universe

MDPI

Article

Inverting the Quark Self-Energy from Chiral Low-Energy Constants: A Possible L_8 Tension in the Leading-N_c Rainbow Framework

Xiao Xiao * and De-Kai Kong

Key Laboratory for Relativistic Astrophysics, School of Physical Science and Technology, Guangxi University, Nanning 530004, China; m18843108872@163.com

* Correspondence: xx15367480116@163.com

Abstract

We pose and study the inverse problem of the chiral low-energy constants (LECs): with our re-implemented and checked Yang large-N_c forward operator (reproducing all eight LECs of Yang's Table I within 30%, four within 15%), we use the experimental LECs of Pan et al. and the lattice quark mass function of Oliveira et al. as data to invert for the quark self-energy $\Sigma(p^2)$. (1) Treating Σ as a free function results in a multimodal posterior; two independent analyses locate the LEC information window at ∼0.5–1.75 GeV and distinguish the profile-sensitive L_1–L_3, L_5 from the comparatively profile-rigid L_7, L_8 under the tested changes in Σ. (2) Parametrizing Σ through the kernel of the gap equation (Maris–Tandy/Qin–Chang) collapses the inversion to three kernel parameters (D, ω, m): the resulting posterior is unimodal, the lattice data and LECs are fitted simultaneously, and the uncertainties in Σ at representative momenta are reduced by factors of 2–4. (3) The fitted result contains four apparent deviations from Pan/FLAG. A proper-time regulator scan examines their response to the forward prescription. The apparent tensions in L_5, L_7 and the condensate are not particularly stable under the checks performed, whereas L_8 remains above the Pan central value throughout the scan. It also varies little across the two kernel realizations and the exploratory $M = B/A$ profile proxy. At the representative $\Lambda_{\rm reg} = 2\,\text{GeV}$ point, $10^3 L_8 = 0.99$ differs by 2.6σ from the quoted Pan value $10^3 L_8^{\rm Pan} = 0.60 \pm 0.15$; the condensate-based B_0 comparison and the scalar resonance-saturation estimate show the same tendency. These results support a possible tension associated with the scalar-channel LEC L_8 within the tested leading-N_c rainbow construction.

Keywords: chiral low-energy constants; quark self-energy inversion; large-N_c framework; kernel parametrization; information window; L_8 scalar saturation

Check for updates

Academic Editors: David Blaschke, Oleksii Ivanytskyi and Janusz Miśkiewicz

Received: 8 June 2026
Revised: 5 August 2026
Accepted: 6 August 2026
Published: 7 August 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

The low-energy constants (LECs) of chiral perturbation theory (ChPT) encode the dynamics of QCD in the pseudoscalar-meson sector [1–5]. A long-standing line of research connects them to the underlying theory through the quark self-energy $\Sigma(p^2)$: Yang, Wang, Kuang and Lu [6] worked in the large-N_c limit [7,8] within an improved ladder approximation, expressing the $O(p^4)$ LECs as a set of momentum integrals over $\Sigma(p^2)$ and its derivatives $(\kappa_1, \ldots, \kappa_{15})$, and obtained values broadly consistent with experiment by solving the Schwinger–Dyson equation [9–11]. The forward map $\Sigma \mapsto \{L_i\}$ is thereby well defined; yet its *inverse problem*—what the measured set of LECs implies for the self-energy, and what that comparison in turn reveals about the framework itself—has not been studied.

This paper poses and studies that inverse problem directly. Using the experimentally determined LECs of Pan et al. [12] together with the lattice quark mass function of Oliveira et al. [13] as data, we ask which self-energies are compatible with them under the Yang forward operator. We re-implement the operator in full (15 κ integrals, six auxiliary polynomials, the explicit constant-Σ subtraction, and the LEC assembly) as a differentiable PyTorch [14] map, and check it against Yang's Table I (all eight LECs within 30%, four within 15%; Section 2), including an audit of the remaining numerical discrepancies.

Two features make the inverse problem at once nontrivial and ultimately informative:

(i) It is underdetermined and multimodal. Six data points (or thirteen when the lattice pins are included) constrain a $\sim 10^3$-point discretization of $\Sigma(p)$. Treating Σ as a free function, the posterior is severely underdetermined. We show (Section 3) that the LEC data carry information about Σ only in a mid-momentum window $\sim$0.5–1.75 GeV (located by two independent analyses); moreover, different LECs respond differently to the tested changes in Σ—the profile-sensitive L_1–L_3, L_5 versus the comparatively profile-rigid L_7, L_8 (Section 5)—a classification that guides the subsequent diagnostics.

(ii) It can be constrained by physical priors. Rather than fitting a free Σ, we parametrize it through the kernel of the gap equation: replacing Yang's bare interaction with the Maris–Tandy/Qin–Chang effective interaction [15,16], we solve the (rainbow, $A = 1$) gap equation, generating a self-energy with the expected ultraviolet tail, monotonicity and dynamical chiral symmetry breaking built in. The inverse problem then collapses to three kernel parameters (D, ω, m); the resulting posterior is unimodal (Section 4).

These features enable a diagnostic comparison of four apparent deviations from Pan/FLAG. We use a proper-time regulator scan to examine their response to variations in the forward prescription. Within the checks performed, the apparent tensions in L_5, L_7 and the condensate are not particularly stable, while the independent η'-resonance estimate for L_7 is close to Pan [17,18]. L_8 shows a different pattern: it remains above the Pan central value throughout the regulator scan, responds only weakly to the tested changes in the Σ profile, and varies little across the two kernel realizations and the exploratory $M = B/A$ profile proxy. The condensate-based B_0 comparison and the scalar resonance-saturation estimate point in the same direction, supporting a possible L_8 tension within the tested leading-N_c rainbow construction (Section 5).

The paper is organized as follows. Section 2 presents the forward operator and its implementation check; Section 3 establishes the multimodality diagnosis of the free-Σ inversion and the LEC information window; Section 4 carries out the kernel-parametrized inversion; Section 5 examines the four apparent deviations using the LEC profile split, proper-time regulator scan, B_0 and resonance-saturation comparisons, and the two tested kernel realizations; Section 6 summarizes the conclusions and discusses them.

2. The Forward Operator

Following Yang et al. [6,19], in the large-N_c limit the $O(p^4)$ low-energy constants $L_1, \ldots, L_{10}, H_1, H_2$ are given by linear combinations of 15 κ integrals (their Equation (25)). Five of the κ vanish identically ($\kappa_5 = \kappa_6 = \kappa_8 = \kappa_{10} = \kappa_{12} = 0$), so the large-$N_c$ prediction is $L_4 = L_6 = 0$; the gap to the nonzero experimental values is the known $1/N_c$ deviation, and we exclude L_4, L_6 from the likelihood, keeping them as a model-bias check.

The remaining 10 nonzero κ integrals (Yang Equation (36)) involve the self-energy $\Sigma(p^2)$ and its first and second derivatives Σ', Σ'' together with six auxiliary polynomials A_p–F_p (Yang Equation (38)). The Wick-rotated integration measure is

$$\int d\tilde{p}\, F = -\frac{N_c}{16\pi^2} \int dt\, e^{2t} \exp\big[-(x + \Sigma^2)/\Lambda_{\mathrm{reg}}^2\big]\, F\big|_{\mathrm{Wick}}, \tag{1}$$

where $x = e^t$ is the Euclidean p_E^2 (dimensionless) and $\Lambda_{\rm reg}$ is the Schwinger regulator cutoff; the Wick rotation takes $p^2 \to -x$, $\Sigma_p' \to -d\Sigma/dx$, $X_p \equiv 1/(p^2 - \Sigma^2) \to -1/(x + \Sigma^2)$. After inserting the vanishing κ, the six LECs entering the likelihood are assembled from the κ (Yang Equation (25), $N_f = 3$) as

$$
\begin{aligned}
L_1 &= \frac{\kappa_4 + 2\kappa_{13} - \kappa_{14}}{32}, & L_2 &= \frac{\kappa_4 + 2\kappa_{13} - \kappa_{14}}{16}, \\
L_3 &= \frac{\kappa_3 - 2\kappa_4 - 6\kappa_{13} + 3\kappa_{14}}{16}, & L_5 &= \frac{\kappa_{11}}{16 B_0}, \\
L_7 &= -\frac{\kappa_1}{16 N_f} - \frac{\kappa_{15}}{16 B_0 N_f}, & L_8 &= \frac{\kappa_1 + (\kappa_7 - \kappa_9)/B_0^2 + \kappa_{15}/B_0}{16},
\end{aligned}
\tag{2}
$$

where B_0 is determined self-consistently by $F_0^2 B_0 = 4 \int d\tilde{p}\, \Sigma_p X_p$ (Yang Equation (34)). All operations are implemented as PyTorch tensors, ensuring automatic differentiability; a single forward pass takes ∼30 ms. The forward pass requires a dense grid for the explicit constant-Σ subtraction, for which we take $N = 1000$ points. Throughout this paper, numerical values quoted for the L_i without an explicit power of ten are given in units of 10^{-3}.

As an implementation check, we digitize the Model A self-energy $\Sigma_A(p^2)$ of Yang's Figure 2 [6], interpolate it onto the grid as forward input, and compare it against the first row of Yang's Table I (Table 1): all eight LECs agree within ±30%, four of them within ±15%, and the F_0^2 self-consistency residual is 1.3%.

Table 1. Implementation check of the forward operator on the Yang Model A baseline. $L_4 = L_6 = 0$ is the exact large-N_c prediction, independent of the fit; F_0^2 is checked self-consistently through the independent Pagels–Stokar relation [20].

LEC	This Work ($\times 10^3$)	Yang Table I ($\times 10^3$)	Deviation
F_0^2 (Pagels–Stokar)	0.0364	0.0369	−1.3%
B_0 (MeV)	914	∼1000	reasonable
L_1	+0.85	+1.10	−22.6%
L_2	+1.70	+2.20	−22.6%
L_3	−6.09	−7.82	+22.1%
L_5	+1.50	+1.62	−7.4%
L_7	−0.67	−0.70	+4.5%
L_8	+1.59	+1.75	−9.4%
L_9	+4.57	+5.07	−9.9%
L_{10}	−5.52	−7.06	+21.8%

To examine the likely origin of the 20%-level implementation-check residuals, we varied the quadrature density and domain, floating-point precision, ultraviolet continuation of the digitized self-energy, an independent digitization of the published curve, and the finite numerical realization of the constant-Σ subtraction. Grid and precision effects are below about 0.01%; across the converged integration domains, the reported LECs change by at most about 1.5% (and L_8 by at most 0.15%). The ultraviolet continuation and re-digitization produce observable-dependent shifts of several percent—about 4–5% for L_1, L_3 and 6–10% for L_7, L_8. The finite subtraction realization is more consequential: the tested variations are at the tens-of-percent level for L_1, L_3, below about 4% for L_5, about 10% for L_7, and about 10–14% for L_8. Because the numerical profile and complete implementation details were not published, the origins of these discrepancies cannot be determined uniquely. They do not alter the inverse-problem formulation or the differentiation, optimization and sampling procedures, but they do affect the quantitative interpretation of the fitted outputs; this point is taken up in Section 5.

3. Free-Σ Inversion: Multimodality and the LEC Information Window

3.1. Single-Chain Fit

Bayesian setup. We discretize $\Sigma(p^2)$ into $N = 1000$ points in the logarithmic momentum variable $t = \ln(p^2/\Lambda_{\rm QCD}^2)$, with the posterior $p(\boldsymbol{\Sigma} \mid \text{data}) \propto e^{-\chi^2/2}\, p(\boldsymbol{\Sigma})$. The likelihood takes the Gaussian form

$$\chi^2 = (\boldsymbol{L} - \boldsymbol{L}_{\rm Pan})^\top \Sigma_L^{-1} (\boldsymbol{L} - \boldsymbol{L}_{\rm Pan}) + \sum_{k=1}^{7} \left(\frac{\Sigma(p_k) - M_k}{\sigma_k} \right)^2, \tag{3}$$

where the first term contains Pan's six LECs ($L_1, L_2, L_3, L_5, L_7, L_8$), with the corresponding 6×6 block of Pan's covariance Σ_L [12]; the large-N_c-vanishing L_4, L_6 are excluded. $\boldsymbol{L} = \boldsymbol{L}(\boldsymbol{\Sigma})$ is assembled via Equation (2), and the second term contains Oliveira's 7 lattice pins [13], giving about 13 observations in total. The prior is a zero-mean Gaussian process [21] with RBF covariance $K_{ij} = \sigma_f^2 \exp[-(t_i - t_j)^2/2\ell^2] + \eta\delta_{ij}$ ($\sigma_f = 1.5$, $\ell = 1.0$, $\eta = 10^{-6}$). This prior is chosen because it encodes only the weakest physical constraint—smoothness of the self-energy: in the strongly underdetermined inversion of a 10^3-dimensional curve from $\sim$13 observations, it supplies the necessary physical prior without imposing excessive constraints.

Single-chain result: We find the MAP with L-BFGS [22], then sample a single chain with Hessian-preconditioned Metropolis [23,24]. The single chain already achieves a good fit ($\chi^2_{\rm med} = 11.0$): the posterior $\Sigma(p^2)$ band passes through all 7 Oliveira lattice points and reproduces Pan's LECs (Figure 1). As far as fitting the data themselves is concerned, the free-Σ scheme can therefore reproduce the modern LEC data within the Yang framework.

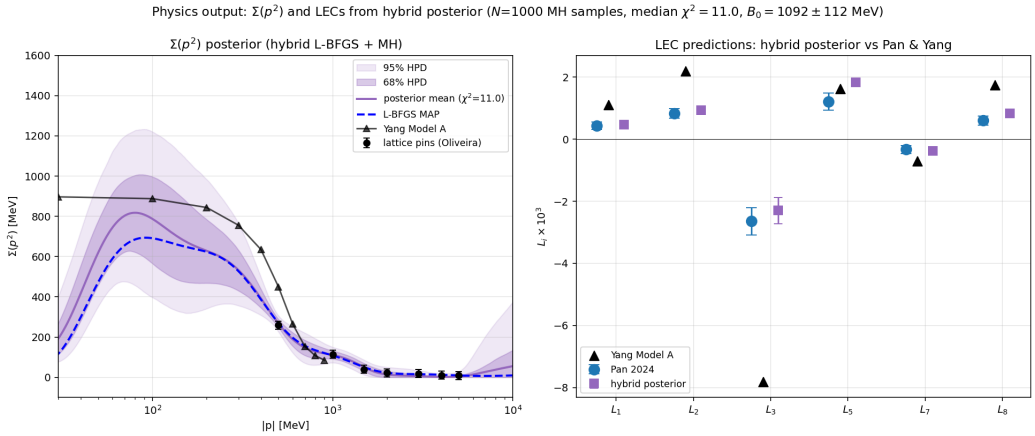


Figure 1. Posterior of the free-Σ inversion (L-BFGS MAP + Hessian-MH). **Left**: 68%/95% HPD band of $\Sigma(p^2)$, posterior mean ($\chi^2 = 11.0$), L-BFGS MAP, Yang Model A and Oliveira lattice points; **right**: LEC predictions (posterior vs. Pan et al. [12] and Yang). The band passes through all lattice points and reproduces the LECs; the single chain already reaches $\chi^2 \approx 11$.

3.2. Multiple Chains Reveal Multimodality

A good single-chain fit does not imply that the posterior is unique. We repeat the sampling with 12 independent chains, each with 6×10^4 steps, and compute the Gelman–Rubin $\hat{R}$ [25] (Table 2). Quantities constrained by the data agree across chains ($\hat{R} = 1.03$, ESS = 611 for χ^2), but the Σ directions not constrained by the data do not converge over the sampled runs (median componentwise $\hat{R} = 1.05$, max 1.21, with $\hat{R}$ not decreasing to 1.01 even after 6×10^4 total steps): different chains stabilize on different Σ shapes, i.e., the posterior is *multimodal*. Correspondingly, the deep-IR region appears unconstrained ($\Sigma(p \to 0) = 706 \pm 223$ MeV).

In other words, although the free-Σ inversion can fit the data, it does not yield a unique, well-constrained Σ: 13 data points are insufficient to pin down a 1000-dimensional curve, and the excess degrees of freedom manifest as multimodality. At the same time, the contrast between the good fit in the mid-momentum region and the unconstrained deep IR suggests that the information on Σ in the deep IR is insufficient. We therefore perform an LEC-sensitivity test in both the forward and the inverse directions to identify the cause of the multimodality of the free-Σ posterior. The posterior and diagnostics in Table 2 show the multimodality across chains.

Table 2. Posterior and convergence diagnostics of the free-Σ inversion with multiple chains (12 chains, each with 6×10^4 steps). Quantities constrained by data agree across chains; the $\hat{R}$ of unconstrained directions stays high, indicating posterior multimodality.

Quantity	Posterior / Diagnostic	Note
χ^2_{med}	11.0	Goodness-of-fit
$\Sigma(0.5\,\text{GeV})$	$268 \pm 20\,\text{MeV}$	In window, constrained
$\Sigma(p \to 0)$	$706 \pm 223\,\text{MeV}$	Outside window, essentially unconstrained
B_0	$1092 \pm 112\,\text{MeV}$	Condensate normalization
$\hat{R}(\chi^2)$	1.03	Data-constrained quantities agree across chains
$\hat{R}(\Sigma_{\text{dim}})$ median/max	1.05/1.21	Unconstrained directions do not converge (multimodal)
$\text{ESS}(\chi^2)$	611	Effective sample size

3.3. Forward/Inverse Scans Locate the Information Window

Forward line (sensitivity kernel). We perturb the baseline Σ (L-BFGS MAP) with a Gaussian bump whose width matches the prior length scale, scan it across momentum, and record the response $|\Delta L_i|/\sigma_{\text{Pan},i}$. The region in which the maximum response over the six LECs exceeds 1 is 0.62–1.78 GeV. **Inverse line (posterior width).** Propagating the posterior width from the MAP via the LEC-only Gauss–Newton Hessian, the ratio of posterior to prior width drops below 0.5 over 0.45–1.75 GeV, with the strongest constraint at 1.22 GeV. The two windows coincide at ∼0.5–1.75 GeV (Figure 2). Outside the window, the data carry little information, so the deep IR is weakly constrained. This reflects the limited sensitivity of the LECs to Σ in that region.

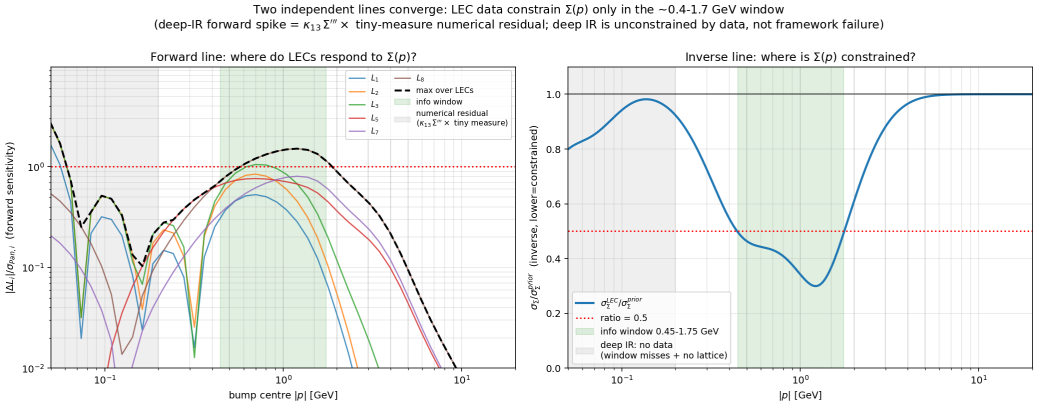


Figure 2. Two independent lines close on the ∼0.5–1.75 GeV information window. **Left**: forward sensitivity kernel, where the red dotted line marks the unit-sensitivity threshold; **right**: inverse posterior-to-prior width ratio, where the black horizontal line marks equality of the posterior and prior widths. The small forward peak in the deep IR is an unphysical numerical residual arising from the product of $\kappa_{13}\,\Sigma'''$ and a tiny measure.

3.4. Cross-Validation and Summary

As a cross-validation of the LECs' predictive power for Σ, we remove the two lattice pins at 0.5 and 1.0 GeV and refit using only the 6 LECs plus the 5 outer pins. The held-out prediction $\Sigma(1.0\,\mathrm{GeV}) = 137 \pm 71\,\mathrm{MeV}$ agrees with the removed lattice value $115 \pm 5\,\mathrm{MeV}$ within 0.47σ, with a prediction width of only 30% of the prior, indicating a constraint from the retained data rather than direct use of the removed pins.

Summary: Together, the two tests characterize the free-Σ inversion. First, the information-window analysis from two independent approaches shows that the LECs are insensitive to Σ outside the window (especially the deep IR); second, the cross-validation above shows that within the sensitive region the LECs and lattice data provide constraints supported by the held-out comparison (0.47σ agreement, with a width only 30% of the prior). In the mid-momentum region where the data carry information, the free-Σ inversion therefore gives a data-constrained Σ. The multimodality exposed in Section 3.2 is associated with **underdetermination in the deep IR**: the Σ there is only weakly constrained, and the excess degrees of freedom manifest as multimodality. This information deficit outside the window is the starting point for Section 4, where a physical kernel supplies the deep-IR degrees of freedom.

4. Kernel-Parametrized Inversion

4.1. From Fitting Σ to Fitting the Kernel

The free-Σ multimodality arises from using 13 data points to constrain 10^3 degrees of freedom. Rather than fitting Σ itself, we fit the kernel of the gap equation that produces Σ, compressing the curve's degrees of freedom to a few physical parameters. The large number of parameters allows the function to vary widely where deep-IR information is limited, whereas with only a few parameters, constraints from the LEC information window can propagate into the deep IR through the functional form. The kernel replacement also addresses a weakness of the Yang framework: Yang's bare interaction is $(3N_c/2)\,\alpha_s(k)/k^2$, whose running coupling $\alpha_s(k)$ lacks a first-principles definition in the infrared and therefore requires an explicit modeling and regularization prescription. We replace it with the Maris–Tandy/Qin–Chang effective interaction [15,16] $\mathcal{K}(s) = \mathcal{G}(s)/s$,

$$\mathcal{K}(s) = \frac{8\pi^2}{\omega^4} D\, e^{-s/\omega^2} + \frac{8\pi^2 \gamma_m}{\ln[\tau + (1 + s/\Lambda_{\mathrm{QCD}}^2)^2]} \cdot \frac{1 - e^{-s/4m_t^2}}{s}, \tag{4}$$

which characterizes the infrared strength by an infrared-finite Gaussian enhancement (strength D, width ω) and incorporates the expected ultraviolet asymptotic behavior through the logarithmic term. Its canonical parametrization is calibrated against hadron phenomenology and has an explicit infrared form. This replaces the infrared prescription of Yang's bare α_s by an explicitly parametrized interaction for the subsequent inversion. We retain the $A = 1$ truncation to keep the Yang forward operator (built on $X_p = 1/(p^2 - \Sigma^2)$) usable as is—the forward operator, the data layer and the κ assembly are all reused, and only the source of Σ is replaced. The self-energy is constrained by the gap equation, and the dimensionality collapses from 1000 to 3.

4.2. Gap Equation and Ultraviolet-Anchor Validation

The Euclidean, Landau-gauge, rainbow, $A = 1$ scalar equation [26] is

$$B(p^2) = m + \frac{1}{2\pi^3} \int_0^\infty dq^2\, q^2\, W_B(p^2, q^2) \frac{B(q^2)}{q^2 + B^2(q^2)}, \quad W_B = \int_0^\pi d\theta \sin^2\theta\, \mathcal{K}(p^2 + q^2 - 2pq\cos\theta), \tag{5}$$

with $\Sigma = B = M$. The two-dimensional angular integral must be evaluated explicitly (Yang's angular approximation is exact only for the bare $1/k^2$ kernel and fails when the kernel is replaced); for the bare kernel our numerical W_B agrees with the analytic expression to 0.2%. The overall constant is calibrated by physical anchors: a UV anchor ($\mathcal{G} \to 4\pi\alpha$, validated analytically) plus the asymptotic tail (the chiral-solution UV logarithmic slope $d\ln M/d\ln p$ measured as -2.42 over $[3, 10]$ GeV, consistent with Yang Equation (46)). Before using the solver in the inverse analysis, we also check it at the canonical MT parameter point $D = 1.024, \omega = 0.5$: the coupled $A \neq 1$ equation gives $M(0) = 504$ MeV, $A(0) = 1.70$, and $|\langle\bar{q}q\rangle|^{1/3} = 285$ MeV, at the expected constituent-mass scale quoted in the MT/FLAG literature [15,27,28]. This comparison is used as a reasonable implementation check for the coupled gap solver.

4.3. Three-Parameter Fit Results

The posterior $\propto e^{-\chi^2/2}$, where χ^2 combines contributions from the Pan LEC data, the Oliveira lattice data, and the f_π anchor, is computed directly on a three-dimensional grid in (D, ω, m), without MCMC sampling ($\sim$0.6 s per point). The joint fit gives $D = 0.385 \pm 0.042$, $\omega = 0.627 \pm 0.045$, $m = 5.9 \pm 2.5$ MeV; all 7 lattice points are fitted simultaneously ($\chi^2_{\text{pin}} = 5.7/7$), yielding $\Sigma(0.5\,\text{GeV}) = 261 \pm 9$ MeV, $\Sigma(1.0\,\text{GeV}) = 119 \pm 5$ MeV, with posterior uncertainties 2–4 times smaller than those from the 1000-dimensional free inversion reported in Table 2 (Figure 3). Because there are only three parameters and the posterior is integrated directly on the grid, the evaluated posterior is unimodal on the tested grid: the multimodality and non-converging $\hat{R}$ of the 1000-dimensional free inversion in Section 3.2 are not observed here. This supports the interpretation that the excess deep-IR degrees of freedom drive the multimodality of the free inversion; compressing Σ to three kernel parameters removes those degrees of freedom in the present fit. The resulting LECs are listed in Table 3, with small uncertainties ($\sigma \lesssim 0.07$), providing a baseline for the framework diagnosis of Section 5.

The fitted D differs from the canonical value quoted above, mainly because the latter comes from a calibration with coupled $A(p^2), B(p^2)$ dynamics and meson Bethe–Salpeter calculations [15,16], whereas the present inverse fit fixes $A = 1$ to reuse the Yang forward operator and refits D within this reduced equation. At the fitted point, the Pagels–Stokar estimate [20] of the pion decay constant is $F_{\text{PS}} = 92.4$ MeV, in close agreement with the input value $F_0 = 93$ MeV. This agreement provides a basic pion-sector consistency check and supports the physical plausibility of the present fit within the $A = 1$ setup.

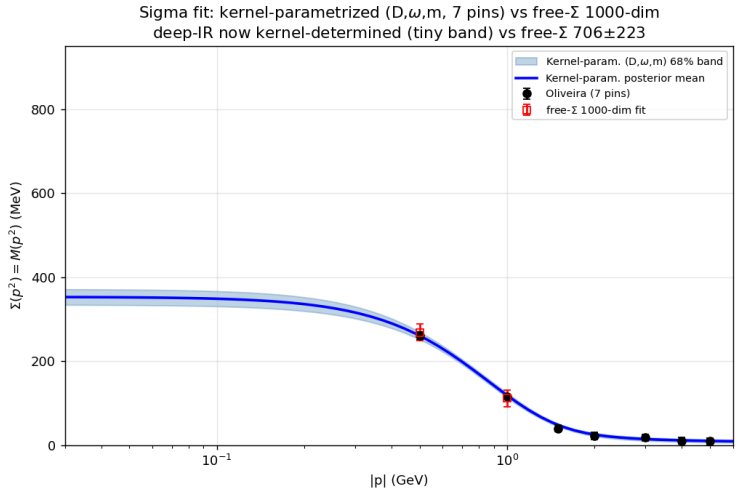


Figure 3. Posterior $\Sigma(p^2)$ band (blue) of the kernel-parametrized inversion (D, ω, m), passing through all 7 Oliveira lattice points (black).

Table 3. LECs from the three-parameter kernel-parametrized fit ($\Lambda_{\text{reg}} = 1.1\,\text{GeV}$; central value from the best self-energy, uncertainties from the three-parameter posterior standard deviations; units $\times 10^3$, pull relative to Pan [12]). The posterior is unimodal and the uncertainties are small; the higher L_5, L_7, L_8 values relative to Pan are analyzed in Section 5.

LEC	Kernel-Parametrized	Pan 2024	Pull
L_1	$+0.41 \pm 0.01$	$+0.43 \pm 0.12$	-0.2
L_2	$+0.81 \pm 0.01$	$+0.83 \pm 0.15$	-0.1
L_3	-2.96 ± 0.07	-2.64 ± 0.44	-0.7
L_5	$+1.96 \pm 0.03$	$+1.21 \pm 0.27$	$+2.8$
L_7	-0.61 ± 0.01	-0.33 ± 0.13	-2.2
L_8	$+1.54 \pm 0.03$	$+0.60 \pm 0.15$	$+6.2$

5. A Possible L_8 Tension in the Leading-N_c Rainbow Framework

A comparison of the framework forward predictions (kernel-parametrized LECs, Table 3) with Pan/FLAG [12,28] reveals four apparent deviations: L_5, L_7, the chiral condensate and L_8. This section compares their responses to identifiable regulator and Σ-shape variations. Within the variations examined here, the apparent deviations in L_5, L_7 and the condensate are less persistent, whereas the L_8 discrepancy is more persistent. The comparison rests on the response split of the LECs to the Σ profile (Section 5.1) and the proper-time regulator response (Section 5.2), followed by further diagnostics of L_8 (Sections 5.3 and 5.4).

5.1. Different Responses of the LECs to the Σ Profile

The LECs' sensitivity to Σ is concentrated in the mid-momentum information window (Section 3), but different LECs are affected differently by Σ. When the kernel replacement changes the overall Σ profile substantially ($\Sigma(0) : 890 \to 350\,\text{MeV}$), L_1–L_3 move by $\sim$50% and L_5 by 31% (profile-sensitive), whereas L_7, L_8 move by only 9% and 3% (profile-rigid, Figure 4); the sensitivity-peak positions also differ (L_1–L_3, L_5 at 0.7 GeV, L_7, L_8 at 1.2 GeV). This split originates in the $\kappa \to L$ assembly structure.

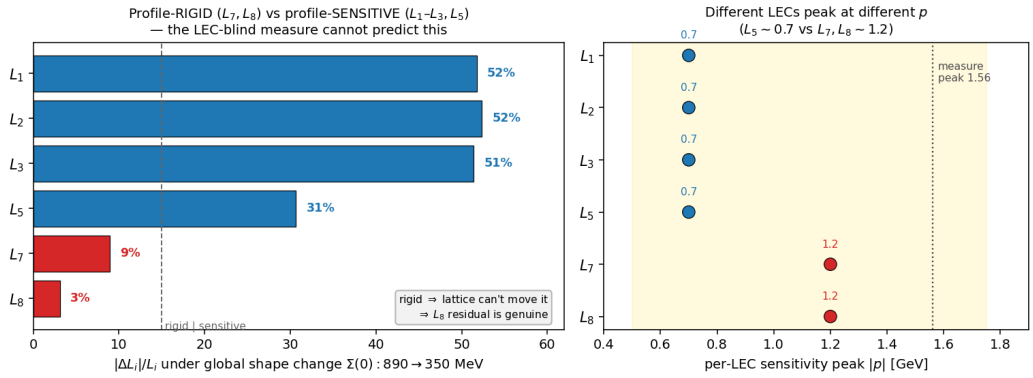


Figure 4. Comparatively profile-rigid (L_7, L_8) versus profile-sensitive (L_1–L_3, L_5) responses under the tested global profile change. **Left**: fractional shift of each LEC; **right**: position of each LEC's sensitivity peak.

This comparison shows that the profile-sensitive L_1–L_3, L_5 are strongly affected by the choice of Σ/kernel, so replacing the kernel (Section 4) can move them away from the Yang bare-kernel prediction toward the modern values; L_7, L_8 are comparatively insensitive to this particular profile change at fixed forward prescription.

The fixed-prescription qualifier is important. The implementation audit in Section 2 finds a subtraction response below about 4% for L_5, about 10% for L_7, and about 10–14% for L_8. Although not negligible, the subtraction responses of L_7 and L_8 are moderate compared

with the tens-of-percent responses of L_1 and L_3. In the present inverse fit and kernel-replacement comparison, the change in the shape of Σ is the primary variation of interest. Combining this consideration with the moderate subtraction responses of L_7 and L_8, we give priority in the subsequent analysis to these two comparatively profile-rigid LECs.

5.2. Proper-Time Regulator Sensitivity of B_0, L_5 and L_7

The three LECs carrying B_0 (Yang Equation (25)) are

$$L_5 = \frac{\kappa_{11}}{16B_0}, \quad L_7 = -\frac{\kappa_1}{16N_f} - \frac{\kappa_{15}}{16B_0N_f}, \quad L_8 = \frac{1}{16}\left(\kappa_1 + \frac{\kappa_7 - \kappa_9}{B_0^2} + \frac{\kappa_{15}}{B_0}\right). \tag{6}$$

The relation $B_0 = -\langle\bar{q}q\rangle / F_0^2$ [3] enters together with the κ integrals, and both are recomputed self-consistently for the same self-energy at each proper-time regulator value Λ_{reg}. The ChPT renormalization scale μ_χ enters directly in the definition and phenomenological comparison of the renormalized LECs $L_i^r(\mu_\chi)$, but incorporating this scale dependence consistently into the inverse fit is nontrivial. For example, the one-loop ChPT property $\Gamma_7 = 0$ [3,18] provides an additional constraint on the interpretation of the L_7 response. As a practical preliminary diagnostic, we therefore examine the response of the apparent deviations in B_0, L_5, L_7 and L_8 to a scan of the forward-model regulator Λ_{reg}. The aim is to identify which deviations remain comparatively persistent under this controlled variation and therefore merit priority in the subsequent analysis.

We perform a multipoint calculation in which the same self-energy is evaluated at 14 regulator values from 1.1 to 2.25 GeV, with all κ_i and B_0 recomputed at every point (Figure 5). Table 4 lists two representative points. Between $\Lambda_{\mathrm{reg}} = 1.1$ and 2.0 GeV, B_0 rises by 86%, while the κ numerators rise by only 16–27%; consequently, $L_5 : 1.96 \to 1.34$, $L_7 : -0.61 \to -0.36$, and $L_8 : 1.54 \to 0.99$.

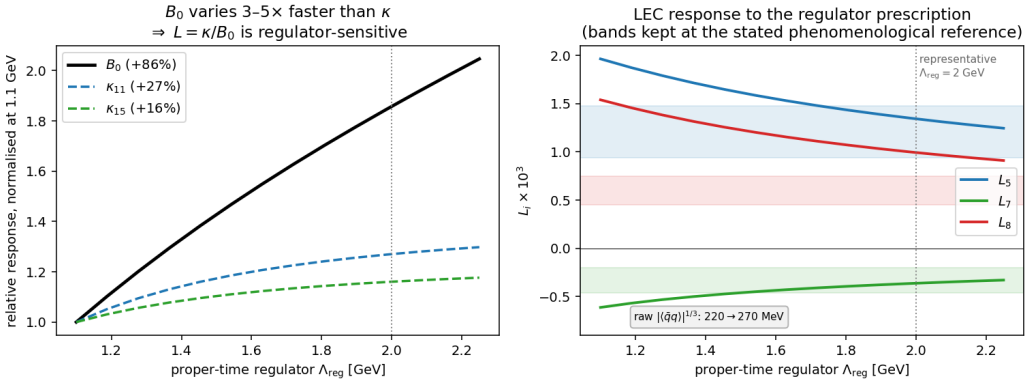


Figure 5. The 14-point proper-time regulator scan. **Left**: relative responses of B_0, κ_{11} and κ_{15}. **Right**: L_5, L_7, L_8 as functions of the proper-time regulator; the phenomenological bands are held fixed at their stated reference.

Over the scanned interval, L_5, L_7 and the condensate estimate all move toward the quoted Pan/FLAG values. For L_7, an additional comparison independent of the present inversion is provided by the resonance-saturation η'-channel value [17,18] $L_7^{\eta'} = -0.3 \times 10^{-3}$, which is close to both the framework value -0.36×10^{-3} and the Pan value -0.33×10^{-3} [12]. The condensate estimate changes from about 220 to 271 MeV, approaching the quoted FLAG value [28]. Across these diagnostic checks, the apparent tensions in L_5, L_7, and the condensate are not particularly stable; the following subsection therefore focuses on the more persistent candidate.

L_8 behaves differently in this diagnostic. Inverting the assembly relation $L_i(B_0) = L_i^{\mathrm{Pan}}$ with the baseline κ_i held fixed gives the effective B_0 required to reproduce each Pan

central value: L_5 requires 2.01 GeV and L_7 requires 2.17 GeV, whereas L_8 requires 3.13 GeV, substantially higher than the values required for the former two. We therefore examine L_8 further in the following subsection.

Table 4. Representative points from the multipoint proper-time regulator scan (kernel-parametrized best self-energy).

Λ_{reg} (GeV)	B_0 (MeV)	κ_{11}	κ_{15}	L_5	L_7	L_8	$\lvert\langle\bar{q}q\rangle\rvert^{1/3}$
1.1	1240	0.0805	0.0807	+1.96	−0.61	+1.54	220
2.0	2303	0.1022	0.0936	+1.34	−0.36	+0.99	271
Change	+86%	+27%	+16%	−32%	—	—	—
Pan/FLAG	—	—	—	1.21 ± 0.27	-0.33 ± 0.13	0.60 ± 0.15	272 ± 5

5.3. Further Diagnostics of L_8

At the representative $\Lambda_{\mathrm{reg}} = 2$ GeV point, $L_8^{\mathrm{framework}} = 0.99$. Comparison with the quoted Pan value $L_8^{\mathrm{Pan}} = 0.60 \pm 0.15$ [12] gives a numerical difference of 2.6σ. Across the 14-point scan over 1.1–2.25 GeV, L_8 exhibits a different pattern from L_5, L_7, and the condensate: the latter three reach their quoted Pan/FLAG ranges, whereas L_8 remains above the Pan central value throughout. This contrast draws attention to the L_8 discrepancy.

The condensate provides a reference independent of the regulator scan. Using the Yang normalization $F_0 = 93$ MeV [6] and the quoted FLAG value $\lvert\langle\bar{q}q\rangle\rvert^{1/3} = 272 \pm 5$ MeV [28], the relation $B_0 = -\langle\bar{q}q\rangle / F_0^2$ gives $B_0^{\mathrm{FLAG/Yang}} \simeq 2.33$ GeV. The algebraic values obtained above with the baseline κ_i held fixed, 2.01 GeV for L_5 and 2.17 GeV for L_7, lie near this condensate-based reference, whereas L_8 requires the substantially larger value 3.13 GeV.

A further external comparison is provided by resonance saturation. In the standard resonance-saturation decomposition [17,18], L_8 is saturated purely by the scalar channel (vector, axial-vector and η_1 contributions are all zero), with the large-N_c scalar value $L_8^S = 0.9 \times 10^{-3}$. This agrees with the older phenomenological value $L_8^r(M_\rho) = 0.9 \pm 0.3$, but is above the modern NLO value 0.60 ± 0.15. By contrast, L_7 is saturated by the $\eta_1(\eta')$ channel, whose estimate near -0.3×10^{-3} is consistent with Pan. The regulator response, the condensate-based B_0 comparison and the resonance-saturation pattern all direct attention to L_8 within the leading large-N_c framework. Omitted subleading scalar contributions provide a possible origin of this pattern [29].

5.4. Effects of Rainbow Truncation and Kernel Choice

To examine the sensitivity of L_8 to the mass-function profile, we solve the coupled rainbow gap equations for $A(p^2)$ and $B(p^2)$ [15,16] and insert the resulting mass function $M = B/A$ into the existing Yang $A = 1$ forward operator ($A(0) = 1.30$; Table 5, Figure 6). Three input self-energies of very different profiles, with $M(0.5\,\mathrm{GeV})$ ranging from 255 to 449 MeV, give $10^3 L_8 = 1.54$–1.59 in this common assembly, dominated by the scalar term κ_{15}/B_0 (128–135%). Across the tested coupled (A, B) parameter grid, the lattice-compatible proxies with $\chi^2_{\mathrm{pin},7} \leq 7$ give $10^3 L_8 = 1.529$–1.614. Thus, the mass-function substitution does not remove the L_8 discrepancy within the tested grid.

The $M = B/A$ result remains a profile proxy rather than a complete $A \neq 1$ calculation of the LECs. The quoted range therefore characterizes the profile sensitivity within the unchanged assembly, rather than the uncertainty of a complete $A \neq 1$ forward operator. Yang's own Table I [6] gives $10^3 L_8 = 1.75$, which itself lies above the modern Pan value, independently of the kernel replacement and inverse fit introduced in this work. Yang's result, the MT/QC-based $A = 1$ result and the exploratory proxy all yield high values across the tested profiles and the two kernel realizations, further strengthening the indication of a possible L_8 tension.

Table 5. Results after the kernel replacement: L_8 varies only over a narrow range.

Self-Energy	$M(0.5\,\mathrm{GeV})$ (MeV)	L_7	L_8
Yang Model A ($A = 1$, digitized)	449	-0.67	$+1.59$
Kernel-parametrized ($A = 1$ fit)	265	-0.61	$+1.54$
$M = B/A$ proxy in the $A = 1$ operator	255	-0.63	$+1.59$
Yang Table I (published)	—	-0.70	$+1.75$
Pan 2024 (experiment)	260	-0.33	$+0.60$

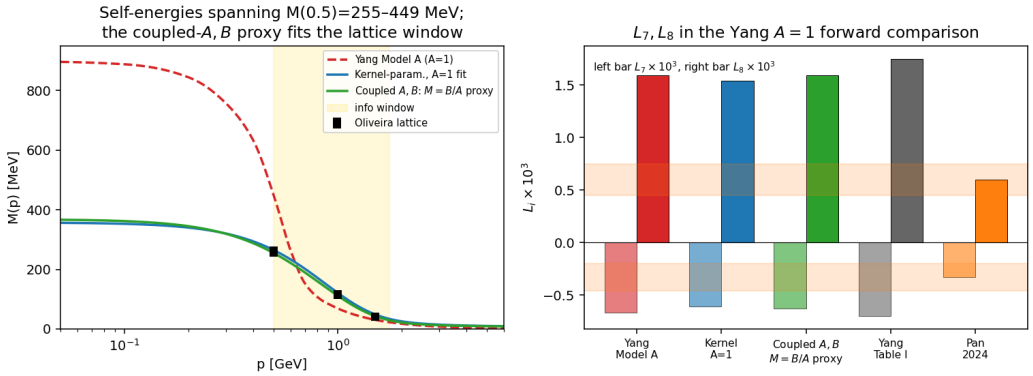


Figure 6. Left: the three scalar input profiles coincide in the window and differ in the deep infrared. **Right**: the predictions for L_7, L_8 are ∼2–3 times the corresponding Pan et al. [12] values in the configurations shown. The orange horizontal bands denote the corresponding quoted Pan uncertainty intervals.

5.5. Summary

Of the four apparent deviations in the fitted result, those in L_5, L_7 and the condensate are not persistent under the relevant diagnostic variations. L_5 and L_7 move substantially under the proper-time regulator scan, while the condensate comparison remains prescription-dependent. For L_7, the independent η'-resonance estimate is also close to the Pan value. These results therefore leave only weak indications of tension in these three quantities.

L_8 displays a different pattern. It remains above the Pan central value throughout the proper-time regulator scan and remains high across the tested self-energy profiles, the two kernel realizations and the exploratory $M = B/A$ proxy. The effective B_0 required to reproduce the Pan central value, 3.13 GeV, is also substantially higher than the condensate-based reference 2.33 GeV and the values required for L_5 and L_7. Moreover, the scalar resonance-saturation value $10^3 L_8^S \simeq 0.9$ itself lies above the modern Pan value. Taken together, these observations strengthen the indication of a possible tension localized in the scalar-channel LEC L_8 within the tested leading-N_c rainbow construction.

6. Conclusions

We have turned the LEC-to-self-energy inverse problem into an end-to-end quantitative diagnosis. The main conclusions are as follows.

Method and diagnosis. We re-implemented the Yang large-N_c forward operator and checked it against Yang's Table I, reproducing all eight LECs within 30% and four within 15%, with F_0^2 self-consistency at the 1.3% level. The implementation audit finds negligible grid and precision effects, percent-level responses to the digitized profile and its ultraviolet continuation, and larger observable-dependent responses to the subtraction prescription. Since the original numerical profile and full implementation details are not available, the origins of the remaining discrepancies cannot be determined uniquely, and these sensitivities are included in the interpretation of the LEC comparisons. Treating Σ as

a free function gives a multimodal posterior; two independent analyses locate the LEC information window at $\sim$0.5–1.75 GeV, while the held-out prediction agrees with the removed lattice point within 0.47σ and has a width equal to 30% of the prior. Parametrizing the gap equation by three kernel parameters yields a unimodal posterior on the tested grid, fits the lattice and LEC data simultaneously, and reduces the uncertainty in Σ by a factor of 2–4.

Phenomenological indication: a possible L_8 tension in the tested leading-N_c rainbow framework. The fitted result contains four apparent deviations from Pan/FLAG. We use the proper-time regulator scan to examine how these four apparent tensions respond to variations in the forward prescription. Within the checks performed, the apparent tensions in L_5, L_7 and the condensate are not particularly stable, and the independent η'-resonance estimate for L_7 is close to the Pan value. L_8 shows a different pattern. Although it changes under the regulator and subtraction prescriptions, it remains above the Pan central value throughout the regulator scan. It also varies little across the two kernel realizations and the exploratory $M = B/A$ profile proxy. At the representative $\Lambda_{\rm reg} = 2\,\text{GeV}$ point, $10^3 L_8 = 0.99$ differs by 2.6σ from the quoted Pan value 0.60 ± 0.15. The effective B_0 required to reproduce the Pan central value is 3.13 GeV, above the condensate-based reference 2.33 GeV, while the scalar resonance-saturation estimate $10^3 L_8^S \simeq 0.9$ is also above Pan. These observations support a possible tension associated with the scalar-channel LEC L_8 within the construction tested here.

Limitations and outlook. The present kernel comparison covers two interaction realizations, and the $M = B/A$ substitution is a profile-level diagnostic rather than a complete $A \neq 1$ calculation of the LECs. A broader kernel comparison and a full $A \neq 1$ forward operator would therefore be needed to assess general kernel and truncation dependence. Cross-validation of the kernel-parametrized inversion and a treatment beyond rainbow with a consistent Bethe–Salpeter equation [30–32] may be addressed in future work. More broadly, the method applies to inverse problems of similar structure (high-dimensional parameters, few observations, and a physically regularizable forward operator).

Author Contributions: Conceptualization, X.X. and D.-K.K.; methodology, X.X. and D.-K.K.; software, X.X.; validation, X.X.; formal analysis, X.X.; investigation, X.X.; data curation, X.X.; visualization, X.X.; writing—original draft preparation, X.X.; writing—review and editing, X.X. and D.-K.K.; resources, D.-K.K.; supervision, D.-K.K. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

ChPT	Chiral perturbation theory
LEC	Low-energy constant
QCD	Quantum chromodynamics
DSE	Dyson–Schwinger equation
MT	Maris–Tandy
QC	Qin–Chang

IR	Infrared
UV	Ultraviolet
MAP	Maximum a posteriori
HPD	Highest posterior density
ESS	Effective sample size
GP	Gaussian process
RBF	Radial basis function

References

1. Weinberg, S. Phenomenological Lagrangians. *Physica A* **1979**, *96*, 327–340. [CrossRef]
2. Gasser, J.; Leutwyler, H. Chiral Perturbation Theory to One Loop. *Ann. Phys.* **1984**, *158*, 142–210. [CrossRef]
3. Gasser, J.; Leutwyler, H. Chiral Perturbation Theory: Expansions in the Mass of the Strange Quark. *Nucl. Phys. B* **1985**, *250*, 465–516. [CrossRef]
4. Scherer, S. Introduction to Chiral Perturbation Theory. *Adv. Nucl. Phys.* **2003**, *27*, 277–538.
5. Bijnens, J.; Ecker, G. Mesonic Low-Energy Constants. *Annu. Rev. Nucl. Part. Sci.* **2014**, *64*, 149–174. [CrossRef]
6. Yang, H.; Wang, Q.; Kuang, Y.P.; Lu, Q. Calculation of the Chiral Lagrangian Coefficients from the Underlying Theory of QCD: A Simple Approach. *Phys. Rev. D* **2002**, *66*, 014019. [CrossRef]
7. 't Hooft, G. A Planar Diagram Theory for Strong Interactions. *Nucl. Phys. B* **1974**, *72*, 461–473. [CrossRef]
8. Kaiser, R.; Leutwyler, H. Large N_c in Chiral Perturbation Theory. *Eur. Phys. J. C* **2000**, *17*, 623–649. [CrossRef]
9. Roberts, C.D.; Williams, A.G. Dyson-Schwinger Equations and Their Application to Hadronic Physics. *Prog. Part. Nucl. Phys.* **1994**, *33*, 477–575. [CrossRef]
10. Alkofer, R.; von Smekal, L. The Infrared Behaviour of QCD Green's Functions: Confinement, Dynamical Symmetry Breaking, and Hadrons as Relativistic Bound States. *Phys. Rep.* **2001**, *353*, 281–465. [CrossRef]
11. Maris, P.; Roberts, C.D. Dyson-Schwinger Equations: A Tool for Hadron Physics. *Int. J. Mod. Phys. E* **2003**, *12*, 297–365. [CrossRef]
12. Pan, H.X.; Kong, D.K.; Wen, Q.Y.; Jiang, S.Z. Bayesian Method for Fitting the Low-Energy Constants in Chiral Perturbation Theory. *Front. Phys.* **2024**, *19*, 64203. [CrossRef]
13. Oliveira, O.; Kızılersü, A.; Silva, P.J.; Skullerud, J.I.; Sternbeck, A.; Williams, A.G. Lattice Landau Gauge Quark Propagator and the Quark-Gluon Vertex. *Acta Phys. Pol. B Proc. Suppl.* **2016**, *9*, 363–368. [CrossRef]
14. Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Proceedings of the Advances in Neural Information Processing Systems*; Neural Information Processing Systems Foundation, Inc.: San Diego, CA, USA,2019; Volume 32, pp. 8024–8035.
15. Maris, P.; Tandy, P.C. Bethe-Salpeter Study of Vector Meson Masses and Decay Constants. *Phys. Rev. C* **1999**, *60*, 055214. [CrossRef]
16. Qin, S.X.; Chang, L.; Liu, Y.X.; Roberts, C.D.; Wilson, D.J. Interaction Model for the Gap Equation. *Phys. Rev. C* **2011**, *84*, 042202. [CrossRef]
17. Ecker, G.; Gasser, J.; Pich, A.; de Rafael, E. The Role of Resonances in Chiral Perturbation Theory. *Nucl. Phys. B* **1989**, *321*, 311–342. [CrossRef]
18. Pich, A. Chiral Perturbation Theory. *Rep. Prog. Phys.* **1995**, *58*, 563–609. [CrossRef]
19. Lu, Q.; Yang, H.; Wang, Q. Fermion Determinant with Dynamical Chiral Symmetry Breaking. *Commun. Theor. Phys.* **2002**, *38*, 185–188. [CrossRef]
20. Pagels, H.; Stokar, S. Pion Decay Constant, Electromagnetic Form Factor, and Quark Electromagnetic Self-Energy in Quantum Chromodynamics. *Phys. Rev. D* **1979**, *20*, 2947–2952. [CrossRef]
21. Rasmussen, C.E.; Williams, C.K.I. *Gaussian Processes for Machine Learning*; MIT Press: Cambridge, MA, USA, 2006.
22. Liu, D.C.; Nocedal, J. On the Limited Memory BFGS Method for Large Scale Optimization. *Math. Program.* **1989**, *45*, 503–528. [CrossRef]
23. Metropolis, N.; Rosenbluth, A.W.; Rosenbluth, M.N.; Teller, A.H.; Teller, E. Equation of State Calculations by Fast Computing Machines. *J. Chem. Phys.* **1953**, *21*, 1087–1092. [CrossRef]
24. Hastings, W.K. Monte Carlo Sampling Methods Using Markov Chains and Their Applications. *Biometrika* **1970**, *57*, 97–109. [CrossRef]
25. Gelman, A.; Rubin, D.B. Inference from Iterative Simulation Using Multiple Sequences. *Stat. Sci.* **1992**, *7*, 457–472. [CrossRef]
26. Higashijima, K. Dynamical Chiral-Symmetry Breaking. *Phys. Rev. D* **1984**, *29*, 1228–1232. [CrossRef]
27. Bhagwat, M.S.; Pichowsky, M.A.; Roberts, C.D.; Tandy, P.C. Analysis of a Quenched Lattice-QCD Dressed-Quark Propagator. *Phys. Rev. C* **2003**, *68*, 015203. [CrossRef]
28. Aoki, Y. et al. [Flavour Lattice Averaging Group (FLAG)] FLAG Review 2021. *Eur. Phys. J. C* **2022**, *82*, 869. [CrossRef]

29. Sanz-Cillero, J.J. Scalar and Pseudoscalar Correlators in Resonance Chiral Theory. *Nucl. Phys. B Proc. Suppl.* **2010**, *207–208*, 236–239. [CrossRef]
30. Chang, L.; Roberts, C.D. Sketching the Bethe–Salpeter Kernel. *Phys. Rev. Lett.* **2009**, *103*, 081601. [CrossRef] [PubMed]
31. Chang, L.; Roberts, C.D. Tracing Masses of Ground-State Light-Quark Mesons. *Phys. Rev. C* **2012**, *85*, 052201. [CrossRef]
32. Williams, R.; Fischer, C.S.; Heupel, W. Light Mesons in QCD and Unquenching Effects from the 3PI Effective Action. *Phys. Rev. D* **2016**, *93*, 034026. [CrossRef]

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