Identifiability and calibration limits of a single dispersion branch in a two-field effective theory

Consolidated version 1.0 — 23 September 2026

Abstract

We study what one dispersion branch can determine about a two-field quadratic effective theory, and which additional assumptions are required to turn exact reconstruction into a reliable inference from data. In a fixed metric and field normalization, three distinct positive-momentum samples determine the three pole parameters away from an exceptional linear branch. We characterize the complete exceptional family and show how a calibrated response residue resolves it. Exact finite-error identities, including jointly uncertain momentum and frequency, connect this structural result to conditional error certificates. An explicit classical Gaussian heavy-field extension supplies a checkable matching example, while also showing why positivity alone does not bound the displacement of a heavy pole. We then examine published plasmon-polariton data through the separate zero-gradient boundary of the model. The full pilot contains 167 peak positions in five particle-size groups; a subsequent layer-resolved comparison uses five interior points with matched endpoint fitting. Its energy RMSE is 20.89 meV, compared with 25.79 meV for a two-parameter saturation curve, without a statistical significance claim. We derive an exact family under unknown common momentum calibration and certify nominal energy-tolerance floors for four layer groups. An audit of the archived spectra identifies missing uncertainty information and limits on reproducing the original peak extraction. The result is a conditional inverse framework and a reproducible empirical audit, not an independently certified measurement of the unobserved upper gap.

1 Question, scope, and relation to established results

A measured low-frequency branch may depend on parameters associated with a second, inaccessible branch. This dependence motivates a precise inverse question: when do finitely many samples determine those parameters, and when do finite errors or calibration freedom invalidate the inference? Structural identifiability concerns exact observations in a specified model; practical identifiability additionally depends on an observation protocol and its uncertainty [1, 2]. These questions must be separated from establishing that the model describes a physical system.

The forward two-field dispersion relation used below belongs to established coupled-field theory. Closely related kernels and branches occur in the effective-field-theory literature on heavy and light modes [3–6]. We claim neither a new forward propagator nor a demonstrated priority for the general inverse methodology. Bounded-error set inversion, nonlinear eigenvalue localization, and correlated effective-theory errors also have established frameworks [7–10]. The purpose here is to give explicit, mutually compatible statements for one model and to test what they actually support on an existing dataset.

The consolidation contains three distinct levels of result. First, exact algebra characterizes a minimal reconstruction and its complete exceptional family. Second, finite-error and classical matching statements provide conditional certificates whose hypotheses must be supplied independently. Third, a retrospective audit of published plasmon-polariton measurements tests the usefulness and limitations of the boundary model. No claim about a broader physical framework, a quantum ultraviolet completion, or an experimentally certified unobserved pole follows from combining these levels.

The four accompanying supplements preserve the detailed record. S1 contains the original exact inverse analysis, S2 the robustness and common-window analysis, and S3 the explicit classical matching calculation. Those three PDFs are unchanged archival components. S4 integrates the boundary, calibration, and finite-data arguments needed for the empirical audit. Five complete computational subpackages accompany them. The present article states the consolidated claims and the limitations governing their interpretation.

2 Reference theory and two different inverse problems

We use a fixed nondynamical metric of signature $(+---)$ and fixed field normalization. Define

$$Y_\sigma = (\partial\sigma)^2, \qquad Y_\theta = (\partial\theta)^2, \qquad Z = \partial\sigma \cdot \partial\theta. \tag{1}$$

The reference Lagrangian is

$$\mathcal{L}_C = \tfrac{1}{2}Y_\sigma + \tfrac{1}{2}(1+g\sigma)Y_\theta - \tfrac{1}{2}gv^2\sigma - \tfrac{1}{2}m^2\sigma^2 + C(Z^2 - Y_\sigma Y_\theta). \tag{2}$$

The rolling solution $\sigma_0 = 0$, $\theta_0 = vt$ has $\rho = p = v^2/2$. This is a solution in a fixed metric, not a claim that a nonzero stress tensor supports Minkowski space in dynamical gravity. We restrict the local field patch to $1 + g\sigma > 0$ and take $m^2 > 0$ and $gv \neq 0$.

For perturbations $\sigma = s$ and $\theta = vt + \pi$, introduce

$$D = gv, \quad M = m^2, \quad B = M + D^2, \quad r = M/B, \quad a = 1 - 2Cv^2, \quad \chi = 1 - r. \tag{3}$$

The strict parameter domain in Sections 3–5 is $a > 0$ and $B > M > 0$, so $0 < r < 1$. The quadratic Lagrangian is

$$\mathcal{L}_2 = \tfrac{1}{2}(\dot{s}^2 + \dot{\pi}^2) - \tfrac{1}{2}\{a(\nabla s)^2 + (\nabla\pi)^2\} - \tfrac{1}{2}Ms^2 + Ds\dot{\pi}. \tag{4}$$

Its Hamiltonian is the sum of $p_s^2/2$, $(p_\pi - Ds)^2/2$, the two gradient terms, and $Ms^2/2$; it is positive at nonzero momentum in this domain.

There are two different comparisons. If only C varies while m, g, v stay fixed, the operator $Z^2 - Y_\sigma Y_\theta$ and its first field and metric variations vanish for collinear gradients. The full homogeneous action and stress tensor are unchanged, while a can change. This is a restricted homogeneous degeneracy. If instead (a, B, r) are all unknown, changing B at fixed r, v changes m and $|g|$. Such theories can share the selected rolling solution and its stress tensor, but their complete homogeneous actions need not agree. Conflating these two comparisons would overstate the result. Pole positions also leave the sign of D undetermined.

The homogeneous kernel statement is similarly restricted: in the first-derivative quartic basis $\{Y_\sigma^2, Y_\theta^2, Y_\sigma Y_\theta, Z^2\}$, its null direction is proportional to $(0, 0, -1, 1)$. It is not a classification of operators involving second derivatives or curvature. Field redefinitions must transform the source and observable along with the fields [11]; they do not remove a distinction between fixed, differently normalized observation protocols.

3 Exact reconstruction and the complete exception

Write $x = k^2$ and $y = \omega^2$. A Fourier convention gives the kernel and determinant

$$P_0 = \begin{pmatrix} y - ax - M & -iD\omega \\ iD\omega & y - x \end{pmatrix}, \tag{5}$$

$$F(y, x) = y^2 - [B + (a+1)x]y + ax^2 + Mx. \tag{6}$$

The two branches are

$$y_\pm(x) = \tfrac{1}{2}\big[B + (a+1)x \pm \sqrt{\Delta(x)}\big], \quad \Delta = B^2 + 2B(a+1-2r)x + (a-1)^2x^2. \tag{7}$$

At $x > 0$, $F(0,x) > 0$ and $F(x,x) = -D^2x < 0$, hence

$$0 < y_-(x) < x < y_+(x). \tag{8}$$

For $z(x) = y_-(x)/x$ the pole equation becomes

$$x(1-z)(a-z) = B(z-r), \qquad z' = \frac{(1-z)(a-z)}{B + x(a+1-2z)}. \tag{9}$$

The denominator is the positive branch separation. Thus z is strictly increasing if $a > r$, strictly decreasing if $a < r$, and constant if $a = r$.

Theorem 1 (Three samples and the exceptional family)**.** *For three distinct positive, exactly known momenta x_i, exact samples of the lower branch identify (a, B, M) uniquely whenever $a \neq r$. If $a = r$, the entire lower branch is $y_-(x) = rx$ and its complete admissible parameter family is $(a, B, r) = (r, B, r)$ with arbitrary $B > 0$.*

Proof. Set $z_i = y_i/x_i$. The samples give a linear system with row

$$\big[x_i(1-z_i), -z_i, 1\big] \begin{pmatrix} a \\ B \\ M \end{pmatrix} = x_i z_i (1-z_i). \tag{10}$$

Using (9) in the coefficient determinant yields

$$\det A = B(a-r)\frac{\prod_{i<j}(z_j - z_i)}{\prod_i (a - z_i)}. \tag{11}$$

For $a \neq r$, strict monotonicity makes the numerator nonzero, and (9) excludes $a = z_i$. This proves uniqueness. At $a = r$,

$$F = (y - rx)(y - B - x), \tag{12}$$

so B is invisible in the lower branch. Conversely, substituting the observed line $y = rx$ into (6) makes the coefficients of x and x^2 vanish, forcing $M = rB$ and $a = r$ for $0 < r < 1$. □

Two samples give a rank-two system off the exceptional set and therefore a local one-dimensional admissible family. This establishes minimality in the unrestricted three-parameter problem; it does not preclude reconstruction with fewer samples when some parameters are known. For example, known B, r and one positive-momentum pole give

$$a = \frac{y^2 - (B+x)y + rBx}{x(y-x)}. \tag{13}$$

If r alone is known, define $h_i = (z_i - r)/[x_i(1-z_i)]$. Two suitable samples give $B = (z_2 - z_1)/(h_1 - h_2)$ and $a = z_1 + Bh_1$.

A low-momentum expansion contains the same singularity. Writing

$$y_- = rx + b_2x^2 + b_3x^3 + \cdots, \quad b_2 = \frac{\chi(a-r)}{B}, \quad b_3 = \frac{b_2(2r-a-1)}{B}, \tag{14}$$

one obtains $d = b_2^2 + \chi b_3 = -b_2\chi^2/B$ and, off the exceptional set, $B = -b_2\chi^2/d$ and $a = r + Bb_2/\chi$. The exceptional condition is $b_2 = 0$, not $b_3 = 0$ alone. S1 includes the finite-jet remainder treatment.

3.1 A calibrated response resolves the linear exception

With a specified source term $-J_\theta\pi$, the response function in the chosen convention is

$$G_{\theta\theta}(y,x) = \frac{y - ax - M}{F(y,x)}. \tag{15}$$

At $a = r$, the residue with respect to y at $y = rx$ is

$$R_\theta(x) = \frac{rB}{B + \chi x}, \qquad B = \frac{\chi x R_\theta}{r - R_\theta}. \tag{16}$$

Thus one positive-momentum calibrated residue fixes B. A residue with respect to frequency is instead $R_\theta/(2\omega_-)$, and an unknown source amplitude introduces an additional unknown. The relative condition factor $r/(r - R_\theta) = 1 + B/(\chi x)$ also shows why an exact formula need not be a stable practical measurement at small momentum.

4 Finite errors require a joint error set and a common validity window

Theorem 1 contains no error model. A physical pole can differ from a reference-model pole, and measured frequency and momentum can both differ from their physical values. Those contributions must be combined in a justified joint set, including shared calibration errors. A covariance matrix alone is not a hard bound; a probability statement requires an independently justified coverage statement.

Here is an exact finite-error identity that avoids treating momentum as error-free. Choose a fixed scale μ^2 and write $u = x/\mu^2$, $v = y/\mu^2$ and $p = (a, \beta, m) = (a, B/\mu^2, M/\mu^2)$. In this paragraph m is a dimensionless mass-squared parameter, not the mass in (2). For each row,

$$A(u,v) = [u(u - v), -v, u], \qquad f(u,v) = v(u - v). \tag{17}$$

Let the recorded values be U_i, V_i, the true values $U_i + \xi_i, V_i + \eta_i$, and define $\widehat{p} = \widehat{A}^{-1}\widehat{f}$ and $T = \widehat{A}^{-1}$. With $C = A - \widehat{A}$, its rows are

$$C_i = [(2U_i - V_i)\xi_i - U_i\eta_i + \xi_i^2 - \xi_i\eta_i, -\eta_i, \xi_i]. \tag{18}$$

Define

$$c_i = (1 + \widehat{a})V_i - 2\widehat{a}U_i - \widehat{m}, \tag{19}$$

$$d_i = (1 + \widehat{a})U_i + \widehat{\beta} - 2V_i, \tag{20}$$

$$\psi_i = c_i\xi_i + d_i\eta_i - \widehat{a}\xi_i^2 + (1 + \widehat{a})\xi_i\eta_i - \eta_i^2. \tag{21}$$

Direct subtraction of the two pole systems gives the exact identity

$$(I + TC)(p - \widehat{p}) = T\psi. \tag{22}$$

If componentwise nonnegative bounds $K \geq \sup|TC|$ and $g \geq \sup|T\psi|$ hold over the same joint error set and the spectral radius satisfies $\rho(K) < 1$, then

$$|p - \widehat{p}| \leq (I - K)^{-1}g. \tag{23}$$

This follows by a nonnegative Neumann-series comparison. Correlations can be retained before taking the suprema. Parameter-domain margins must still be checked. Failure of this sufficient test is inconclusive; the identity does not supply the needed error set. S4 gives the complete statement and the two-parameter $a = 0$ specialization.

4.1 Model remainders and pole continuation

A justified model remainder must control the kernel in a region, not just its value at a fitted pole. Let a reference root $y_0 > 0$ be separated from the other root by $g_0 > 0$. On a complex disk of radius $\varrho < \min(y_0, g_0/2)$, assume an analytic determinant perturbation satisfying Schwarz reflection and a uniform bound $\varepsilon_F < \varrho(g_0 - \varrho)$. Rouché's theorem gives one real, positive, simple continued pole and

$$|\delta y| \le \frac{2\varepsilon_F}{g_0 + \sqrt{g_0^2 - 4\varepsilon_F}} < \varrho. \tag{24}$$

The statement is local to this pole. It does not establish global absence of extra modes or positivity of all source residues. A heavy-pole interpretation needs a separate disk and a justified frequency range extending to that pole. All sample nodes, inverse bounds, parameter margins, momentum ranges, and frequency ranges must fit a single common validity window; S2 spells out these intersections.

Operator dimension alone does not determine a harmless high power of momentum on a rolling background. For example,

$$\delta\mathcal{L} = \frac{c_t Y_\theta^2 + c_{st} Y_\sigma Y_\theta + c_Z Z^2}{\Lambda^4} + \frac{\lambda\sigma^2 Y_\theta}{2\Lambda^2} \tag{25}$$

already changes quadratic kinetic and mass terms. For the last term alone, putting $d = \lambda v^2/\Lambda^2$ gives $M' = M - d$, $B' = B - d$, and

$$r' - r = -\frac{d(B - M)}{B(B - d)}. \tag{26}$$

The induced lower-pole difference starts at order x. Neither unknown Wilson coefficients nor unknown loop contributions acquire a numerical bound merely by writing a cutoff symbol. S2 also proves nonuniformity near $a = r$ and in shrinking momentum windows: fixed absolute observation errors cannot support a uniform accuracy for B there.

5 An explicit classical matching example

To make a remainder source concrete, introduce a heavy field S with $L = M_S^2 > 0$ and

$$\mathcal{L} = \mathcal{L}_C + \tfrac{1}{2}(\partial S)^2 - \tfrac{1}{2}LS^2 + SJ - \frac{J^2}{2L}, \qquad J = \tfrac{1}{2}(h_0 + h_1\sigma)Y_\theta. \tag{27}$$

The subtraction is an explicit finite matching choice. This is a selected classical extension, not a derivation of C or a complete quantum theory. The rolling solution extends to $S_0 = h_0 v^2/(2L)$ with the same selected rolling stress tensor; equality of all homogeneous histories is not implied.

Write $A = h_1 v^2/2$, $b = h_0 v$, $c = 1 - b^2/L$ and $d = D - Ab/L$. The full quadratic Hamiltonian is positive for $a, M, L > 0$ and $b^2 < L$, with arbitrary A. Eliminating the heavy fluctuation gives the exact Schur kernel

$$P_{\text{eff}} = P_0 + fW, \quad f = \frac{y - x}{L(L + x - y)}, \quad W = \begin{pmatrix} A^2 & -iAb\omega \\ iAb\omega & b^2 y \end{pmatrix}. \tag{28}$$

Writing $F_{\text{ref}} = \det P_0 = F$ for the reference determinant (6), its determinant is $F_{\text{ref}} + fQ$, where

$$Q = A^2(y - x) + b^2 y(y - ax - M) - 2DAb\, y. \tag{29}$$

The determinant of the full three-field kernel is

$$\det P_3 = (y - x - L)F_{\text{ref}} - (y - x)Q/L, \tag{30}$$

which is finite at the apparent elimination pole. Positivity and $\det P_3(x,x) = LD^2x > 0$ yield a unique physical root in $0 < y < x$.

On a disk around y_0, let $Y = y_0 + \varrho$, $T_0 = |y_0 - x| + \varrho$, $U_0 = |y_0 - ax - M| + \varrho$ and $\eta_H = |L + x - y_0| - \varrho > 0$. An explicit determinant remainder bound is

$$\varepsilon_F = \frac{T_0}{L\eta_H}\left[A^2T_0 + b^2YU_0 + 2|DAb|Y\right]. \tag{31}$$

When the remaining hypotheses hold, it can be inserted into (24). Geometric truncation of f is also controlled on $|y - x| \le \zeta L < L$: truncating its series through power n leaves at most $\zeta^{n+1}/[L(1-\zeta)]$, for $n \ge 1$, before multiplication by W or Q.

Matching conventions matter. Absorbing the induced two-derivative s normalization gives $Z_s = 1 + A^2/L^2$, $a_c = (a + A^2/L^2)/Z_s$, $M_c = M/Z_s$, $D_c = D/\sqrt{Z_s}$, $B_c = B/Z_s$, and $r_c = r$. The bounds above retain the explicitly subtracted reference convention (27). In that convention the leading additional light-jet coefficient is

$$\delta b_2 = \frac{(1-r)(DA + Mb)^2}{B^2L^2}. \tag{32}$$

5.1 Positivity does not guarantee a nearby heavy pole

At zero momentum the two massive roots obey

$$cy^2 - T_H y + LB = 0, \qquad T_H = L + cB + K_H/L, \quad K_H = (A - Db)^2. \tag{33}$$

For $L > B$, write the smaller massive root in stable form as

$$Y_H = \frac{2LB}{T_H + \sqrt{T_H^2 - 4cLB}}, \qquad Y_S = \frac{T_H + \sqrt{T_H^2 - 4cLB}}{2c}. \tag{34}$$

Then $0 < Y_H \le B < L \le L/c \le Y_S$. Its displacement $\delta_H = B - Y_H$ satisfies

$$\delta_H(L - cY_H) = K_H Y_H/L, \qquad \delta_H \le \frac{K_H B}{L(L - cB)}. \tag{35}$$

For $0 < \tau < B$, the exact target condition is

$$\delta_H \le \tau \quad \Longleftrightarrow \quad K_H \le \frac{L\tau[L - c(B - \tau)]}{B - \tau}. \tag{36}$$

Taking $|A| \to \infty$ while preserving positivity gives $Y_H \to 0$ and $Y_S \to \infty$. A coupling bound is therefore indispensable for a uniform statement that the heavy pole remains close to B.

S3 supplies the full quadratic derivation, an explicitly declared synthetic application, and its checks. The special alignment $a = r$, $DA + Mb = 0$ preserves the exact linear light branch, but does not imply a free B when the additional couplings are held separately fixed. A light-independent pure-heavy Gaussian determinant also does not remove mixed heavy–light quantum loops. These facts prevent interpreting this example as a completed ultraviolet matching calculation.

6 Existing data: a distinct zero-gradient boundary

We use the published experimental peak positions and archived spectra of plasmonic nanoparticle crystals from Mueller et al. [12, 13]. Their Hopfield-type description maps to

$$F_0 = y^2 - (B + x)y + Mx = 0, \qquad B > M > 0, \tag{37}$$

Table 1: Full nominal-calibration pilot. Gaps are inferred values, not independent upper-branch observations. The last column is the largest absolute energy error among the internal held-out points.

Diameter (nm)	Points	$\sqrt{B}$, two points (eV)	$\sqrt{B}$, all points (eV)	Max. error (meV)
25	13	4.192	4.414	93.5
30	75	5.866	5.679	133.3
41	16	5.677	5.426	197.4
46	5	5.316	5.318	31.2
59	58	6.954	6.781	63.9

which is the boundary $a = 0$. It is treated separately from the strict $a > 0$ theorem. Here $M = (\hbar\omega_{\mathrm{pl}}^{t})^2$ is a collective transverse plasmon scale, not an isolated-particle resonance; $B = M + 4(\hbar\Omega_R)^2$. The implied zero-momentum upper gap is $\sqrt{B}$, and the inferred coupling is $\hbar\Omega_R = \sqrt{B - M}/2$. The forward physical model was already fitted in the source study; applying an algebraic inverse to it is not a discovery of that model.

For the boundary, the lower root obeys $0 < y_- < \min(x, M)$ and has decreasing y_-/x. Two exact points suffice:

$$z_i = y_i/x_i, \quad t_i = y_i(1 - z_i), \quad B = \frac{t_1 - t_2}{z_2 - z_1}, \quad M = t_1 + Bz_1. \tag{38}$$

For numerical evaluation we use

$$y_-(x) = \frac{2Mx}{B + x + \sqrt{(B + x)^2 - 4Mx}}. \tag{39}$$

This boundary reconstruction and its finite-error specialization are proved in S4.

6.1 Nominal calibration and the full pilot

The source assigns standing-wave coordinates $\kappa = k/\Gamma_L = j/n_\ell$ using mode index j and layer count n_ℓ. These are geometric assignments rather than independent momentum readings. The nominal conversion is

$$x = \left(\frac{\hbar c\,\Gamma_L}{n_{\mathrm{eff}}}\right)^2 \kappa^2, \quad \Gamma_L = \frac{\pi}{\sqrt{2/3}(d_{\mathrm{NP}} + g_{\mathrm{NP}})}, \tag{40}$$

with $f_{\mathrm{vol}} = 0.74[d_{\mathrm{NP}}/(d_{\mathrm{NP}} + g_{\mathrm{NP}})]^3$ and $n_{\mathrm{eff}} = 1.4(1 - f_{\mathrm{vol}}) + f_{\mathrm{vol}}$. The five diameter/gap pairs in nm are $(25, 3), (30, 4), (41, 3), (46, 2), (59, 1)$. Published particle-size spreads do not constitute a joint uncertainty model for geometry, index, boundary assignment, and extracted energy.

The pilot contains all 167 selected experimental peak positions, cross-checked against the publisher's experimental cells; fitted dispersion curves are not substituted for measurements. Before this pilot calculation, the two outer κ points in each size group were assigned to reconstruction and the remaining points to an internal check. This gives ten endpoint observations and 157 held-out observations, all from the same published experiment. Table 1 retains the full pilot, including its least favorable group. Implicit all-point least squares is reported descriptively; uncertainty in both axes prevents treating it as an optimal likelihood without additional assumptions.

Pooling different layer counts produces distinct energies at identical nominal κ. Half the largest such span is an unavoidable maximum-error floor for any single-valued pooled curve: respectively 8.87, 53.80, 93.79, 24.33, and 41.20 meV for the five diameter groups. For example, the 41-nm source gives 1.48969180741 eV for three layers and 1.30211169490 eV for six layers at $\kappa = 1/3$, a 187.5801 meV difference. This indicates unresolved layer dependence or calibration/model variation; it is neither an instrument error estimate nor a rejection of the underlying experiment.

Table 2: Retrospective comparison on the same five interior points. All rows use two endpoint-fitted parameters per eligible layer. No statistical significance is assigned.

Model	Energy RMSE (meV)	Maximum error (meV)
Boundary Hopfield curve, nominal calibration	20.893636	35.65239
Saturation, $y = \alpha_0\kappa^2/(\beta_0 + \kappa^2)$	25.793157	41.13277
Linear energy versus κ	177.110576	256.24986

6.2 Layer-resolved follow-up and a matched comparison

The 41-nm group contains all 16 source-labelled observations across layer counts two through seven. Two layers provide one point, and three layers provide only two, so neither contributes an interior check. Endpoint fits for four, five, six, and seven layers leave five interior points in total. This follow-up was chosen after seeing the pooled pilot and uses more parameters; it is retrospective, not a new independent validation sample.

Under the nominal common calibration, those five points have RMSE 20.893636 meV and maximum absolute error 35.65239 meV (maximum relative error 2.2091%). The pooled fit evaluated on the same five points has RMSE 30.350962 meV. Table 2 compares models fitted to the same endpoints with two fitted parameters per layer. The Hopfield curve also uses the external nominal momentum calibration; equal fitted-parameter counts do not make all external information identical.

The 4.899521 meV RMSE advantage over the saturation curve is a descriptive result. Five retrospective points with shared calibration and unknown extraction errors do not establish statistical superiority. In addition, the saturation curve arises as an exact calibration limit of the boundary family, which makes this comparison especially relevant to identifiability rather than to a broad model-selection claim.

7 What an unknown common calibration permits

Set $q_i = \kappa_i^2$ and $x = sq$ with an unknown common scale s. In this section q is a source-assigned geometric coordinate; the earlier symbol for $1 - r$ is χ. Hold two endpoint energies exact, with $q_1 < q_2$, $y_1 < y_2$, and $h_1 = y_1/q_1 > h_2 = y_2/q_2$. Define

$$\alpha = \frac{y_2 - y_1}{h_1 - h_2}, \quad \beta = \frac{y_1 h_1 - y_2 h_2}{h_1 - h_2}, \quad m_\infty = y_1 + \alpha h_1, \quad \gamma = \alpha h_1 h_2. \tag{41}$$

Every endpoint-compatible parameter pair is

$$B(s) = \alpha s + \beta, \qquad M(s) = m_\infty - \gamma/s. \tag{42}$$

The observed endpoints lie on the lower branch precisely for $s > h_1$. In particular,

$$B - M = \frac{\alpha(s - h_1)(s - h_2)}{s}. \tag{43}$$

A second positive-parameter interval can place the same observations on the upper branch and must be excluded; positivity by itself does not identify the observed branch.

As $s \to \infty$, the lower branch tends to

$$y_\infty(q) = \frac{m_\infty q}{\alpha + q}, \qquad \sqrt{B(s)} \to \infty. \tag{44}$$

For an interior $q_1 < q < q_2$, its energy decreases strictly with s towards $E_\infty = \sqrt{y_\infty}$. One way to see this is to write the pole equation as $H(y, q) + sG(y, q) = 0$, where

$$H = y^2 - \beta y - \gamma q, \qquad G = m_\infty q - (\alpha + q)y, \tag{45}$$

Table 3: Six-layer endpoint-compatible calibration family. Residual means predicted minus published interior energy. Arbitrarily high implied gaps can coexist with a small limiting interior residual if no upper calibration bound is supplied.

Calibration	Implied $\sqrt{B}$ (eV)	Interior residual (meV)
s_0	5.837382	+0.042767
$4s_0$	11.799622	−3.229262
$16s_0$	23.661262	−3.992106
$s \to \infty$	unbounded	−4.241847

and evaluate

$$H(y_\infty, q) = -\frac{\gamma q(q-q_1)(q-q_2)}{(\alpha+q)^2} > 0. \tag{46}$$

The upper limit as $s \downarrow h_1$ is $\sqrt{\min(h_1 q, y_2)}$. S4 gives the complete monotonicity and branch proof.

A compatible third exact interior point generically identifies $s = -H/G$. Thus the conclusion is not structural nonidentifiability with three exact points. Rather, a finite interval for that third energy gives a finite upper bound on s only if its lower endpoint lies strictly above E_∞, assuming compatibility elsewhere. Intervals containing the limit from above can admit arbitrarily large s. These statements hold with the two endpoints fixed; uncertain endpoints require the joint finite-error treatment.

For the 41-nm nominal conversion, $s_0 = 221.08533593187062\,\mathrm{eV}^2$. In the six-layer subgroup, the published interior energy is just 4.241847 meV above its saturation limit. This is a required discrimination scale under the stated endpoint assumptions, not an available experimental error bar. Table 3 displays the exact family. The large changes in s are mathematical stress cases; no claim is made that they are physically plausible after independent geometric constraints are imposed.

The five interior points separately imply exact scales of approximately 39.529980, 62.742127, 74.549441, 223.154429, and $56.460611\,\mathrm{eV}^2$. Taken literally, they do not select a single common exact calibration. Without an error model this is an algebraic incompatibility of exact nominal observations, not a statistical rejection. With all layer endpoints fixed and one common free s, the minimum possible maximum absolute interior residual is certified within

$$[16.5679541142, 16.5679541196]\,\mathrm{meV}, \tag{47}$$

at a scale in $[58.24533607, 58.24533608]\,\mathrm{eV}^2$. This minimax problem differs from allowing all endpoint energies to move at fixed nominal calibration, considered next.

Independent bounds on s would directly limit $\sqrt{B}$ through (42). For a target gap interval $[E_L, E_U]$, the required scale interval is

$$\left[\frac{E_L^2-\beta}{\alpha}, \frac{E_U^2-\beta}{\alpha}\right] \cap (h_1, \infty). \tag{48}$$

For a symmetric relative scale radius δ_s around s_0 and a target relative gap radius $0 < \rho < 1$, a sufficient tight lower-side condition is $\delta_s \le (2\rho-\rho^2)B(s_0)/(\alpha s_0)$, together with the lower-branch domain margin. This translates a desired conclusion into a calibration requirement.

8 Certified nominal energy-tolerance floors

At the fixed nominal x_i, suppose every energy in one layer group, including endpoints, may move by at most a common δ. Set $L_i = (E_i - \delta)^2$ and $U_i = (E_i + \delta)^2$, assuming $E_i > \delta$ and

Table 4: Certified minimum uniform energy tolerance at fixed nominal calibration. Each layer group has its own B, M. These are compatibility thresholds, not confidence intervals or measurement uncertainties.

Layers	Lower bracket (meV)	Upper bracket (meV)
4	17.0999288558	17.0999348164
5	11.8911147117	11.8911206723
6	0.0193774700	0.0193834305
7	4.5780837535	4.5780897141

$0 < L_i \le U_i < x_i$. The lower root belongs to this interval if and only if

$$L_i B - x_i M \le L_i^2 - x_i L_i, \tag{49}$$

$$-U_i B + x_i M \le x_i U_i - U_i^2, \tag{50}$$

with $B > M > 0$. This equivalence follows from the sign of the upward-opening quadratic below the upper root, since $y_+ > x_i$. It turns the boundary compatibility problem into a finite half-plane intersection.

Exact rational arithmetic applied to the stored numerical inputs gives the brackets in Table 4. Upper endpoints have feasible parameter witnesses and direct lower-root checks; lower endpoints have independently checked infeasibility certificates. The independent review also eliminates M by Fourier–Motzkin reduction and supplies Farkas certificates. Decimal digits here certify computation on the recorded inputs, not the precision of the experiment.

For all four groups with separate parameters but the same tolerance, the largest threshold governs, approximately 17.10 meV. For example, 17.0 meV is insufficient under these nominal assumptions. The very small six-layer threshold expresses agreement of three stored points with a two-parameter curve; it does not imply microelectronvolt experimental precision. The physically relevant question remains whether an independently justified combined measurement, calibration, and model budget is consistent with these requirements.

9 What the archived spectra can and cannot establish

We additionally audited all 15 archived absorbance, reflectance, and transmittance CSV files across the five particle sizes and checked their repository MD5 values. For the 41-nm specimen the archive has 441 values from 3.0 to 0.8 eV in 0.005 eV steps, with columns for two through seven layers. The stored relation $A = 100 - R - T$ holds to numerical precision; these are not three independent observations. The plotted vertical offsets are not offsets to subtract from the raw archived columns.

Archive grid spacing is not instrument resolution or a joint energy error bound. The source describes substrate and silver-mirror referencing and the combination of visible and near-infrared spectral ranges, but the archive does not provide the full original fitting windows, weights, baseline conventions, covariance, reference repetitions, or a joint calibration budget. The earlier qualitative reproducibility discussion concerns a different mono-/bilayer study [14] and does not supply these missing quantities for the 2020 data. A further processing difference is explicit for the 30-nm group: the repository series contains 226 points over 0.75–3.0 eV, whereas the source plotting series has 1000 points over 0.8–3.0 eV; the processing relation is not documented in the recovered material.

One essential endpoint is outside the archived spectral interval: the six-layer first published peak, 0.780680325 eV, lies below 0.8 eV. The reason for this mismatch is not established. It prevents checking a two-sided peak neighborhood in this archive and does not justify attributing extrapolation or an error to the original authors. Some stored absorbance/transmittance values

are negative and some reflectance values exceed 100%; these were preserved, with their origin left unresolved.

9.1 A retrospective extraction sensitivity audit

To test sensitivity, we fitted six declared recipes per eligible layer group: quadratic or cubic baselines and upper energy limits 1.95, 2.05, or 2.15 eV, all starting at 0.8 eV. The model uses positive Lorentzian amplitudes with bounds 0–500% and full widths between 0.01 and 0.6 eV. Mode counts, initial center labels, and midpoint constraints come from the published peaks. These are constrained retrospective fits, not blind recovery of physical modes.

Three width starts (0.06, 0.12, 0.24 eV) were used for each recipe. A diagnostic caught cubic-baseline solutions with worse residual sums of squares than their nested quadratic models. Adding a deterministic fourth start from the quadratic solution with zero cubic coefficient for the twelve cubic recipes yields 84 starts in total and passes that nesting check. This repair does not establish global optimality. All earlier diagnostics and the final outputs remain in the reproduction package.

Under the declared primary recipe (cubic baseline, upper limit 2.05 eV), the first fitted four-layer component is 1.343839569 eV, compared with the published assigned peak 1.090908118 eV. Absorbance RMSE is 1.274005 percentage points and residual lag-one correlation is 0.4161. A low aggregate residual does not establish that a fitted component represents the same physical mode. The first-center spread across the six final four-layer recipes is 22.1104 meV.

Other groups expose different limits: the primary five-layer first component is at the 0.8 eV boundary, compared with the published 0.9170358 eV, and five of six recipes have at least one active bound. The six-layer primary first component is 0.8127015 eV, while its published endpoint lies outside the archive. For seven layers the primary first component is 1.2342346 eV, compared with 1.1812854 eV, and all six recipes reach the first-component width upper bound. A singleton passing a descriptive filter is not a zero-width uncertainty estimate.

These diagnostics cannot replace the original mode assignments or reproduce the original extraction protocol, whose exact settings are unavailable. Correlated residuals, boundary solutions, spectral joins, and baseline dependence preclude turning these fits into a validated confidence interval by assumption. Equally, a mismatch under a new protocol does not refute the original study. The audit's positive result is to identify which missing inputs prevent the exact inverse from becoming a certified empirical upper-gap measurement.

10 Conclusions and closure criteria

The strict two-field reference problem has a complete exact answer: three positive-momentum lower-branch samples determine the pole parameters away from $a = r$, while the exceptional line leaves a one-parameter family that a calibrated residue can resolve. Finite-error reconstruction is possible when a justified joint error set and a common model-validity window are supplied. The explicit heavy-field example makes one classical remainder source calculable and demonstrates why positivity cannot replace a bound on matching couplings.

The published plasmonic data provide a concrete use of the separate $a = 0$ boundary. They yield a fully recorded 167-point pilot, a modest retrospective advantage on five layer-resolved interior points, exact calibration families, and certified nominal compatibility floors. These are measurable and reproducible outcomes. They do not furnish an independent observation of the inferred upper branch: the source's gold interband region around 2.3 eV obstructs the relevant upper spectral range, while the extracted peaks and calibration lack a recovered joint uncertainty certificate.

The scope can therefore be closed as a mathematical reconstruction framework with a conditional empirical audit. Strengthening the physical inference would require new infor-

mation rather than another unconstrained numerical example: a justified common calibration/measurement/model error budget, recoverable peak-extraction settings or independent spectral replication, and preferably a separately calibrated upper-branch or response observable. Whether these inputs are available is an experimental question. Until then, the gap values remain conditional model inferences.

Data, code, supplements, and administrative status

The external experimental data are cited by their source article and repository DOI [12, 13]. The English publication package includes S1–S4 and editable sources. Its complete research-archive companion retains five unchanged computational archives with inventories, provenance, internal review records, and checksums; historical German working reports in that companion are internal evidence, not publication documents. The English reproduction guide gives a portable extraction and verification path. No new public repository DOI or external deposit is claimed by this version. All fresh empirical analyses are retrospective; the distinction between exploratory follow-up and the initial pilot is retained.

Author names and order, affiliations, contact details, contributions and approval, funding, conflicts of interest, manuscript license, and destination-specific declarations are not established by the technical record. They remain explicit fields in the administrative checklist and must be resolved before external submission. The internal computational and editorial checks are not external peer review. This version has not been uploaded to Academia or another publication platform by this workflow.

References

[1] Richard Bellman and Karl Johan Åström. On structural identifiability. *Mathematical Biosciences*, 7:329–339, 1970. doi: 10.1016/0025-5564(70)90132-X. URL https://www.lunduniversity.lu.se/publication/e5be7c83-a0d2-4500-826f-6e1b414d1733.

[2] 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.

[3] Toshifumi Noumi, Masahide Yamaguchi, and Daisuke Yokoyama. Effective field theory approach to quasi-single field inflation and effects of heavy fields. *JHEP*, 06:051, 2013. doi: 10.1007/JHEP06(2013)051. URL https://arxiv.org/abs/1211.1624v3.

[4] 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/abs/1205.0710v2.

[5] Rhiannon Gwyn, Gonzalo A. Palma, Mairi Sakellariadou, and Spyros Sypsas. Effective field theory of weakly coupled inflationary models. *JCAP*, 04:004, 2013. doi: 10.1088/1475-7516/2013/04/004. URL https://arxiv.org/abs/1210.3020v2.

[6] Lucas Pinol, Sébastien Renaux-Petel, and Denis Werth. The cosmological flow: A systematic approach to primordial correlators, 2023. URL https://arxiv.org/abs/2312.06559v1.

[7] 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.

[8] 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.

[9] 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.

[10] 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.

[11] 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/abs/1811.09413v2.

[12] Niclas S. Mueller et al. Deep strong light–matter coupling in plasmonic nanoparticle crystals. *Nature*, 583:780–784, 2020. doi: 10.1038/s41586-020-2508-1.

[13] Stephanie Reich, Niclas S. Mueller, Yu Okamura, Sabrina Juergensen, Bruno G. M. Vieira, Holger Lange, Eduardo B. Barros, and Florian Schulz. Data repository: Deep strong light–matter coupling in nanoparticle crystals. Refubium data repository, 2020. URL https://doi.org/10.17169/refubium-26819.

[14] Niclas S. Mueller, Bruno G. M. Vieira, Florian Schulz, Patryk Kusch, Valerio Oddone, Eduardo B. Barros, Holger Lange, and Stephanie Reich. Dark interlayer plasmons in colloidal gold nanoparticle bi- and few-layers. *ACS Photonics*, 5:3962–3969, 2018. doi: 10.1021/acsphotonics.8b00898. URL https://doi.org/10.1021/acsphotonics.8b00898.