

The Exact Matter-Contraction Non-Gaussian Amplitude: Four-Vertex Derivation and Conditional Large-Scale-Structure Mapping

Houston Golden^{1,*}

¹*Independent Researcher, Los Angeles, California, USA*

(Dated: July 24, 2026, 18:35 PDT)

A matter-dominated contracting phase gives a local-type non-Gaussian amplitude $f_{\text{NL}}^{\text{local}} = -35/16 = -2.1875$ before the nonsingular transition. We derive this contraction-phase coefficient for the stated $\epsilon = 3/2$ background and cubic action by re-summing all four cubic vertices, re-expand the result in the ordered symmetric basis, and obtain the unique coefficients $(3, 1, -9, 5, -33, 9)$; independent checks use Cai *et al.*'s order-grouped expressions and Li *et al.*'s general- c_s formula. The result corrects the unreproduced printed $-35/8$ literature value. For orientation only, and conditional on faithful cubic-order transmission through a specified bounce completion, we map the published Heinrich *et al.* SPHEREx multi-tracer bispectrum sensitivity through the exact shape. Its flat-grid amplitude recovery is $r = 0.8354$ and shape cosine is $r_{\text{cos}} = 0.9817$; the corresponding arithmetic map, recomputed with the adopted $r = 0.84$ convention, is 2.63σ before additional nuisance marginalization. A channel-native surrogate-covariance check spans 3.5σ with nuisances fixed, 3.1σ after marginalizing the relativistic-projection amplitude A_{GR} , 2.3σ with an explicit 30% theory prior on the PNG bias-response coefficient b_ϕ , and 0.4σ when b_ϕ is free. These values are illustrative conditional diagnostics, not an observational headline, a new joint-covariance forecast, or a detection forecast. The primary contribution is the exact contraction-phase amplitude derivation.

I. INTRODUCTION

Scope and conventions. The paper's primary contribution is an exact four-vertex contraction-phase amplitude derivation. The observational numbers are illustrative conditional mappings, not a unified or independent survey forecast. They map the published Heinrich *et al.* [1] $\sigma(f_{\text{NL}}^{\text{local}}) \approx 0.7$ baseline through Eq. (9), use the exact-shape overlap, and test nuisance sensitivity with an in-house surrogate covariance. The latter does not replace the unpublished external per-triangle covariance. Cubic transmission through the bounce is a load-bearing model assumption: it is established at linear order for the specified completion, while the nonlinear extension is motivated by single-clock conservation and is not claimed here as a completed third-order calculation. We therefore report conditional diagnostics, not a derived measurement precision. Here r denotes amplitude recovery, r_{eff} the survey-weighted recovery check, r_{cos} the shape cosine, and ρ a Fisher correlation. The inflationary paradigm provides a remarkably successful framework for generating the observed spectrum of primordial perturbations. Standard single-field slow-roll inflation gives a nearly scale-invariant, nearly Gaussian spectrum with the conventional primordial squeezed-template coefficient $f_{\text{NL}} \approx (5/12)(1 - n_s) \approx 0.015$ from the Maldacena consistency relation [2]. This global-template coefficient is not by itself an independently observable scale-dependent-bias signal: local-observer consistency-relation and projection treatments absorb the leading coordinate contribution, leaving only slow-roll and projection residuals [3, 4]. We therefore use 0.015 only as a theory-

normalization benchmark; the bounce-vs-inflation contrast remains $|f_{\text{NL}}^{\text{bounce}}| \gg |f_{\text{NL}}^{\text{inf}}|$ after the observable projection treatment.

Bouncing cosmology offers an alternative origin for primordial perturbations: modes exit the Hubble radius during a contracting phase and re-enter after a nonsingular bounce. In particular, a matter-dominated contraction ($w \approx 0$) produces a scale-invariant scalar spectrum through the growth of the curvature perturbation ζ on superhorizon scales [5, 6].

A distinctive pre-bounce result is the negative local amplitude $f_{\text{NL}} = -35/16 = -2.1875$ [7–9]. Appendix B derives it from the exact four-vertex sum and documents why the separately printed $-35/8$ is not reproduced. The leading amplitude is fixed once the matter-contraction background and cubic action are specified; quasi-dust corrections remain model dependent. Applying this pre-bounce amplitude to late-time observables additionally requires faithful nonlinear transfer across the chosen nonsingular transition. That transfer is verified only at linear order in the cited Wilson–Ewing construction, so all observational statements in this paper are explicitly conditional on assumption (d) of Sec. II C. The result is not asserted to be model independent across bounce completions, post-bounce inflationary histories, or fermion-sourced torsion sectors. SPHEREx [10] can test local-type non-Gaussianity through scale-dependent bias and the galaxy bispectrum [1, 11]. We recast its published local-template sensitivity onto the exact matter-bounce shape, itemize the assumptions and nuisance dependence, and keep the model-comparison exercise subordinate to its prior dependence. A proposed future spectroscopic facility could provide a complementary test, but such projections are speculative and are not part of the headline result.

* houston@hubify.com; ORCID: 0009-0008-5616-5994

II. THE MATTER-BOUNCE BISPECTRUM BENCHMARK

A. The Prediction

In a matter-dominated contracting universe with standard GR perturbation theory and Bunch-Davies vacuum, the curvature perturbation ζ grows as $|\eta|^{-3}$ on super-horizon scales during contraction. The cubic interactions, governed by the Maldacena action [2] specialized to $\epsilon = 3/2$, produce a bispectrum with shape function [7]:

$$A_T(k_1, k_2, k_3) = \frac{3}{256 k_1^2 k_2^2 k_3^2} \mathcal{K}_9(k_1, k_2, k_3), \quad (1)$$

where $\mathcal{K}_9 \equiv \mathcal{K}_9(k_1, k_2, k_3)$ is a degree-nine homogeneous momentum polynomial, deliberately distinguished from the power spectra $P_\zeta(k)$ and $P_\Phi(k)$. The configuration-dependent nonlinearity amplitude, whose squeezed limit is the local nonlinearity parameter, is:

$$B_{\text{NL}} = \frac{10}{3} \frac{A_T}{\sum_i k_i^3} \rightarrow -\frac{35}{16} \quad \text{as } k_1/k \rightarrow 0, \quad (2)$$

where $k \equiv k_2 \approx k_3$ denotes the hard-mode scale and $k_1 \ll k$ is the long mode. Here A_T is the exact vertex-sum shape. Appendix B documents the discrepancy between that sum and Cai *et al.*'s transcribed final polynomial and shows why the separately stated $-35/8$ is not reproduced. Every overlap and Fisher calculation below uses the exact vertex-sum polynomial of Eq. (3), while $-35/16$ agrees independently with Li *et al.* [8] at $c_s = 1$. Here $\sum_i k_i^3 \equiv k_1^3 + k_2^3 + k_3^3$, and B_{NL} is dimensionless by construction: $\mathcal{K}_9$ has degree nine, the prefactor of Eq. (1) removes degree six, and the $\sum_i k_i^3$ denominator removes the remaining degree three. Note that B_{NL} retains its full dependence on the coefficients ($c_1, \dots, c_6$) through $\mathcal{K}_9$ via A_T ; no cancellation of $\mathcal{K}_9$ occurs between Eqs. (1) and (2).

We evaluate the vertex-consistent shape in the ordered symmetric degree-nine basis. At first use, $\sum_{i \neq j}$ denotes the six ordered pairs and $\sum_{i \neq j \neq l}$ the six ordered all-distinct triples:

$$\begin{aligned} \mathcal{K}_9 = & 3 \sum_i k_i^9 + \sum_{i \neq j} k_i^7 k_j^2 - 9 \sum_{i \neq j} k_i^6 k_j^3 + 5 \sum_{i \neq j} k_i^5 k_j^4 \\ & - 33 \sum_{i \neq j \neq l} k_i^5 k_j^2 k_l^2 + 9 \sum_{i \neq j \neq l} k_i^4 k_j^3 k_l^2. \end{aligned} \quad (3)$$

The coefficient vector is therefore uniquely

$$(c_1, c_2, c_3, c_4, c_5, c_6) = (3, 1, -9, 5, -33, 9). \quad (4)$$

This is a re-expansion of the exact sum of all four cubic vertices, not a fit to three benchmark configurations. In particular, the ordered (5, 2, 2) orbit has six index permutations but only three distinct monomials because the two exponent-2 factors are identical. Each distinct monomial appears twice, so the ordered-basis coefficient -33

expands to coefficient -66 on each distinct (5, 2, 2) monomial. The earlier three-benchmark fit mistook this book-keeping redundancy for a physical coefficient uncertainty; no such uncertainty is propagated here.

On the fixed 23,098-triangle ratio grid, the exact shape gives amplitude recovery $r = 0.83542294$ and shape cosine $r_{\text{cos}} = 0.98167825$ against the local template. The exact benchmark values are $B_{\text{NL}} = -35/16$ (squeezed), $-255/128$ (equilateral), and $-9/8$ (folded). The adopted noise-weighted recast convention $r = 0.84 \pm 0.02$ is unchanged because the exact flat-grid value lies inside that weighting-scheme interval; the interval describes estimator-weight variation, not uncertainty in the polynomial coefficients. The local/equilateral/orthogonal decomposition and the independent C13–C15 Fisher checks are regenerated from Eq. (4).




FIG. 1. Exact four-vertex matter-bounce shape $B_{\text{NL}}(k_L, k_S, k_S)$ as a function of squeeze ratio. The curve is evaluated directly from Eq. (3) and approaches $-35/16$; the marker identifies the equilateral value $-255/128$.

Configuration	Exact B_{NL}
Squeezed ($k_1 \rightarrow 0$)	$-35/16 = -2.1875$
Equilateral ($k_1 = k_2 = k_3$)	$-255/128 = -1.9922$
Folded ($k_1=2k_2=2k_3$) ^a	$-9/8 = -1.125$

TABLE I. Confirmation of the matter-bounce shape function at three benchmark momentum configurations, quoted at the corrected amplitude $f_{\text{NL}}^{\text{local}} = -35/16$ (Appendix B): the values are exactly one-half the amplitudes published in Cai *et al.* [7] ($-35/8, -255/64, -9/4$), whose transcribed printed polynomial carries a $-(99/128) \sum_i k_i^3$ discrepancy relative to the exact vertex sum (Appendix B), and agree with the general- c_s result of Li *et al.* [8] at $c_s = 1$. ^aThe folded row sits on the degenerate boundary $k_1 = k_2 + k_3$ and is evaluated as the limit of the sequence $k_1 = 2k, k_2 = k_3 = k$.

B. Contraction-Phase Coefficient and the Cubic-Transfer Gate

The exact $-35/16$ coefficient is a contraction-phase result under: (a) matter-dominated contraction ($w \approx 0$, $\epsilon \approx 3/2$), (b) the stated standard-GR cubic action during contraction, and (c) Bunch–Davies vacuum initial conditions. It is not asserted to be independent of a UV completion. Relating that coefficient to a late-time observable additionally requires faithful third-order transfer through a specified nonsingular transition (assumption (d), Sec. II C); only linear transfer is verified for the cited Wilson–Ewing construction [12]. Thus the algebraic coefficient and the cubic-transfer gate are separate claims: the first is certified here, while the second remains a completion-dependent condition.

C. Assumptions

The contraction-phase amplitude and the observational recast require distinct assumptions. The amplitude calculation assumes: (a) matter domination, $w = 0$ and $\epsilon = 3/2$; (b) standard GR perturbation theory during contraction; and (c) Bunch–Davies initial conditions. The recast additionally assumes: (d) faithful transmission of the bispectrum through the nonsingular transition; (e) no prolonged post-bounce inflation that erases the signal; and (f) fermion-sourced torsion negligible during contraction and the bounce—now bounded explicitly below rather than assumed.

Assumption (d) is the load-bearing limitation. Wilson–Ewing [12] establishes linear propagation for the specified effective LQC construction. Single-clock separate-universe reasoning makes small nonlinear corrections plausible for deeply superhorizon modes, but it does not substitute for an explicit third-order evolution through the bounce, particularly across quantization schemes with different effective sound speeds. We now compute the linear transmission directly in one concrete scheme. In the *dressed-metric* geometric prescription ($\tilde{z} \propto a$, bounded a''/a , $c_s = 1$; the Agullo–Ashtekar–Nelson quantum-mass term is *not* included, see below), the Weinberg constant (conserved) curvature mode continues smoothly—and, for the symmetric bounce, doubly parity-protected—through the bounded transition with unit transmission, $T_c(k) = 1$ identically: the measured departure is $|T_c - 1| = 2.7 \times 10^{-10}$ to 2.2×10^{-5} across $k\eta_B = 0.002$ – 0.03 (with η_B the conformal bounce timescale; the growth toward large k is k^6 gradient-basis truncation, not physics), so the scheme-specific transmission correction to the $-35/16$ amplitude is $|\delta f_{\text{NL}}| \leq 6.8 \times 10^{-8}$ at the observable scale $k\eta_B = 10^{-2}$ —more than four orders of magnitude below the earlier order-of-magnitude residual (`research/cubic_bounce_transmission/g1_dressedmetric_ic_close.py`, `research/cubic_bounce_transmission/g1_dressedmetric_ic_close.json`). This is explicitly

a scheme-specific statement, not a scheme-independent theorem: in the effective-fluid variable ($z \propto a/c_s$) the corresponding gradient-expansion basis integral $K = \int z^2 d\eta$ diverges as $d_{\text{cut}}^{-1/2}$ at the $H = 0$ pole (fitted exponent -0.4998 against the analytic $-1/2$), so no finite transmission coefficient exists in that scheme—the concrete mechanism behind the model’s absence of a scheme-independent transfer coefficient. We therefore still do not claim a closed cubic-transfer theorem valid across all schemes: the exact Agullo–Ashtekar–Nelson quantum-mass term $U(\eta)$ for a quasi-dust (rather than scalar-field) background is not implemented (any U even about the symmetric bounce cannot alter the parity result, so only an odd-parity or asymmetric completion could contribute), the deformed-algebra scheme carries a different subleading structure, and a fully nonlinear third-order branch calculation remains open. Every late-time sensitivity in this paper accordingly remains conditional on faithful cubic transmission.

Quasi-dust corrections are also kept separate from the exact matter-limit polynomial. Existing scaling arguments suggest percent-level changes, but their coefficient requires evaluating the four cubic integrals with the quasi-dust mode functions. Likewise, assumption (f) is automatic only in the scalar-only sector; for models with appreciable fermion populations the Einstein–Cartan four-fermion operator is estimated explicitly in the next paragraph. These limitations affect model applicability, not the algebraic uniqueness of Eq. (4).

Assumption (f) is now supported by an explicit order-of-magnitude parametric estimate. Integrating out the algebraic Einstein–Cartan torsion in the presence of fermions leaves an axial four-fermion contact term $-(3\kappa/16)[\gamma^2/(1+\gamma^2)](\bar{\psi}\gamma^\mu\gamma^5\psi)^2$, with $\kappa = 8\pi G$ and γ the Barbero–Immirzi parameter; its coefficient, finite-Holst factor, and dimensional benchmark follow the standard Einstein–Cartan–Holst result [13] as convention-audited in the companion paper [14] (reproduced here to 0.1%). Estimating the torsion-sourced correction by its fractional energy density (an asserted energy-density transfer proxy, $|\delta f_{\text{NL}}^{\text{tor}}| \lesssim |f_{\text{NL}}| \rho_{\text{tor}}/\rho$, not a derived in-in propagation of the four-fermion operator) $\rho_{\text{tor}}/\rho = (3/16)[\gamma^2/(1+\gamma^2)]\kappa\langle J_5^2 \rangle/\rho$ with the assumed spin-coherent proxy $\langle J_5^2 \rangle \lesssim n_\psi^2$ (thermal or incoherent spin distributions only reduce this), and using the committed quasi-dust dilution $n_\psi \propto a^{-3} \propto \rho$ (which makes the fraction maximal at the bounce), gives

$$|\delta f_{\text{NL}}^{\text{tor}}| \lesssim \frac{35}{16} \cdot \frac{3}{16} \frac{\gamma^2}{1+\gamma^2} \frac{\kappa n_{\psi,c}^2}{\rho_c}, \quad (5)$$

with $n_{\psi,c}$ the (free) bounce-epoch fermion number density and ρ_c the critical density; $n_{\psi,c}$ is not fixed by any committed artifact and is carried symbolically. Writing $x_\psi \equiv \kappa n_{\psi,c}^2/\rho_c$, the effective bounce description holds only for $x_\psi < 1$, and the estimate is largest at the $x_\psi \rightarrow 1$ edge of that domain, where it reaches the prefactor $(35/16)(3/16)\gamma^2/(1+\gamma^2)$, i.e. from 0.022 ($\gamma = 0.2375$)

to 0.21 ($\gamma = 1$): within this parametric estimate, torsion contributes at most ~ 1 –10% of the $-35/16$ amplitude wherever the effective bounce description is valid. The correction reaches the earlier $\delta f_{\text{NL}} \sim 10^{-3}$ order-of-magnitude reference level only for $n_{\psi,c} \sim 3 \times 10^{-2} M_{\text{Pl}}^3$ and becomes comparable to the full $-35/16$ amplitude only at near-Planckian $n_{\psi,c} \sim M_{\text{Pl}}^3$, where the semiclassical bounce itself breaks down; for any sub-Planckian abundance $n_{\psi,c} \ll M_{\text{Pl}}^3$ —as required for a semiclassical bounce—one has $|\delta f_{\text{NL}}^{\text{tor}}| \ll 10^{-3}$ and assumption (f) holds. This supports assumption (f) with an explicit parametric estimate in place of a bare assertion, with $n_{\psi,c}$ the explicit model parameter (`research/cubic_bounce_transmission/g3_torsion_fourfermion_bound.py`).

D. The Viable Model

The Wilson-Ewing Λ CDM quasi-dust model [12] supplies one concrete setting for the conditional recast: $w = -0.003$ fits $n_s \simeq 0.964$, and the construction reports strong tensor suppression. The exact contraction-phase cubic amplitude is $f_{\text{NL}} = -35/16$ at the matter limit, with quasi-dust corrections and nonlinear bounce transmission kept as separate theoretical systematics. We do not elevate this package to a complete fit or a model-independent forecast.

III. OBSERVABLE MAPPING TO LARGE-SCALE STRUCTURE

A. Scale-Dependent Bias

The primordial and late-time conventions are connected explicitly. We define $\zeta = \zeta_g + (3/5)f_{\text{NL}}(\zeta_g^2 - \langle \zeta_g^2 \rangle)$, hence $B_\zeta^{\text{loc}} = (6/5)f_{\text{NL}}[P_\zeta P_\zeta + \text{cyc.}]$. During matter domination $\Phi = (3/5)\zeta$, so $P_\Phi = (3/5)^2 P_\zeta$ and $B_\Phi^{\text{loc}} = 2f_{\text{NL}}[P_\Phi P_\Phi + \text{cyc.}]$. The same dimensionless f_{NL} therefore enters the large-scale-structure kernel below; no additional factor of $3/5$ is applied.

Primordial local non-Gaussianity induces a scale-dependent correction to galaxy bias [11, 15]:

$$\Delta b(k, z) = \frac{f_{\text{NL}} b_\phi}{\mathcal{M}(k, z)}, \quad (6)$$

$$b_\phi^{\text{UMF}} = 2 \delta_c (b_1 - 1). \quad (7)$$

The second line is the universal-mass-function specialization used by the numerical surrogate; substituting it in the first line recovers $\Delta b = 2f_{\text{NL}}(b_1 - 1)\delta_c/\mathcal{M}$. More general tracer responses can instead keep b_ϕ free, as in the final row of Table III. The Poisson–Newtonian transfer kernel is

$$\mathcal{M}(k, z) = \frac{2k^2 T(k) D(z)}{3\Omega_m H_0^2}, \quad (8)$$

where $T(k)$ is the matter transfer function (normalized to $T(k) \rightarrow 1$ as $k \rightarrow 0$), $D(z)$ is the linear growth factor (normalized to $D(0) = 1$), wavenumbers k are comoving and quoted in $h \text{Mpc}^{-1}$ throughout (we work in units $c = 1$ and likewise express H_0 in $h \text{Mpc}^{-1}$, so that $\mathcal{M}$ is dimensionless and Δb carries no residual units), $\delta_c \approx 1.686$ is the spherical-collapse threshold, and b_1 is the linear Eulerian galaxy bias. Since $\Delta b \propto 1/\mathcal{M}$ and $\mathcal{M} \propto k^2$ on ultra-large scales (where $T(k) \rightarrow 1$), the signal grows as $\Delta b \propto 1/k^2$ as $k \rightarrow 0$ [11, 15]; this $1/k^2$ enhancement on the largest scales is what makes scale-dependent bias the most sensitive single channel for local-type f_{NL} . All downstream Fisher weightings, plots, and forecasts that invoke the SDB kernel in this paper use Eqs. (6)–(8) as the canonical definition.

B. Template Projection and Amplitude Recovery

The exact matter-bounce shape is not purely local: B_{NL} changes from $-35/16$ in the squeezed limit to $-9/8$ at the folded boundary. A local-template estimator therefore recovers a weighted fraction r of a bounce-shaped signal:

$$f_{\text{NL}}^{\text{measured}} = r f_{\text{NL}}^{\text{bounce}}, \quad \sigma(f_{\text{NL}}^{\text{bounce}}) = \sigma(f_{\text{NL}}^{\text{local}})/r. \quad (9)$$

Direct evaluation of the unique exact shape gives flat-grid $r = 0.83542294$ and $r_{\text{cos}} = 0.98167825$. The adopted $r = 0.84 \pm 0.02$ is a weighting-scheme envelope around this exact-shape value. It is not a coefficient posterior or polynomial uncertainty. The independent survey-weighted Fisher gives r_{eff} as a benchmark statistic under its adopted covariance; no value from one metric is silently substituted into another.

Standard-basis template decomposition. Projecting the exact physical bispectrum onto the local, equilateral, and orthogonal templates [16] gives uniform-weight local cosine 0.9817. A joint three-template projection increases the recovered Fisher-norm fraction from 0.9637 to 0.9680, corresponding to only $\delta r = 0.0022$. The large raw orthogonal cosine reflects template collinearity and is not independent signal. This is a geometry-only check; the true SPHEREx estimator variance still requires its survey covariance. Across the stated survey-motivated weighting choices, the exact shape is summarized by

$$r = 0.84 \pm 0.02, \quad (10)$$

with the flat-grid exact value 0.8354 and the previously defined signal-only endpoint 0.876. We use 0.84 for the recast and keep the ± 0.02 spread as a weighting-choice sensitivity only. A shape-matched estimator could approach unit recovery, but constructing that estimator with the actual SPHEREx covariance is outside this recast. For the published SPHEREx bispectrum baseline $\sigma(f_{\text{NL}}^{\text{local}}) = 0.7$, direct substitution gives 2.63σ with the adopted $r = 0.84$ template map, reduced from the naive ratio 3.13σ . This is an illustrative arithmetic mapping

before additional nuisance marginalization, not an observational headline. The channel-native ladder in Sec. VII reports those nuisance assumptions explicitly, including the 0.4σ free- b_ϕ surrogate limit.

C. Galaxy Bispectrum

The galaxy bispectrum provides an independent measurement channel that accesses information at shorter wavelengths, reducing the dependence on ultra-large-scale modes [1]. This makes bispectrum-based constraints more robust to large-scale systematics than power-spectrum-based scale-dependent bias alone.

IV. ILLUSTRATIVE SPHEREX MAPPING

We use the published Heinrich *et al.* [1] multi-tracer bispectrum baseline $\sigma(f_{\text{NL}}^{\text{local}}) \simeq 0.7$ only for an illustrative mapping to the exact matter-bounce template. With $r = 0.8354$, direct substitution gives $|f_{\text{NL}}|r/0.7 = 2.61\sigma$; recomputing with the adopted $r = 0.84$ convention gives 2.63σ (the rounding is in r : $0.8354 \rightarrow 0.84$, not a rounding of 2.61σ). This conditional arithmetic map is not a paper headline, an independently reproduced SPHEREx likelihood, or a detection forecast.

Two in-house Fisher calculations test the mapping under a leading-order Gaussian multi-tracer covariance. In real space, the local-template uncertainty is 0.626 with bias fixed and 0.687 after the implemented bias marginalization; the exact bounce-template values are 0.631 and 0.688 , with survey-weighted Fisher recovery $r_{\text{eff}} = 0.9929$ and 0.9986 , respectively (a covariance-weighted information cosine, *not* the bare σ ratio, e.g. $0.626/0.631 = 0.992$). These correspond to 3.5σ and 3.2σ before the additional nuisance axes discussed below. The survey weighting is squeezed dominated, explaining why this recovery ratio is closer to unity than the flat-grid amplitude recovery $r = 0.8354$; the two quantities use different inner products and are not interchangeable.

A redshift-space extension numerically integrates the line-of-sight orientation and includes the tree-level Kaiser and Z_2 kernels. It gives bounce-template uncertainties 0.417 (bias fixed) and 0.449 (bias marginalized), with full-signal recovery ratios $r_{\text{eff}} = 0.9953$ and 0.9991 . A separate primordial-transfer-only run gives $r_{\text{eff}} = 0.9981$, which is reported separately because the full result includes the standard density-bispectrum contribution. A targeted orientation-grid convergence test changes marginalized uncertainties by at most 1.4×10^{-6} fractionally between 8×8 and 12×12 quadrature. In that reduced-grid test the primordial-only marginalized response ratio is 1.0037 ; because marginalization changes the two template responses differently, this statistic is not a bounded Fisher cosine and is not interpreted as an overlap exceeding unity. The approximately 34.7% bias-marginalized gain over the real-space calculation

($0.688 \rightarrow 0.449$, the bounce-template bias-marginalized baseline) is materially larger than the roughly 18% gain quoted for an already-redshift-space monopole-to-multipole comparison; differences in baseline, nuisance set, and signal model prevent treating that agreement as validation. Fingers-of-God damping and non-Gaussian covariance are omitted, so the absolute redshift-space significance is optimistic. The recovery ratio is the more stable diagnostic, but it is not assigned an uncertainty band.

Finally, a channel-native $(f_{\text{NL}}, b_\phi, A_{\text{GR}})$ surrogate Fisher makes the nuisance dependence explicit. Here b_ϕ is the coefficient of the local-PNG scale-dependent-bias response and A_{GR} scales the modeled relativistic-projection contribution. At $f_{\text{NL}} = -35/16$ it gives 3.5σ with both nuisances fixed, 3.1σ after marginalizing A_{GR} , 2.3σ with A_{GR} marginalized and a declared 30% Gaussian theory prior on b_ϕ (an illustrative benchmark bracketing typical assembly-bias departures from the universal-mass-function b_ϕ , not a first-principles bound), and 0.4σ when b_ϕ is free. The exact values remain in Table III and its generating artifact. The 30% result is conditional on that prior; the free- b_ϕ result is the unconstrained limit of this surrogate model. Neither replaces the missing external per-triangle covariance, and no endpoint is presented as a guaranteed floor.

V. SPECULATIVE FUTURE-SURVEY OUTLOOK

A proposed high-redshift spectroscopic facility such as MegaMapper [17] could eventually test the same shape with different tracer and redshift leverage. We do not transfer the SPHEREx systematic budget or quote a headline significance for an unapproved design: the high-redshift GR projection, tracer selection, and covariance require a dedicated calculation. This is an outlook only and is not used in the abstract, conclusion, or evidence claims.

VI. INFLATIONARY ALTERNATIVES

A. Can Inflation Reproduce the Signal?

Standard single-field slow-roll inflation gives $f_{\text{NL}} = (5/12)(1 - n_s) \approx +0.015$ [2] in the standard global primordial-template convention. It is not itself an independent on-sky scale-dependent-bias signal: in a local-observer or conformal-Fermi description the leading single-clock consistency-relation contribution is absorbed by the coordinate/projection treatment, leaving parametrically smaller slow-roll and projection residuals [3, 4]. We retain $|f_{\text{NL}}^{\text{bounce}}|/0.015 = 2.1875/0.015 \approx 146$ only as a conventional template-normalization benchmark, not as a survey prediction; after the observable

projection treatment the bounce-vs-inflation contrast remains qualitatively $|f_{\text{NL}}^{\text{bounce}}| \gg |f_{\text{NL}}^{\text{inf}}|$ and opposite in sign. Non-canonical single-field models (DBI, etc.) produce equilateral-shape f_{NL} , not local. Table II makes the convention distinction explicit.

TABLE II. Expected squeezed-limit local f_{NL} for the quasi-dust matter bounce and single-field slow-roll inflation under two comparison conventions. The global primordial-template row is a like-for-like theory-normalization convention; it does not assert that a survey measures the single-clock consistency-relation term as an independent scale-dependent-bias signal. The local-observer, projection-aware row [3, 4] removes or absorbs the leading single-clock coordinate contribution. This distinction primarily changes the inflation entry and does not alter the conditional bounce recast.

convention	$f_{\text{NL}}^{\text{bounce}}$	$f_{\text{NL}}^{\text{inf}}$	$ f_{\text{NL}}^{\text{bounce}} / f_{\text{NL}}^{\text{inf}} $
primordial template	-2.1875	+0.015	≈ 146
local observer	-2.1875	$\rightarrow 0^+$	$\gg 146$

Non-attractor single-field inflation gives $f_{\text{NL}} = +5/2$ in its standard limit [18], while curvaton and curved-field-space constructions can access negative local amplitudes with additional parameters. The amplitude alone therefore does not uniquely identify a bounce; distinguishing these classes requires their full shape, scale dependence, nuisance structure, and physically derived priors.

To keep estimator coordinates explicit, Fig. 2 displays the imported Heinrich $\sigma(f_{\text{NL}}^{\text{local}}) = 0.7$ baseline in bounce-amplitude coordinates using Eq. (9): the bar is centered at $f_{\text{NL}}^{\text{bounce}} = -35/16$ with half-width $\sigma(f_{\text{NL}}^{\text{bounce}}) = 0.7/0.84 = 0.83$. The separate shape-matched surrogate result $\sigma(f_{\text{NL}}^{\text{bounce}}) \simeq 0.688$ is not the blue bar.




FIG. 2. f_{NL} landscape: matter bounce vs. inflationary alternatives. The blue bar is the imported SPHEREx local-template uncertainty transformed to bounce-amplitude coordinates by Eq. (9): $\sigma(f_{\text{NL}}^{\text{bounce}}) = 0.7/0.84 = 0.83$, centered on $-35/16$. It is not the separate shape-matched surrogate-covariance uncertainty.

B. The Kinematic vs. Parametric Asymmetry

The bounce predicts $f_{\text{NL}} \approx -35/16$ *kinematically* in the exact matter limit. The quasi-dust correction is presently unquantified because its coefficient requires evaluating all four cubic integrals with quasi-dust mode functions (Sec. II C). Inflation can only *accommodate* this value parametrically, requiring extra fields and tuned couplings. This asymmetry—one tightly determined matter-limit prediction versus multiple tunable parameters—is qualitative and does not by itself establish a Bayesian preference.

VII. SYSTEMATICS AND ROBUSTNESS

The illustrative observational calculation has two deliberately separate layers. The published-baseline mapping uses $\sigma(f_{\text{NL}}^{\text{local}}) \simeq 0.7$ and the exact-shape recovery $r = 0.8354$, giving 2.61σ , or 2.63σ when the adopted $r = 0.84$ convention is used (the rounding is in r : $0.8354 \rightarrow 0.84$). It does not contain a new covariance. The in-house channel-native calculation instead uses a leading-order Gaussian multi-tracer covariance and is a sensitivity test of nuisance directions, not a replacement for the published likelihood.

The channel-native ladder is summarized in Table III. The corresponding exact-shape uncertainties are 0.631, 0.697, 0.941, and 5.173. The native correlations are $\rho(f_{\text{NL}}, A_{\text{GR}}) = -0.4264$ in the two-parameter block, -0.5025 in the three-parameter block, and $\rho(f_{\text{NL}}, b_\phi) = +0.9909$. The 30% row is conditional on a theory prior; it is neither an empirical calibration nor a lower bound. The free- b_ϕ row is the genuinely unconstrained limit of this surrogate model and shows that the result is prior sensitive.

The dominant unmodeled effects are the unavailable external per-triangle covariance, non-Gaussian covariance terms, higher-order bias marginalization, photometric-redshift outliers, relativistic light-cone terms beyond the amplitude surrogate, and fingers-of-God damping in the redshift-space check. Omitting fingers-of-God damping makes the absolute C14 information gain optimistic. Common modeling errors can cancel in a local-to-bounce recovery ratio, but this cancellation is not guaranteed and no uncertainty band is invented for r_{eff} .

VIII. CURRENT DATA AND CONSISTENCY RELATION

Current CMB and large-scale-structure constraints are not yet decisive. Planck PR4/NPIPE reports $f_{\text{NL}} = -0.1 \pm 5.0$ [19]; using the CMB-weighted recovery $r = 0.876$ gives $f_{\text{NL}}^{\text{bounce}} = -0.11 \pm 5.71$, only 0.36σ from the predicted $-35/16$. DESI DR1 local-PNG constraints have uncertainties near 9–10 [20, 21]. These data are

TABLE III. Channel-native surrogate-covariance nuisance ladder. The 30% row is explicitly theory-prior conditional (an illustrative benchmark, not a first-principles bound); the free- b_ϕ row is the unconstrained limit of this surrogate.

Nuisance treatment	Significance at $f_{\text{NL}} = -35/16$
b_ϕ and A_{GR} fixed	3.47σ
A_{GR} marginalized	3.14σ
A_{GR} marginalized; 30% Gaussian b_ϕ theory prior	2.32σ
A_{GR} marginalized; b_ϕ free	0.42σ

consistent with both the matter-bounce amplitude and zero.

The Wilson–Ewing quasi-dust construction relates the measured tilt to a small departure from matter contraction. That departure can also perturb the cubic amplitude, but the required four-vertex calculation with quasi-dust mode functions has not been completed here. We therefore do not assign a derived correction coefficient or propagate an invented joint (n_s, f_{NL}) covariance. The exact $w = 0$ result remains the benchmark, and departures from it are a named theory uncertainty.

IX. DISCUSSION

The strongest result of this work is algebraic: the sum of the four cubic vertices has the unique ordered-basis coefficients $(3, 1, -9, 5, -33, 9)$ and yields $f_{\text{NL}} = -35/16$. The ordered $(5, 2, 2)$ orbit contains six permutations but only three distinct monomials, so each distinct monomial is counted twice. This resolves the earlier artificial coefficient freedom without assigning a statistical uncertainty to an exact basis conversion.

The observational material is an illustration, not a second headline result. Mapping the published SPHEREx local-template sensitivity through the exact-shape recovery gives the arithmetic value 2.63σ before additional nuisance marginalization. The channel-native surrogate shows why this must remain conditional: the significance ranges from 3.5σ with nuisances fixed to 0.4σ with b_ϕ free, and the intermediate 2.3σ value requires a declared 30% theory prior. An eventual analysis must use the survey’s full per-triangle covariance and a justified b_ϕ calibration.

Faithful cubic transmission through the bounce is the leading theoretical condition. Linear propagation is established for the specified construction [12], and the linear transmission of the conserved mode is now computed to be unity in the dressed-metric scheme ($|\delta f_{\text{NL}}| \leq 6.8 \times 10^{-8}$ at $k\eta_B = 10^{-2}$; Sec. II C), as a scheme-specific rather than scheme-independent statement; single-clock conservation motivates the nonlinear extension but does not replace an explicit third-order calculation through the transition. Consequently a future null measurement would test the combination of the contraction-phase amplitude and the stated transmission assumptions, not every nonsingular cosmology.

A proposed future high-redshift facility may offer com-

plementary leverage, but no such projection enters the headline, conclusion, or readiness assessment. A nonzero negative local amplitude would also not uniquely identify a bounce because multifield inflation can generate similar amplitudes. Shape information, scale dependence, nuisance control, and model-derived priors remain necessary for model discrimination.

X. CONCLUSION

The contraction-phase calculation predicts $f_{\text{NL}}^{\text{local}} = -35/16$. Re-summing the four cubic vertices and expanding the result in the ordered symmetric basis gives the unique coefficients $(3, 1, -9, 5, -33, 9)$, with the $(5, 2, 2)$ permutation multiplicity treated explicitly. The result agrees with the independent general- c_s expression at $c_s = 1$ and replaces the unreproduced printed $-35/8$ value.

Conditional on faithful cubic transmission through the chosen nonsingular completion, direct substitution into the published SPHEREx baseline gives an illustrative 2.63σ arithmetic map before additional nuisance marginalization. The channel-native surrogate ranges from 3.5σ with nuisances fixed to 0.4σ with b_ϕ free; its 2.3σ intermediate result requires an explicit 30% b_ϕ theory prior. None of these values is an observational headline, a unified forecast, or a guaranteed discovery floor. The missing external covariance, higher-order nuisance treatment, and direct third-order bounce evolution are the principal work required before a survey-level claim.

DATA AND CODE AVAILABILITY

Analysis code and numerical artifacts are available in the public project repository at <https://github.com/Hubify-Projects/bigbounce>. An immutable versioned archival deposit of the reviewed release—the exact manuscript PDF, the L^AT_EX source, the arXiv source bundle, a tracked-provenance archive, and SHA256SUMS—is published on Zenodo under the minted version DOI [10.5281/zenodo.21461881](https://doi.org/10.5281/zenodo.21461881); the concept DOI [10.5281/zenodo.21461880](https://doi.org/10.5281/zenodo.21461880) resolves to the latest deposited version. That deposit archives the reviewed v1.7.125 release PDF and source (the exact bytes reviewed at that release); the present manuscript is v1.7.130, 5 patch re-

leases ahead, and will be added to the same Zenodo record as a new version on the next re-stage. The exact vertex certification is implemented in `scripts/p2_vertex_check.py`, the exact shape and overlap artifact is generated by `scripts/exact_shape_analysis.py`, the real- and redshift-space Fisher checks are implemented in `scripts/c13_independent_bounce_fisher.py` and `scripts/c14_rsd_multipole_fisher.py`, and the native nuisance ladder is generated by `scripts/c15_channel_native_fisher.py`. The orientation convergence artifact is stored with the other JSON outputs. The retired benchmark-fit coefficient scan is not used or distributed as an uncertainty model.

The per-triangle covariance of Heinrich *et al.* [1] is external to this project and unavailable here; only its published scalar $\sigma(f_{\text{NL}}^{\text{local}}) \simeq 0.7$ is used in the illustrative mapping. No new observational data are introduced.

AI USAGE DISCLOSURE

This manuscript and the analysis behind it were produced with substantial and substantive AI assistance, and the disclosure below is deliberately specific rather than pro forma. The work was carried out with an agentic AI research pipeline operated under the author’s direction. Anthropic Claude (Opus 4 family, 2026 releases) was used for pipeline orchestration, generation of the analysis and certification scripts named in the Data and Code Availability section, symbolic manipulation with SYMPY, literature triage, and drafting and revision of the text. xAI Grok (Grok-4 and Grok-4.3), Google Gemini (2.5 Pro and 3.1 Pro), and OpenAI ChatGPT (GPT-5-class models, in review rounds through 2026 July) were used as independent cross-checking and adversarial referee-style reviewers of the compiled manuscript, each bound to the exact PDF under review; their reports and the resulting dispositions are retained in the public project repository.

What remained with the author is the scientific content. The author set the scope and the conventions, chose the vertex basis and the squeezed and equilateral limits, decided which results are reported as exact and which are labelled illustrative or conditional, and made every judgment about what this paper does and does not claim. Every analytic result reported here is fixed by exact-fraction symbolic computation in committed scripts—principally `scripts/p2_vertex_check.py`—whose source, inputs, and printed output the author read and verified line by line, and which any reader can rerun; the four-vertex sum $f_{\text{NL}}^{\text{local}} = -35/16$, its per-vertex decomposition in Table V, and its agreement with the independent general- c_s expression at $c_s = 1$ are each certified this way rather than by an AI assertion. Findings raised by the AI reviewers were adjudicated individually against the source and the artifacts, not adopted on the reviewer’s authority. The author takes full responsibility for the contents, including all material produced with AI assistance. No AI system

is an author, and none is credited as one.

Appendix A: Prior-Volume Reproducibility Note

For completeness, the committed reproducibility scans include a one-dimensional Gaussian summary likelihood under several bounce and multifield amplitude priors. This is not a forecast or model-selection result: in the broad-competitor limit $\text{BF} \propto W/\sqrt{\sigma_{\text{eff}}^2 + \sigma_{\text{theory}}^2}$, so the evidence ratio changes directly with the chosen competitor width W . A defensible comparison would require a joint likelihood and physically derived priors for every competing class; the scans are therefore retained only as provenance artifacts.

Appendix B: Resolution of the Cai–Li Factor of Two: The Correct Matter-Bounce Local Amplitude is $-35/16$

The historical factor-of-two discrepancy between $f_{\text{NL}} = -35/8$ (Cai *et al.* [7]) and $f_{\text{NL}} = -35/16$ (Li *et al.* [8], evaluated at $c_s = 1$) is resolved at the level needed here by an exact re-summation of Cai *et al.*’s own four cubic vertices. That sum gives $-35/16$. The transcribed final polynomial does not reproduce either the vertex sum or Cai *et al.*’s separately stated $-35/8$, so we identify the polynomial discrepancy but do not claim a complete reconstruction of how the published $-35/8$ arose.

Trusted-expression provenance. The adopted chain is deliberately bounded: Cai *et al.*’s four source-level vertex expressions and their ϵ -order-grouped intermediates are trusted inputs; their exact symbolic sum fixes Eqs. (3)–(4) and $f_{\text{NL}} = -35/16$; Li *et al.*’s independent Eq. (5.1) agrees at $c_s = 1$. Cai’s Eq. (37) and Li’s Eq. (4.19) share the printed polynomial whose $-305/64$ squeezed reduction conflicts with that chain. We disclose that tension, do not use the printed polynomial downstream, and do not claim to have reconstructed how the separately stated $-35/8$ arose.

The decisive computation. Cai *et al.*’s cubic action (Maldacena form, $\epsilon = 3/2$) gives four vertex contributions to the shape function A : the field-redefinition term, the $\zeta\dot{\zeta}^2$ term, the $\dot{\zeta}\partial\zeta\partial\chi$ term, and the $\zeta(\partial_i\partial_j\chi)^2$ term (their own expressions, χ defined by $\partial^2\chi = a\epsilon\dot{\zeta}$). Summing these four exactly and taking the squeezed limit $k_1 \ll k_2 = k_3 = k$ gives

$$f_{\text{NL}}(k_1 \ll k_2 = k_3) = -\frac{35}{16} + \frac{35}{64} \frac{k_1^2}{k^2} + \cdots \xrightarrow{k_1 \rightarrow 0} -\frac{35}{16}, \quad (\text{B1})$$

a clean $k_1 \rightarrow 0$ limit with no cancellation-sensitive subleading ambiguity (equilateral: $f_{\text{NL}}^{\text{equil}} = -255/128$). Two independent cross-checks fix this value. (1) Cai *et al.*’s own ϵ -order-grouped intermediate expressions (their $A^\epsilon, A^{\epsilon^2}, A^{\epsilon^3}$) sum to the per-vertex result and give

$-35/16$. (2) Li *et al.*'s Eq. (5.1), $f_{\text{NL}}^{\text{local}} = -165/16 + 65/(8c_s^2)$, gives $-35/16$ at $c_s = 1$ by an independent general- c_s in-in method. By contrast, Cai *et al.*'s transcribed final polynomial A_T (their Eq. 37) does not equal the sum of their own vertices:

$$A_T^{\text{printed}} - \sum(\text{vertices}) = -\frac{99}{128} \sum_i k_i^3 \neq 0, \quad (\text{B2})$$

i.e. a spurious $-(99/128) \sum_i k_i^3$ local-shaped term is present in the transcribed printed polynomial A_T relative to the exact sum of Cai *et al.*'s own four vertices (the sign is fixed by the committed symbolic check `scripts/cailli_certification/cai_vertices.py`, which evaluates $\sum(\text{vertices}) - A_T^{\text{printed}} = +(99/128) \sum_i k_i^3$). Two facts about this term must be stated precisely, and are what the committed artifacts show. *First*, under the six-Wick-permutation convention used consistently in Cai *et al.*'s vertex equations and Eq. (B8), the transcribed printed polynomial A_T does *not* itself squeezed-reduce to Cai's stated $-35/8$: it reduces to $-305/64 = -4.766$ (`scripts/cailli_certification/cai_conv.py`, `scripts/cailli_certification/cai_shape.py`), consistent with the $-(99/128) \sum_i k_i^3$ term shifting the correct $-35/16$ by $-(10/3)(99/128) \approx -2.58$ to $-305/64$. Treating the repeated (5,2,2) orbit as three distinct monomials while retaining six (4,3,2) permutations is not this source convention: it drops a factor of two from only one orbit and breaks the equal forms of Cai *et al.*'s vertex expressions. *Second*, therefore, this single term is *one identified discrepancy* between the transcribed printed polynomial and the exact vertex sum — *not* by itself the full mechanism of Cai's separately-published $-35/8$ (which the transcribed printed coefficients do not reproduce). The certified result is the vertex-sum squeezed limit $-35/16$ (Eq. B1), which is unaffected by either issue: it is fixed independently by the exact four-vertex sum, by Cai's own ϵ -order-grouped intermediates (Eq. B5), and by Li *et al.*'s general- c_s formula. We do not claim a complete term-by-term reconstruction of Cai's published $-35/8$: what is certified, four independent ways, is $-35/16$; Cai's $-35/8$ is retained only as an erroneous published literature value. For clarity: Li *et al.*'s printed total polynomial A_{tot} (their Eq. 4.19) agrees with Cai *et al.*'s transcribed polynomial A_T (their Eq. 37) coefficient-for-coefficient at $c_s = 1$, so the two transcribed polynomials share the same $(-305/64)$ squeezed reduction; the corrected amplitude $-35/16$ comes from re-summing Cai's four cubic-vertex contributions directly (the step that bypasses the

shared printed polynomial), and Li's closed-form general- c_s formula (their Eq. 5.1) independently gives the same $-35/16$ at $c_s = 1$.

Vertex-by-vertex certification. We certify the $-35/16$ correction explicitly at the level of Cai *et al.*'s four individual cubic vertices. Reading their own shape-function contributions verbatim from the arXiv source of Ref. [7] (`matterbounceng2.tex`, arXiv:0903.0631v2, source retrieved July 14, 2026), the four source-level vertices contribute to the shape function A with their general ϵ dependence as listed in Table IV; $\epsilon = 3/2$ is substituted only when evaluating the stated matter-contraction limits. Summing the four rows of Table IV exactly (symbolic fractions) and forming $f_{\text{NL}} = (10/3)A / \sum_i k_i^3$, the squeezed limit $k_1 \ll k_2 = k_3 = k$ is

$$f_{\text{NL}}^{\text{vertex sum}}(k_1 \ll k) = -\frac{35}{16} + \frac{35}{64} \frac{k_1^2}{k^2} + \dots \xrightarrow{k_1 \rightarrow 0} -\frac{35}{16}, \quad (\text{B3})$$

identical to Eq. (B1).

Term-by-term vertex walkthrough. To make the re-summation fully auditable at referee-report granularity we tabulate the squeezed and equilateral local-amplitude contribution of *each individual* vertex of Table IV, formed by inserting that vertex's shape-function row alone into $f_{\text{NL}} = (10/3)A / \sum_i k_i^3$ and taking the two limits (squeezed $k_1 \ll k_2 = k_3 = k$, then $k_1 \rightarrow 0$; equilateral $k_1 = k_2 = k_3$). Table V lists the four values; both columns sum, exact-fraction, to the certified benchmarks $-35/16$ (squeezed) and $-255/128$ (equilateral) of Table I. Each per-vertex entry is the exact squeezed/equilateral limit of the corresponding single row of Table IV (formed by substituting that vertex's shape-function contribution alone into $f_{\text{NL}} = (10/3)A_v / \sum_i k_i^3$); the four rows are built from the same exact-fraction vertex expressions that the committed SYMPY certification `scripts/p2_vertex_check.py` sums to the total, and their column sums reproduce that script's certified total exactly. No value is fitted or rounded.

Displayed intermediate algebra. For full referee/reader verifiability we display the two intermediate objects that the squeezed reduction Eq. (B3) collapses — both transcribed verbatim from the committed symbolic-certification script (`scripts/p2_vertex_check.py`, exact-fraction SYMPY). First, the collapsed exact vertex sum. Multiplying the four Table IV rows through by the common denominator $\Pi k^2 \equiv k_1^2 k_2^2 k_3^2$, the total shape function at $\epsilon = 3/2$ is the degree-9 polynomial

$$\begin{aligned} 256 \Pi k^2 A = & 9 \sum_i k_i^9 + 3 \sum_{i \neq j} k_i^7 k_j^2 - 27 \sum_{i \neq j} k_i^6 k_j^3 + 15 \sum_{i \neq j} k_i^5 k_j^4 \\ & - 198 \sum_i^{\text{dist}} k_i^5 k_j^2 k_l^2 + 27 \sum_{i \neq j \neq l} k_i^4 k_j^3 k_l^2, \end{aligned} \quad (\text{B4})$$

where $\sum_{i \neq j}$ runs over the six ordered pairs $i \neq j$ (six dis-

tinct monomials), $\sum_{i \neq j \neq l}$ over the six all-distinct triples,

TABLE IV. The four source-level cubic-vertex contributions to Cai *et al.*'s [7] shape function A , retaining their general ϵ dependence, read from the arXiv source (`matterbounceng2.tex`, arXiv:0903.0631v2; retrieved July 14, 2026). Here Σ_{ij} denotes the six ordered pairs, Σ_{ijl} the six ordered all-distinct triples, and $\Pi k^2 \equiv k_1^2 k_2^2 k_3^2$. The matter-contraction limits stated below are evaluated only after setting $\epsilon = 3/2$. Their exact sum has squeezed limit $-35/16$ (Eq. B3); Cai's transcribed printed polynomial (their Eq. 37) differs from this sum by the spurious $-(99/128) \sum_i k_i^3$ term of Eq. (B2) and squeezed-reduces to $-305/64$.

vertex	contribution to A (general ϵ)
field redef. ($\mathcal{L}_{\text{redef}}$)	$-\frac{\epsilon}{2} \sum k^3 - \frac{\epsilon^2}{32 \Pi k^2} \{ \Sigma_{ij} k^7 k^2 + \Sigma_{ij} k^6 k^3 - 2 \Sigma_{ij} k^5 k^4 - 2 \Sigma_{ijl} k^5 k^2 k^2 - \Sigma_{ijl} k^4 k^3 k^2 \}$
$\mathcal{L}_{\zeta \zeta^2}$	$\left(-\frac{\epsilon^2}{12} + \frac{\epsilon^3}{24} \right) \sum k^3$
$\mathcal{L}_{\zeta \partial \zeta \partial \chi}$	$\frac{\epsilon^2}{24 \Pi k^2} \{ 2 \Sigma_{ij} k^7 k^2 - 2 \Sigma_{ij} k^5 k^4 - \Sigma_{ijl} k^5 k^2 k^2 \}$
$\mathcal{L}_{\zeta(\partial_i \partial_j \chi)^2}$	$\frac{\epsilon^3}{96 \Pi k^2} \{ \sum k^9 - 3 \Sigma_{ij} k^7 k^2 - \Sigma_{ij} k^6 k^3 + 3 \Sigma_{ij} k^5 k^4 - \Sigma_{ijl} k^5 k^2 k^2 + \Sigma_{ijl} k^4 k^3 k^2 \}$

TABLE V. Per-vertex squeezed- and equilateral-limit local-amplitude contribution of each of Cai *et al.*'s [7] four cubic vertices (rows of Table IV), at $\epsilon = 3/2$, formed as $f_{\text{NL}} = (10/3)A_v / \sum_i k_i^3$ for each vertex v alone. Column sums are the certified benchmarks $f_{\text{NL}}^{\text{sq}} = -35/16$ and $f_{\text{NL}}^{\text{eq}} = -255/128$ (Table I). The single-vertex $\mathcal{L}_{\zeta \partial \zeta \partial \chi}$ squeezed contribution vanishes; the net $-35/16$ is carried by the field-redefinition and highest-order bulk vertices. Values are the exact per-row limits of the same exact-fraction vertex expressions summed by `scripts/p2_vertex_check.py` (SYMPY), whose column sums reproduce that script's certified total; no fitting or rounding.

vertex	squeezed $f_{\text{NL}}^{\text{sq}}$	equilateral $f_{\text{NL}}^{\text{eq}}$
field redef. ($\mathcal{L}_{\text{redef}}$)	$-\frac{25}{16}$	$-\frac{35}{32}$
$\mathcal{L}_{\zeta \zeta^2}$	$-\frac{5}{32}$	$-\frac{5}{32}$
$\mathcal{L}_{\zeta \partial \zeta \partial \chi}$	0	$-\frac{5}{8}$
$\mathcal{L}_{\zeta(\partial_i \partial_j \chi)^2}$	$-\frac{15}{32}$	$-\frac{15}{128}$
sum (all four)	$-\frac{35}{16}$	$-\frac{255}{128}$

and $\Sigma_i^{\text{dist}} k_i^5 k_j^2 k_l^2 \equiv k_1^5 k_2^2 k_3^2 + k_2^5 k_1^2 k_3^2 + k_3^5 k_1^2 k_2^2$ over the three distinct $(5, 2, 2)$ monomials (so, e.g., the leading squeezed piece $k_1 \rightarrow 0$ isolates the $9k_2^9 + 9k_3^9 + 3k_2^7 k_3^2 + \dots$ collinear block that reduces to $-35/16$ after forming $f_{\text{NL}} = (10/3)A / \sum_i k_i^3$). Second, the ϵ -order-grouped decomposition. Writing $A = \epsilon A^\epsilon + \epsilon^2 A^{\epsilon^2} + \epsilon^3 A^{\epsilon^3}$ (Cai *et al.*'s own intermediate grouping) and forming the squeezed f_{NL} -contribution of each order separately at $\epsilon = 3/2$ gives

$$f_{\text{NL}}|_{\epsilon^1} = -\frac{5}{2}, \quad f_{\text{NL}}|_{\epsilon^2} = +\frac{5}{16}, \quad f_{\text{NL}}|_{\epsilon^3} = 0, \quad (\text{B5})$$

whose sum $-\frac{5}{2} + \frac{5}{16} + 0 = -\frac{35}{16}$ reproduces Eq. (B3) term-by-term (the A^ϵ piece is the pure $-\frac{\epsilon}{2} \sum_i k_i^3$ local shape; the A^{ϵ^2} and A^{ϵ^3} pieces carry the non-local Πk^2 -denominator structure of Eq. B4). The order-grouped values Eq. (B5) are the ϵ -decomposition Cai *et al.* describe, made explicit.

Two independent certifications bracket this result. First, Cai *et al.*'s own ϵ -order-grouped intermediate expressions ($A^\epsilon, A^{\epsilon^2}, A^{\epsilon^3}$; Eq. B5) sum to the Table IV total *exactly* (difference = 0 symbolically) and likewise give $-35/16$, so the correction is internal to Cai's own paper. Second, subtracting the vertex sum from Cai's transcribed printed polynomial A_T (their Eq. 37) leaves exactly the spurious local term of Eq. (B2), $A_T^{\text{printed}} - \sum(\text{vertices}) = -(99/128) \sum_i k_i^3$, one traceable term by

which the transcribed printed polynomial differs from the exact vertex sum (and by which its squeezed reduction departs from $-35/16$ to $-305/64$, not to Cai's stated $-35/8$; not a naive additive shift; cf. Eq. B2 discussion above). Independently, Li *et al.*'s [8] general- c_s formula (their Eq. 5.1) $f_{\text{NL}}^{\text{local}} = -165/16 + 65/(8c_s^2)$ evaluates to $-35/16$ at $c_s = 1$ by a wholly separate in-in method. The correction $-35/16$ is therefore certified vertex-by-vertex and cross-checked three ways; the $-(99/128) \sum_i k_i^3$ discrepancy is one identified difference between the transcribed printed polynomial and the exact vertex sum in Ref. [7].

Interpretation. The discrepancy is not a local-template normalization change: Cai *et al.* and Li *et al.* use the same $f_{\text{NL}} = 10A/(3 \sum_i k_i^3)$ convention and the same squeezed limit. Nor can the commutator sign identity, used consistently within either the Hamiltonian or interaction-Lagrangian convention, adjudicate it, because both calculations use the full in-in structure. What is established is narrower and sufficient: the exact four-vertex sum, Cai *et al.*'s order-grouped intermediates, and Li *et al.*'s general- c_s formula all give $-35/16$ at $c_s = 1$. The transcribed total polynomial differs from that vertex sum by Eq. (B2) and squeezed-reduces to $-305/64$, not to the separately published $-35/8$. We therefore adopt $-35/16$ and do not assign a time-ordering mechanism to

the historical value.

A.1 Wick contraction and the in-in commutator

At first order in the cubic interaction Hamiltonian,

$$\langle \zeta^3(\eta_*) \rangle = -i \int_{-\infty}^{\eta_*} d\eta \langle [\zeta^3(\eta_*), H_{\text{int}}(\eta)] \rangle. \quad (\text{B6})$$

For Hermitian H_{int} , define $I_{v,H}^{(\sigma)} \equiv \int_{-\infty}^{\eta_*} d\eta \langle \zeta^3(\eta_*) H_{\text{int},v}(\eta) \rangle_c^{(\sigma)}$. The two terms in the commutator are complex conjugates, so the Hamiltonian convention gives

$$-i \langle [\zeta^3, H_{\text{int}}] \rangle = +2 \text{Im} \langle \zeta^3 H_{\text{int}} \rangle. \quad (\text{B7})$$

For each cubic vertex v , Wick's theorem expands the connected time-ordered correlator as

$$\langle \zeta_{\mathbf{k}_1} \zeta_{\mathbf{k}_2} \zeta_{\mathbf{k}_3} \mathcal{O}_v \rangle_c = \sum_{\sigma \in S_3} \prod_{a=1}^3 \langle \zeta_{\mathbf{k}_{\sigma(a)}} \Phi_{v,a} \rangle, \quad (\text{B8})$$

with all six external-leg permutations included. Combining the four vertices and their symmetry factors gives

$$B_\zeta = +2 \text{Im} \sum_{v=1}^4 \sum_{\sigma \in S_3} \frac{I_{v,H}^{(\sigma)}}{\mathcal{S}_v}. \quad (\text{B9})$$

Equivalently, an interaction-Lagrangian convention may define $I_{v,L}^{(\sigma)} \equiv -I_{v,H}^{(\sigma)}$ and write the same result as $B_\zeta = -2 \text{Im} \sum I_{v,L}^{(\sigma)} / \mathcal{S}_v$; the integral definition and its sign must not be mixed across conventions. Equation (B7) is universal and does not by itself decide between the reported local amplitudes. The adopted result instead follows from the four explicit, already signed vertex rows in Table V, whose squeezed contributions are $-25/16$, $-5/32$, 0 , and $-15/32$, summing to $-35/16$; their equilateral contributions likewise sum to $-255/128$. The accompanying symbolic checker prints all four rows and the two exact totals.

ACKNOWLEDGMENTS

The author thanks the developers of the open-source numerical and symbolic tools used in this work. The author designed the study, made all scientific and editorial judgments, and takes full responsibility for the content. No external funding was received.

-
- [1] C. Heinrich, O. Dore, and E. Krause, Measuring f_{nl} with the spherex multi-tracer redshift space bispectrum, *Phys. Rev. D* **109**, 123511 (2024), [arXiv:2311.13082 \[astro-ph.CO\]](#).
 - [2] J. Maldacena, Non-gaussian features of primordial fluctuations in single field inflationary models, *JHEP* **0305**, 013, [arXiv:astro-ph/0210603](#).
 - [3] E. Pajer, F. Schmidt, and M. Zaldarriaga, The Observed Squeezed Limit of Cosmological Three-Point Functions, *Phys. Rev. D* **88**, 083502 (2013), [arXiv:1305.0824 \[astro-ph.CO\]](#).
 - [4] T. Tanaka and Y. Urakawa, Dominance of Gauge Artifact in the Consistency Relation for the Primordial Bispectrum, *JCAP* **1105**, 014, [arXiv:1103.1251 \[astro-ph.CO\]](#).
 - [5] D. Wands, Duality invariance of cosmological perturbation spectra, *Phys. Rev. D* **60**, 023507 (1999), [arXiv:gr-qc/9809062](#).
 - [6] F. Finelli and R. Brandenberger, On the generation of a scale-invariant spectrum of adiabatic fluctuations in cosmological models with a contracting phase, *Phys. Rev. D* **65**, 103522 (2002), [arXiv:hep-th/0112249](#).
 - [7] Y.-F. Cai, W. Xue, R. Brandenberger, and X. Zhang, Non-gaussianity in a matter bounce, *JCAP* **0905**, 011, [arXiv:0903.0631 \[astro-ph.CO\]](#).
 - [8] Y.-B. Li, J. Quintin, D.-G. Wang, and Y.-F. Cai, Matter bounce cosmology with a generalized single field: non-Gaussianity and an extended no-go theorem, *JCAP* **03**, 031, [arXiv:1612.02036 \[astro-ph.CO\]](#).
 - [9] D. Wands, Local non-Gaussianity from inflation, *Class. Quant. Grav.* **27**, 124002 (2010), [arXiv:1004.0818](#).
 - [10] O. Dore *et al.*, Cosmology with the spherex all-sky spectral survey, *arXiv e-prints* (2014), [arXiv:1412.4872](#).
 - [11] N. Dalal, O. Dore, D. Huterer, and A. Shirokov, The imprints of primordial non-gaussianities on large-scale structure: scale-dependent bias and abundance of virialized objects, *Phys. Rev. D* **77**, 123514 (2008), [arXiv:0710.4560](#).
 - [12] E. Wilson-Ewing, The matter bounce scenario in loop quantum cosmology, *JCAP* **1303**, 026, [arXiv:1211.6269](#).
 - [13] L. Freidel, D. Minic, and T. Takeuchi, Quantum gravity, torsion, parity violation and all that, *Phys. Rev. D* **72**, 104002 (2005), [arXiv:hep-th/0507253](#).
 - [14] H. Golden, Algebraic Cartan Elimination in Minimal Einstein–Cartan–Holst Gravity: Spin-Sourced Contact and Zero-Spin Scalar Branches, *Zenodo* [10.5281/zenodo.21481838](#) (2026), Companion paper (Paper I A), publicly archived on Zenodo under version DOI 10.5281/zenodo.21481838 (concept DOI 10.5281/zenodo.21481837, resolving to the latest deposited version), deposited 2026 July 22 under CC-BY-4.0. The record is a public permanent archive of the manuscript and its exact source; it is not an arXiv preprint and the paper is not peer reviewed.
 - [15] A. Slosar, C. Hirata, U. Seljak, S. Ho, and N. Padmanabhan, Constraints on local primordial non-gaussianity from large scale structure, *JCAP* **2008** (08), 031,

- [arXiv:0805.3580](#).
- [16] L. Senatore, K. M. Smith, and M. Zaldarriaga, Non-Gaussianities in Single Field Inflation and their Optimal Limits from the WMAP 5-year Data, *JCAP* **01**, 028, [arXiv:0905.3746 \[astro-ph.CO\]](#).
 - [17] D. J. Schlegel *et al.*, The MegaMapper: A stage-5 spectroscopic instrument concept for the study of inflation and dark energy, arXiv e-prints (2022), [arXiv:2209.04322](#).
 - [18] Y.-F. Cai, X. Chen, M. H. Namjoo, M. Sasaki, D.-G. Wang, and Z. Wang, Revisiting non-gaussianity from non-attractor inflation models, *JCAP* **2018** (05), 012, [arXiv:1712.09998 \[astro-ph.CO\]](#).
 - [19] G. Jung, M. Citran, B. van Tent, L. Dumilly, and N. Aghanim, Constraints on primordial non-Gaussianity from Planck PR4 data, *Astronomy & Astrophysics* **702**, A204 (2025), [arXiv:2504.00884 \[astro-ph.CO\]](#).
 - [20] E. Chaussidon *et al.*, Constraining Primordial Non-Gaussianity with DESI 2024 LRGs and QSOs, *J. Cosmol. Astropart. Phys.* (2024), dESI DR1 LRG combined PNG constraint; $f_{\text{NL}}^{\text{loc}} = -3.6_{-9.1}^{+9.0}$; $\sigma \approx 9-10$, [arXiv:2411.17623](#).
 - [21] A. Fondi *et al.*, Assembly bias and local Primordial non-Gaussianity from DESI DR1 Quasars, arXiv e-prints (2025), dESI QSO assembly-bias PNG analysis; $f_{\text{NL}}^{\text{loc}} = -3.3 \pm 9.2$; distinct from Chaussidon et al. LRG+QSO combined result, [arXiv:2602.12357](#).