cern/`.
Relevance. The public-results companion to the diphoton paper, cited in the channel-by-channel exclusion ledger as the traceable diphoton input for a geometry-aware reinterpretation. First cited in §30.

CMS Collaboration (2024). Search for resonant Higgs-boson pair production in the $b\bar{b}b\bar{b}$ final state at $\sqrt{s} = 13$ TeV.
arXiv:2407.13872 [hep-ex].

Relevance. The hh channel that probes the spin-2 graviton's hh decay and the IR-sensitive Higgs sector of the cluster. We cite it as the best scalar-corroborator search and as a template for the boosted-hh leg of a combined-channel GG-6 analysis. First cited in §30.

CMS Collaboration (2024). Model-agnostic search for dijet resonances with anomalous jet substructure at $\sqrt{s} = 13$ TeV.
arXiv:2412.03747 [hep-ex].
Relevance. The boosted-dijet / substructure search most sensitive to the leptophobic, color-octet side of the cluster. We cite it as the channel where the $\mathrm{U}(1)_B$ admixture and KK-gluon partner become visible, and as a natural host for a cluster-aware (multi-pole) template. First cited in §10.4.

CMS Collaboration (2024). The CMS Upgrade Week: knowledge sharing within the collaboration.
Official CMS news item. `https://cms.cern/`.
Relevance. The CMS Phase-2 upgrade program — new tracker, MIP timing detector, high-granularity calorimetry, and the Level-1 scouting stream — as presented in the collaboration's upgrade-week material. Cited for the detector capabilities that the trigger and reduced-stream proposals of this paper (Table 23) build on, and that keep the boosted-top performance assumptions of Part VIII conservative. First cited in §32.

CMS Collaboration (2025). How CMS records all 40 million events per second (Level-1 scouting).
Official CMS news item. `https://cms.cern/`.
Relevance. The CMS data-scouting capability cited in the trigger roadmap as the mechanism that would let a dedicated reduced-format stream capture the clustered high-mass signal without saturating the standard trigger budget. First cited in §35.

Connes, A., & Lott, J. (1991). Particle models and noncommutative geometry. *Nuclear Physics B (Proceedings Supplements)*, *18*(2), 29–47.
Relevance. The noncommutative-geometry Standard Model: gauge and Higgs structure encoded in a genuinely non-spatial internal geometry. Cited in the genealogy of §5.5 as the furthest prior step toward non-spatial internal dimensions — GG-6 differs by making its two internal directions literal bulk coordinates, forced by Theorem 0 of Arisaka (2026, A2-GG6). First cited in §5.5.

Cowan, G. (1998). *Statistical Data Analysis*. Oxford University Press.
Relevance. The standard first statistics text for particle physicists; the recommended companion to the discovery-statistics course of Appendix L and the formalism of Part VIII. First cited in Part XII.

Cowan, G., Cranmer, K., Gross, E., & Vitells, O. (2011). Asymptotic formulae for likelihood-based tests of new physics. *European Physical Journal C*, *71*, 1554. arXiv:1007.1727.
Relevance. The standard machinery of discovery statistics — profile-likelihood test statistics, Wilks asymptotics, the Asimov dataset, and the significance formula with background uncertainty — derived step by step and used for every projection in Part VIII. First cited in §40.

Davoudiasl, H., Hewett, J. L., & Rizzo, T. G. (2000). Phenomenology of the Randall–Sundrum gauge hierarchy model. *Physical Review Letters*, *84*, 2080–2083. arXiv:hep-ph/9909255.
Relevance. The standard normalization of warped-graviton couplings and production cross sections in terms of $k/\bar{M}_{\mathrm{Pl}}$. We use it in *Read-out II only* (§19, assumption B1) as the bulk-graviton template against which $\kappa_{\mathrm{eff}} \approx 1/\sqrt{K}$ is calibrated — an external, model-level input, not a GG-6-A result. First cited in §19.

Dobrescu, B. A., & Yu, F. (2013). Coupling–mass mapping of dijet peak searches. *Physical Review D*, *88*, 035021. arXiv:1306.2629.

Relevance. The standard widths and coupling–mass mapping for leptophobic baryonic Z' states; §31 uses its conventions for the g_B^2 scaling of the Z'_B rates and for the octet width cross-check. First cited in §31.2.

Ellis, R. K., Stirling, W. J., & Webber, B. R. (1996). *QCD and Collider Physics*. Cambridge University Press.
Relevance. The definitive treatment of factorization, parton luminosities, and resonance production at hadron colliders; the primary companion to Appendix K, whose master formulas it derives at full depth. First cited in Part XII.

European Strategy Group (2025). European Strategy for Particle Physics: 2026 update milestone — FCC-ee recommended as the preferred next flagship collider.
Official CERN press release, December 2025. `https://home.cern/`.
Relevance. Fixes the post-LHC calendar assumed in §32.7: FCC-ee as the next machine (indirect at the lock scale), with the first direct probe of 3.85 TeV deferred to FCC-hh. First cited in §32.7.

European Strategy Group (2026). The European Strategy for Particle Physics — 2026 Update. Adopted unanimously by the CERN Council, Budapest, 22 May 2026.
Relevance. Names the FCC-ee the preferred option for the next flagship collider at CERN, with a Council decision expected by 2028, and prioritizes muon-collider R&D. The post-LHC planning baseline of Part IX. First cited in §47.

FCC Collaboration (2018). FCC-hh: the Hadron Collider — Future Circular Collider Conceptual Design Report.
CERN. `https://fcc.web.cern.ch/`.
Relevance. The 100 TeV machine that would turn the lock-cluster corridor into a fully mappable sector: the first cluster, the radial $n = 2$ (~ 7 TeV) and $n = 3$ (~ 10 TeV) families, and the colored partners would all be in reach. Cited in the long-term-prospects section. First cited in §10.9.

FCC Collaboration (2025). Future Circular Collider Feasibility Study Report, Volume 1.
arXiv:2505.00272 [physics.acc-ph].
Relevance. The staged FCC scenario (FCC-ee physics ∼2047+, FCC-hh ∼2074) that sets the right edge of Figure 10 and the date at which $\sigma(G_1) \approx 205$ fb becomes directly testable. First cited in §32.7.

Franceschini, R., et al. (2016). What is the $\gamma\gamma$ resonance at 750 GeV? *Journal of High Energy Physics*, *03*, 144. arXiv:1512.04933.
Relevance. The reference narrow-width master formula and parton-luminosity coefficients used to validate the workbook of §31 (reproduced there to 2–5%), together with its quoted NLO K-factors. First cited in §31.2.

Fukuda, Y. et al. (Super-Kamiokande Collaboration) (1998). Evidence for oscillation of atmospheric neutrinos. *Physical Review Letters*, *81*, 1562–1567.
Relevance. The discovery of atmospheric-neutrino oscillation — the lone established crack in the Standard-Model edifice of the past fifty years, and pointedly *not* a collider result. Cited in §1 as the exception that sharpens the rule: no collider has ever added a particle to the Standard Model's list, and the one confirmed beyond-Standard-Model fact prescribes no collider scale — the vacuum this paper's derived target fills. First cited in Part I.

Goldberger, W. D., & Wise, M. B. (1999). Modulus stabilization with bulk fields. *Physical Review Letters*, *83*, 4922–4925. arXiv:hep-ph/9907447.

Relevance. The canonical stabilization mechanism for the Randall–Sundrum radion: a bulk scalar whose potential fixes the brane separation. Cited in Table 2 as the emblem of what every *spatial* extra dimension owes and what GG-6 does not: the scale direction's length is the integer $K = 35$, so there is no modulus to stabilize and no radion — an absence this paper promotes to a falsifiable Higgs-precision prediction (Part IX). First cited in §5.5.

Gross, E., & Vitells, O. (2010). Trial factors for the look elsewhere effect in high energy physics. *European Physical Journal C*, *70*, 525–530. arXiv:1005.1891.
Relevance. The rigorous (up-crossings) treatment of the look-elsewhere effect; §44.1 uses the resolution-cell estimate it validates to quantify the $\sim$75-fold trials factor that the pre-registered GG-6 search does not pay. First cited in §44.1.

Halzen, F., & Martin, A. D. (1984). *Quarks and Leptons: An Introductory Course in Modern Particle Physics*. Wiley.
Relevance. The classic accessible source for cross-section and decay estimates; the gentlest complete route into the worked examples of Appendix K. First cited in Part XII.

Han, T., Lykken, J. D., & Zhang, R.-J. (1999). On Kaluza–Klein states from large extra dimensions. *Physical Review D*, *59*, 105006. arXiv:hep-ph/9811350.
Relevance. The canonical spin-2 partial-width formulas used for the G_1 rates in §31 (with the massive-vector polynomial taken from Allanach et al., 2000). First cited in §31.2.

Hosotani, Y. (1983). Dynamical mass generation by compact extra dimensions. *Physics Letters B*, *126*(5), 309–313.
Relevance. The Hosotani mechanism: gauge-symmetry breaking by the Wilson-line holonomy of a compact direction. In GG-6 this is not one model-building option among many but the *operative* electroweak mechanism — the phase direction S^1_θ supplies the Wilson line θ_H whose selected value gives $v = M_{\text{lock}}(2/5)^3$ in Arisaka (2026, A9-SM) — so the comparison row of Table 2 credits the mechanism at its source. First cited in §5.5.

Accettura, C., et al. (International Muon Collider Collaboration) (2024). Interim report for the International Muon Collider Collaboration. arXiv:2407.12450.
Relevance. The staged 3 $\to$ 10 TeV design, the ionization-cooling demonstrator roadmap, and the reference schedule (stage-1 near 2050, 10 TeV near 2070) assumed in Part IX; the leptophobia null test of §50.3 is defined against this machine. First cited in §47.

Kaluza, T. (1921). Zum Unitätsproblem der Physik. *Sitzungsberichte der Preußischen Akademie der Wissenschaften zu Berlin*, 966–972.
Relevance. The historical origin of the idea that gauge structure can be read geometrically from an extra coordinate. We cite it to situate the *non-spatial* GG-6 reading: the lineage of "KK" is retained for discoverability even though the GG-6 directions are scale and phase, not space. First cited in §5.

Klein, O. (1926). Quantentheorie und fünfdimensionale Relativitätstheorie. *Zeitschrift für Physik*, *37*, 895–906.
Relevance. Introduced compactification and the discrete momentum tower that gives "Kaluza–Klein" its spectral meaning. We cite it for the standard mechanism by which a compact direction yields a discrete spectrum — the precedent we deliberately contrast with the RG-scale-throat origin of the GG-6 radial cluster. First cited in §5.

Lillie, B., Randall, L., & Wang, L.-T. (2007). The bulk RS KK-gluon at the LHC. *Journal of High Energy Physics*, *0709*, 074. arXiv:hep-ph/0701166.
Relevance. The reference computation of the warped KK-gluon: its top-rich branching pattern, broad width, and light-quark-suppressed production. We borrow its coupling ratios in *Read-out II only* (assumption B4) to convert the GG-6 color-octet structure into the $\sigma \times \mathrm{BR}(t\bar{t})$ ballpark; the result is model-level. First cited in §19.

Maldacena, J. (1998). The large-N limit of superconformal field theories and supergravity. *Advances in Theoretical and Mathematical Physics*, *2*, 231–252. arXiv:hep-th/9711200.
Relevance. The AdS/CFT correspondence: the renormalization-group scale of a boundary theory realized as a bulk radial coordinate — *as a duality*. Cited in §5.5 because GG-6 takes the same identification literally rather than dually: the throat resonances of this paper *are* the composite states of the boundary theory at its infrared scale, which is why the holographic language of "composite resonances" describes the lock cluster exactly. First cited in §5.5.

Mathews, P., & Ravindran, V. (2005). NLO-QCD corrections to dilepton production in the Randall–Sundrum model. arXiv:hep-ph/0506158.
Relevance. NLO K-factors for spin-2 (RS graviton) production at the LHC ($K \simeq 1.5$–1.6); source of the $K_{gg} \simeq 1.6$ applied to the G_1 rates in Part VI. First cited in §21.

Minkowski, H. (1908). Raum und Zeit. *Physikalische Zeitschrift*, *10*, 75–88.
Relevance. The first non-spatial geometric dimension: time, fused into spacetime. Cited in §5.5 as the precedent that legitimizes classifying dimensions by their symmetry generators rather than by spatial intuition — the classification under which GG-6's scale and phase directions are dimensions in exactly Minkowski's sense, and spatial in no sense. First cited in §5.5.

F. W. J. Olver, D. W. Lozier, R. F. Boisvert, and C. W. Clark (eds.). *NIST Handbook of Mathematical Functions.* Cambridge University Press, 2010. Companion online release: the NIST Digital Library of Mathematical Functions, `https://dlmf.nist.gov/`.
Relevance. The standard modern reference for the Bessel functions J_n, Y_n and their zeros x_n, which carry the entire radial spectrum of this paper — the lock mass, the tower ratios x_2/x_1 and x_3/x_1, and the mode profiles of the overlap integrals; cited at the boundary-value course of Appendix J as the recommended companion for a reader meeting the machinery for the first time. First cited in §J.

Particle Data Group (Navas, S., et al.) (2024). Review of Particle Physics. *Physical Review D*, *110*, 030001.
Relevance. The canonical compilation of measured Standard-Model parameters. Cited for the world-average strong coupling $\alpha_s(M_Z) = 0.1180 \pm 0.0009$ against which the GG-6 structural value $1/(\pi e) = 0.11710$ is compared in Part IX — the single number FCC-ee will turn into a decisive test — and as the source of the top mass and electroweak inputs used throughout. First cited in §50.2.

Peskin, M. E., & Schroeder, D. V. (1995). *An Introduction to Quantum Field Theory.* Westview Press.
Relevance. The standard graduate field-theory text; supplies the mode-decomposition and effective-Lagrangian machinery assumed in the KK course of Appendix J. First cited in Part XII.

Pumplin, J., Stump, D. R., Huston, J., Lai, H.-L., Nadolsky, P., & Tung, W.-K. (2002). New generation of parton distributions with uncertainties from global QCD analysis. *Journal of High Energy Physics*, *0207*, 012. arXiv:hep-ph/0201195.
Relevance. The CTEQ6 parton distributions used for the parton luminosities of §31; the CTEQ6M/CTEQ6L1 spread is part of the quoted uncertainty band. First cited in §31.2.

Randall, L., & Sundrum, R. (1999). A large mass hierarchy from a small extra dimension. *Physical Review Letters*, *83*, 3370–3373. arXiv:hep-ph/9905221.
Relevance. The warped two-brane construction whose Bessel-root spectrum and TeV-scale resonances are the closest phenomenological analogue of the GG-6 lock cluster. We cite it throughout the scorecard: GG-6 borrows the warped-throat discreteness but differs decisively in that its infrared scale is *derived* (the integer $K = 35$) rather than a fitted curvature, and its cluster is multi-field rather than a single graviton pole. First cited in §5.

Thomson, M. (2013). *Modern Particle Physics*. Cambridge University Press.
Relevance. The experiment-aware modern survey; the recommended first pass over every sector touched by the appendix course, written from the detector side of the field — the perspective this paper's Preface addresses. First cited in Part XII.

Weyl, H. (1918). Gravitation und Elektrizität. *Sitzungsberichte der Preussischen Akademie der Wissenschaften*, 465–480.
Relevance. Weyl's original *Eichinvarianz*: the first geometrization of *scale* as a connection. It failed as a theory of electromagnetism (Einstein's second-clock objection), but it is the ur-precedent for the GG-6 scale dimension I_r — the 1918 half of the genealogy traced in §5.5, completed when GG-6 promotes scale to a forced, non-spatial dimension of the bulk. First cited in §5.5.

Weyl, H. (1929). Elektron und Gravitation. I. *Zeitschrift für Physik*, *56*, 330–352.
Relevance. The birth of modern gauge theory: Weyl's reinterpretation of his 1918 idea as a local *phase* ($U(1)$) invariance. The GG-6 phase dimension S^1_θ is this insight taken literally — the $U(1)$ fiber promoted to a coordinate — making Weyl the single ancestor of both non-spatial directions (§5.5). First cited in §5.5.