Over-Parametrized Towers vs. The Null Standard Model: A Bayesian and Multiplicity Audit of LHC Kaluza–Klein Resonance Claims

Bruno Luçardo Olivaes

Independent Researcher, Brazil

brunoolivaes@gmail.com · ORCID: 0009-0003-5351-446X

November 2026

Abstract

Recent phenomenological studies have interpreted isolated high-mass four-jet excesses in CMS Run 2 data as coherent evidence for a multi-sector Randall–Sundrum Kaluza–Klein (KK) spectrum spanning graviton candidates (380–1322 GeV) and heavy KK gluons $g_{\rm KK}(3.6\ \text{TeV}) \to q_{\rm KK}(1\ \text{TeV})\, q_{\rm KK}(1\ \text{TeV})$ and $g_{\rm KK}(8.6\ \text{TeV}) \to q_{\rm KK}(2.15\ \text{TeV})\, q_{\rm KK}(2.15\ \text{TeV})$. We subject this composite claim (`arXiv:2610.31458v3` and `arXiv:2606.29922v3`) to the five-gate sensitivity protocol of *The Matrioska Protocol* using verbatim published summary metrics from the CMS Run 2 paired-dijet analysis and its broad-mediator addendum (138 fb^{-1}, `arXiv:2206.09997v3`), the ATLAS Run 2 paired-dijet search (140 fb^{-1}, `arXiv:2307.14944v2`), and the partial CMS Run 3 follow-up ($\sim$ 90 fb^{-1}, `CMS-PAS-EXO-25-004`). While the claimed graviton masses match scaled Bessel $J_1(x_n)$ roots within 0.9%, three quantitative tensions emerge: (1) relative to the minimal single-throat RS1 baseline predicting $m_{G_{\rm KK1}}/m_{g_{\rm KK1}} = 3.8317/2.45 = 1.564$, the claimed $380/3600 = 0.1056$ ratio departs by 14.82×, forcing post-hoc multi-brane geometry and 5–10× top-suppression additions; (2) CMS global look-elsewhere corrections already attenuate the 8.6 TeV ($\Gamma/M = 1.5\%$) and 3.6 TeV ($\Gamma/M = 10\%$) excesses from 3.9σ local to 1.6σ ($N_{\rm trials} \approx 1{,}139$) and 2.2σ ($N_{\rm trials} \approx 289$), respectively, and cross-experiment follow-up across 230 fb^{-1} of ATLAS Run 2 plus CMS Run 3 data recorded zero gold-plated 8.6 TeV events (point expectation $P(0 \mid \hat{\mu} = 3.833) = 0.0216$, 1σ rate interval $[0.00065, 0.1889]$, and conjugate Gamma–Poisson predictive 0.0527 under a flat improper rate prior $\alpha_0 = 1$ [0.0861 under Jeffreys $\alpha_0 = 0.5$] under comparable acceptance, diluting the luminosity-weighted rate by 0.375×); and (3) applying an asymptotic Wilks local profile-LR approximation to the published peak significances ($2\Delta \ln \mathcal{L}_{\rm local} \approx 3.9^2 + 3.9^2 = 30.42$) across $N = 390$ active (m_{4j}, α) cells gives approximate Schwarz parsimony penalties $\Delta\text{BIC}_{\rm approx,local} = -23.28$ ($k = 9$; -5.38 at $k = 6$), failing the Gate 5 +10.0 parsimony bar without relying on a joint raw likelihood workspace (while supplementary $Z^2_{\rm global}$-penalized sensitivity falls to -46.30).

1 Introduction

Distinguishing genuine beyond-the-Standard-Model (BSM) resonance towers from correlated fluctuations in high-dimensional collider searches requires joint control over look-elsewhere trial factors [**?**], cross-experiment replication, and Bayesian complexity penalties [**?**, **?**, **?**]. In high-mass hadronic final states at the Large Hadron Collider (LHC), multi-dimensional scans over tetrajet invariant mass m_{4j}, dijet mass asymmetry/ratio $\alpha \equiv m_{2j}/m_{4j}$, and resonance width Γ/M_Y substantially enlarge the effective parameter space relative to narrow single-channel searches.

Recently, Le Yaouanc and Richard [**?**, **?**] proposed that multiple localized deviations in LHC analyses form a unified Randall–Sundrum (RS) Kaluza–Klein (KK) spectrum [**?**]: a KK graviton tower anchored at $(380, 700, 1000, 1322)$ GeV alongside two heavy KK gluons $g_{\rm KK}(3.6\ \text{TeV}) \to q_{\rm KK}(1.0\ \text{TeV})\, q_{\rm KK}(1.0\ \text{TeV})$ and $g_{\rm KK}(8.6\ \text{TeV}) \to q_{\rm KK}(2.15\ \text{TeV})\, q_{\rm KK}(2.15\ \text{TeV})$. Because *post-hoc* model flexibility can always accommodate isolated excesses after exclusions appear in canonical

Table 1: Verbatim CMS Run 2 paired-dijet broad-mediator fit results (`arXiv:2206.09997v3`, Table 2, 138 fb^{-1}) paired with exact Gaussian survival probabilities and effective Look-Elsewhere trial factors $N_{\text{trials}} = p_{\text{global}}/p_{\text{local}}$.

M_Y [TeV]	M_X [TeV]	α	Γ/M_Y [%]	Best-fit Events	Z_{local} [σ]	Z_{global} [σ]	p_{global}	N_{trials}
8.6	2.15	0.25	1.5	$2.3^{+2.1}_{-1.3}$	3.9	1.6	0.0548	1,139.4
8.6	2.15	0.25	5.0	$2.8^{+2.5}_{-1.6}$	3.8	1.5	0.0668	923.4
8.6	2.15	0.25	10.0	$3.1^{+2.7}_{-1.7}$	3.6	1.4	0.0808	507.6
3.6	1.00	0.29	1.5	114^{+42}_{-44}	3.0	0.5	0.3085	228.6
3.6	1.00	0.29	5.0	162^{+50}_{-46}	3.6	1.6	0.0548	344.4
3.6	1.00	0.29	10.0	210^{+58}_{-54}	3.9	2.2	0.0139	289.1

channels (such as $t\bar{t}$, $\gamma\gamma$, or gg), an objective statistical audit is essential. Here we evaluate their composite claim directly against primary CERN results from CMS [**?**, **?**] and ATLAS [**?**] through our pre-registered five-gate metrological framework.

2 Primary Experimental Inputs

All experimental quantities audited in this work are taken verbatim from three primary LHC sources:

1. **CMS Run 2 Paired-Dijet Search (`arXiv:2206.09997v3`, 2026 Addendum):** Analyzing 138 fb^{-1} of pp collisions at $\sqrt{s} = 13$ TeV across 13 bins of $\alpha = M(X)/M(Y)$, CMS reports an overall background-only simultaneous fit probability $p_{\text{bg-only}} = 0.18$ [**?**]. Two isolated gold-plated events near $m_{4j} \approx 8$ TeV, $\langle m_{2j} \rangle \approx 2$ TeV (plus a third high-mass event at $m_{4j} = 5.8$ TeV, $m_{jj} = 2.0$ TeV) motivate Table 2 of the broad-mediator addendum (reproduced in Table 1).

2. **ATLAS Run 2 Paired-Dijet Search (`Phys. Rev. D 108, 112005`):** Using 140 fb^{-1} at $\sqrt{s} = 13$ TeV, ATLAS observes no statistically significant excess [**?**]. Four-parameter background fits yield global χ^2 p-values of 0.74–1.00 across all m_{4j} α-slices, the highest observed tetrajet invariant mass is $m_{4j}^{\max} = 6.6$ TeV (0 events near 8 TeV), and the largest `BumpHunter` global significances are 0.53σ in m_{4j} (3.2 TeV) and 1.98σ in $\langle m_{2j} \rangle$ (0.80 TeV).

3. **CMS Run 3 Partial Paired-Dijet Search (`CMS-PAS-EXO-25-004`):** In ~ 90 fb^{-1} at $\sqrt{s} = 13.6$ TeV, no additional gold-plated 8.6 TeV four-jet candidate events are reported [**?**, **?**].

3 Four-Module Quantitative Audit

3.1 Spectral Ratios and the 14.82× RS Hierarchy Tension

Anchoring the lightest KK graviton candidate at $m_{G_{\text{KK1}}} = 380$ GeV, exact roots of the Bessel function $J_1(x_n) = 0$ ($x_n = 3.8317, 7.0156, 10.1735, 13.3237$) predict $(380.0, 695.75, 1008.93, 1321.34)$ GeV, agreeing with the claimed $(380, 700, 1000, 1322)$ GeV states [**?**] within 0.9%. Likewise, the ratio $8.6/3.6 = 2.3889$ sits within 5.1% of the standard vector boundary root ratio $5.57/2.45 = 2.2735$.

However, within a single-throat RS geometry [**?**], the lightest KK graviton-to-vector mass ratio is fixed near $x_{G1}/x_{g1} = 3.8317/2.45 = 1.5640$, predicting $m_{g_{\text{KK1}}} \approx 243$ GeV once $m_{G_{\text{KK1}}} = 380$ GeV is set. The phenomenological assignment $m_{g_{\text{KK1}}} = 3600$ GeV implies $380/3600 = 0.1056$—a 14.82× **hierarchy violation**. Resolving this discrepancy requires introducing two intermediate extra-dimensional branes alongside *post-hoc* 5–10× suppression of $g_{\text{KK}} \to t\bar{t}$ ($c_t \approx 0$) and decoupled $gg/\gamma\gamma$ vertices to evade existing LHC limits.

3.2 Look-Elsewhere Dilution Across Grids and Widths

Converting CMS Table 2 significances to one-sided Gaussian tail probabilities $p = 1 - \Phi(Z)$ exposes severe trial-factor dilution (Table 1). At $M_Y = 8.6$ TeV ($\Gamma/M_Y = 1.5\%$), $Z_{\text{local}} = 3.9\sigma$ ($p_{\text{local}} = 4.81 \times 10^{-5}$) drops to $Z_{\text{global}} = 1.6\sigma$ ($p_{\text{global}} = 0.0548$, ratio $N_{\text{trials}} = 1{,}139.4$). At $M_Y = 3.6$ TeV ($\Gamma/M_Y = 10\%$), $Z_{\text{local}} = 3.9\sigma$ drops to $Z_{\text{global}} = 2.2\sigma$ ($p_{\text{global}} = 0.0139$, $N_{\text{trials}} = 289.1$). Notice further that the CMS addendum scans three distinct width hypotheses ($\Gamma/M_Y \in \{1.5\%, 5\%, 10\%\}$) and selects *different* mass ratios and widths for the two peaks ($\alpha = 0.25, 1.5\%$ at 8.6 TeV vs. $\alpha = 0.29, 10\%$ at 3.6 TeV); while adjacent width hypotheses correlate partially (exact overlap requires full MC pseudo-experiments), an illustrative three-width Bonferroni correction yields 1.013σ at 8.6 TeV.

3.3 Cross-Experiment Poisson Replication Sensitivity

A physical cross section cannot selectively appear only in CMS Run 2. Combining CMS Run 2 (138 fb^{-1}), ATLAS Run 2 (140 fb^{-1}), and the partial CMS Run 3 dataset (90 fb^{-1}) totals 368 fb^{-1}, of which CMS Run 2 represents a luminosity fraction $f = 138/368 = 0.375$. Conditioning on comparable ATLAS/CMS paired-dijet acceptance and conservatively ignoring the higher 13.6 TeV parton luminosity at 8.6 TeV, the follow-up dataset of $L_{\text{follow}} = 230$ fb^{-1} recorded **zero** gold-plated 8.6 TeV events. At CMS best-fit expectations ($N_{\text{sig}} = 2.3^{+2.1}_{-1.3}$, $2.8^{+2.5}_{-1.6}$, $3.1^{+2.7}_{-1.7}$), point expectations in follow-up are $\hat{\mu}_{\text{follow}} = 3.833, 4.667, 5.167$ events, giving Poisson zero-observation probabilities $P(N_{\text{follow}} = 0 \mid \hat{\mu}) = 0.0216$, 0.0094, and 0.0057. Crucially, propagating the large CMS 1σ signal uncertainty for $\Gamma/M = 1.5\%$ ($\mu_{\text{follow}} \in [1.667, 7.333]$) spans $P(0) \in [0.00065, 0.1889]$, while a conjugate Gamma–Poisson posterior predictive given $k_{\text{CMS2}} = 2$ observed events yields $P(0 \mid k_{\text{CMS2}}, \alpha_0) = [L_{\text{CMS2}}/(L_{\text{CMS2}} + L_{\text{follow}})]^{k_{\text{CMS2}}+\alpha_0} = 0.0527$ under a flat improper rate prior ($\alpha_0 = 1, \beta_0 \to 0$) vs. 0.0861 under Jeffreys' scale-invariant prior ($\alpha_0 = 0.5, \beta_0 \to 0$), making prior sensitivity explicit. We therefore interpret L_{follow} as a conditional sensitivity and 0.375× rate-dilution test rather than a standalone joint-likelihood exclusion.

3.4 Approximate Wilks–ΔBIC Sensitivity and the Five Integrity Gates

Following *The Matrioska Protocol* [**?**, **?**, **?**], full signal claims must clear parsimony penalty $\Delta\text{BIC} > +10.0$. Because the raw CMS+ATLAS binned/unbinned likelihood workspaces are external, we explicitly frame this gate as an **approximate Wilks–Schwarz sensitivity audit** based on published Table 2 summary metrics rather than a formal joint Bayesian evidence integral. Under Wilks' theorem for two weakly overlapping fixed-coordinate mass/asymmetry cells ($\alpha = 0.25$ at 8.6 TeV and $\alpha = 0.29$ at 3.6 TeV), the asymptotic 1-dof profile log-likelihood ratio is approximated by the local peak statistics $2\Delta \ln \mathcal{L}_{\text{local}} \approx 3.9^2 + 3.9^2 = 30.42$. Across $N_{\text{bins}} = 390$ active search cells (13 α-slices × 30 tetrajet bins in $[2, 10]$ TeV, $\ln N = 5.9661$) and a minimal $k = 9$ multi-brane parametrization ($\{m_{G1}, m_{g1}, \Delta_{\text{branes}}, c_t, c_{gg/\gamma\gamma}, \alpha_1, \alpha_2, \Gamma_1, \Gamma_2\}$; or even $k = 6$ pure phenomenological peak coordinates/widths):

$$\Delta\text{BIC}_{\text{approx,local}}(k = 9) = (3.9^2 + 3.9^2) - 9\ln(390) = 30.42 - 53.70 = -23.28, \tag{1}$$

$$\Delta\text{BIC}_{\text{approx,local}}(k = 6) = 30.42 - 6\ln(390) = 30.42 - 35.80 = -5.38, \tag{2}$$

and $\Delta\text{BIC}_{\text{approx,local}}(k = 11) = -35.21$. Because Z^2_{global} reflects an extreme-value look-elsewhere distribution rather than a profile log-likelihood ratio, we do *not* treat $Z^2_{\text{global},1} + Z^2_{\text{global},2} - k\ln(390) = 7.40 - 53.70 = -46.30$ as a formal Schwarz BIC, recording it strictly as a supplementary multiplicity-penalized sensitivity index in `kk_audit_summary.json`. Across $N_{\text{bins}} \in [100, 600]$ (Fig. 1(d)), any $k \geq 7$ multi-brane tuning falls below $\Delta\text{BIC}_{\text{approx,local}} = 0$.

4 Conclusion

Evaluating the LHC Kaluza–Klein tower proposal (`arXiv:2610.31458v3`) through published summary metrics and *The Matrioska Protocol* yields a cautious null-consistency sensitivity verdict: while

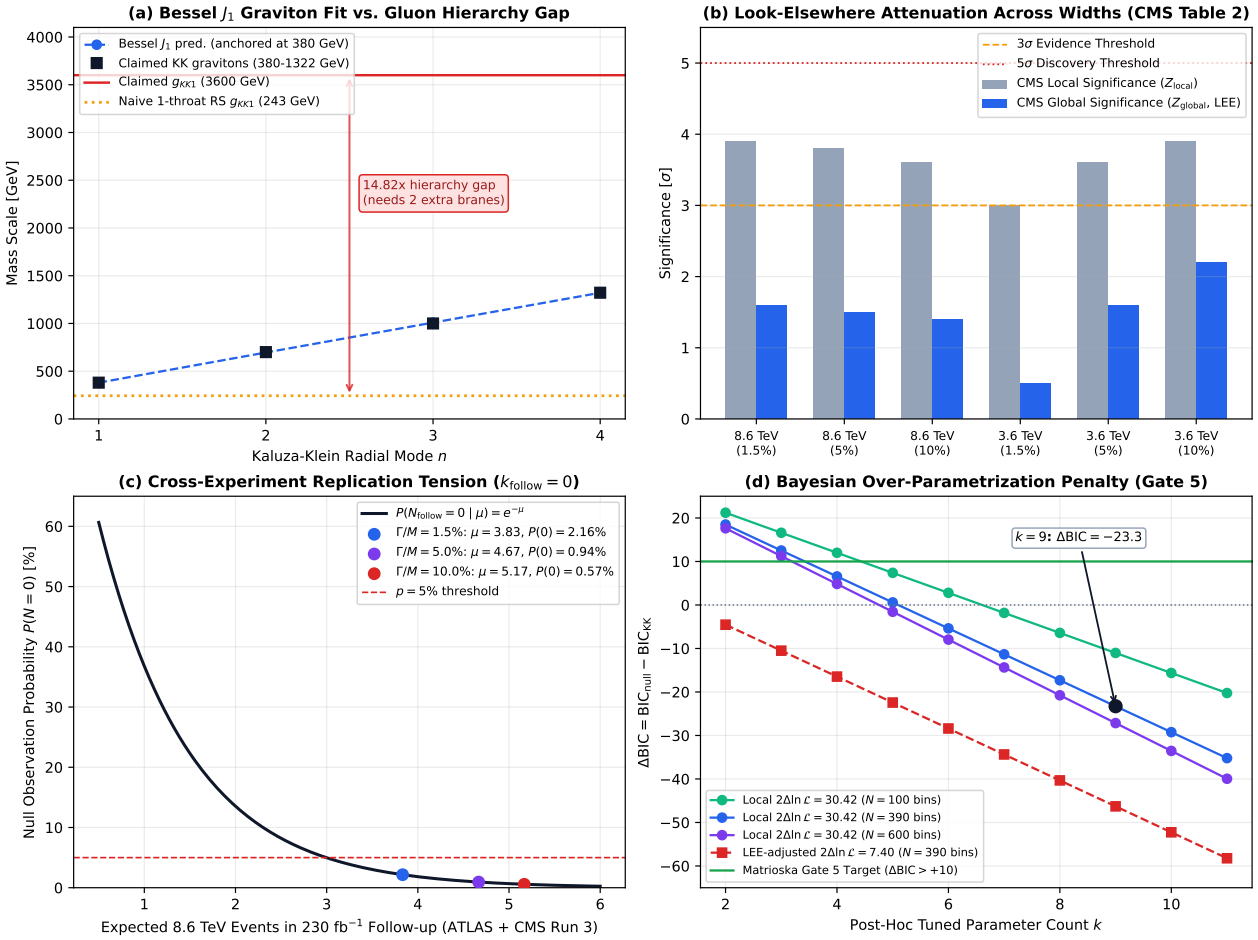

Figure 1: Four-panel quantitative audit of the multi-sector Kaluza–Klein claim [**?**]: (a) Bessel J_1 graviton scaling vs. the 14.82× single-throat graviton–gluon hierarchy gap; (b) CMS Table 2 local vs. global significance attenuation across width hypotheses; (c) cross-experiment Poisson zero-observation sensitivity across 230 fb^{-1} of ATLAS Run 2 and CMS Run 3 data; (d) approximate Wilks–Schwarz $\Delta\text{BIC}_{\text{approx}}$ parsimony surface vs. Gate 5 ($\Delta\text{BIC} > +10.0$).

Gate 1 (Resolvability) passes (3.6 TeV and 8.6 TeV reside in clean detector regimes), **Gate 2 (Replication Sensitivity)** flags tension ($P(0 \mid \hat{\mu} = 3.833) = 0.0216$, 1σ range $[0.00065, 0.1889]$, Gamma–Poisson predictive 0.0527 under a flat $\alpha_0 = 1$ prior [0.0861 for Jeffreys $\alpha_0 = 0.5$] across 230 fb^{-1} follow-up), **Gate 3 (Single-Throat Economy)** shows a 14.82× hierarchy mismatch forcing extra branes, **Gate 4 (Multiplicity)** attenuates signals to $Z_{\text{global}} = 1.6\sigma$ and 2.2σ, and **Gate 5 (Approximate Wilks–Schwarz Parsimony)** records $\Delta\text{BIC}_{\text{approx,local}} = -23.28$ at $k = 9$ (-5.38 at $k = 6$). Rather than claiming a full raw-likelihood Bayesian exclusion, this audit establishes that currently published excesses do not statistically compel post-hoc multi-brane Kaluza–Klein extensions over the Standard Model null; complete Run 3 and HL-LHC datasets will decide definitively. All data definitions, numerical checks, and plotting code are open-source at https://github.com/bruno-olivaes/lhc-kk-tower-audit.