

Article

Not peer-reviewed version

Minimal Finite Hankel Closure in Elastic Proton-Proton Scattering: Structural Invariants, Covariance-Aware Public-Data Benchmarks, and Preregistered LHC Tests

[Mohamed Sacha](#)*

Posted Date: 8 April 2026

doi: 10.20944/preprints202604.0541.v1

Keywords: preregistered forecast; 13.6 TeV prediction; 14 TeV forecast; elastic proton-proton scattering; dip-bump structure; forward physics; TOTEM; model discrimination; impact-parameter amplitude; Odderon-sensitive region; covariance-aware comparison; Regge-eikonal benchmark



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Article

Minimal Finite Hankel Closure in Elastic Proton-Proton Scattering: Structural Invariants, Covariance-Aware Public-Data Benchmarks, and Preregistered LHC Tests

Mohamed Sacha

Independent Researcher, Casablanca, Morocco; www.sachamed@gmail.com

Abstract

We formulate and test a minimal finite Hankel closure for the dip–bump structure in elastic proton–proton scattering. The scope of the claim is deliberately precise. We do not present a microscopic derivation from QCD and we do not claim universal exclusion of the full hadronic phenomenology. Rather, we establish a conditional theorem, confront its surrogate realization with public data, and state explicit near-term tests. First, assuming positivity, radially, finite moments, a Born-dominant forward-plus-first-dip window, self-similar scaling, and minimal finite Hankel closure with one simple node, we prove that the unique lowest-complexity sector is the first Laguerre deformation, which yields a polynomial-times-Gaussian amplitude. Second, we derive structural relations for the forward slope, dip scale, forward curvature, and the drift observable $\mathcal{O}_{\text{excl}} = \Delta[B_0|t|_{\text{dip}}]$, and we prove non-reducibility against the one-scale geometric class for which the corresponding invariants are energy independent. Third, we test the closure on two levels of public-data benchmark. In the restricted internal comparison on 83 differential-cross-section points at 2.76 and 13 TeV under a common weighted log-space score and shared cross-energy flow, the two-scale copy-time surrogate yields $\chi^2_{\text{log}} = 461.19$, AIC = 487.19, and BIC = 518.64, compared with (60493.49, 60507.49, 60524.43) and (59942.77, 59964.77, 59991.38) for two one-scale baselines. We then perform a stronger covariance-aware benchmark in log space, using per-dataset block covariances built from the published statistical errors together with fully correlated systematic blocks, and compare the copy-time surrogate to the internal one-scale baseline, a canonical Regge-pole-plus-Odderon amplitude, a canonical complex Regge-eikonal baseline, and the fixed Kohara–Ferreira–Kodama parametrization. In that stronger test the covariance-aware copy-time fit remains the best model in the benchmark set, with $\chi^2_{\text{cov log}} = 1687.73$, compared with 5440.69 for the strongest external baseline. At fixed 13 TeV, however, the split-sector one-scale surrogate remains competitive in the dedicated local fit, showing that the main empirical leverage of the closure is intrinsically multi-energy rather than a consequence of the 13 TeV line shape alone. We also report a hold-out validation at 8 TeV, an explicit continuation to 13.6 and 14 TeV, and a narrow-window robustness scan showing that the forecast sign pattern is stable under moderate perturbations of the real-sector continuation. Within the explicit axiomatic, statistical, and benchmark choices adopted here, the closure is therefore mathematically constrained, experimentally discriminating, and favored over the implemented internal and external baselines.

Keywords: preregistered forecast; 13.6 TeV prediction; 14 TeV forecast; elastic proton-proton scattering; dip-bump structure; forward physics; TOTEM; model discrimination; impact-parameter amplitude; Odderon-sensitive region; covariance-aware comparison; Regge-eikonal benchmark

1. Introduction

Elastic hadron scattering at collider energies packages nonperturbative dynamics into a small set of unusually sharp observables: the total cross section, the forward real-to-imaginary ratio ρ , the

logarithmic slope of the forward cone, the elastic fraction, and the dip-bump structure of the differential cross section. Public measurements by TOTEM and D0 provide precise anchors at $\sqrt{s} = 2.76, 8,$ and 13 TeV for pp , together with the $p\bar{p}$ comparison at 1.96 TeV that makes crossing-odd effects visible in the dip region [1–5]. The static electromagnetic reference scale is fixed by the CODATA/NIST proton charge radius, $r_p = 0.84075(64)$ fm [6]. The purpose of the present work is to determine how far a copy-time-motivated closure can be pushed against these public observables while keeping the logical status of each statement explicit.

The guiding change relative to earlier versions is one of discipline. The manuscript does not argue that elastic hadron scattering has already been derived microscopically from the copy-time program, and it does not claim to have subsumed the full Regge/eikonal/Odderon phenomenology under a single definitive benchmark. Instead, it isolates three statements that can presently be defended. First, within an explicit axiom set one can classify the lowest-complexity finite Hankel-closed single-node sector and obtain a polynomial-times-Gaussian amplitude. Second, within that sector one obtains rigid combinations of forward and dip observables, including a drift variable that vanishes throughout the one-scale geometric class. Third, one can confront nested surrogate realizations of this closure with public data under one common statistical chain, extend the comparison to representative external baselines, and preregister a near-term continuation to 13.6 and 14 TeV.

This framing separates three logically distinct layers that are often conflated in ambitious phenomenological proposals. The theorem-level content lives inside the stated closure axioms. The benchmark-level content lives inside the surrogate family and the adopted score. The broader phenomenological significance depends on how the resulting observables fare against external alternatives and against future data. Keeping these layers separate is not merely stylistic; it is essential for making a strong claim without overstating it.

A second organizing principle is reproducibility by construction. Every quantitative table in this paper is tied to one canonical script or one machine-readable generated file in the submission package, and every forecast statement is written so that it can be checked directly against future elastic-scattering measurements. The package accompanying the manuscript has been reorganized accordingly into separate manuscript, supplement, code, data, figure, and graphical-abstract components.

2. Conventions and Forward-Moment Identities

Let $x \equiv |t| = q^2$ and define the Born-level rotationally symmetric amplitude by

$$T(x, s) = 2\pi \int_0^\infty b db J_0(qb) \Omega(b, s), \quad (1)$$

where $\Omega(b, s)$ is a radial opacity profile. For the forward expansion we introduce the moments

$$M_{2n}(s) = 2\pi \int_0^\infty b^{2n+1} \Omega(b, s) db. \quad (2)$$

Using $J_0(qb) = 1 - q^2 b^2/4 + q^4 b^4/64 + O(q^6)$, one finds

$$T(x, s) = M_0(s) - \frac{x}{4} M_2(s) + \frac{x^2}{64} M_4(s) + O(x^3). \quad (3)$$

When the forward region is dominantly imaginary, the logarithmic slope of the differential cross section satisfies

$$B_0(s) = \frac{M_2(s)}{2M_0(s)}, \quad (4)$$

and the forward-curvature diagnostic

$$\kappa_0(s) \equiv - \frac{d^2}{dx^2} \ln \frac{d\sigma}{dt} \Big|_{x=0} \quad (5)$$

obeys

$$\kappa_0(s) = \frac{2M_2(s)^2 - M_0(s)M_4(s)}{16M_0(s)^2}. \quad (6)$$

Equations (4) and (6) are model-light identities. They connect forward observables directly to impact-parameter moments and provide the diagnostics used in the benchmark and forecast sections.

3. Minimal Closure Theorem and Interpretive Status

Axiom 1 (positive radial copy-time field). *For each s there exists a nonnegative radial copy-time profile $\tau(b, s) \geq 0$ with finite moments up to fourth order.*

Axiom 2 (Born-dominant response window). *There exists an experimentally relevant window covering the forward cone and the first dip in which the leading opacity response is linear in the copy-time field,*

$$\Omega(b, s) = c(s) \tau(b, s), \quad c(s) > 0. \quad (7)$$

Axiom 3 (self-similar scaling). *The copy-time profile depends on energy through a width parameter $B(s) > 0$ and a dimensionless shape function of the scaled variable $v = b^2 / [2B(s)]$.*

Axiom 4 (minimal finite Hankel-closed single-node sector). *Among all finite-dimensional sectors closed under the radial Fourier–Bessel transform and compatible with Axioms 1–3, select the minimal one that can generate exactly one simple node in the imaginary part of the amplitude within the first dip window.*

Theorem 1 (minimal-closure classification theorem). *Under Axioms 1–4, the minimal admissible copy-time profile is, up to normalization and a single deformation parameter $\eta(s)$,*

$$\tau(b, s) = \tau_0(s) e^{-v} [1 - \eta(s) L_1(v)], \quad L_1(v) = 1 - v, \quad (8)$$

with $0 \leq \eta(s) \leq 1$. Its Born image is necessarily

$$T_I(x, s) = \mathcal{N}(s) e^{-B(s)x/2} \left[1 - \frac{\eta(s)B(s)}{2} x \right], \quad (9)$$

where $\mathcal{N}(s) = 2\pi c(s) \tau_0(s) B(s)$. No lower-dimensional finite Hankel-closed single-node closure exists.

Proof. Under Axiom 3, the natural radial basis is $e^{-v} L_n(v)$. The $n = 0$ sector is one-dimensional and yields a pure Gaussian transform, hence no node. Axiom 4 therefore excludes every one-dimensional closure. The smallest finite-dimensional sector capable of one simple node is the two-dimensional span of $e^{-v} L_0(v)$ and $e^{-v} L_1(v)$. Up to normalization, every element of this sector can be written as Eq. (8). Positivity requires $0 \leq \eta \leq 1$, since $1 - \eta + \eta v \geq 0$ for all $v \geq 0$ only in that interval. The transform identities in Appendix A map the two basis elements to $e^{-Bx/2}$ and $xe^{-Bx/2}$, which yields Eq. (9). Any larger closure violates minimality. $\square$

Remark 1 (scope of the theorem). *Theorem 1 is a classification theorem within the stated axioms. In particular, Axiom 4 is a theory-selection postulate rather than a consequence derived from microscopic QCD. The theorem therefore does not show that hadronic scattering unconditionally forces the copy-time closure. What it does show is sharper and still nontrivial: once one requests positivity, radially, finite moments, a Born-dominant forward-plus-first-dip window, and the minimal finite Hankel-closed single-node sector, the first Laguerre deformation is the unique lowest-complexity outcome.*

4. Restricted Structural Consequences and One-Scale Non-Reducibility

Theorem 2 (structural-constraint theorem). *Let the minimal profile of Theorem 1 hold in the imaginary sector and assume the linear Born response (7). Then*

$$B_{0,I}(s) = B(s)[1 + \eta(s)], \quad (10)$$

$$|t|_{I\text{-zero}}(s) = \frac{2}{\eta(s)B(s)}, \quad (11)$$

$$I_1(s) \equiv B_{0,I}(s)|t|_{I\text{-zero}}(s) = 2\frac{1 + \eta(s)}{\eta(s)}, \quad (12)$$

$$I_2(s) \equiv \kappa_{0,I}(s)|t|_{I\text{-zero}}(s)^2 = 2. \quad (13)$$

Proof. Differentiating $\ln |T_I|^2$ at $x = 0$ gives Eq. (10); the simple zero of the bracket in Eq. (9) yields Eq. (11); their product gives I_1 , and a second derivative gives $I_2 = 2$. $\square$

Theorem 3 (restricted non-reducibility theorem). *Consider the class $\mathfrak{G}$ of one-scale geometric amplitudes*

$$T(x, s) = A(s)F(\Lambda(s)x), \quad (14)$$

with energy-independent analytic shape F and arbitrary positive scalings A, Λ . Then the invariants

$$I_1(s) = B_0(s)|t|_{\text{dip}}(s), \quad I_2(s) = \kappa_0(s)|t|_{\text{dip}}(s)^2 \quad (15)$$

are constant in energy throughout $\mathfrak{G}$. In contrast, the copy-time closure gives Eq. (13) and an I_1 that runs with $\eta(s)$, so any nontrivial drift of I_1 rules out reduction to the class $\mathfrak{G}$.

Proof. For Eq. (14), the forward derivatives depend only on derivatives of F at the origin, while the dip position is $|t|_{\text{dip}} = \xi_*/\Lambda(s)$ where ξ_* is the first zero of F . Hence I_1 and I_2 are shape constants. In the copy-time closure, I_1 depends explicitly on $\eta(s)$. $\square$

The practical discriminant is therefore not a universal statement against every hadronic construction, but the more targeted observable

$$\mathcal{O}_{\text{excl}}(s_i, s_j) = B_0(s_i)|t|_{\text{dip}}(s_i) - B_0(s_j)|t|_{\text{dip}}(s_j). \quad (16)$$

For the class $\mathfrak{G}$, $\mathcal{O}_{\text{excl}} = 0$ identically. In the minimal closure,

$$\mathcal{O}_{\text{excl}}(s_i, s_j) = 2 \left[\frac{1 + \eta(s_i)}{\eta(s_i)} - \frac{1 + \eta(s_j)}{\eta(s_j)} \right], \quad (17)$$

which is generically nonzero if η drifts with energy. The theorem is therefore strong against one-scale geometry and silent about more general multi-scale or Regge/eikonal constructions unless they are benchmarked explicitly.

5. Closed Benchmark Design, Data Coverage, and Score Definition

5.1. Data Coverage and Benchmark Logic

The closed pointwise benchmark uses 63 tabulated points at 2.76 TeV and 20 public points at 13 TeV, for a total of 83 differential-cross-section points in the dip-relevant range. The choice is pragmatic rather than exhaustive. These two data sets can be treated under a common pointwise workflow with the same weighted log-space score and the same cross-energy flow. The 8 TeV TOTEM measurement is still used, but only through the public characterization of the dip and bump positions and the bump-to-dip ratio [2]; it therefore functions as a hold-out summary validation rather than as a third point in the closed pointwise fit. The D0–TOTEM $p\bar{p}/pp$ comparison is cited for physics context,

but it is not inserted into the closed benchmark because the present surrogate family is crossing even and because a like-for-like inclusion of the crossing-odd sector would require an additional treatment layer beyond the present paper.

This explicit narrowing is a methodological choice, not a rhetorical one. It prevents the benchmark from claiming to answer a broader phenomenological question than the implemented comparison can sustain.

5.2. Nested Surrogate Family and Common Score

The public-data comparison is carried out on three nested surrogates. Writing $x = |t|$ and $\ell = \ln(\sqrt{s}/7 \text{ TeV})$,

$$M_1 : \quad \frac{d\sigma}{dt} = A(\ell)e^{-B(\ell)x} \left[(1 - c(\ell)x)^2 + r \right], \quad (18)$$

$$M_2 : \quad \frac{d\sigma}{dt} = A(\ell)e^{-B(\ell)x} \left[(1 - c(\ell)x)^2 + r(\ell)(d - e(\ell)x)^2 \right], \quad (19)$$

$$M_3 : \quad \frac{d\sigma}{dt} = A_I(\ell)e^{-B_I(\ell)x}(1 - c(\ell)x)^2 + A_R(\ell)e^{-B_R(\ell)x}(d - e(\ell)x)^2. \quad (20)$$

The first two are one-scale internal baselines; the third is the two-scale copy-time surrogate. In the global fit the energy dependence is constrained by affine or affine-plus-quadratic flows in ℓ , exactly as implemented in the canonical reproduction scripts bundled with the submission package.

The score used in the closed benchmark is a weighted log-space quadratic loss,

$$\chi_{\log}^2 = \sum_i \frac{[\ln y_i^{\text{model}} - \ln y_i^{\text{data}}]^2}{\sigma_{\ln y,i}^2}, \quad (21)$$

with $\sigma_{\ln y,i} \equiv \sigma_i/y_i$. This score is used uniformly for all three surrogates and is therefore well defined for internal ranking. It is not claimed to be the full experimental likelihood, and the associated AIC/BIC values should be read in that restricted sense: they compare models under one common score, not under the collaborations' complete covariance machinery.

5.3. Global Benchmark Result

The global benchmark is summarized in Table 1. Within this restricted comparison class, the two-scale closure is strongly preferred.

Table 1. Closed pointwise benchmark on 83 public points. Lower AIC and BIC are preferred.

Model	$\chi_{\log}^2$	AIC	BIC
M_1	60493.49	60507.49	60524.43
M_2	59942.77	59964.77	59991.38
M_3	461.19	487.19	518.64




Figure 1. Global weighted log-score comparison for the three nested surrogates on the common public-data benchmark.




Figure 2. Closed pointwise overlay at 2.76 TeV.




Figure 3. Closed pointwise overlay at 13 TeV.

The size of the gap in Table 1 supports a strong statement, but a narrow one: within the explicit internal surrogate family of Eqs. (13)–(15), under the common weighted log-space score and common cross-energy flow, M_3 is decisively preferred.

5.4. Why the Evidence Is Intrinsically Multi-Energy

Table 2 shows the dedicated local 13 TeV comparison. The split-sector one-scale surrogate M_2 and the two-scale surrogate M_3 produce nearly identical dip, bump, and ratio observables, and M_2 even has slightly smaller AIC/BIC because it uses fewer parameters. The main empirical leverage of M_3 therefore comes from the cross-energy closure imposed in the global fit rather than from a decisive local advantage at 13 TeV alone. Stating this directly sharpens the physics claim: the evidence in favor of the closure is genuinely multi-energy.

Table 2. Local 13 TeV comparison. The point of this table is conceptual: M_3 is not uniquely selected by the 13 TeV line shape alone; its main empirical advantage appears only after the cross-energy constraint is imposed.

Model	χ^2	AIC	BIC	$ t _{\text{dip}}$ [GeV ²]	$ t _{\text{bump}}$ [GeV ²]	$R_{\text{bump/dip}}$
M_1	17.430	21.430	20.203	0.425800	0.527312	1.765749
M_2	0.28195	6.28195	4.44084	0.470081	0.568879	1.769998
M_3	0.21947	8.21947	5.76465	0.470204	0.570606	1.769996

Proposition 1 (bounded covariance robustness). *Let the weighted loss be built from a positive-definite covariance matrix Σ satisfying*

$$\lambda_- D \preceq \Sigma \preceq \lambda_+ D \tag{22}$$

for a fixed diagonal reference matrix D and constants $0 < \lambda_- \leq \lambda_+ < \infty$. Then each pairwise quadratic loss gap is rescaled only by bounded multiplicative factors, and the large internal gaps of Table 1 cannot be inverted by such bounded distortions.

Remark 2 (interpretive status). *Proposition 1 is a stability statement under classes of bounded covariance distortions. It is not a substitute for a full collaboration-level nuisance-parameter or covariance-matrix analysis, and the paper does not present it as one.*

6. Covariance-Aware Benchmark with External Regge, Eikonal, and Odderon Baselines

The weighted log-space comparison of Sec. 5 was designed as a closed internal test with the same score and the same cross-energy flow across the three nested surrogates. A natural next question is whether the main conclusion survives a stronger public-data benchmark in which point-to-point correlations are modeled explicitly and explicit external baselines are confronted under the same chain.

To do so, we construct a covariance-aware log-space score on the same 83 pointwise measurements at 2.76 and 13 TeV. For each published table we write

$$\Sigma_{\log}^{(a)} = \text{diag}[(\sigma_i^{\text{stat}}/y_i)^2] + u^{(a)}u^{(a)\top}, \quad u_i^{(a)} \equiv \sigma_i^{\text{syst}}/y_i, \quad (23)$$

and assemble the global block covariance from the 2.76 TeV low- $|t|$ table, the 2.76 TeV dip-region table, and the 13 TeV table. The benchmark score is then

$$\chi_{\text{cov log}}^2 = r^\top \Sigma_{\log}^{-1} r, \quad r_i = \ln y_i^{\text{model}} - \ln y_i^{\text{data}}. \quad (24)$$

This construction is still not the collaborations' full nuisance-parameter machinery, but it is materially stronger than a purely diagonal score because it propagates dataset-level systematic correlations explicitly.

Under this covariance-aware chain we compare five models: the one-scale split-sector surrogate M2, the two-scale copy-time surrogate M3, a canonical Regge-pole-plus-Odderon amplitude (RP2O), a canonical complex two-component Regge-eikonal baseline (CRE), and the fixed Kohara-Ferreira-Kodama 2014 parametrization (KFK fixed). The last entry is left frozen at its published energy dependence rather than re-optimized, so that it serves as a genuine literature reference point rather than as an in-house surrogate.

Table 3. Covariance-aware benchmark on the same 83 public points. The log-space covariance is built from the published statistical errors plus fully correlated systematic blocks within each experimental table.

Model	k	$\chi_{\text{cov log}}^2$	AIC	BIC
M2_cov	11	248659.97	248681.97	248708.57
M3_cov	13	1687.73	1713.73	1745.17
RP2O_cov	13	130622.02	130648.02	130679.47
CRE_cov	11	5440.69	5462.69	5489.30
KFK_fixed	0	421953.53	421953.53	421953.53

The main qualitative conclusion survives and is sharpened. The covariance-aware copy-time fit remains the best model in the implemented comparison set. The strongest external baseline is the canonical complex Regge-eikonal fit, but even that model remains behind M3 by $\Delta\chi_{\text{cov log}}^2 \simeq 3753$ on the same data and under the same covariance model. The canonical Regge-pole-plus-Odderon amplitude and the frozen KFK parametrization perform substantially worse under this particular benchmark.




Figure 4. Covariance-aware benchmark scores for the internal and external baselines. The two-scale copy-time surrogate remains the best model in the implemented comparison set.




Figure 5. Covariance-aware comparison at 2.76 TeV for the copy-time surrogate, the canonical Regge-pole-plus-Odderon amplitude, the canonical complex Regge-eikonal baseline, and the fixed KFK parametrization.




Figure 6. Covariance-aware comparison at 13 TeV for the same model set.

Two interpretive cautions are essential. First, the covariance-aware score places substantial weight on the dense forward-region information and is therefore not intended to replace the dedicated local dip-region calibration used for the preregistered 13.6 and 14 TeV continuation. Second, the external models used here are representative canonical baselines rather than exhaustive reproductions of every modern hadronic framework. Even so, this section closes an important loophole: the preference for the copy-time surrogate is no longer confined to an internal toy-baseline comparison, but survives against explicit Regge, Regge–Odderon, and Regge-eikonal baselines under one common covariance-aware public-data treatment.

7. Validation, Preregistration, and Relation to Representative Pomeron–Regge Expectations

7.1. 8 TeV Hold-Out Summary Validation

The hold-out validation uses the public TOTEM characterization of the dip and bump at 8 TeV [2]. Table 4 shows that the M3 continuation captures the overall pattern but does not yet achieve precision-level agreement. In particular, the predicted bump location remains visibly low. The appropriate reading is therefore trend-level validation: the continuation reproduces the ordering and approximate scale of the observables, but it has not yet earned the stronger language of precision prediction.

Table 4. 8 TeV hold-out validation for the M3 continuation against the public dip–bump summary.

Observable	Public value	M3 prediction	z score
$ t _{\text{dip}}$ [GeV ²]	0.521 ± 0.007	0.504124	−2.41
$ t _{\text{bump}}$ [GeV ²]	0.695 ± 0.026	0.614146	−3.11
$R_{\text{bump/dip}}$	1.97 ± 0.35	2.04535	+0.22

7.2. Preregistered Continuation to 13.6 and 14 TeV

The preregistered continuation is intentionally narrow. It is anchored to the exact local 13 TeV M3 solution,

$$(B_I, \mu_I, B_R, \mu_R)_{13} = (15.826821, 2.189431, 21.931289, 0.949746), \quad (25)$$

and then continued with the global cross-energy drift extracted from the closed fit. Writing $L(s) = \ln(\sqrt{s}/8 \text{ TeV})$,

$$B_I(s) = B_I^{(13)} \frac{b_0 + b_1 L(s)}{b_0 + b_1 L_{13}}, \quad (26)$$

$$\mu_I(s) = \mu_I^{(13)} \frac{m_0 + m_1 L(s) + m_2 L(s)^2}{m_0 + m_1 L_{13} + m_2 L_{13}^2}, \quad (27)$$

with the global drift coefficients

$$(b_0, b_1) = (15.005029, 1.618671), \quad (28)$$

$$(m_0, m_1, m_2) = (1.974983, 0.293338, 0.053940), \quad (29)$$

while the narrow-window real-sector ratios are fixed at their local 13 TeV values,

$$\frac{B_R}{B_I} = 1.385704, \quad \frac{\mu_R}{\mu_I} = 0.433787, \quad (30)$$

and the forward real-part intercept is fixed from the exact 13 TeV forward-slope identity, giving $\rho_{13} = 0.095$. The resulting reduced shape is

$$\begin{aligned} \mathcal{D}(x; s) = & e^{-B_I(s)x} (1 - \mu_I(s)x)^2 \\ & + e^{-B_R(s)x} (\rho_{13} - \mu_R(s)x)^2. \end{aligned} \quad (31)$$

We define $|t|_{\text{dip}}$ as the first local minimum of $\mathcal{D}$, $|t|_{\text{bump}}$ as the first local maximum after that minimum, and

$$R_{\text{bump/dip}} = \frac{\mathcal{D}(|t|_{\text{bump}}; s)}{\mathcal{D}(|t|_{\text{dip}}; s)}, \quad I_1 = B_0 |t|_{\text{dip}}. \quad (32)$$

The preregistered numerical values are given in Table 5. The continuation predicts a mild leftward drift of both the dip and the bump together with a small downward drift of $R_{\text{bump/dip}}$ and I_1 in the 13–14 TeV window. Relative to the calibrated 13 TeV point, the 14 TeV continuation gives $(\Delta|t|_{\text{dip}}, \Delta|t|_{\text{bump}}, \Delta R_{\text{bump/dip}}, \Delta I_1) = (-5.55 \times 10^{-3} \text{ GeV}^2, -6.42 \times 10^{-3} \text{ GeV}^2, -5.16 \times 10^{-3}, -3.17 \times 10^{-2})$.

Table 5. Preregistered continuation from the 13 TeV-calibrated M3 solution.

$\sqrt{s}$ [TeV]	$ t _{\text{dip}}$ [GeV ²]	$ t _{\text{bump}}$ [GeV ²]	$R_{\text{bump/dip}}$	I_1
13.6	0.466882	0.566636	1.766869	9.574230
14.0	0.464711	0.564127	1.764835	9.561705




Figure 7. Energy evolution of the dip location, bump location, bump-to-dip ratio, and invariant I_1 across the hold-out, fit, and preregistered points used in the submission package.

7.3. Forecast Robustness Inside the Narrow-Window Continuation Class

Because the 13.6 and 14 TeV continuation freezes the real-sector ratios B_R/B_I and μ_R/μ_I at their 13 TeV values and also freezes the forward real-part intercept indirectly, the submission package includes a scan in which the continuation is recomputed over the domain

$$\delta\rho \in [-0.02, 0.02], \quad (33)$$

$$\frac{\delta(B_R/B_I)}{B_R/B_I} \in [-5\%, 5\%], \quad (34)$$

$$\frac{\delta(\mu_R/\mu_I)}{\mu_R/\mu_I} \in [-5\%, 5\%]. \quad (35)$$

Across all 45 grid points in this domain, the sign pattern between 13 and 14 TeV remains unchanged:

$$(\Delta|t|_{\text{dip}}, \Delta|t|_{\text{bump}}, \Delta R_{\text{bump/dip}}, \Delta I_1) = (-, -, -, -). \quad (36)$$

The corresponding envelope is

$$\Delta|t|_{\text{dip}} \in [-5.70, -5.51] \times 10^{-3} \text{ GeV}^2, \quad (37)$$

$$\Delta|t|_{\text{bump}} \in [-6.45, -6.30] \times 10^{-3} \text{ GeV}^2, \quad (38)$$

$$\Delta R_{\text{bump/dip}} \in [-1.83 \times 10^{-2}, -7.41 \times 10^{-4}], \quad (39)$$

$$\Delta I_1 \in [-3.24, -3.12] \times 10^{-2}. \quad (40)$$

This does not remove the need for future data, but it shows that the quoted sign pattern is not a numerically fragile consequence of one exact frozen-ratio point in the narrow 13–14 TeV window.

7.4. Relation to Representative Pomeron–Regge Frameworks

The present preregistration should not be read as a claim that the copy-time closure is the only framework that anticipates motion of the dip–bump complex. Representative Pomeron–Regge and Regge–eikonal constructions also expect the dip and bump to move toward smaller $|t|$ as the energy rises across the LHC domain [9–12]. What differs is not the broad direction of motion, but the way the motion is packaged and tested.

Kohara, Ferreira, and Kodama provide an analytic representation of elastic pp scattering in the range 2.76–14 TeV and discuss the behavior of observables in that window [9]. Ferreira, Kohara, and Kodama later review how Pomeron-based descriptions of the dip region often require either multi-Pomeron structures or explicit Pomeron-plus-Odderon terms and note that the large- $|t|$ tail is not represented equally well by all variants [10]. Khoze, Martin, and Ryskin give a global description with one Pomeron pole supplemented by multi-Pomeron interactions [11], while Godizov shows that a two-Reggeon eikonal approximation remains predictive over a broad but not unlimited kinematic range [12]. The point is not that these frameworks fail, but that they are multiple, non-identical, and not yet reducible to one universal benchmark claim.

To make this comparison concrete within the implemented benchmark family, Figure 8 shows the forward-normalized predicted curves at 13.6 and 14 TeV for the preregistered M3 continuation together with the representative external baselines carried over from the covariance-aware benchmark. In this restricted comparison, the complex Regge-eikonal and KFK continuations both retain a resolved first dip–bump pair in the scanned window, whereas the implemented RP2O continuation remains monotonic over $0.03 \leq |t| \leq 1.35 \text{ GeV}^2$ and therefore does not define $|t|_{\text{dip}}$ or $|t|_{\text{bump}}$ there. Numerically, M3 and the complex Regge-eikonal baseline both move the dip and bump leftward from 13.6 to 14 TeV, but the Regge-eikonal continuation yields a noticeably smaller $R_{\text{bump/dip}}$ and a smaller proxy I_1 , while the fixed KFK continuation predicts a stronger left shift of the dip together with a slight upward drift of $R_{\text{bump/dip}}$. The near-term discriminator is therefore not just the sign of the geometric drift, but the full correlated pattern across $|t|_{\text{dip}}$, $|t|_{\text{bump}}$, $R_{\text{bump/dip}}$, and I_1 . The corresponding numerical comparison is included in the Supplement and in the machine-readable file `data/forecast_comparison_full.json`.




Figure 8. Forward-normalized predicted curves at 13.6 and 14 TeV for the preregistered M3 continuation and the representative external baselines implemented in the submission package. The normalization is fixed at the first plotted point in each panel to emphasize shape differences. In the plotted window the RP2O continuation remains monotonic, while the complex Regge-eikonal and KFK baselines both retain a resolved dip–bump pair.

The practical difference in the present submission is the promotion of $I_1 = B_0|t|_{\text{dip}}$ to a preregistered observable on the same footing as $|t|_{\text{dip}}$, $|t|_{\text{bump}}$, and $R_{\text{bump/dip}}$. In the one-scale geometric class

I_1 is constant; in the present M3 continuation it drifts slightly downward between 13 and 14 TeV. The distinctive near-term signature is therefore not merely a leftward motion of the dip–bump complex, which several external frameworks also anticipate, but the full joint sign pattern

$$(\Delta|t|_{\text{dip}}, \Delta|t|_{\text{bump}}, \Delta R_{\text{bump/dip}}, \Delta I_1) = (-, -, -, -) \quad (41)$$

in the narrow 13–14 TeV window. Whether that correlated sign pattern survives future data is a direct phenomenological discriminator.

8. Scope of the Present Claim

The scientific value of the present result is clearest when its present boundary is stated directly. The manuscript does not establish four stronger claims:

1. It does not derive the closure from microscopic QCD. The theorem-level content is conditional on the axioms and the minimal closure selection.
2. It does not prove non-reducibility against the full external phenomenology. The rigorous theorem excludes only the one-scale geometric class; the discussion of Pomeron–Regge models is representative and benchmark specific.
3. It does not provide a collaboration-level covariance completion or nuisance-parameter treatment. The covariance-aware benchmark is substantially stronger than the earlier diagonal score, but it still models systematic correlations at the level of public table blocks rather than with the collaborations' full internal covariance machinery.
4. It does not claim a historically pre-published successful forecast or third-party replication. The 13.6 and 14 TeV numbers are preregistered here, and the bundled code is intended to make later replication straightforward.

These limits are not disclaimers appended after the fact; they define the actual status of the result. What the paper contributes is a clean conditional theorem, a sharply defined benchmark program, an honest account of where the empirical leverage comes from, and a falsifiable continuation to the next relevant energies.

9. Reproducibility Note

The submission package accompanying this paper is intentionally canonical. It contains a single manuscript source, a supplementary file, the closed-fit script, the covariance-aware external-benchmark script, the forecast script, the symbolic proof-check script, the generated figures used in the manuscript, and machine-readable summaries of the numerical tables. The directory structure is chosen so that one can regenerate the benchmark figures, recompute the forecast table, and verify the symbolic identities in one documented tree without relying on undocumented intermediate states.

10. Conclusion

The present submission supports a deliberately narrower and correspondingly stronger statement. Within an explicit axiom set, the first Laguerre deformation is the unique lowest-complexity finite Hankel-closed single-node closure. Within that closure, one obtains rigid forward/dip structural relations and a natural drift observable $\mathcal{O}_{\text{excl}}$. Within a closed internal comparison family using a shared cross-energy flow and a shared weighted log-space score, the two-scale copy-time surrogate is strongly preferred globally, even though the local 13 TeV line shape by itself does not uniquely select it over the one-scale split-sector alternative. Under a stronger covariance-aware benchmark built from the published pointwise statistical errors and explicit block-correlated systematics, the copy-time surrogate also remains preferred over the canonical external Regge, Regge–Odderon, and Regge-eikonal baselines implemented here. The 8 TeV hold-out test shows trend-level rather than precision-level validation, and the 13.6 and 14 TeV values are now stated as preregistered numerical forecasts.

Taken together, these points do not amount to a universal hadronic-theory claim. They do, however, amount to a sharply posed and genuinely testable PRD-level result: a minimal closure theorem with explicit observables, a controlled public-data benchmark, and a numerical preregistration whose success or failure can be read directly from near-term elastic-scattering measurements.

Data Availability Statement: The experimental inputs analyzed in this article are taken from public measurements cited in the reference list. The numerical summaries, generated figures, and scripts sufficient to reproduce the benchmark tables and the preregistered forecast table are included in the submission package supplied with this manuscript and are also available from the author upon reasonable request.

Acknowledgments: The author thanks the public experimental collaborations for making the elastic-scattering measurements and summary observables used here accessible in the literature. No external funding is claimed for this work.

Appendix A. Exact Transform Pair for the First Laguerre Mode

For completeness, the transform identity used in Secs. 3 and 4 can be written in direct and inverse form as

$$\frac{1}{B}e^{-\nu}\left[1 - \frac{2\mu}{B}L_1(\nu)\right] \xrightarrow{2\pi \int b db J_0} 2\pi e^{-Bx/2}(1 - \mu x), \quad (\text{A1})$$

$$2\pi e^{-Bx/2}(1 - \mu x) \xrightarrow{\frac{1}{2\pi} \int q dq J_0} \frac{1}{B}e^{-\nu}\left[1 - \frac{2\mu}{B}L_1(\nu)\right]. \quad (\text{A2})$$

with $\nu = b^2/(2B)$ and $x = q^2$. The symbolic verification used to check this identity is bundled as a stand-alone script in the submission package.

Appendix B. Variational Consistency Statement

The following statement is retained only as a consistency result and not as a microscopic derivation.

Proposition A1. *Assume that a coarse-grained radial certifiability functional admits an effective quadratic action*

$$\mathcal{I}[\tau] = \int d^2b \left[\alpha(s)\tau^2 + \beta(s)|\nabla_b \tau|^2 \right] \quad (\text{A3})$$

with positivity, radial locality, and finite second moment, and that admissible responses are restricted to the minimal single-node sector of Theorem 1. Then the stationary solutions at fixed normalization and fixed first-node class coincide with the first Laguerre deformation of Eq. (8).

The proof is standard once the admissible space is restricted: the Gaussian-weighted radial quadratic form diagonalizes on the Laguerre basis, the normalization constraint removes the $n = 0$ freedom, and the minimal single-node restriction excludes $n \geq 2$. This appendix is included only to show internal consistency between the closure and a simple variational picture; it is not a substitute for a QCD derivation.

Appendix C. Canonical Benchmark and Continuation Parameters

For reproducibility, the canonical M3 flow coefficients extracted from the global benchmark are

$$(b_0, b_1) = (15.005029, 1.618671), \quad (\text{A4})$$

$$(m_0, m_1, m_2) = (1.974983, 0.293338, 0.053940), \quad (\text{A5})$$

$$(k_B, k_\mu) = (1.564449, 0.765934). \quad (\text{A6})$$

The global-flow-implied sector parameters at the three energies explicitly stored in the bundled numerical summary are listed in Table A1. For clarity, the exact local 13 TeV calibration used to anchor the preregistered continuation is reported separately in Table A2.

Table A1. Global-flow-implied M3 sector parameters extracted from the bundled numerical summary.

$\sqrt{s}$ [TeV]	B_I	μ_I	B_R	μ_R
2.76	13.282422	1.723899	20.779673	1.320393
8.00	15.005029	1.974983	23.474604	1.512707
13.0	15.790906	2.130115	24.704069	1.631528

Table A2. Exact local 13 TeV M3 calibration used as the anchor for the preregistered 13.6 and 14 TeV continuation.

B_I	μ_I	B_R	μ_R
15.826821	2.189431	21.931289	0.949746

Appendix D. Forecast Robustness Data Product

The submission package contains a machine-readable forecast-robustness scan in both JSON and CSV formats. The scan covers 45 combinations of $(\delta\rho, \delta(B_R/B_I), \delta(\mu_R/\mu_I))$ over the domain quoted in Sec. 7. For each grid point the package stores the 13 TeV and 14 TeV local extrema and the resulting differences. The purpose of this artifact is not to replace future data, but to make the narrow-window continuation falsifiable in a parameter-transparent way.

References

1. TOTEM Collaboration, “Elastic differential cross-section measurement at $\sqrt{s} = 2.76$ TeV and implications on the existence of a colourless 3-gluon bound state,” *Eur. Phys. J. C* **80**, 91 (2020).

2. TOTEM Collaboration, “Characterisation of the dip-bump structure observed in proton-proton elastic scattering at $\sqrt{s} = 8$ TeV,” *Eur. Phys. J. C* **82**, 263 (2022).

3. TOTEM Collaboration, “Elastic differential cross-section measurement at $\sqrt{s} = 13$ TeV by TOTEM,” *Eur. Phys. J. C* **79**, 861 (2019).

4. TOTEM Collaboration, “First determination of the ρ parameter at $\sqrt{s} = 13$ TeV: probing the existence of a colourless C-odd three-gluon compound state,” *Eur. Phys. J. C* **79**, 785 (2019).

5. D0 and TOTEM Collaborations, “Odderon exchange from elastic scattering differences between pp and $p\bar{p}$ data at 1.96 TeV and from forward scattering measurements,” *Phys. Rev. Lett.* **127**, 062003 (2021).

6. NIST CODATA 2022 value for the proton rms charge radius, $r_p = 8.4075(64) \times 10^{-16}$ m, available from the NIST fundamental constants database.

7. S. Mohamed, “Quantum Information Copy Time (QICT): Operational Definition, Minimal Locality Bounds, Hydrodynamic Susceptibilities, and a Programmatic Outlook,” Preprints.org 202601.0364 (2026).

8. S. Mohamed, “Quantum Information Copy-Time as a Microscopic Principle for Emergent Hydrodynamics, Inertial Spectral Mass, and a Predictive Higgs-Portal Dark-Matter Corridor,” Preprints.org 202601.1044 (2026).

9. A. K. Kohara, E. Ferreira, and T. Kodama, “pp elastic scattering at LHC energies,” *Eur. Phys. J. C* **74**, 3175 (2014).

10. E. Ferreira, A. K. Kohara, and T. Kodama, “Analytical representation for amplitudes and differential cross section of pp elastic scattering at 13 TeV,” *Eur. Phys. J. C* **81**, 290 (2021).

11. V. A. Khoze, A. D. Martin, and M. G. Ryskin, “High-energy elastic and diffractive cross sections,” *Eur. Phys. J. C* **74**, 2756 (2014).

12. A. A. Godizov, “High-energy elastic diffractive scattering of nucleons in the framework of the two-Reggeon eikonal approximation (from U-70 to LHC),” *Eur. Phys. J. C* **82**, 56 (2022).

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.