Conditional robust reconstruction from a single two-field dispersion branch

Separate supplement to P01, 23 September 2026

Abstract

We give a conditional finite-error reconstruction certificate for the three-parameter two-field inverse problem of P01. A controlled quadratic-kernel remainder is converted into an isolated-pole error and combined with measurement uncertainty, including shared systematics. An exact structured residual yields componentwise bounds for the reference parameters and heavy scale. Explicit inequalities connect a total error budget, requested reconstruction accuracy and one common momentum/frequency validity domain. Indistinguishable-model bounds show why fixed absolute errors obstruct uniform heavy-scale recovery near the exceptional locus and in shrinking observation windows. Allowed background operators can alter the leading light slope, so engineering dimension alone cannot justify a high-power momentum remainder. The results specialize established inverse and spectral-perturbation methods; they supply neither missing Wilson/loop bounds nor a numerical physical cutoff. A physical heavy-gap interpretation requires a separate heavy-pole certificate. The original exact P01 classification is unchanged.

1 Scope and fixed model

The original P01 manuscript is unchanged. It classifies exact light-pole data in the classical domain

$$a > 0, \qquad B > 0, \qquad 0 < r < 1, \qquad q = 1 - r, \qquad M = rB, \quad N = qB.$$

Writing $x = k^2$ and $y = \omega^2$, its spectral polynomial is

$$F(y, x) = y^2 - [B + (a + 1)x]y + ax^2 + rBx. \tag{1}$$

The lower root satisfies $0 < y_-(x) < x$ for $x > 0$, and its squared heavy gap is $B = y_+(0)$. Three distinct positive known momenta identify (a, B, r) off $a = r$. At $a = r$ the entire light branch is rx for every $B > 0$. These facts are prior results of P01, not new claims of this supplement.

The scalar realization uses $\sigma_0 = 0$, $\theta_0 = vt$ and

$$\mathcal{L} = \tfrac{1}{2}(\partial\sigma)^2 + \tfrac{1}{2}(1 + g\sigma)(\partial\theta)^2 - \tfrac{1}{2}gv^2\sigma - \tfrac{1}{2}m^2\sigma^2 + C\big[(\partial\sigma \cdot \partial\theta)^2 - (\partial\sigma)^2(\partial\theta)^2\big],$$

with $M = m^2$, $N = g^2v^2$ and $a = 1 - 2Cv^2$. We use $Y_s = (\partial\sigma)^2$, $Y_t = (\partial\theta)^2$ and $Z = \partial\sigma \cdot \partial\theta$ below. The reference metric, background convention, field and source normalization are fixed. Gravity is nondynamical. An EFT interpretation adds assumptions about matching, omitted operators, loops and a validity domain; the displayed classical action does not supply those assumptions.

2 A conditional interface for omitted EFT terms

The retained P01 pole and its Taylor remainder do not bound omitted operators. For fixed physical frequency, reference metric, field normalization, background and matching convention, write the corrected quadratic kernel as

$$K = P_0 + E, \qquad P_0 = \begin{pmatrix} y - ax - rB & -iD\omega \\ iD\omega & y - x \end{pmatrix}, \quad y = \omega^2, \quad x = k^2, \quad D^2 = (1-r)B. \tag{2}$$

Here $a > 0$, $B > 0$, $0 < r < 1$. The existence and control of K are additional physical premises. No numerical Wilson-coefficient, loop, background-error or cutoff bounds are supplied by the original action.

2.1 Background insertions and matching

Engineering dimension does not determine the momentum order of a quadratic correction on a rolling background. For example, the allowed operators

$$\delta\mathcal{L} = \frac{c_t Y_t^2 + c_{st} Y_s Y_t + c_Z Z^2}{\Lambda^4} + \frac{\lambda\sigma^2 Y_t}{2\Lambda^2} \tag{3}$$

give, on $\sigma = s$, $\theta = vt + \pi$, with $h = v^2/\Lambda^4$,

$$\delta\mathcal{L}_2 = h\left\{c_t[6\dot{\pi}^2 - 2(\nabla\pi)^2] + c_{st}[\dot{s}^2 - (\nabla s)^2] + c_Z\dot{s}^2\right\} + \frac{\lambda v^2 s^2}{2\Lambda^2}, \tag{4}$$

$$E_{11} = 2h[c_{st}(y-x) + c_Z y] + \lambda v^2/\Lambda^2, \qquad E_{22} = 4hc_t(3y - x), \qquad E_{12} = E_{21} = 0. \tag{5}$$

For the last operator alone, put $d = \lambda v^2/\Lambda^2$ and suppose $M - d > 0$, where $M = rB$. The corrected reference parameters are $M' = M - d$, $B' = B - d$, and

$$r' - r = -\frac{d(B-M)}{B(B-d)}, \qquad y'_-(x) - y_-(x) = (r' - r)x + O(x^2). \tag{6}$$

For nonzero allowed d, this rules out a remainder bounded by a finite constant times x^p, $p > 1$, on an interval approaching zero. No numerical value of λ is assumed.

One may either keep the original reference parameters and include all such terms in E, or define matched parameters by explicit renormalized conditions. Terms absorbed by those conditions must not be counted twice. A generic corrected quadratic action need not lie in the same three-parameter family: a changed intrinsic π speed cannot be reset to one by field normalization in the fixed reference metric. Field redefinitions require the induced higher-order terms to be retained or bounded [1, eqs. (5.1)–(5.3)].

2.2 A hard kernel envelope requires a domain contract

Let an omitted operator of dimension d_j have coefficient $c_j\Lambda^{4-d_j}$ and quadratic Fourier Hessian $Q_{j,ab}$. Suppose independently justified bounds $|c_j| \le C_j$ are given, and the kernel discrepancy equals an absolutely uniformly convergent operator sum plus a remainder L. On a declared domain T,

$$h_{ab}(T) = \sum_j C_j\Lambda^{4-d_j} \sup_T |Q_{j,ab}| + \sup_T |L_{ab}| \quad \Longrightarrow \quad \sup_T |E_{ab}| \le h_{ab}(T). \tag{7}$$

Operator multiplicities and Hessian factors belong in this sum. L includes any unrepresented loop, matching or background error; setting it to zero requires justification.

A useful conditional implementation is $E_{ab} = \Lambda^2 \sum_\alpha d_{ab,\alpha} t^\alpha + L_{ab}$, with $|d_{ab,\alpha}| \le W_{ab,\alpha}$ and $|t_j| \le q_j$. Then $h_{ab} \le \Lambda^2 \sum_\alpha W_{ab,\alpha} q^\alpha + \ell_{ab}$. If the omitted total degrees start at n_0, all $q_j \le q$, and the aggregate shell satisfies $\sum_{|\alpha|=n} W_{ab,\alpha} \le A_{ab} R^n$, then

$$h_{ab} \le \Lambda^2 A_{ab} \frac{(Rq)^{n_0}}{1 - Rq} + \ell_{ab}, \qquad Rq < 1. \tag{8}$$

This assumes an actual convergent representation and a hard coefficient envelope. It does not follow from power counting or naturalness. Background ratios can carry total degree without adding external momentum powers.

The domain contract must specify a stationary corrected background, local field normalization, matching/renormalization convention, analytic sheet, thresholds, and quantitative control of momentum, frequency, masses and background ratios. With declared domain sizes K_* and Ω_*, every root tube used below must satisfy

$$\sqrt{x} \le K_*, \qquad \sqrt{y_-(x) + \rho(x)} \le \Omega_*. \tag{9}$$

The ratios K_*/Λ, Ω_*/Λ, m/Λ, $|D|/\Lambda$ and $|v|/\Lambda^2$ must also lie in the independently proved expansion domain. Merely declaring them small is not a remainder theorem. In particular, massless thresholds or dissipative loop corrections may obstruct the real isolated-pole assumptions below.

2.3 From determinant error to an isolated pole

Fix $x > 0$ and a parameter candidate. Set $y_0 = y_-(x)$, $g = y_+(x) - y_0 = \sqrt{\Delta(x)} > 0$, and $F(y, x) = (y - y_0)(y - y_0 - g)$. Choose $0 < \rho < \min(y_0, g/2)$ and use the analytic square-root branch positive at y_0 on $T = \{|y - y_0| \le \rho\}$. Define

$$p_{11} = |y_0 - ax - rB| + \rho, \qquad p_{22} = |y_0 - x| + \rho, \tag{10}$$

$$p_{12} = p_{21} = |D|\sqrt{y_0 + \rho}. \tag{11}$$

For entry bounds $|E_{ab}| \le h_{ab}$ on T, the determinant identity gives

$$\begin{aligned} R &= \det(P_0 + E) - F \\ &= P_{22}E_{11} + P_{11}E_{22} - P_{12}E_{21} - P_{21}E_{12} + E_{11}E_{22} - E_{12}E_{21}, \qquad (12) \\ |R| &\le \epsilon_F := p_{22}h_{11} + p_{11}h_{22} + p_{12}h_{21} + p_{21}h_{12} + h_{11}h_{22} + h_{12}h_{21}. \qquad (13) \end{aligned}$$

Here h_{ab} has mass dimension two and ϵ_F has dimension four.

Conditional pole certificate. Suppose $G = F + R$ is holomorphic on a neighborhood of T and

$$\sup_T |R| \le \epsilon_F < \rho(g - \rho). \tag{14}$$

Then G has exactly one zero in T, counted with multiplicity, and

$$|\widetilde{y} - y_0| \le e_{\mathrm{EFT}}(x; p) := \frac{2\epsilon_F}{g + \sqrt{g^2 - 4\epsilon_F}} < \rho. \tag{15}$$

If $G(\overline{y}) = \overline{G(y)}$, the zero is real, positive and simple. Indeed, on the boundary $|F| \ge \rho(g - \rho) > |R|$, so Rouché's theorem preserves the single zero. At the perturbed zero, $t(g - t) \le \epsilon_F$, where $t = |\widetilde{y} - y_0| < \rho < g/2$. Monotonicity of $t(g - t)$ yields (15). Conjugation and uniqueness give reality; $\rho < y_0$ gives positivity. This is an explicit specialization of standard nonlinear spectral localization [2, Lemma 2.1 and Theorem 4.4].

Alternatively, a real C^1 determinant on $[y_0 - \rho, y_0 + \rho]$, the same supremum bound, and $\sup |\partial_y R| < g - 2\rho$ imply the same unique-root conclusion: the endpoint signs differ and $\partial_y G < 0$ throughout. A bound only on the unperturbed shell is insufficient. For example, $R = -F + \varepsilon$ has arbitrarily small $R(y_0) = \varepsilon \neq 0$ but $G = \varepsilon$ has no zero.

The enclosed zero is the continuation of the P01 light pole. To identify it as the lowest positive pole of a more general kernel, exclude additional zeros and kernel singularities below the tube; one sufficient real test is $|R| < F$ on $0 \leq y \leq y_0 - \rho$. The local theorem is not a global ghost-freedom statement. Source observability and any positive spectral-weight claim require additional premises or certificates: $K = -P_0$ has the same determinant but reverses response residues. Without Schwarz reality it encloses a complex zero, whose relationship to a real spectral peak needs another observation model.

At $a = r$, the forward gap is $g = B + (1 - r)x > 0$, despite the inverse ambiguity in B. Thus the forward gap cannot replace separation from the inverse degeneracy. The light disk theorem itself is restricted to $x > 0$.

2.4 Heavy gap and joint inverse inputs

The reconstructed B is a matched reference parameter. To interpret it as the corrected heavy squared gap, require a separate tube centered at B at $x = 0$. If $0 < \rho_H < B/2$, the analogous hypotheses and $\epsilon_H < \rho_H(B - \rho_H)$ hold, then

$$|\widetilde{B} - B| \leq e_H = \frac{2\epsilon_H}{B + \sqrt{B^2 - 4\epsilon_H}}. \tag{16}$$

This requires domain control up to $|\omega| \leq \sqrt{B + \rho_H}$ and a justified heavy-mode label. It does not follow from a light-tube bound. An interval $B \in [B_L, B_U]$ and a uniform $e_H^* < B_L$ give a heavy frequency interval $[\sqrt{B_L - e_H^*}, \sqrt{B_U + e_H^*}]$. The mixing-enhanced gap is already present in the forward literature [3, eqs. (7)–(11)]; a single-field reduction scale is not automatically the retained two-field UV cutoff.

At exact momenta, bounded squared-frequency measurement error combines as

$$|\widehat{y}_i - y_-(x_i; p)| \leq e_{\mathrm{meas},i} + e_{\mathrm{EFT}}(x_i; p). \tag{17}$$

For P01's fixed mass normalization μ, divide this sum by μ^2 before using the dimensionless inverse certificate. If the measured error is $|\omega - \widehat{\omega}| \leq \eta$, first use $|\omega^2 - \widehat{\omega}^2| \leq 2|\widehat{\omega}|\eta + \eta^2$. Parameter-dependent envelopes require a proved supremum over a justified candidate set or a candidate-dependent feasibility calculation. Evaluation only at an unvalidated fit is circular.

Correlated deterministic measurement and model sets combine by a Minkowski sum; pointwise boxes are outer enclosures. Set-membership inference is established methodology [4, eqs. (1)–(4)], as is the distinction between structural and practical identifiability [5, Sec. 2.3]. Statistical EFT truncation models [6, 7] require their own priors; their credible intervals do not supply hard support bounds for P01.

The unequal-speed forward spectrum also predates P01 [8, eqs. (4.17), (4.21)]. The present contribution is a conditional specialization connecting controlled kernel errors to this inverse problem. The same parameter set, data nodes and error budgets must satisfy the expansion-domain, pole-gap, inverse-conditioning, admissibility and desired heavy-scale-accuracy conditions. The available P01 inputs do not establish a nonempty physical window satisfying them, nor a comprehensive priority claim for this composition.

3 Finite reconstruction with joint observation and EFT pole errors

The following result is conditional on a declared observation model and a separately certified omitted-EFT pole remainder. It supplies neither physical Wilson-coefficient bounds nor a cutoff, and makes no novelty claim about general matrix perturbation methods. The original P01 exact three-point identifiability result is an input.

3.1 Units, error sets, and the EFT interface

Let $x_i = k_i^2$ be three distinct positive, exactly known squared momenta. Fix a known reference mass $\mu > 0$, independently of the inferred parameters, and put

$$u_i = x_i/\mu^2, \quad v_i = y_-(x_i)/\mu^2, \quad \widehat{v}_i = \widehat{y}_i/\mu^2, \quad p = (a, \beta, m)^T = (a, B/\mu^2, M/\mu^2)^T. \tag{18}$$

The domain is $a > 0$, $\beta > m > 0$; $r = m/\beta$ and $N = B - M$. Here $y = \omega^2$. The observation contract is

$$\widehat{y}_i = y_{\mathrm{phys}}(x_i) + e_i^{\mathrm{obs}}, \qquad y_{\mathrm{phys}}(x_i) = y_-(x_i; p_*) + e_i^{\mathrm{EFT}}. \tag{19}$$

For a stated joint set $\mathcal{J}$ of measurement errors and EFT pole shifts, define the dimensionless total error set

$$\delta = v - \widehat{v} \in \mathcal{E} = -\mu^{-2}\{e^{\mathrm{obs}} + e^{\mathrm{EFT}} : (e^{\mathrm{obs}}, e^{\mathrm{EFT}}) \in \mathcal{J}\}. \tag{20}$$

Marginal deterministic sets may be combined by a Minkowski sum, without an independence assumption. Known shared calibration variables and joint restrictions should be retained. Component envelopes $|\delta_i| \le e_i = (\epsilon_i^{\mathrm{obs}} + \epsilon_i^{\mathrm{EFT}})/\mu^2$ are sufficient but can lose correlation information. Momentum calibration errors are outside this theorem, including an unknown common momentum scale.

The EFT input must bound the displacement of the relevant continuation pole in squared frequency. For example, a separately proved root-tube certificate may give, with $g = y_+ - y_- > 0$,

$$\epsilon^{\mathrm{EFT}}(x; p) = \frac{2\epsilon_F(x; p)}{g(x; p) + \sqrt{g(x; p)^2 - 4\epsilon_F(x; p)}}. \tag{21}$$

Its hypotheses include control of the determinant discrepancy throughout a valid tube $|y - y_-| \le \rho$, with $0 < \rho < \min(y_-, g/2)$ and $\epsilon_F < \rho(g - \rho)$, and the relevant analyticity and reality hypotheses, or an appropriate real derivative alternative. Those hypotheses are supplied by the separate EFT argument; they are not consequences of reconstruction. A tube certificate identifies the continuation pole, without by itself excluding additional lower poles of a more general kernel.

If the majorant depends on p, a uniform error set requires a supremum over an independently justified candidate set containing p_*. Alternatively, retain $\delta(p) \in \mathcal{E}(p)$ in a set-membership computation. A majorant evaluated only at the fitted center is insufficient. Using the resulting error box to justify its own EFT majorant needs a separate localization or exclusion argument. Naturalness and power counting alone provide no hard numerical error bound. For frequency errors the finite conversion is $|\omega_{\mathrm{phys}}^2 - \widehat{\omega}^2| \le 2|\widehat{\omega}|\epsilon_\omega + \epsilon_\omega^2$; correlations must also be transformed. Complex or damped poles require a different observation model.

3.2 An exact structured residual and a finite theorem

For $z_i = v_i/u_i$, define

$$A(v)_{i\cdot} = [u_i(1 - z_i), -z_i, 1], \qquad f(v)_i = u_i z_i(1 - z_i). \tag{22}$$

Every matching reference model obeys $A(v)p = f(v)$. If $\widehat{A} = A(\widehat{v})$ is invertible, form

$$T = \widehat{A}^{-1}, \qquad \widehat{p} = Tf(\widehat{v}), \qquad d_i = \widehat{a} + 1 - 2\widehat{z}_i + \widehat{\beta}/u_i. \tag{23}$$

Put

$$L(\delta)_{i\cdot} = [\delta_i, \delta_i/u_i, 0], \qquad \psi_i(\delta) = d_i\delta_i - \delta_i^2/u_i. \tag{24}$$

Direct substitution yields the exact identity

$$(I - TL(\delta))(p - \widehat{p}) = T\psi(\delta). \tag{25}$$

It preserves the dependence between data-matrix and right-hand-side errors.

Finite reconstruction theorem. Let $\mathcal{E}$ be nonempty and bounded. Choose certified nonnegative bounds

$$K_{j1} \geq \sup_{\delta\in\mathcal{E}} \left|\sum_i T_{ji}\delta_i\right|, \qquad K_{j2} \geq \sup_{\delta\in\mathcal{E}} \left|\sum_i T_{ji}\delta_i/u_i\right|, \qquad K_{j3} = 0, \tag{26}$$

$$g_j \geq \sup_{\delta\in\mathcal{E}} |(T\psi(\delta))_j|. \tag{27}$$

If the spectral radius $\rho(K) < 1$, every admissible reference model consistent with the declared observation/error contract obeys

$$|p - \widehat{p}| \leq R = (I - K)^{-1}g \tag{28}$$

componentwise. Moreover, $A(\widehat{v} + \delta)$ is nonsingular for every $\delta \in \mathcal{E}$.

Proof. Equation (25) implies $|p - \widehat{p}| \leq g + K|p - \widehat{p}|$. The inverse $(I - K)^{-1} = \sum_{n\geq 0} K^n$ exists and is nonnegative, proving the bound. If $(I - TL)w = 0$, then $|w| \leq K|w|$ forces $w = 0$ by the same argument. This proves the nonsingularity statement. All data-error powers in the original quadratic equations are retained. □

Only a two-dimensional test is needed:

$$K_{11} < 1, \quad K_{22} < 1, \quad \mathcal{D} = (1 - K_{11})(1 - K_{22}) - K_{12}K_{21} > 0. \tag{29}$$

For this nonnegative K with zero third column, these conditions are equivalent to $\rho(K) < 1$. They give

$$R_a = \frac{(1 - K_{22})g_1 + K_{12}g_2}{\mathcal{D}}, \qquad R_\beta = \frac{K_{21}g_1 + (1 - K_{11})g_2}{\mathcal{D}}, \tag{30}$$

$$R_m = g_3 + K_{31}R_a + K_{32}R_\beta. \tag{31}$$

Rational arithmetic or outward-rounded interval bounds can verify the strict inequalities. Failure is inconclusive, not a physical instability or a general impossibility result.

For a box $|\delta_i| \leq e_i$, one valid choice is

$$K_{j1} = \sum_i |T_{ji}|e_i, \qquad K_{j2} = \sum_i |T_{ji}|e_i/u_i, \qquad g_j = \sum_i |T_{ji}|(|d_i|e_i + e_i^2/u_i). \tag{32}$$

For correlated errors $\delta = H\xi$, $|\xi_l| \leq 1$, define for a row w

$$k_1(w) = \sum_l |\sum_i w_i H_{il}|, \qquad k_2(w) = \sum_l |\sum_i w_i H_{il}/u_i|, \tag{33}$$

$$G(w) = \sum_l |\sum_i w_i d_i H_{il}| + \sum_l |\sum_i w_i H_{il}^2/u_i| + 2\sum_{l<h} |\sum_i w_i H_{il}H_{ih}/u_i|. \tag{34}$$

Use $K_{j\cdot} = (k_1(T_{j\cdot}), k_2(T_{j\cdot}), 0)$ and $g_j = G(T_{j\cdot})$. Linear supports are exact; G is a valid quadratic triangle bound. For a single shared nuisance $\delta = st$, $|s| \le \epsilon$, the exact residual support is

$$g_j = \epsilon |\sum_i T_{ji} d_i t_i| + \epsilon^2 |\sum_i T_{ji} t_i^2 / u_i|. \tag{35}$$

Other error geometries use (27) directly.

3.3 Output parameters and prescribed target accuracies

For any fixed row h, form bounds g_h, K_{h1}, K_{h2} by replacing the row $T_{j\cdot}$ by hT in (27). Then

$$|h(p - \widehat{p})| \le E_h = \min\{\sum_j |h_j| R_j,\ g_h + K_{h1} R_a + K_{h2} R_\beta\}. \tag{36}$$

Write $R_n = E_{(0,1,-1)}$, $\widehat{r} = \widehat{m}/\widehat{\beta}$. Sufficient admissibility margins for the certified enclosure, including its n projection, are

$$\widehat{a} > R_a, \quad \widehat{\beta} > R_\beta, \quad \widehat{m} > R_m, \quad \widehat{\beta} - \widehat{m} > R_n. \tag{37}$$

The finite output bounds are

$$|a - \widehat{a}| \le R_a, \qquad |B - \widehat{B}| \le \mu^2 R_\beta, \tag{38}$$

$$|M - \widehat{M}| \le \mu^2 R_m, \qquad |N - \widehat{N}| \le \mu^2 R_n, \tag{39}$$

$$|r - \widehat{r}| \le R_r = \frac{E_{(0,-\widehat{r},1)}}{\widehat{\beta} - R_\beta}, \qquad |\sqrt{B} - \sqrt{\widehat{B}}| \le R_H = \frac{\mu R_\beta}{\sqrt{\widehat{\beta}} + \sqrt{\widehat{\beta} - R_\beta}}. \tag{40}$$

The ratio identity is $r - \widehat{r} = (\Delta m - \widehat{r}\Delta\beta)/(\widehat{\beta} + \Delta\beta)$. The square-root expression is the largest deviation on the certified symmetric β interval. A true-relative B error is bounded by $R_\beta/(\widehat{\beta} - R_\beta)$. Given absolute targets $\tau_a, \tau_B, \tau_r, \tau_M, \tau_N, \tau_H$, check (29), (37), and

$$(R_a, \mu^2 R_\beta, R_r, \mu^2 R_m, \mu^2 R_n, R_H) \le (\tau_a, \tau_B, \tau_r, \tau_M, \tau_N, \tau_H). \tag{41}$$

Also $|\widehat{a} - \widehat{r}| > R_a + R_r$ is a convenient explicit separation margin from the exceptional locus.

Here $\sqrt{B}$ is a matched reference heavy-gap parameter. A separately certified physical heavy root at $x = 0$ with $|Y_H^{\text{phys}} - B| \le \eta_H$ permits, with $T_H = \mu^2 R_\beta + \eta_H < \widehat{B}$,

$$|\sqrt{Y_H^{\text{phys}}} - \sqrt{\widehat{B}}| \le \frac{T_H}{\sqrt{\widehat{B}} + \sqrt{\widehat{B} - T_H}}. \tag{42}$$

A light-root bound does not provide that heavy-root certificate or ensure that the required heavy-energy tube is within the EFT domain.

The theorem bounds every compatible model; an empty compatibility set is not successful physical inference. An admissible three-point algebraic fit does not establish the EFT assumptions or adequacy over an entire dispersion curve. Existence/model adequacy must be supplied separately for a statement about the physical system; further observations can test it.

3.4 Comparison with the original P01 bound

Let $E_A = \max_i e_i(1 + 1/u_i)$, $E_f = \max_i(|1 - 2\widehat{z}_i|e_i + e_i^2/u_i)$ and $\alpha = \|T\|_\infty E_A$. For (32), with $\iota = (1, 1, 1)^T$,

$$K\iota \le \alpha\iota, \qquad g \le \|T\|_\infty (E_f + E_A \|\widehat{p}\|_\infty)\iota. \tag{43}$$

Therefore whenever the original P01 condition $\alpha < 1$ holds, the new condition also holds, and each new coordinate radius is no greater than the original scalar bound. This is a dominance result for the displayed certificates, not an optimality theorem. Correlation can improve it further. An exact rational witness uses $a = 4/5$, $\beta = 2$, $r = 1/3$ and $u = (5/9, 20/9, 20/3)$. For shared additive error $\delta = s(1, 1, 1)$, $|s| \le 1/150$, the original test has $\alpha = 1.12$, whereas the structured test passes with $(R_a, R_\beta, R_m) \simeq (0.021643, 0.398095, 0.187692)$ and $R_r \simeq 0.034330$. These are synthetic arithmetic witnesses, not measurements or a physical error model.

3.5 Local conditioning and the limits of uniform recovery

At exact admissible data, put $D = \operatorname{diag}(a + 1 - 2z_i + \beta/u_i)$. Its entries equal the positive dimensionless pole gap divided by u_i, and implicit differentiation gives

$$\mathrm{d}p = A(v)^{-1} D \, \mathrm{d}v. \tag{44}$$

For a fixed interior model with $a \neq r$ and distinct positive nodes this is finite. The finite theorem gives a nonzero error neighborhood. On compact parameter sets separated from $a = r$ and the domain boundaries, with β bounded above and below and nodes in a compact positive interval with a positive separation, continuity gives uniform finite conditioning and a sufficiently small uniform error radius. It does not guarantee experimentally useful tolerances.

At $a = r$, the entire reference light branch is $v(u) = ru$, independently of β. If the declared error set contains such data, their matrix is singular. Consequently no sound K satisfying the theorem's test can certify that set: apparent rank generated by noise cannot restore the missing information.

For a quantitative approach to this locus, fix $r \in (0, 1)$, distinct $\beta_1, \beta_2 > 0$, and $a = r + \varepsilon > 0$. The P01 identity $u(1 - z)(a - z) = \beta(z - r)$ and monotonicity imply

$$|v_j(u) - ru| \le |\varepsilon| u^2/\beta_j, \qquad \sup_{0 \le u \le U} |v_1(u) - v_2(u)| \le |\varepsilon| U^2 (1/\beta_1 + 1/\beta_2). \tag{45}$$

The parameter separation stays fixed while the complete light-pole traces approach each other. If a deterministic sup-norm error radius is at least half this bound, a midpoint trace is compatible with both candidates. No estimator using only this experiment can guarantee absolute β error less than $|\beta_1 - \beta_2|/2$ for both. This assertion requires error sets containing the two deviations; it does not apply automatically to every restricted correlation geometry, other observables, or additional priors.

There is a separate low-momentum obstruction. Fix r and nonzero b, and choose two admissible models with $a_j = r + \beta_j b/(1 - r)$. They share

$$v_j(u) = ru + bu^2 + c_j u^3 + O(u^4), \qquad c_j = -b(1 - r)/\beta_j - b^2/(1 - r). \tag{46}$$

If $K_{3j} \ge \sup_{[0,U]} |v_j'''|$, Taylor's theorem gives the finite bound

$$\sup_{[0,U]} |v_1 - v_2| \le (K_{31} + K_{32}) U^3/6. \tag{47}$$

For $a_j \ge r$, one explicit bound is

$$K_{3j} = \frac{3(1 - r)|a_j - r|}{\beta_j^3} [2\beta_j (a_j + 1 - 2r) + 2(a_j - 1)^2 U], \tag{48}$$

using $\Delta \ge \beta_j^2$. The same midpoint argument applies. For nodes $u_i = ht_i$ of fixed distinct positive ratios, the exact determinant has leading term

$$\det A = \frac{b(1 - r)^2}{\beta} h^3 \prod_{i<j} (t_j - t_i) + O(h^4). \tag{49}$$

The finite bound above, rather than an asymptotic estimate alone, supplies the obstruction at fixed error tolerance.

A narrow band around $u_0 > 0$ is also distinct from a low absolute momentum. Fix admissible z_0 and $s = z'(u_0)$ from a nonexceptional model. A local family sharing these data satisfies

$$[su_0 - (1 - z_0)]a + s\beta = -(1 - z_0)z_0 - su_0(1 - 2z_0), \tag{50}$$

$$m = \beta z_0 - u_0(1 - z_0)(a - z_0). \tag{51}$$

The coefficient $su_0 - (1 - z_0)$ equals $-(1 - z_0)[\beta + u_0(1 - z_0)]/[\beta + u_0(a + 1 - 2z_0)] < 0$. Thus β can vary locally while admissibility persists and both $v(u_0), v'(u_0)$ remain fixed. If $|v_j''| \le K_{2j}$ on $|u - u_0| \le w$, then

$$\sup |v_1 - v_2| \le (K_{21} + K_{22})w^2/2. \tag{52}$$

For $a_j \ge r_j$, $K_{2j} = 2(1 - r_j)|a_j - r_j|/\beta_j$ suffices. This is an experiment-specific obstruction, not a proof that all methods or additional response data must fail.

3.6 A common momentum/energy window and probabilistic inputs

A useful window must simultaneously meet the following requirements. Three distinct, exactly calibrated nodes satisfy $0 < x_i \le K_{\max}^2$. For every relevant candidate, the complete EFT root tube lies in the separately proved momentum, frequency, and background domain; in particular $y_-(x_i; p) + \rho_i(p) \le \Omega_{\max}^2$. The determinant and root-tube hypotheses hold uniformly, or remain explicit in candidate-dependent membership. The combined error set then passes (29), (37), and the target test (41). Where the EFT assumptions require the explicit heavy scale to be below a specified energy limit, the upper enclosure $\mu\sqrt{\widehat{\beta} + R_\beta}$ must also respect that limit; $\sqrt{M}, \sqrt{N} \le \sqrt{B}$ in the admissible reference model. A physical-heavy-gap claim needs its separate root tube. Since $0 < y_- < x$, a conservative reference-frequency envelope is $\Omega_{\max}^2 \ge K_{\max}^2 + \sup \rho$, but this does not itself prove EFT validity. There is no unconditional assertion that such a common window is nonempty. An upper EFT limit can conflict with the resolution requirements just established.

A covariance matrix alone is not a deterministic error bound. For an explicitly mean-zero measurement vector with known finite covariance Σ of rank d, the set $e^T\Sigma^\dagger e \le d/\alpha$, restricted to the covariance range, has coverage at least $1 - \alpha$ by Markov's inequality applied to a nonnegative variable of mean d. Under the stronger centered-Gaussian assumption, the threshold is the corresponding chi-square quantile. Unknown bias needs another nuisance set; estimated covariance needs a separate coverage argument. Composing a measurement confidence set with a uniform deterministic EFT set preserves its coverage. If both input sets are probabilistic, a union bound gives coverage at least $1 - \alpha_{\text{obs}} - \alpha_{\text{EFT}}$ without independence. These probability statements depend on the declared assumptions; the deterministic inverse theorem supplies no confidence level by itself.

Reproducibility. The accompanying standard-library Python code and rational input file verify the exact residual, finite output bounds, correlation-preserving supports, four finite nuisance grids with 114 points in total, the old-bound comparison, exceptional-set failure, and the three window witnesses. Directed rational root enclosures are used for square roots and forward poles. Finite sampling tests the implementation; the arguments above prove the continuous-set guarantees. The examples are synthetic and have no asserted physical realism. No physical data or EFT coefficient bounds are manufactured.

4 One joint applicability test

The preceding results must be applied to the same parameter set, physical domain and observations. Separate measurement and EFT windows do not suffice. The following conservative version uses the original P01 matrix certificate; the structured certificate above can replace it when its hypotheses hold.

Theorem 1 (Conditional joint applicability)**.** *Let Θ be an independently justified set containing one common admissible reference triple $p_* = (a, B, r)$ for all observations. Fix the physical reference mass $\mu > 0$ and three distinct, exactly known $x_i > 0$. Suppose:*

1. *The corrected kernel, background and observation conventions satisfy the EFT domain and pole hypotheses at every candidate in Θ. The observed real modes are the certified continuation poles, with nonzero response in the measured channel. Every pole tube and all required background, retained-mass and mixing scales lie in the independently justified validity domain below its applicable cutoff. Stability and physical mode identification are assumed where not established by the local root theorem.*

2. *The EFT pole errors have valid uniform bounds $\tau_i \geq \sup_{p\in\Theta} e_{\mathrm{EFT}}(x_i; p)$, and the measurement errors in squared frequency obey a simultaneous event $|n_i| \leq \varepsilon_i$. Thus $e_i = (\varepsilon_i + \tau_i)/\mu^2$ bounds the total dimensionless error. A justified joint uncertainty set may replace this box.*

3. *Put $u_i = x_i/\mu^2$, $\widehat{z}_i = \widehat{y}_i/x_i$, and form*

$$\widehat{A}_{i\cdot} = [u_i(1-\widehat{z}_i), -\widehat{z}_i, 1], \quad \widehat{f}_i = u_i\widehat{z}_i(1-\widehat{z}_i), \quad \widehat{p} = \widehat{A}^{-1}\widehat{f} = (\widehat{a}, \widehat{\beta}, \widehat{m})^T.$$

   Here $\beta = B/\mu^2$ and $m = M/\mu^2$; this dimensionless m is not the original dimensional bare mass. Require $\widehat{A}$ nonsingular and define

$$E_A = \max_i e_i(1 + 1/u_i), \tag{53}$$

$$E_f = \max_i \left(|1 - 2\widehat{z}_i|e_i + e_i^2/u_i\right), \tag{54}$$

$$\alpha = \|\widehat{A}^{-1}\|_\infty E_A < 1, \tag{55}$$

$$E_p = \frac{\|\widehat{A}^{-1}\|_\infty(E_f + E_A\|\widehat{p}\|_\infty)}{1-\alpha}. \tag{56}$$

4. *The sufficient domain margins hold:*

$$\widehat{a} > E_p, \qquad \widehat{m} > E_p, \qquad \widehat{\beta} - \widehat{m} > 2E_p.$$

   For a requested true-relative accuracy $\delta_B > 0$, also require

$$\frac{E_p}{\widehat{\beta} - E_p} \leq \delta_B. \tag{57}$$

Then on the stated simultaneous event the reference parameters satisfy

$$|a - \widehat{a}| \leq E_p, \quad |B - \mu^2\widehat{\beta}| \leq \mu^2 E_p, \quad |M - \mu^2\widehat{m}| \leq \mu^2 E_p,$$

and $|\mu^2\widehat{\beta}/B - 1| \leq \delta_B$. In particular,

$$\sqrt{B} \in \mu\,[\sqrt{\widehat{\beta} - E_p}, \sqrt{\widehat{\beta} + E_p}]. \tag{58}$$

Writing $\widehat{r} = \widehat{m}/\widehat{\beta}$, a conservative ratio bound is

$$|r - \widehat{r}| \le \frac{E_p(1 + |\widehat{r}|)}{\widehat{\beta} - E_p}.$$

The analogous bound for $N = B - M$ is $2\mu^2 E_p$; the structured projections give sharper alternatives. These are reference-parameter statements. A physical heavy-gap interval additionally requires the separate heavy-pole remainder certificate.

Proof. The pole theorem and the observation assumption put the same reference model's three pole values within the total error box. The triangle inequality needs no statistical independence. Applying P01's finite matrix perturbation bound yields (56). The domain margins put the entire parameter box in the open admissible domain. Since $B/\mu^2 \ge \widehat{\beta} - E_p > 0$, division gives (57); monotonicity of the square root gives (58). The ratio identity

$$r - \widehat{r} = \frac{(m - \widehat{m}) - \widehat{r}(\beta - \widehat{\beta})}{\beta}$$

gives the displayed ratio bound. No step establishes the existence of the common model or the physical majorants from the measured fit. □

If a simultaneous measurement event has probability at least $1 - \gamma$, the conclusion holds on that event with the same coverage conditional on the deterministic EFT assumptions. Covariance alone does not define such a bounded event. A probabilistic EFT model needs its own joint event and coverage accounting. Unknown systematic offsets cannot be silently set to zero.

4.1 A quantitative allocation from kernel control to inverse accuracy

For a selected node design, one can first choose total squared-pole budgets that pass the finite inverse and target-accuracy inequalities, and then check whether measurement and EFT inputs fit them. This is an acceptance test for supplied bounds, not a prescription to invent them.

For clarity let the rigorous entry envelope have the form $h_{ab} \le t\,\bar{h}_{ab}$ on each fixed tube, uniformly over Θ, where the dimensionless $t \ge 0$ is a certified bound. The determinant formula yields

$$\epsilon_F(x_i; p) \le tL_i + t^2 Q_i,$$

using uniform bounds L_i on the sum of the four linear products and Q_i on the two quadratic products. Let $G_i > 0$ be a uniform lower bound on the squared-pole gap. Choose an allowed EFT pole error $0 < \tau_i < \inf_{p\in\Theta} \rho_i(p)$ with $\tau_i < G_i/2$. Then

$$tL_i + t^2 Q_i \le C_i := \tau_i(G_i - \tau_i) \tag{59}$$

implies the strict tube test and $e_{\mathrm{EFT}}(x_i; p) \le \tau_i$ for every candidate. A concrete sufficient threshold, if $L_i + Q_i > 0$, is

$$t \le \frac{2C_i}{L_i + \sqrt{L_i^2 + 4Q_iC_i}}. \tag{60}$$

If $L_i = Q_i = 0$, the determinant bound imposes no additional restriction on t. The common t must pass all nodes and all independently required expansion and background conditions. Known contributions that do not scale with t, including loops or matching uncertainty, remain in the budget and must be included before using this simplified allocation.

To see (59), $s(G - s)$ is increasing for $0 \le s < G/2$. Thus $\tau_i < G_i/2$ and $\tau_i < \rho_i(p) < G(p)/2$ give

$$\epsilon_F \le \tau_i(G_i - \tau_i) \le \tau_i(G(p) - \tau_i) < \rho_i(p)(G(p) - \rho_i(p)).$$

The pole theorem then gives the desired error. Solving the scalar quadratic inequality gives (60).

4.2 What a nonempty window means

Define the acceptable design set to be all ordered triples of distinct positive nodes inside the same validated momentum/frequency domain for which the total-error certificate, domain margins and requested accuracy pass. It is not necessarily an interval in momentum. A continuous-window statement additionally requires the EFT/domain bounds for every momentum in that interval; three samples alone do not establish that statement.

At a fixed interior off-exception model and three fixed distinct nodes with strict domain and pole-tube margins, sufficiently small measurement errors and sufficiently small controlled kernel discrepancies pass any fixed positive parameter-accuracy targets: the exact data matrix is invertible, the bounds vanish continuously with the two errors, and the strict margins persist. Equations (56)–(60) supply an explicit finite test rather than a universal numerical threshold. This local nonvacuity does not survive uniformly as the design or model approaches an inverse degeneracy, nor does it prove that physical uncertainties can be made sufficiently small.

The supplied P01 materials provide no numerical physical cutoff, Wilson/loop majorant or measurement-error event. Accordingly no numerical physical momentum or energy window is certified here. Failure of a sufficient test is inconclusive; the separately proved indistinguishable-model constructions establish impossibility only for their named observation and error classes.

5 Relation to established methods and remaining physical inputs

The full targeted source comparison and exact inspected passages are documented in the accompanying literature report. The mixed two-field forward propagator and its poles are prior results: the determinant in Pinol, Renaux-Petel and Werth [8, eqs. (4.17)–(4.21)] maps algebraically to (1). This formula comparison is not a change of the calibrated reference metric or an assertion of identical phenomenological assumptions. Achucarro et al. [3, eqs. (7)–(11)] already distinguish the mixing-enhanced gap and the scale for a single-field reduction. That reduction scale does not by itself supply the UV validity cutoff of the retained two-field theory.

Bounded-error parameter inference by set inversion is established [4]; structural versus practical identifiability is likewise an established distinction [5]. Adding measurement and theory discrepancies and retaining correlated EFT remainders also precede this supplement [7]. The naturalness priors of EFT uncertainty analyses [6, 7] do not imply deterministic bounds for every P01 completion. Analytic matrix perturbation and boundary homotopy provide the general zero-localization framework [2]. The elementary pole-radius formula used here specializes that framework to a quadratic baseline determinant. Consistent higher-order field redefinitions and matching are addressed by Criado and Perez-Victoria [1].

The contribution is consequently a model-specific composition: a declared operator/kernel envelope is propagated to isolated pole errors and then to a finite inverse certificate, with explicit observation, matching, heavy-gap and domain qualifications. The symbolic background examples identify why high operator dimension alone does not force a high external-momentum remainder. No new general inverse or perturbation method, exhaustive priority result, radiative protection or guaranteed publication novelty is asserted.

The physical gap in the supplied inputs is specific: neither the P01 action nor the primary passages examined provide this model's hard Wilson/loop/background envelope, quantitative cutoff or experimental uncertainty event. The conditional theorem can be applied when those inputs are independently supplied. This is a limit of the present information, not a proof that no UV completion or useful experiment can ever exist.

Reproducibility and status

The accompanying package contains editable sources, deterministic verification inputs and outputs, independent review records, build provenance and visual PDF inspection. Generated witnesses are mathematical examples and are not observations. Internal checks are AI-assisted, not external peer review. No external submission, experimental validation, UV completion or overall Apeiron physics approval is claimed.

References

[1] J. C. Criado and M. Pérez-Victoria. Field redefinitions in effective theories at higher orders. *JHEP*, 03:038, 2019. doi: 10.1007/JHEP03(2019)038. URL https://arxiv.org/pdf/1811.09413v2.

[2] David Bindel and Amanda Hood. Localization theorems for nonlinear eigenvalue problems, 2013. URL https://www.cs.cornell.edu/~bindel/papers/2013-simax.pdf. Author manuscript, 22 pages; Lemma 2.1 and Theorem 4.4 inspected.

[3] Ana Achúcarro, Vicente Atal, Sebastián Céspedes, Jinn-Ouk Gong, Gonzalo A. Palma, and Subodh P. Patil. Heavy fields, reduced speeds of sound and decoupling during inflation. *Phys. Rev. D*, 86:121301, 2012. doi: 10.1103/PhysRevD.86.121301. URL https://arxiv.org/pdf/1205.0710v2.

[4] Luc Jaulin and Eric Walter. Set inversion via interval analysis for nonlinear bounded-error estimation. *Automatica*, 29(4):1053–1064, 1993. doi: 10.1016/0005-1098(93)90106-4. URL https://webperso.ensta.fr/Jaulin/paper_automatica93.pdf.

[5] A. Raue, C. Kreutz, T. Maiwald, J. Bachmann, M. Schilling, U. Klingmüller, and J. Timmer. Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood. *Bioinformatics*, 25(15):1923–1929, 2009. doi: 10.1093/bioinformatics/btp358. URL https://academic.oup.com/bioinformatics/article/25/15/1923/213246.

[6] R. J. Furnstahl, N. Klco, D. R. Phillips, and S. Wesolowski. Quantifying truncation errors in effective field theory. *Phys. Rev. C*, 92:024005, 2015. doi: 10.1103/PhysRevC.92.024005. URL https://harvest.aps.org/v2/journals/articles/10.1103/PhysRevC.92.024005/fulltext.

[7] J. A. Melendez, R. J. Furnstahl, D. R. Phillips, M. T. Pratola, and S. Wesolowski. Quantifying correlated truncation errors in effective field theory. *Phys. Rev. C*, 100:044001, 2019. doi: 10.1103/PhysRevC.100.044001. URL https://arxiv.org/pdf/1904.10581v2.

[8] Lucas Pinol, Sébastien Renaux-Petel, and Denis Werth. The cosmological flow: A systematic approach to primordial correlators, 2023. URL https://arxiv.org/pdf/2312.06559v1. Version 1, inspected PDF pp. 36–38.