Hadronic τ decays at higher orders in QCD

What four coefficients can, and cannot, support

Gauhar Abbas

Department of Physics, Indian Institute of Technology (BHU), Varanasi

Workshop on Tau Physics at Belle II and STCF
September 1st, 2026

Based on arXiv:2601.11277 and more recent work. To appear/update on arXiv soon.

Outline

1. Motivation
$\delta^{(0)}$, $c_{5,1}$, and why it limits $\alpha_s(m_\tau)$

2. Formalism
The three-weight test is a tolerance
A no-go for a whole class of methods
Cross-validated renormalon closure (CVRC)

3. Validation
Moments, pole mass, all orders in large-β_0
. . . and the honest negative

4. Global comparison
Every approximant, one code path, one figure

5. Final predictions
$c_{5,1}$, $c_{6,1}$, and the calibrated Borel model

6. Summary and conclusions

One methodological thread runs through all of it: **declare the test before making the forecast**, and report the tests that fail as prominently as those that pass.

Why inclusive τ decay is still the sharpest low-energy probe

and where the theory error actually sits

Optical theorem + analyticity:

$$R_{\tau,V+A} = N_c |V_{ud}|^2 S_{\rm EW} \left[1 + \delta'_{\rm EW} + \delta^{(0)} + \sum\nolimits_{D\geq 2} \delta^{(D)}\right]$$

- $\delta^{(0)}$ is the massless perturbative term — the piece that carries α_s.
- Everything else (EW factors, condensates, duality violations) is *not* touched here.
- $\delta^{(0)}$ is known to $\mathcal{O}(\alpha_s^4)$: four Adler coefficients.

$$\widehat{D}(Q^2) = \sum_{n\geq 1} c_{n,1}\, a^n, \qquad a = \frac{\alpha_s}{\pi}$$

$$c_{1,1} = 1, \quad c_{2,1} = 1.6398, \quad c_{3,1} = 6.3710, \quad c_{4,1} = 49.076$$

The practical question

At fifth order, $c_{5,1}$ fixes the first omitted fixed-order term. It does *not* resolve the asymptotic FOPT–CIPT separation, which also depends on contour running and Borel singularities.

As Belle II and STCF reduce spectral-moment errors, perturbative control becomes increasingly important.

Can four coefficients support a **renormalon-resolved** extrapolation, not merely a locally successful one?

Why the FOPT–CIPT difference matters

Two perturbative organizations

- FOPT expands the running coupling at a fixed reference scale before performing the contour integral.
- CIPT evolves the coupling locally along the complex contour and does not re-expand it at a fixed scale.

At finite order,

$$\delta_W^{\mathrm{FOPT}}(N) \neq \delta_W^{\mathrm{CIPT}}(N), \tag{1}$$

and the difference has traditionally been included as an important theory uncertainty in extractions of α_s from hadronic τ decays.

Why the FOPT–CIPT difference matters

The modern renormalon interpretation

In the presence of infrared renormalons, the Borel representation underlying standard CIPT is not compatible with the usual local-operator form of the OPE. Consequently,

$$\delta_W^{\mathrm{CI,Borel}} - \delta_W^{\mathrm{FO,Borel}} = \Delta_W^{\mathrm{Asymptotic-Separation.}} \neq 0, \tag{2}$$

where Δ_W^{AS} is the asymptotic separation.

Interpretation

- The difference is not merely caused by the unknown coefficient $c_{5,1}$.
- Knowing $c_{5,1}$ reduces the first omitted fixed-order uncertainty, but does not remove the all-order separation.
- A renormalon-subtracted OPE scheme can strongly reduce the numerical FOPT–CIPT discrepancy. See Pich and Rodriguez-Sanchez 2026

Why $c_{5,1}$ is an important missing theory input

a qualitative statement — no α_s extraction is performed in this work

- FOPT and CIPT use the same Adler coefficients but organize contour running differently. Their truncated numerical sums differ, although the CIPT result re-expanded in the reference coupling reproduces FOPT through the stated order.
- $c_{5,1}$ fixes the next unknown fixed-order contribution and is therefore an important stability diagnostic. It does not make the all-order prescription separation vanish.
- The all-order separation involves analytic continuation and the Borel singularity structure. A credible continuation must be *renormalon-aware*, not merely a successful local fit.

Literature values of $c_{5,1}$

283	geometric ratio (not the model!) (Beneke–Jamin 2008)
277 ± 51	Borel–Padé (Boito, Masjuan, Oliani 2018)
278 ± 27	Levin (Abbas, Singh 2026)

The spread is not small, and — as we will see — different extrapolation methods applied to the *same four numbers* give anything from -22 to 3328.

So the interesting question is not "what number?" but "which methods are entitled to answer?"

The rules, declared before any forecast

a validation hierarchy, not a post-hoc justification

Every method in this work must face the same four tests:

1. **Leave-one-order-out on three moments.** Withhold $c_{N,1}$, forecast it, and check the implied $d_N[W]$ for

$$W_\tau = (1-x)^3(1+x), \quad W_{\mathrm{GC}} = (1-x)^3,$$
$$W_{P2} = (1-x)^2(1+x).$$

   Criterion: below 10% errors on *all three*.

2. **Nearby renormalon.** Blind forecast of the four-loop pole–$\overline{\mathrm{MS}}$ mass coefficient ($u = 1/2$).
3. **All orders.** Strictly recursive blind forecasting in the large-β_0 limit, where every coefficient is known exactly.
4. **A non-observable control.** Withhold the known five-loop $\overline{\mathrm{MS}}$ β-function coefficient.

Two ground rules

Accuracy through order. A method must reproduce every supplied coefficient exactly. No fitting away the input.

No withheld coefficient may enter a forecast. Blind means blind, including in the calibration of any free parameter.

W_{GC} and W_{P2} are both $u = 2$-sensitive but weight the Adler series very differently. Passing one and failing the other is diagnostic.

From moments back to a single number

the triangular weight map has unit diagonal

Analyticity turns the weighted spectral integral into a contour integral,

$$\delta_W^{(0)}(s_0) = \frac{1}{2\pi i} \oint_{|x|=1} \frac{dx}{x}\, W(x)\, \widehat{D}(-s_0 x) = \sum_{N\geq 1} d_N[W]\, a^N.$$

In fixed order this is an exact **triangular** map from Adler coefficients to moment coefficients,

$$d_N[W] = \sum_{n=1}^{N} M_{Nn}[W]\, c_{n,1}, \qquad M_{NN}[W] = 1 \quad \text{for every normalised } W.$$

Transfer identity

Because the diagonal is unity, an error in the forecast of $c_{N+1,1}$ passes into *every* moment with unit coefficient:

$$d_{N+1}[W] - \widehat{d}_{N+1}[W] = c_{N+1,1} - \widehat{c}_{N+1,1} \qquad \textit{independently of } W.$$

The three-weight test is really a tolerance on one coefficient

For *any* transformation acting on the common Adler sequence, the declared ten-percent three-weight criterion is exactly

$$\left|\widehat{c}_{N+1,1} - c_{N+1,1}\right| \;\le\; \min_{W\in\{W_\tau, W_{\mathrm{GC}}, W_{P2}\}} 0.1\,\left|d_{N+1}[W]\right|,$$

and with the exact moments,

$$\boxed{\left|\widehat{c}_{3,1} - c_{3,1}\right| \le 1.4523, \qquad \left|\widehat{c}_{4,1} - c_{4,1}\right| \le 2.1509}$$

- The withheld-d_4 "ten-percent" test is in fact a 1.69% requirement on $d_4[W_\tau] = 127.08$.
- The withheld-d_3 test is a 5.51% requirement on $d_3[W_\tau] = 26.37$.

This single reformulation is what makes the rest of the analysis possible: a multi-observable criterion collapses to a one-dimensional window, and the whole problem becomes *"where can* $\widehat{c}_{5,1}$ *be?"*

Not a fixed-order artefact: in CIPT the map is *diagonal*, not triangular, and the error is no longer weight independent — but the collapse survives, and the same criterion is $\sim 35\times$ *weaker* (± 80 vs ± 2.15 at $N = 4$).

A no-go: why most methods cannot possibly work here

Define the one-step Padé/Shanks **defect** $\Delta_n = c_{n,1} - c_{n-1,1}^2/c_{n-2,1}$. From the four known coefficients,

$$(\Delta_2, \Delta_3, \Delta_4) = (+0.640,\ +3.682,\ +24.323).$$

Positive, and growing by a factor of six per order. But the geometric baseline at the next order is $c_{4,1}^2/c_{3,1} = 378.028$, so *any* literature-range $c_{5,1}$ requires $\widehat{\Delta}_5 = \widehat{c}_{5,1} - 378.028 < 0$: a **sign reversal**.

Proposition

Let $\mathcal{T}$ act on the common Adler sequence and be *defect-sign-preserving* (whenever $\Delta_2, \ldots, \Delta_N > 0$ and increasing, it returns $\widehat{\Delta}_{N+1} \geq 0$). Then

(i) at $N = 4$ it returns $\widehat{c}_{5,1} \geq 378.028$;

(ii) the withheld-$c_{4,1}$ three-weight test does *not* exclude it, because the test window $\widehat{\Delta}_4 \in [22.17, 26.47]$ lies entirely at $\widehat{\Delta}_4 > 0$.

Hence: it can pass the weight test and still never return a $c_{5,1}$ in $[229, 305]$. The two are incompatible inside this class.

Ordinary Shanks/Padé, ARTS/PORT with a positive transported defect, and any single positive infrared factorial mode all live inside this class. Their mixed record is **forced, not accidental**.

Escaping the obstruction: declare the Borel structure

In the Borel plane, $u = \beta_0 t$, a singularity $N_p(p-u)^{-\alpha_p}$ contributes

$$c_{n,1}\big|_p = N_p\, p^{-\alpha_p - n + 1}\, \frac{\Gamma(\alpha_p + n - 1)}{\Gamma(\alpha_p)}\, \beta_0^{n-1},$$

and a Borel monomial u^j contributes $\beta_0^j j!$ to $c_{j+1,1}$ *and nothing to any other order.*

The exponents are *declared*, not fitted:

$$\alpha_p = 1 + \frac{p\beta_1}{\beta_0^2}\ (p > 0), \qquad \alpha_{-k} = 2 - \frac{k\beta_1}{\beta_0^2}\ (k > 0)$$

$$\alpha_2 = 2.5802, \quad \alpha_3 = 3.3704, \quad \alpha_{-1} = 1.2099$$

- $\alpha_{-1} = 2 - \beta_1/\beta_0^2$: the large-$\beta_0$ *double* pole at $u = -1$, RG-dressed. This is the Beneke–Jamin exponent 1.21.

Nesting rule at order N

N *constrained* columns, seeded by the leading IR singularity and the analytic term, then alternating:

$$\{u{=}2,\ u^0\},\ \{u{=}3\},\ \{u^1\},\ \{u{=}4\},\ \{u^2\},\ldots$$

plus **one more** column — the closure: the leading UV singularity $u = -1$, amplitude ρ.

$N + 1$ parameters for N constraints — deliberately **one parameter richer than the data**.

Cross-validated renormalon closure (cvrc)

the closure family is exactly one-dimensional and affine

Imposing $c_{1,1}, \ldots, c_{N,1}$ on $N+1$ parameters leaves a one-dimensional affine set. Every member reproduces all supplied coefficients exactly (**accuracy through order for every** ρ), and every forecast is affine:

$$\boxed{\widehat{c}_{N+1,1}(\rho) = A_N + B_N\,\rho}$$

For the physical $n_f = 3$ series:

$$\widehat{c}_{3,1}(\rho) = 6.60484 + 24.500\,\rho$$
$$\widehat{c}_{4,1}(\rho) = 31.6788 - 287.960\,\rho$$
$$\widehat{c}_{5,1}(\rho) = 395.340 + 2439.209\,\rho$$
$$\widehat{c}_{6,1}(\rho) = 3447.716 + 5720.039\,\rho$$

- The slopes are *structural*: they depend on the declared singularity list and β_0, not on the data.
- They **alternate in sign** with the order.

This is exactly the missing degree of freedom

One negative ρ simultaneously *lowers* $\widehat{c}_{3,1}$, *raises* $\widehat{c}_{4,1}$, and *lowers* $\widehat{c}_{5,1}$ relative to the pure infrared model.

A defect-sign-preserving transformation has no such handle. CVRC is not a better fit — it is **outside the class the Proposition forbids**, as is the Beneke–Jamin model, for the same reason. The new element is not the model but the criterion that can reject it.

The escape is the alternating ultraviolet mode at $u = -1$ — a physical statement, not a numerical trick.

Fixing the one free parameter — blind

There is exactly one free direction, so use exactly one leave-one-order-out fold. The model built from $c_{1,1}, \ldots, c_{N-1,1}$ must reproduce the known $c_{N,1}$:

$$\boxed{\rho_\star = \frac{c_{N,1} - A_{N-1}}{B_{N-1}}} \quad \Longrightarrow \quad \boxed{\widehat{c}^{\,\mathrm{CVRC}}_{N+1,1} = A_N + B_N \, \frac{c_{N,1} - A_{N-1}}{B_{N-1}}}$$

A **closed-form nonlinear** transformation — nonlinear because $A_{N-1}, B_{N-1}, A_N, B_N$ themselves depend on the supplied data. No withheld coefficient enters it.

At physical order $N = 4$:

$$\rho_\star = \frac{49.07570 - 31.67880}{-287.96007} = -0.060414$$

- The *negative* sign is a statement about the physical series: the UV mode must enter with the sign that *suppresses* $c_{5,1}$.

Structural reading: closure = null newest column

$\rho_\star$ is the unique value for which the column added at order N carries *zero* amplitude — i.e. the order-N and order-$(N-1)$ models coincide.

Verified numerically: at $N = 4$ the newest column is the monomial u^1 and its fitted amplitude is $\sim 10^{-16}$.

cvrc *is* the Beneke–Jamin model — stated explicitly

same singularities, same exponents, same closure, same data

Identical, ingredient by ingredient

Converting their convention, $\beta_k^{\rm BJ} = 2\beta_{k-1}^{\overline{\rm MS}}$:

$1+\tilde{\gamma}$	BJ (2008)	this work
$u=2$	2.580247	2.580247
$u=3$	3.370370	3.370370
$u=-1$	1.209877	1.209877

- same singularity list $\{u{=}{-1},\ 2,\ 3\}$;
- same parameter count (four);
- their $d_1^{\rm PO}=0$ **is** our closure = null newest column;
- same four coefficients.

Running our leave-one-out closure on *their* columns returns $\rho_\star = -0.016089$ — exactly their fitted $d_1^{\rm UV}$.

The only difference = the whole gap

Identical protocol, one ingredient at a time:

declared content	$\widehat{c}_{5,1}$	shift
pure powers (CVRC)	247.98	—
+ $\tilde{b}_1, \tilde{b}_2$ cuts	257.53	+9.6
+ $c_1[\langle aG^2\rangle]$ (=BJ)	280.12	+22.6

The 32-unit gap is **not** a disagreement about renormalon content, closure, or data. It is 9.6 units of subleading cut structure plus 22.6 of a gluon-condensate Wilson coefficient we do not declare (backup: why not).

283 is *not* their model's output — it is their Eq. (3.10) geometric ratio, used as an **input**. Their own four-coefficient forecast is 280.

We claim **no new model class** — the class that escapes the Proposition is theirs. What is new is the acceptance criterion that can reject a declared structure.

Test 1 — the three-weight leave-one-out test

fold	weight	exact	error
d_3	W_τ	26.366	-4.73%
d_3	W_{GC}	34.484	-3.61%
d_3	W_{P2}	14.523	-8.58%
d_4	W_τ	127.079	0.00%
d_4	W_{GC}	214.140	0.00%
d_4	W_{P2}	21.509	0.00%

All six inside the declared 10% band. No other construction examined in this work does that.

What this actually certifies — read carefully

The three d_4 entries are the **calibration fold** and are exact *by construction*. The three d_3 entries are a self-consistency reconstruction.

So this is a **consistency certificate for the declared structure**, *not* six blind forecasts. We say so in the paper, and we say so here.

The blind tests are the *external* ones — pole mass and large-β_0 — and they are next.

Is the certificate discriminating? — misspecification scan

is the admissible closure set nonempty for a *wrong* declared structure?

declared structure	admissible ρ window	$\widehat{c}_{5,1}$
baseline: $u_{\mathrm{IR}} = 2$, $\alpha_{-1} = 2 - \beta_1/\beta_0^2$	$[-0.0679,\ -0.0529]$	247.98
$\alpha_{-1} = 1 - \beta_1/\beta_0^2$ (value used earlier in this work)	EMPTY — rejected	—
$\alpha_{-1} = 1$ (simple UV pole)	$[-0.0879,\ -0.0710]$	258.39
$\alpha_{-1} = 2$ (undressed double pole)	$[-0.0292,\ -0.0228]$	210.52
UV singularity at $u = -3/2$	EMPTY — rejected	—
UV singularity at $u = -2$	EMPTY — rejected	—
leading IR singularity at $u = 1$	EMPTY — rejected	—
leading IR singularity at $u = 3/2$	$[-0.0585,\ -0.0460]$	290.40
leading IR singularity at $u = 5/2$	EMPTY — rejected	—
leading IR singularity at $u = 3$	EMPTY — rejected	—

Six of ten declared structures are **rejected outright**: no value of ρ passes the three-weight criterion. The test has teeth — it is not satisfiable by construction, and it prefers the renormalization-group exponents.

Tests 2 and 3 — genuinely blind, external

Pole–$\overline{\text{MS}}$ mass, $u = 1/2$ renormalon

Blind forecast of the four-loop coefficient r_4 from the first three:

n_l	exact	CVRC	error
3	1702.7	1742.4	+2.33%
4	1235.6	1306.7	+5.75%
5	839.2	939.0	+11.90%

ARTS/PORT does *better* here (1.60–3.84%) — and that is consistent with the Proposition, because the pole-mass defects do not require a sign reversal.

Large-β_0: every coefficient known, so test all orders

N	CVRC-min	CVRC-full	bal. Padé
4	SINGULAR	−108.3%	−68.7%
5	+23.7%	+264.1%	−11.6%
6	+9.5%	+4.6%	−45.7%
7	+13.0%	+3.4%	+787.0%
8	+9.6%	+0.8%	−22.7%
10	+8.0%	+1.2%	−9.5%
12	+6.6%	+0.7%	−3.7%
14	+5.6%	+0.6%	−1.4%

CVRC-complete: within 10% at every $N \geq 6$, within 1.5% at every $N \geq 8$. The renormalon annihilator with *exact* input needs 11–13 coefficients for the same persistent control.

Two things the method does *not* do

1. At $N = 4$ in large-β_0 it *abstains*

At $N = 4$ in the minimal large-β_0 model, the design matrix is rank deficient. Two or more declared templates produce indistinguishable low-order coefficient patterns, so their amplitudes cannot be separated. The method returns no forecast rather than a pseudoinverse-dependent number.

Below $N = 6$ the procedure is not merely inaccurate — at $N = 4$ it is **not defined**.

2. Run strictly recursively, it misses $c_{4,1}$

Fix ρ one order lower still (from the $c_{3,1}$ fold) and forecast $c_{4,1}$ blind:

ρ from withheld-$c_{3,1}$ fold	-0.0095
ρ from withheld-$c_{4,1}$ fold	-0.0604
closure drift $\lvert\Delta\rho\rvert$	0.0509
blind $\widehat{c}_{4,1}$	34.43
exact $c_{4,1}$	49.08
error	-29.9%

This is reported as prominently as the successes. The closure drift is our **quantitative measure of how far four coefficients are from the asymptotic regime**.

We therefore do *not* claim that four coefficients *determine* $c_{5,1}$.

Every approximant, one set of inputs, one code path

how to read the next slide

Each section of the paper introduced a method and reported its continuation in its own normalisation. This figure recomputes **all of them** from

$$c_{1,1} = 1,\ c_{2,1} = 1.6398,\ c_{3,1} = 6.3710,\ c_{4,1} = 49.076$$

through a *single* inverse renormalization-group map, so the entries are commensurable. No new method, no new assumption.

Status labels :

- CVRC — passes the acceptance criterion
- BASE — conditional baseline
- COND — accuracy through order, admissible as a conditional forecast
- DIAG — diagnostic only
- ABST — the method's own certificate abstains

Reference lines on the figure

- grey band — Borel–Padé literature range 277 ± 51
- dashed — Beneke–Jamin 283
- solid red — the **Shanks bound** $c_{4,1}^2/c_{3,1} = 378.03$ from the Proposition
- blue band — CVRC admissible closure window

The red line is a *prediction of the no-go theorem*: every defect-sign-preserving row must sit on or to the right of it. Watch whether it does.

Only one row is blue.

CVRC, top row:
$c_{5,1} = 248 \pm 70$
(the blue band drawn is the narrower *structural* one)

Red line: 378.03

The Shanks bound $c_{4,1}^2/c_{3,1}$. Every defect-sign-preserving Adler-space row sits on or to the *right* of it — the Proposition, visible.

Grey band: 277 ± 51

Borel–Padé literature range; dashed line is Beneke–Jamin 283.

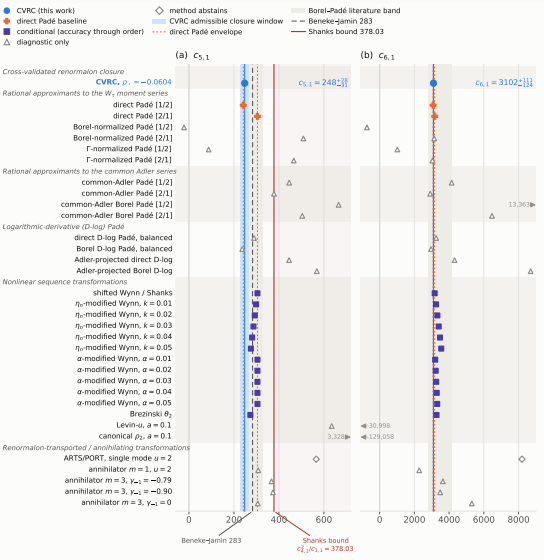

ARTS/PORT → 566

Passes the weight test, fails the range — exactly as the Proposition requires.

The cond block

Sequence transformations cluster tightly at 275–305, but none is validated: accuracy through order is necessary, not sufficient.

The diag scatter

Diagnostic rows run from -21.7 to 3327.8. That spread is the honest measure of how underdetermined four coefficients leave the problem.

Marker shape carries status as well as hue.

What the comparison shows

- The **fourteen** accuracy-through-order rows span

$$241.70 \leq c_{5,1} \leq 304.71.$$

- Admitting the *diagnostic* rows widens this to $[-21.73,\ 3327.76]$.
- So with four coefficients, **method choice dominates the answer** — by two orders of magnitude if one is not selective.

Two nonstatistical summaries, kept separate:

	$c_{5,1}$	$c_{6,1}$
direct-Padé baseline	273 ± 32	3125 ± 46
seq.-transformation grid	289 ± 16	3375 ± 175

Why CVRC is kept *outside* every average

It is the only row that passes the acceptance criterion the others fail. Averaging it with rows that fail that criterion would *destroy* the information the criterion carries.

We also state explicitly which combinations of rows are *not* admissible — e.g. direct Padé and shifted Wynn/Shanks coincide algebraically for the $[N-1/1]$ row and are not independent evidence.

The Proposition is confirmed empirically: every defect-sign-preserving row in the figure returns $c_{5,1} \geq 378.03$.

Final predictions

Physical-order result

$$c_{5,1}^{\mathrm{CVRC}} = 248 \pm \underbrace{30}_{\text{struct}} \pm \underbrace{62}_{\text{method}} = 248 \pm 70$$

The method band is the *demonstrated* one-order-ahead accuracy at this order — not a parameter scan:

large-β_0 at $N = 4$ (the physical order)	undefined
large-β_0 at $N = 5$ (first defined order)	$+23.7\%$
physical recursive test (one order lower)	-29.9%
pole mass (easier: nearby $u = 1/2$)	2.3–11.9%
adopted	$25\% = 62$

Consistent with 277 ± 51, with 273 ± 32, and with 280, the Beneke–Jamin four-coefficient forecast — from which we differ only by subleading cut structure ($+9.6$) and a gluon-condensate Wilson coefficient ($+22.6$) we omit. The criterion *forbids* adding them at this order — backup.

$c_{6,1}$ is only a two-order rolled value (3102^{+111}_{-124}, structural); it is not advanced as a prediction.

. . . and it is a single Borel function

Closure = null newest column, so this is not an N-dependent recipe:

$$\widehat{B}(u) = \frac{6.463246}{(2-u)^{2.580247}} - \frac{26.662923}{(3-u)^{3.370370}} - \frac{0.060414}{(1+u)^{1.209877}} + 0.637084$$

- Positions and exponents: from the renormalization group; amplitudes from $c_{1,1}, \ldots, c_{4,1}$.
- The one direction the data cannot fix — the $u = -1$ amplitude — **is** $\rho_\star = -0.060414$. *Nothing is tuned.*

Summary

The structural result

- The triangular weight map has unit diagonal $\Rightarrow$ the three-weight 10% criterion is exactly $|\widehat{c}_{3,1} - c_{3,1}| \leq 1.4523$, $|\widehat{c}_{4,1} - c_{4,1}| \leq 2.1509$.
- Known defects are positive, but a literature-range $c_{5,1}$ needs $\Delta_5 < 0$.

$\Rightarrow$ **No defect-sign-preserving Adler-space transformation can pass the weight test *and* return a literature-range $c_{5,1}$.** The mixed record of the standard toolkit is forced.

The construction

- Declared Borel singularities with RG exponents, one parameter richer than the data $\Rightarrow$ forecast affine in a single closure amplitude ρ; accuracy through order for all ρ.
- ρ fixed by one leave-one-order-out fold $\Rightarrow$ closed-form map. *Nothing tuned.*

What survived the hierarchy

- All six moment entries pass; six of nine alternatives (six of ten total structures) are rejected outright.
- Blind pole-mass forecast: 2.3–11.9%.
- Blind large-β_0: complete tower $\leq 10\%$ for every $N \geq 6$ and $< 1.5\%$ for every $N \geq 8$; the minimal tower reaches persistent 10% accuracy only for $N \geq 8$.
- Physical order:

$$\boxed{c_{5,1} = 248 \pm 30_{\text{struct}} \pm 62_{\text{meth}} = 248 \pm 70}$$

$c_{6,1}$: rolled, ROLL status — not advanced as a prediction.

What did not

- Strictly recursive at physical order: misses $c_{4,1}$ by -29.9%; closure drift $|\Delta\rho| = 0.0509$.
- At $N = 4$ in large-β_0 the minimal design is rank deficient — the method abstains.
- Every method fails the β-function control.

Conclusions

1. With four coefficients, **method choice dominates the answer**: admissible continuations span $[241.7,\ 304.7]$, diagnostics $[-21.7,\ 3327.8]$; validation is essential.
2. The failure of the standard toolkit is **a theorem, not bad luck**. Escaping it requires a transformation that can reverse the defect sign — physically, an **alternating ultraviolet renormalon**.
3. That class **already exists** (Beneke–Jamin 2008); we claim no new one. CVRC adds an **acceptance criterion**: affine in $\rho \Rightarrow$ the three-weight test selects an interval or *none*, so a declared Borel structure becomes **falsifiable**: six of nine alternatives are rejected. It gives $c_{5,1} = 248 \pm 70$, consistent with 277 ± 51 and the Beneke–Jamin forecast 280.
4. The claim is *weaker* than determination: a single closure value is consistent with everything testable. The -29.9% recursive miss measures the distance still to go.

Outlook for Belle II and STCF

- Higher-precision spectral data can justify a smaller accepted relative tolerance ϵ and a broader weight set. At fixed ϵ and fixed weights, they do not alter the unit-diagonal algebra.
- Weights with different $u = 2$ sensitivity are the most discriminating, and are exactly what a high-statistics spectral function enables.
- A genuine $c_{5,1}$ computation remains the decisive input.

Thank you.

Backup

Rolling the same model forward: $c_{5,1}$ to $c_{12,1}$

n	CVRC $c_{n,1}$	closure window	declared-α_3 range	$\pm$(rel.)	status
5	247.98	[229.76, 266.20]	[235.26, 257.75]	7%	COND
6	3 102.1	[3 059.4, 3 144.9]	[3 021.3, 3 169.9]	2%	ROLL
7	13 003	[10 549, 15 458]	[10 952, 14 644]	19%	ROLL
8	3.7215×10^5	$[3.685, 3.758] \times 10^5$	$[3.638, 3.801] \times 10^5$	2%	ROLL
9	-2.5999×10^5	$[-7.736, 2.537] \times 10^5$	$[-6.555 \times 10^5, 5.7 \times 10^4]$	198%	INDET
10	8.7153×10^7	$[8.233, 9.198] \times 10^7$	$[8.701, 8.783] \times 10^7$	6%	ROLL
11	-9.0648×10^8	$[-1.098, -0.715] \times 10^9$	$[-1.030, -0.810] \times 10^9$	21%	ROLL
12	3.8699×10^{10}	$[3.506, 4.234] \times 10^{10}$	$[3.769, 4.018] \times 10^{10}$	9%	ROLL

- **No controlled values are claimed** for $c_{7,1}$–$c_{12,1}$. These are what the *single calibrated Borel model* returns.
- The roll is a *fixed point* of the closure rule: recomputing ρ at each step returns the same value to $\sim 10^{-17}$.
- Cost of each rolled order, measured *blind* in large-β_0: about half a percent per order, on top of the one-order-ahead error.
- $c_{9,1}$ is flagged sign-indeterminate: its envelope contains zero.

A control that everything fails — and why we report it

Withhold the *known* five-loop $\overline{\mathrm{MS}}$ β-function coefficient at $n_f = 3$ and forecast it from $\beta_0, \ldots, \beta_3$.

- Exact: $\beta_4 = 127.322$.
- Every tested continuation fails the declared ten-percent criterion.
- The least inaccurate member of the predeclared grid is still 65.8% high.
- Direct and Borel Levin-u miss by 155.0% and 177.8%; their higher-order continuations change sign.

Accordingly β_5 and β_6 are retained only as *method diagnostics* and **no average is quoted**.

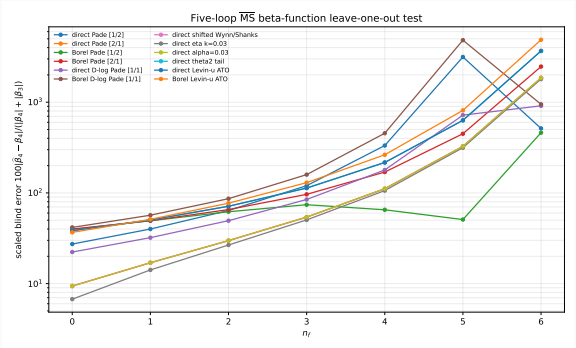

The β function has no renormalon interpretation, so this is a genuine negative control: it catches methods that "work" for reasons unrelated to the physics they claim to capture.

Backup — the closure family and its calibration

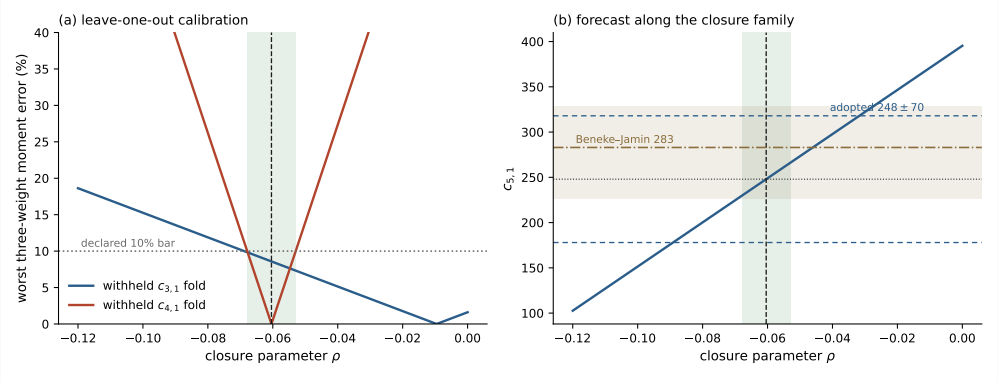

(a) Worst three-weight relative error on each leave-one-order-out fold as a function of ρ; dashed line is $\rho_\star = -0.060414$, shaded band the admissible window $[-0.0679, -0.0529]$. (b) $\widehat{c_{5,1}}(\rho)$ along the same family, with the literature band 277 ± 51, the geometric-ratio value 283, and the adopted result 248 ± 70 as the dashed blue pair — note how much wider the adopted band is than the vertical closure window, which is only its structural part.

Backup — blind large-β_0 forecast error

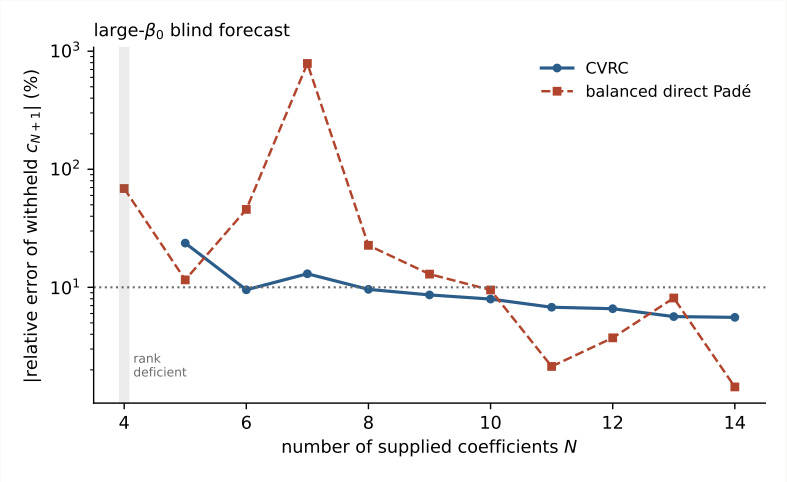

Strictly recursive error for the withheld large-β_0 Adler coefficient: CVRC-minimal versus balanced direct Padé on identical input. The dotted line is the ten-percent criterion; CVRC abstains at the shaded $N = 4$ point.

Backup — why we do *not* add the gluon condensate

The obvious objection: $c_1[\langle aG^2\rangle]$ is known short-distance physics, it is worth $+22.6$ on $c_{5,1}$, and adding it would move $248 \to 280$. Why not declare it?

Because it enters *only* inside the asymptotic bracket of the $u = 2$ cut, and that bracket is not a correction at the orders we actually have:

n	0	1	2	3	4	5	6
bracket	$\mathbf{-2.90}$	$\mathbf{-0.26}$	$\mathbf{+0.27}$	$\mathbf{+0.49}$	$+0.61$	$+0.69$	$+0.74$

Bold = the four orders that are known. It reaches 30% of unity only at $n = 6$, i.e. at $c_{7,1}$.

- At $n = 0$ it *reverses the sign* of the pure power and nearly triples it.
- Cut-corrected structure: admissible set **empty** (gap 0.206), closure drift 0.236 vs 0.051.
- Its blind postdiction of the known $c_{4,1}$: 244.5 against 49.076, an error of $+398\%$.

The criterion rejects a *more complete* structure because the extra physics is not yet numerically operative at $N = 4$. Beneke–Jamin use their model from $c_{6,1}$ upward — exactly where the bracket is meaningful. No disagreement; a statement about the order.