Supplement S4

Boundary-class inversion, joint coordinate errors,
calibration limits and nominal compatibility

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

Version 1 23 September 2026

Abstract

This supplement separates four conditional inference questions: inversion in the known local class $a = 0$; finite propagation of joint momentum and energy errors; inference with an unknown common momentum scale; and compatibility of published peak values with a fixed nominal calibration. Exact polynomial identities give deterministic parameter enclosures for a supplied joint error set. In the local class, two exact lower-branch points leave an unbounded upper-gap family when the scale is unknown, whereas a compatible third exact point identifies the scale. A thirteen-point published-data example illustrates this distinction. Independently checked rational feasibility certificates also bound the minimum uniform energy tolerance at fixed nominal momenta. These mathematical and numerical statements neither supply experimental error bounds nor identify an independently observed upper mode.

S4.1 Scope and notation

The reference spectral polynomial is

$$F(y, x) = y^2 - [B + (a + 1)x]y + ax^2 + Mx, \qquad B > M > 0. \tag{S4.1}$$

Here x is squared momentum after multiplication by a declared reference speed, and y is squared angular frequency. In the energy convention used for the data, both variables are multiplied by $\hbar^2$ and measured in eV^2. The upper zero-momentum energy is then $\sqrt{B}$. The strict model has $a > 0$; the known restriction $a = 0$ is treated separately below.

A pole of a real lossless reference model, a maximum of a measured spectrum, and a fitted component center are distinct objects. Their identification requires a forward observation model and a controlled discrepancy. All error bounds in this supplement are conditional on such a model and on a separately justified error set. Exact arithmetic on printed decimals does not make the physical observations exact. No priority claim is attached to the forward dispersion, matrix perturbation argument or linear feasibility methods used here.

S4.2 The known local class a=0

For $a = 0$, equation (S4.1) becomes

$$F_0(y, x) = y^2 - (B + x)y + Mx. \tag{S4.2}$$

At $x > 0$, $F_0(0, x) = Mx > 0$ and $F_0(x, x) = -(B - M)x < 0$. Both roots are positive and distinct, with

$$0 < y_-(x) < \min(x, M), \qquad y_+(x) > x. \tag{S4.3}$$

Put $z = y_-/x$ and $r = M/B$. Substitution gives

$$-xz(1 - z) = B(z - r), \qquad z' = -\frac{z(1 - z)}{B + x(1 - 2z)} < 0. \tag{S4.4}$$

The denominator is the positive root separation. Thus $0 < z < r$, and distinct positive nodes have distinct phase slopes.

If $a = 0$ is structurally known and x is calibrated, two lower observations determine B, M. Define $z_i = y_i/x_i$ and $t_i = y_i(1 - z_i)$. Then

$$-z_i B + M = t_i, \qquad B = \frac{t_1 - t_2}{z_2 - z_1}, \qquad M = t_1 + Bz_1. \tag{S4.5}$$

Equation (S4.4) makes the denominator nonzero. The reconstructed parameters must still satisfy $B > M > 0$ and the lower-root conditions. One observation instead leaves a local line of admissible pairs. Fitting a as an additional free parameter is a different problem: an interval containing zero does not establish strictly positive spatial stiffness.

A declared physical normalization

For a lossless bulk transverse dielectric channel, suppose

$$\epsilon(\omega, k) = \epsilon_\infty \left(1 + \frac{\Omega_p^2}{\Omega_T^2 + s_T^2 k^2 - \omega^2}\right), \qquad \epsilon_\infty, \Omega_T^2, \Omega_p^2 > 0. \tag{S4.6}$$

Maxwell's equation $c_0^2 k^2 = \epsilon\omega^2$ gives (S4.1) under

$$x = \frac{c_0^2 k^2}{\epsilon_\infty}, \quad y = \omega^2, \quad a = \frac{s_T^2 \epsilon_\infty}{c_0^2}, \quad M = \Omega_T^2, \quad B = M + \Omega_p^2. \tag{S4.7}$$

The local oscillator has $s_T^2 = 0$, hence $a = 0$. The symbol s_T^2 here denotes a curvature coefficient whose sign requires physical justification. Multiplication by the oscillator denominator adds no positive root at its zero, since it would require $\Omega_p^2 y = 0$. This equivalence concerns pole locations; it does not identify microscopic charges, densities, source residues or the full underlying action. Internal bulk momentum, surface momentum and a model-fitted momentum offset cannot be interchanged without a separate observation map.

S4.3 An exact residual for errors on both axes

Choose a fixed known scale $\mu > 0$ and write

$$u_i = x_i/\mu^2, \quad v_i = y_i/\mu^2, \quad p = (a, b, m)^T = (a, B/\mu^2, M/\mu^2)^T. \tag{S4.8}$$

Use b for the normalized gap parameter in this section; the dimensional coefficient β introduced later is different. Let (U_i, V_i) be observed coordinates and express truth-minus-observation errors as

$$u_i = U_i + \xi_i, \quad v_i = V_i + \eta_i, \qquad (\xi, \eta) \in \mathcal{E}. \tag{S4.9}$$

The bounded set $\mathcal{E}$ may correlate nodes, axes and shared calibration variables. No product-box or independence assumption is required.

Equation (S4.1) has the undivided linear form

$$\mathcal{A}(u, v)_{i\cdot} = [u_i(u_i - v_i), -v_i, u_i], \qquad f(u, v)_i = v_i(u_i - v_i), \qquad \mathcal{A}p = f. \tag{S4.10}$$

For nonzero u_i, these rows are fixed-row-scaled versions of the normalized inverse equations. Their polynomial form avoids division by an uncertain coordinate.

For three observations, suppose $\widehat{\mathcal{A}} = \mathcal{A}(U, V)$ is invertible, and set

$$T = \widehat{\mathcal{A}}^{-1}, \quad \widehat{p} = T\widehat{f}, \quad C = \mathcal{A}(u, v) - \widehat{\mathcal{A}}. \tag{S4.11}$$

Direct expansion gives

$$C_{i\cdot} = \left[(2U_i - V_i)\xi_i - U_i\eta_i + \xi_i^2 - \xi_i\eta_i, -\eta_i, \xi_i\right]. \tag{S4.12}$$

Define

$$c_i = (1 + \widehat{a})V_i - 2\widehat{a}U_i - \widehat{m}, \tag{S4.13}$$

$$d_i = (1 + \widehat{a})U_i + \widehat{b} - 2V_i, \tag{S4.14}$$

$$\psi_i = c_i\xi_i + d_i\eta_i - \widehat{a}\xi_i^2 + (1 + \widehat{a})\xi_i\eta_i - \eta_i^2. \tag{S4.15}$$

Then, exactly,

$$f(u, v) - \mathcal{A}(u, v)\widehat{p} = \psi, \qquad \boxed{(I + TC)(p - \widehat{p}) = T\psi.} \tag{S4.16}$$

Indeed $\mathcal{A}(p - \widehat{p}) = f - \mathcal{A}\widehat{p} = \psi$, and $T\mathcal{A} = I + TC$. The squared and mixed terms in (S4.12)–(S4.15) are part of the identity, not a truncation estimate. Setting $\xi = 0$ recovers the exact-momentum specialization after fixed row scaling. In particular, a Jacobian approximation alone is not the finite error statement used here.

S4.4 A finite comparison theorem

Let K and g have nonnegative entries and satisfy certified bounds

$$K_{jk} \geq \sup_{\mathcal{E}} |(TC)_{jk}|, \qquad g_j \geq \sup_{\mathcal{E}} |(T\psi)_j|. \tag{S4.17}$$

Absolute values and inequalities on vectors or matrices are componentwise.

Theorem S4.1. *If the spectral radius $\rho(K) < 1$, every reference-model parameter vector consistent with $\mathcal{E}$ obeys*

$$\boxed{|p - \widehat{p}| \leq R = (I - K)^{-1} g.} \tag{S4.18}$$

Every matrix $\widehat{\mathcal{A}} + C$ generated by the same error set is nonsingular.

Proof. Equation (S4.16) implies $|p - \widehat{p}| \leq g + K|p - \widehat{p}|$. The convergent Neumann series $(I - K)^{-1} = \sum_{n \geq 0} K^n$ is nonnegative, proving (S4.18). A null vector of $I + TC$ would satisfy $|w| \leq K|w|$; the same argument forces $w = 0$. □

A sufficient check is $\max_j \sum_k K_{jk} < 1$. More generally, a certified positive vector w with $Kw < w$ suffices. One constructive attempt forms $w = (I - K)^{-1} e$, $e = (1, 1, 1)^T$, and explicitly checks positivity and the strict inequality. Failure of these tests is inconclusive. Conversely, an empty compatibility set is not a successful parameter determination; the theorem bounds all compatible members but does not prove that one exists.

Explicit component-box bounds

If $|\xi_i| \leq s_i$ and $|\eta_i| \leq t_i$, let

$$H_{i1} = |2U_i - V_i| s_i + |U_i| t_i + s_i^2 + s_i t_i, \tag{S4.19}$$

$$H_{i2} = t_i, \qquad H_{i3} = s_i, \tag{S4.20}$$

$$q_i = |c_i| s_i + |d_i| t_i + |\widehat{a}| s_i^2 + |1 + \widehat{a}| s_i t_i + t_i^2. \tag{S4.21}$$

Then $K = |T| H$ and $g = |T| q$ satisfy (S4.17). The exact residual retains the shared origin of the matrix and right-hand-side errors, although this box bound discards correlations among distinct coordinates.

For correlated affine errors $\xi = H_u t$, $\eta = H_v t$, $|t_\ell| \leq 1$, expand each weighted expression in (S4.17) as a polynomial in t. Collect coefficients across nodes before taking absolute values. The sum of absolute coefficients bounds its magnitude on the box and preserves any exact coefficient cancellations. Polynomial coordinate maps can be treated in the same way, including degree-four terms arising from squaring raw calibration errors. General compact sets instead require certified optimization or outward enclosures; an unverified local maximum is insufficient.

S4.5 Physical margins and the reduced inverse

Write $R = (R_a, R_b, R_m)$. Sufficient strict-domain conditions are

$$\widehat{a} > R_a, \quad \widehat{b} > R_b, \quad \widehat{m} > R_m, \quad \widehat{b} - \widehat{m} > R_b + R_m. \tag{S4.22}$$

With $r = M/B$ and $\widehat{r} = \widehat{m}/\widehat{b}$, useful consequences are

$$|B - \widehat{B}| \le \mu^2 R_b, \qquad |M - \widehat{M}| \le \mu^2 R_m, \tag{S4.23}$$

$$|(B - M) - (\widehat{B} - \widehat{M})| \le \mu^2 (R_b + R_m), \tag{S4.24}$$

$$|r - \widehat{r}| \le \frac{R_m + |\widehat{r}| R_b}{\widehat{b} - R_b}, \tag{S4.25}$$

$$|\sqrt{B} - \sqrt{\widehat{B}}| \le \frac{\mu R_b}{\sqrt{\widehat{b}} + \sqrt{\widehat{b} - R_b}}. \tag{S4.26}$$

Coordinate boxes can be checked with $U_i - s_i > 0$, $V_i - t_i > 0$ and $V_i + t_i < U_i - s_i$. Disjoint ordered u-intervals suffice for distinct nodes; a correlated set may establish their order even if its box envelope overlaps.

When $a = 0$ is known, use two observations and $p_0 = (b, m)^T$:

$$\mathcal{A}_{0,i\cdot} = [-v_i, u_i], \qquad f_i = v_i(u_i - v_i), \tag{S4.27}$$

$$C_{0,i\cdot} = [-\eta_i, \xi_i], \tag{S4.28}$$

$$\psi_{0,i} = (V_i - \widehat{m})\xi_i + (U_i + \widehat{b} - 2V_i)\eta_i + \xi_i\eta_i - \eta_i^2. \tag{S4.29}$$

The same comparison proof applies to this 2×2 system. Require $b > m > 0$ and their finite margins; no positive-a margin belongs to a structurally fixed $a = 0$ inverse.

Constructing the error set

For the Lorentz normalization, a possible exact pushforward is

$$u_i = \frac{c_0^2(\widehat{k}_i + \Delta k_i)^2}{\mu^2 \epsilon_\infty}, \qquad v_i = \frac{(\widehat{\omega}_i + \Delta\omega_i)^2 - \Delta_i^{\text{model}}}{\mu^2}. \tag{S4.30}$$

Here Δ_i^{model} is the separately bounded difference of physical squared frequency from the reference pole at the true momentum. Shared permittivity, thickness, momentum offset, frequency calibration and extraction errors remain shared nuisance variables. An uncertain denominator needs a positive lower bound. For a fixed scale, squaring an error interval gives the exact inequality $|(k + \Delta k)^2 - k^2| \le 2|k|\epsilon_k + \epsilon_k^2$. A model-fitted calibration is not an independent calibration. Standard errors become probabilistic sets only under explicit coverage assumptions; the deterministic theorem then applies on the stated coverage event and supplies no confidence level of its own.

S4.6 Two-point families with an unknown common scale

Now restrict to $a = 0$ but leave the shared calibration unknown:

$$x = sq, \qquad q = \kappa^2, \quad s > 0, \qquad F_s(y, q) = y^2 - (B + sq)y + Msq. \tag{S4.31}$$

For the data convention q is dimensionless and y, B, M, s have units eV^2. Take two exact lower observations with $0 < q_1 < q_2$, $0 < y_1 < y_2$. Define

$$h_i = y_i/q_i, \qquad d = h_1 - h_2 > 0, \tag{S4.32}$$

$$\alpha = \frac{y_2 - y_1}{d}, \qquad \beta = \frac{y_1 h_1 - y_2 h_2}{d}, \tag{S4.33}$$

$$m_\infty = y_1 + \alpha h_1 = y_2 + \alpha h_2, \qquad \gamma = \alpha h_1 h_2. \tag{S4.34}$$

The phase-slope ordering is necessary for two distinct lower points. The coefficient α is dimensionless; β, m_∞ have units eV^2 and γ has units eV^4.

Solving the two data equations gives the complete family

$$\boxed{B(s) = \alpha s + \beta, \qquad M(s) = m_\infty - \gamma/s.} \tag{S4.35}$$

Two useful identities are

$$\beta = m_\infty - \alpha(h_1 + h_2), \qquad B - M = \alpha \frac{(s - h_1)(s - h_2)}{s}. \tag{S4.36}$$

Both observations are lower roots in the strict domain $B > M > 0$ precisely when

$$\boxed{s > h_1.} \tag{S4.37}$$

On this interval, $M > y_2 > 0$, $B > M$, and $y_i < sq_i$ identifies the observed root as the lower one. Conversely, any lower observation requires $s > y_i/q_i$, giving (S4.37). At $s = h_1$, $B = M = y_2$: the coupling combination vanishes and the strict domain is lost.

The algebraic family also has $B > M > 0$ for $\gamma/m_\infty < s < h_2$, but both observations then exceed sq_i and are upper roots. Between h_2 and h_1, $B - M < 0$. These intervals cannot be used as lower-data solutions. On the admissible interval, the inferred gap increases monotonically as $\sqrt{\alpha s + \beta}$ while the two measured endpoints remain fixed. This is a calibration-dependent two-point ambiguity, distinct from any exceptional locus in the strict three-parameter theory.

S4.7 The limiting lower curve and information from a third point

The numerically stable lower root along (S4.35) is

$$y_s(q) = \frac{2M(s)sq}{B(s) + sq + \sqrt{[B(s) + sq]^2 - 4M(s)sq}}. \tag{S4.38}$$

Dividing the polynomial by s proves

$$y_\infty(q) = \lim_{s\to\infty} y_s(q) = \frac{m_\infty q}{\alpha + q}, \qquad \sqrt{B(s)} \sim \sqrt{\alpha s} \longrightarrow \infty. \tag{S4.39}$$

Thus a bounded limiting lower curve need not bound the inferred upper gap. Define

$$H(y, q) = y^2 - \beta y - \gamma q, \qquad G(y, q) = m_\infty q - (\alpha + q)y. \tag{S4.40}$$

The exact equation is $H + sG = 0$, and

$$H_\infty(q) = H(y_\infty(q), q) = -\frac{\gamma q(q - q_1)(q - q_2)}{(\alpha + q)^2}. \tag{S4.41}$$

The factorization follows by multiplying by $(\alpha+q)^2/q$: the resulting quadratic has zeros q_1, q_2 and leading coefficient $-\gamma$. For $q_1 < q < q_2$, $H_\infty > 0$ and

$$y_s > y_\infty, \quad \frac{\partial y_s}{\partial s} = \frac{G(y_s, q)}{B + sq - 2y_s} < 0, \quad \lim_{s\downarrow h_1} y_s = \min(h_1 q, y_2). \tag{S4.42}$$

If $D_s = s(\alpha + q) + \beta - 2m_\infty > 0$, factorization of $F_s(y_\infty)$ yields

$$|y_s - y_\infty| \le \frac{|H_\infty|}{D_s}, \qquad 0 < E_s - E_\infty \le \frac{H_\infty}{2E_\infty D_s} \quad (q_1 < q < q_2). \tag{S4.43}$$

Indeed $y_+ - y_\infty = B + sq - y_s - y_\infty \ge D_s$. In particular $y_s - y_\infty = H_\infty/[s(\alpha + q)] + O(s^{-2})$.

For a third datum, $G_j \neq 0$ determines $s_j = -H_j/G_j$. At an interior node this is admissible precisely when $y_\infty(q_j) < y_j < \min(h_1 q_j, y_2)$. If $G_j = 0$, then $H_j = H_\infty > 0$ and no finite scale fits. In row order $(1, 2, j)$, the linear system for (B, s, Ms) has rows $[-y_i, -q_i y_i, q_i]$ and determinant

$$\det \mathcal{A}_3 = q_1 q_2 (h_1 - h_2) G_j. \tag{S4.44}$$

Three compatible distinct ideal lower points therefore identify a finite model; every further exact point must imply the same scale. Small G_j describes poor conditioning, not a structural ambiguity with arbitrarily many exact points.

S4.8 Finite precision and an independent calibration interval

Keep both endpoints exact. At an interior node the map $s \mapsto E_s(q_j)$ decreases continuously from $E_{\mathrm{cap}} = \sqrt{\min(h_1 q_j, y_2)}$ to $E_\infty(q_j)$. Its inverse on this open energy range is

$$S_j(E) = -\frac{E^4 - \beta E^2 - \gamma q_j}{m_\infty q_j - (\alpha + q_j)E^2}. \tag{S4.45}$$

Intersect a separately justified energy interval $[E_j^-, E_j^+]$ with $(E_\infty, E_{\mathrm{cap}})$. Empty intersection means no finite compatible scale. Otherwise invert its endpoints in reversed order. If the upper energy endpoint reaches E_{cap}, the scale interval has lower limit h_1. Its upper endpoint is finite precisely when

$$E_j^- > E_\infty(q_j). \tag{S4.46}$$

An interval reaching the limiting curve from above permits arbitrarily large scales. For a symmetric radius ϵ_E around $E_j > E_\infty$, condition (S4.46) reads $\epsilon_E < E_j - E_\infty$. This is a required precision, not a supplied measurement error. Multiple points or layers sharing s contribute intersecting intervals together with $s > \max_N h_{1,N}$.

At an exact third-point fit, the local sensitivity is

$$\frac{\mathrm{d}s}{\mathrm{d}y_j} = \frac{B + sq_j - 2y_j}{G_j}, \qquad \left|\frac{\mathrm{d}\sqrt{B}}{\mathrm{d}E_j}\right| = \frac{\alpha E_j (B + sq_j - 2y_j)}{\sqrt{B}\,|G_j|}. \tag{S4.47}$$

The finite interval inversion gives more information than this derivative alone.

If independent calibration supplies $s \in [s_L, s_U]$ inside the admissible domain, then

$$\sqrt{B} \in [\sqrt{\alpha s_L + \beta}, \sqrt{\alpha s_U + \beta}]. \tag{S4.48}$$

To guarantee a positive target interval $[E_L, E_U]$, the calibration interval must lie inside $[(E_L^2 - \beta)/\alpha, (E_U^2 - \beta)/\alpha]$. For $B_0 = \alpha s_0 + \beta$ and a reference-relative target $|\sqrt{B}/\sqrt{B_0} - 1| \le \rho$, $0 < \rho < 1$, the exact scale requirement is

$$\frac{(1-\rho)^2 B_0 - \beta}{\alpha} \le s \le \frac{(1+\rho)^2 B_0 - \beta}{\alpha}. \tag{S4.49}$$

For a symmetric external interval $|s/s_0 - 1| \le \delta_s$, the tight symmetric requirement is $\delta_s \le (2\rho - \rho^2) B_0/(\alpha s_0)$, together with the domain margin. For several layers, intersect their requirements. If endpoints or coordinates are uncertain, all coefficients in (S4.34) must instead be propagated over their justified joint set. The exact-endpoint formulas alone then do not define an empirical confidence interval.

S4.9 Published-data instance and rounding conventions

The example uses all thirteen explicitly $N = 4, 5, 6, 7$ labelled observations for the 41-nm sample in the publisher's source workbook for Ref. [1], sheet *Extended Data Figure 10b*, cells A7:C19. Here N denotes layer count. Each layer's smallest and largest q form the endpoints; all five interior observations are retained. No model-generated workbook column is substituted for an observation.

Table S4.1: Printed coordinates used in the common-scale calculation. E denotes an endpoint and I an interior observation; row numbers identify the source workbook.

N	Row	κ	Energy (eV)	Role
4	7	0.25	1.09090811847	E
4	8	0.5	1.61392209177	I
4	9	0.75	1.75814105275	E
5	10	0.2	0.917035801391	E
5	11	0.4	1.45523366003	I
5	12	0.6	1.66865956004	I
5	13	0.8	1.76287977963	E
6	14	0.166666666666667	0.780680325	E
6	15	0.333333333333333	1.3021116949	I
6	16	0.5	1.58164	E
7	17	0.285714285714286	1.18128537546	E
7	18	0.428571428571429	1.48435309837	I
7	19	0.571428571428571	1.6399843132	E

The common-scale calculation treats the printed decimal strings as exact computational inputs and does not replace κ by an inferred rational mode ratio. Its nominal reference is the printed scalar $s_0 = 221.08533593187062\ \mathrm{eV}^2$. Separately, the fixed-nominal-momentum compatibility calculation uses the explicit high-precision x_i strings in its input record. Both represent the same nominal physical normalization at displayed precision; they are not asserted to be identical rational input arrays.

The nominal normalization is

$$s = \left(\frac{\hbar c\,\Gamma_L}{n_{\mathrm{eff}}}\right)^2, \quad \Gamma_L = \frac{\pi}{\sqrt{2/3}(d+g)}, \quad n_{\mathrm{eff}} = (1-F)\sqrt{\epsilon_m} + F, \tag{S4.50}$$

with $F = 0.74[d/(d+g)]^3$, $d = 41$ nm, gap $g = 3$ nm and $\epsilon_m = 1.96$. These nominal quantities do not supply a joint uncertainty interval. In particular, a reported particle-diameter spread alone does not specify uncertainty in the effective index, gap, boundary conditions or common momentum assignment. Extra decimal digits in the calculations express numerical precision, not physical accuracy.

S4.10 Conditional calibration results for the thirteen points

Each interior observation, together with its two layer endpoints, gives an admissible exact scale. The scales are different, including the two within $N = 5$:

Table S4.2: Exact-input scale candidates and distance above the limiting curve. Values are rounded for display.

N	Interior row	s_j (eV^2)	$E_j - E_\infty(q_j)$ (meV)
4	8	39.529980	41.132770
5	11	62.742127	34.249056
5	12	74.549441	17.402247
6	15	223.154429	4.241847
7	18	56.460611	11.861726

Thus the literal data do not fit one common scale with a separate (B, M) per layer. This is exact-input incompatibility, not statistical rejection. At $N = 6$, increasing the scale changes the inferred upper gap strongly while leaving the interior lower prediction near a finite limit:

Table S4.3: Mathematical scale stress cases for $N = 6$ with exact endpoints. The residual is prediction minus the row-15 observation.

Scale	$\sqrt{B_6}$ (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

These scale factors are not demonstrated metrological alternatives for the sample.

With all eight endpoints exact and one common free scale, define

$$R(s) = \max_j |E_s(q_j) - E_j|, \qquad s > \max_N h_{1,N}, \tag{S4.51}$$

where j runs over the same five interior points. Every signed residual decreases strictly. Their maximum and minimum therefore balance at the unique minimax point. Directed rational root enclosures establish

$$58.24533607 < s_* < 58.24533608 \quad \text{eV}^2, \tag{S4.52}$$

$$16.56795411423598 < R(s_*) < 16.56795411950704 \quad \text{meV}. \tag{S4.53}$$

Rows 8 and 15 provide the opposite active residuals. A uniform interior tolerance below this floor is incompatible with these exact endpoints. The limiting maximum is 41.132770 meV; at or above it, arbitrarily large common scales are allowed. Intermediate feasible tolerances bound the scale. These are conditional feasibility thresholds, not assigned experimental error bars.

S4.11 Fixed nominal momenta with all energies movable

A different problem fixes every nominal x_i and allows *all* energies, including both former endpoints, to move. For each layer separately, seek the smallest uniform tolerance δ admitting $B > M > 0$ with

$$E_-(x_i; B, M) \in [E_i - \delta, E_i + \delta] \quad \text{for every point in that layer.} \tag{S4.54}$$

Assume $0 \le \delta < E_i$ for every observation before squaring the energy bands. Let $L_i = (E_i - \delta)^2$, $U_i = (E_i + \delta)^2$, and suppose

$$0 < L_i \le U_i < x_i. \tag{S4.55}$$

Since the upper root exceeds x_i, the lower root lies in $[L_i, U_i]$ exactly when $F_0(L_i, x_i) \ge 0$ and $F_0(U_i, x_i) \le 0$. Rearrangement gives two linear inequalities in (B, M):

$$L_i B - x_i M \le L_i^2 - x_i L_i, \tag{S4.56}$$

$$-U_i B + x_i M \le x_i U_i - U_i^2. \tag{S4.57}$$

Thus the nonlinear-looking lower-root compatibility test is an exact intersection of half-planes, supplemented initially by $M \ge 0$ and $B \ge M$.

One construction enumerates pairwise intersections and tests all original inequalities. In the present tolerance bracket, the smallest and largest nodes satisfy $L_1/x_1 > U_2/x_2$, which gives a finite upper bound on B by combining the corresponding inequalities. The closed feasible polygon is therefore bounded. Feasible vertices can supply a strict witness, for example by an interior convex combination followed by direct checks of $M > 0$ and $B > M$.

An independent elimination expresses (S4.56)–(S4.57) as

$$\underbrace{\frac{L_i}{x_i} B + L_i - \frac{L_i^2}{x_i}}_{\ell_i(B)} \le M \le \underbrace{\frac{U_i}{x_i} B + U_i - \frac{U_i^2}{x_i}}_{u_i(B)}. \tag{S4.58}$$

Include 0 among the lower affine bounds and B among the upper bounds. Requiring every lower bound to be no larger than every upper bound eliminates M and leaves an exact one-dimensional interval for B. This procedure checks global feasibility without reproducing the polygon-vertex calculation.

The tolerance search uses rational inputs and brackets within $[0, 0.05]$ eV, where (S4.55) holds throughout for the selected data. $M = 0$ cannot satisfy positive lower-energy bands. At $B = M$ the lower root is $\min(B, x)$; because $U_i < x_i$, all bands would need to contain one common constant energy. They do not overlap. Boundary-only solutions therefore do not hide a physical solution in this data instance. Calibration errors remain outside this fixed-x calculation.

S4.12 Exact infeasibility certificates and nominal tolerance floors

Write the half-plane system as $Az \leq b$ with $z = (B, M)^T$. At each reported lower tolerance, a nonnegative vector λ obeys

$$\lambda^T A = 0, \qquad \lambda^T b < 0. \tag{S4.59}$$

Adding the original inequalities with these weights would give $0 \leq \lambda^T b < 0$. This is a direct global infeasibility certificate. At the upper tolerance, a strict physical witness satisfies every original inequality with nonnegative exact slack; direct lower-root predictions lie in all energy bands.

The checked certificates use three energy-band half-planes in each layer. "Lower" and "upper" in Table S4.4 refer to (S4.56) and (S4.57), respectively. The numerical bounds are brackets for the fixed-input compatibility floor, not confidence bounds.

Table S4.4: Fixed-nominal-momentum compatibility brackets. All energy values, including former endpoints, may move. Decimal entries are displays; exact fractions are given below.

N	Lower (meV)	Upper (meV)	Active certificate rows
4	17.099928855896	17.099934816361	7 upper; 8 lower; 9 upper
5	11.891114711761	11.891120672226	10 upper; 11 lower; 13 upper
6	0.019377470016	0.019383430481	14 lower; 15 upper; 16 lower
7	4.578083753586	4.578089714050	17 upper; 18 lower; 19 upper

The exact pairs (δ_-, δ_+) in eV are

$$N = 4: \quad (717223/41943040,\ 2868893/167772160), \tag{S4.60}$$
$$N = 5: \quad (997499/83886080,\ 1994999/167772160), \tag{S4.61}$$
$$N = 6: \quad (3251/167772160,\ 813/41943040), \tag{S4.62}$$
$$N = 7: \quad (153615/33554432,\ 192019/41943040). \tag{S4.63}$$

Every bracket has width $0.05/2^{23}$ eV, below 10^{-8} eV. Exact rational weights, cancelled coefficients, negative right-hand sides, feasible witnesses and their input bindings are retained in the accompanying machine-readable certificate record. The independent calculation also checked all supplied vertices for feasibility and extremality.

Rounding must not strengthen exclusion: for example, 17.099935 meV is already feasible for $N = 4$ although it is below the rounded number 17.10 meV. Rigorous exclusion uses the exact infeasible lower bracket. A tolerance below that bracket cannot accommodate these nominally calibrated printed data within the local model. A feasible upper bracket establishes compatibility only; it neither measures actual error nor validates the physical model. The approximately 16.568-meV result of the preceding common-scale problem has different fixed and free quantities and cannot replace these per-layer floors.

S4.13 Interpretation, source support and reproducibility

The thirteen published observations are reused data, not new independent measurements. At fixed nominal calibration, the lower-branch prediction on the five interior observations has RMSE 20.893636 meV. A two-parameter saturation curve $y = Aq/(C + q)$ fitted to the same two endpoints in each layer gives 25.793157 meV; an energy curve affine in κ gives 177.110576 meV. The 4.899521-meV advantage over saturation is descriptive. The alternatives were examined retrospectively, the observations share spectra and calibration, and no significance level follows from the five residuals. A saturation fit has no inferred upper-gap parameter.

The separately deposited 41-nm spectrum begins at 0.800000000000049 eV, while the published $N = 6$ lowest center is 0.780680325 eV. Its original extraction provenance is unresolved; the archive does not permit a two-sided reconstruction of that target. This does not establish how the original authors obtained it. The point remains in the declared published-data calculations above, making the $N = 6$ numerical conclusions conditional on that input.

Reference-normalized reflectance and transmittance, with absorbance formed as $100 - R - T$, do not constitute three independent repeated measurements. Recipe sensitivity, correlated fit residuals and optimization diagnostics are not quantitative coverage statements. A physical application of the joint-error theorem would require an adequate observation map, branch identification, independent calibration information and a justified correlated set covering extraction and model discrepancy. No such empirical certificate or independently observed upper mode is supplied here.

Computational records

The accompanying records preserve the literal input coordinates, nominal normalization, exact calibration-family results, compatibility brackets and independent certificate checks. The calibration identities were checked with rational arithmetic; the minimax bracket was checked by two distinct directed rational root-enclosure constructions. The nominal feasibility calculation was checked independently by affine elimination and exact certificates (S4.59). These computations support the derivations; they do not substitute for the stated assumptions or supply uncertainty sets. This supplement integrates those existing records without a new fitting or optimization run.

References

[1] N. 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. Publisher source data: Extended Data Figure 10, sheet 10b.

[2] S. Reich et al., original data deposit associated with Ref. [1], Refubium. doi:10.17169/refubium-26819. The repository declares CC BY-NC-SA 4.0; this does not assign that license to the separately hosted publisher workbook.