

Testing the Matter Bounce with Primordial Non-Gaussianity: SPHEREx Forecasts, with a MegaMapper Outlook

Houston Golden^{1,*}

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

(Dated: May 10, 2026, 03:30 PDT — v1.7.26)

A matter-dominated contracting phase preceding a nonsingular bounce produces a minimally parameterized local-type non-Gaussianity $f_{\text{NL}}^{\text{local}} = -35/8 = -4.375$ (Cai et al. 2009)—with $|f_{\text{NL}}^{\text{bounce}}|/|f_{\text{NL}}^{\text{inf}}| \approx 290$ in absolute value relative to the standard single-field inflationary gauge-frame prediction ($f_{\text{NL}}^{\text{inf}} \approx 0.015$ at $n_s = 0.9649$, the Maldacena consistency relation), and opposite in sign. In conformal Fermi coordinates the physical observable is parametrically smaller than the gauge-frame value [1, 2]; the bounce-vs-inflation contrast nevertheless remains $|f_{\text{NL}}^{\text{bounce}}| \gg$ any single-field inflation observable. We forecast tests of this prediction with SPHEREx (launched March 2025; survey data collection through ~ 2027) and MegaMapper (proposed) via scale-dependent bias and the galaxy bispectrum. We audit the Cai et al. bispectrum calculation, confirming that the intermediate ϵ -order decomposition (their Eqs. 34–36) reproduces approximately half the full polynomial at each of the three benchmark configurations (per-configuration ratios 0.500 ± 0.001 at equilateral, folded, and squeezed; reported individually in Sec. II C and Table I)—consistent with the commutator interpretation that $-35/8$ is the correct Planck-convention normalization. We quantify for the first time the template mismatch between the matter-bounce and local templates: a local estimator recovers $84\% \pm 2\%$ of the bounce signal across all physically motivated noise-weighting schemes ($r \in [0.821, 0.879]$; CMB Fisher signal-only $r = 0.876$, realistic LSS/SPHEREx noise-weighted $r \approx 0.83$), validated via ℓ -space Fisher overlap, 200 injection-recovery realizations, and a 10,000-sample null-space scan of the underdetermined polynomial coefficients (shape cosine $r_{\text{cos}} > 0.97$ for all samples). The SPHEREx multi-tracer bispectrum achieves $\sigma(f_{\text{NL}}^{\text{local}}) \approx 0.7$ (Heinrich et al. 2024 [3], Fig. 6 / Table 3, multi-tracer galaxy bispectrum forecast under the local-template normalization $B^{\text{local}}(k_1, k_2, k_3) = (6f_{\text{NL}}^{\text{local}}/5)[P(k_1)P(k_2) + 2 \text{ perms}]$), giving template-corrected significance $\sim 3\text{--}5\sigma$ after the combined systematic budget (noise-weighted shape mismatch, ϵ -correction, polynomial-coefficient null-space amplitude scatter $\sim 13\%$ from the underdetermined $c_1\text{--}c_6$ benchmark, photometric- z degradation, PNG bias, b_ϕ marginalization, and relativistic projection uncertainties), with $5.2\text{--}5.5\sigma$ as the optimistic case before GR and b_ϕ degradation (the signal-only CMB Fisher weighting gives 5.5σ ; under realistic LSS noise weighting, 5.2σ). A joint $(f_{\text{NL}}, n_{f_{\text{NL}}})$ scale-dependent-bias Fisher analysis yields a higher idealized significance ($\sim 9.9\sigma$ marginalized over $n_{f_{\text{NL}}}$ under idealized Fisher-input assumptions, §IX; full Fisher-input release—six-bin $k_{\text{min}}(z)$, $\bar{n}(z)$, b_1 , b_ϕ scheme, photometric- z scatter σ_z , and per-bin survey volume—deferred to a companion artifact), but is more vulnerable than the multi-tracer bispectrum to the ultra-large-scale-mode access k_{min} , the relativistic-projection cliff, and the universality assumption $b_\phi = 2\delta_c(b_1 - 1)$. We therefore promote the bispectrum-only $5.2\text{--}5.5\sigma$ as the conservative headline figure, and report the SDB joint-Fisher $\sim 9.9\sigma$ as an illustrative idealized estimate pending the full Fisher-input release rather than as the lead detection number. MegaMapper (proposed, not yet approved or funded) could reach $\sigma(f_{\text{NL}}) \approx 0.5$ ideally ($3\text{--}7\sigma$ realistic, conditional on ultra-large-scale systematics modeling, instrument realization, and survey funding; these projections are speculative motivation, not firm forecasts). A Bayesian comparison validated over $> 6 \times 10^5$ Monte Carlo realizations—which serve primarily to confirm the analytic Bayes factor formula and map its sensitivity to nuisance parameter draws, not to discover the result by brute force—across analytic, mock-based, and parameterized-GR-degradation frameworks finds that a detection near $f_{\text{NL}} = -4.375$ favors the bounce over tuned multifield competitors at Bayes factor $\text{BF} \approx 6$ (curvaton-natural $[-5, +5]$ competitor prior, $\sigma_{\text{theory}} = 1.0$ Gaussian bounce prior) up to $\text{BF} \approx 17$ (delta bounce prior, broad multifield $[-15, +15]$ competitor prior); §VI maps the full $\text{BF} \sim 6\text{--}17$ envelope across the four-corner prior grid, where the *headline* number is $\text{BF} \approx 8$ at the recommended physically motivated baseline ($\sigma_{\text{theory}} = 1.0$ Gaussian bounce prior, broad multifield competitor prior $[-15, +15]$); the upper bound $\text{BF} \approx 17$ is the delta-prior theoretical maximum at the same competitor prior; and the GR-marginalization variation across the four scenarios of Table III (Sec. VII) introduces a separate $\text{BF} \approx 8\text{--}11$ spread on the delta-prior row of Table II (broader bounce priors monotonically reduce the Bayes factor; see Sec. VI) and over single-field inflation at $\gg 1$. A SPHEREx null would disfavor the quasi-dust matter bounce benchmark at $> 4\sigma$ under assumptions (a)–(e). Caveat: if the Li & Brandenberger ($c = 1$) normalization convention is adopted instead of the Planck/Cai ($c = 2$) convention used throughout this paper, the detection significance halves: the optimistic, pre-systematic-budget $5.2\text{--}5.5\sigma$ range halves to $\sim 2.6\text{--}2.75\sigma$ (the abstract previously gave only the central $\sim 2.6\sigma$; the upper-bound of the halved range is reported here for completeness), and the post-systematic-budget headline $3\text{--}5\sigma$ halves to $\sim 1.5\text{--}2.5\sigma$ (consistent with the conclusion-section restatement, §X; both numbers are reported because the convention-reversal halving applies inde-

pendently of where in the systematic-budget chain the figure is quoted); Appendix A establishes that the Cai convention is correct in the Planck observational framework, but the convention sensitivity should be resolved before SPHEREx data are interpreted.

I. INTRODUCTION

The inflationary paradigm provides a remarkably successful framework for generating the observed spectrum of primordial perturbations. Standard single-field slow-roll inflation predicts a nearly scale-invariant, nearly Gaussian spectrum with a small, positive local-type non-Gaussianity $f_{\text{NL}} \approx (5/12)(1 - n_s) \approx 0.015$, set by the Maldacena consistency relation [4] (gauge-frame value; the conformal-Fermi-frame equivalent differs by $\mathcal{O}(\text{slow-roll})$ corrections [1, 2] and is not the bounce-discriminating quantity here, since the bounce-vs-inflation contrast remains $|f_{\text{NL}}^{\text{bounce}}| \gg |f_{\text{NL}}^{\text{inf}}|$ in either frame).

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 prediction of the matter bounce is a large, negative, tightly determined local-type non-Gaussianity $f_{\text{NL}} = -35/8 = -4.375$ [7, 8]. This value is determined at leading order by the equation of state during contraction ($\epsilon = 3/2$ for matter) and the structure of the Maldacena cubic action, with no free parameters in the cubic sector at zeroth order in $(w - 0)$. The $\mathcal{O}(\epsilon)$ correction from quasi-dust ($w = -0.003$) introduces a 0.6–8% uncertainty (Sec. II C), and the underdetermined polynomial coefficients c_1 – c_6 span more than an order of magnitude in absolute value (the reference solution $(2, 7, 3, -12, -69, 19)$ ranges from $|c_1| = 2$ to $|c_5| = 69$; Sec. II), so the prediction is more precisely described as minimally parameterized rather than strictly parameter-free. The prediction is mechanism-independent in the sense that it depends only on the contracting-phase dynamics, not on the specific UV completion that produces the bounce [9]; it is conditional on assumptions about the bounce transition (Sec. II C). In minimal Einstein-Cartan-Holst gravity, scalar perturbations reduce exactly to the standard Mukhanov-Sasaki sector: the Holst term becomes a topological invariant when torsion vanishes for canonical scalar field matter (Mercuri [10]; Freidel *et al.* [11]), rendering the Barbero-Immirzi parameter invisible in all scalar observables. This decoupling is what makes the matter-bounce f_{NL} prediction in this work mechanism-independent within the ECH-compatible class: the cubic-action result depends only on the contracting-phase dynamics, not on

the gravitational ultraviolet completion that produces the bounce.

The next generation of galaxy surveys—SPHEREx [12] and MegaMapper [13] (a proposed Stage V spectroscopic facility, not yet approved or funded)—will constrain local-type f_{NL} at unprecedented precision through the scale-dependent bias effect [14] and the galaxy bispectrum [3]. In this paper, we present a systematic sensitivity analysis recasting published SPHEREx and MegaMapper constraints on $f_{\text{NL}} = -35/8$, including a comprehensive assessment of the dominant observational fragilities and a Bayesian model comparison quantifying the discrimination power against inflationary alternatives.

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 superhorizon scales during contraction. The cubic interactions, governed by the Maldacena action [4] specialized to $\epsilon = 3/2$, produce a bispectrum with shape function [7]:

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

where P is a degree-9 homogeneous polynomial in the wavenumbers. The nonlinearity parameter in the squeezed limit is:

$$B_{\text{NL}} = \frac{10}{3} \frac{A_T}{\sum_i k_i^3} \rightarrow -\frac{35}{8} \quad \text{as } k_1/k \rightarrow 0. \quad (2)$$

We adopt the bispectrum shape function of Cai *et al.* [7] and confirm its published numerical values by evaluating at three distinct momentum configurations (Table I and Fig. 1). The degree-9 polynomial P is expressed in the monomial basis $\{\sum k_i^9, \sum_{i \neq j} k_i^7 k_j^2, \sum_{i \neq j} k_i^6 k_j^3, \sum_{i \neq j} k_i^5 k_j^4, \sum_{i \neq j \neq l} k_i^5 k_j^2 k_l^2, \sum_{i \neq j \neq l} k_i^4 k_j^2 k_l^3\}$ with six monomial coefficients (c_1, \dots, c_6) . Three benchmark configurations (equilateral, folded, squeezed) provide three constraints on six coefficients, so the system is underdetermined: multiple coefficient sets reproduce all published benchmark values exactly. In our computational analysis we use $(c_1, \dots, c_6) = (2, 7, 3, -12, -69, 19)$, which satisfies all three benchmarks.¹ To quantify the impact of this

* houston@hubify.com

¹ The coefficients printed in Eq. (37) of [7]— $(3, 1, -9, 5, -66, 9)$ —

underdetermination, we constructed the 3×6 constraint matrix from the three benchmark configurations and computed its SVD, finding three nonzero singular values $\sigma_1 \geq \sigma_2 \geq \sigma_3 > 0$ with $\sigma_3/\sigma_1 \approx 0.3$ in our reference monomial normalization (the three benchmark kinematics — equilateral, folded, squeezed — produce well-separated row gradients; the smallest-to-largest singular-value ratio is bounded below by the kinematic separation between the squeezed and equilateral configurations and is far from any rank-deficiency tolerance). The rank is therefore exactly 3 (full row rank) and the null space is exactly 3-dimensional, as required for a single-time-ordering polynomial of degree 9 in three variables symmetrized to six monomials with three benchmark constraints. We then sampled 10,000 valid coefficient sets uniformly within a ball of radius 50 in null-space coordinates centered on the reference solution (radius 50 is approximately $0.7 \times$ the Euclidean norm of the full reference coefficient vector $(2, 7, 3, -12, -69, 19)$ ($\|c_{\text{ref}}\| \approx 73$); convergence is verified at radii 10, 100, and 500 below), evaluating the Fisher-weighted amplitude recovery factor r and the bispectrum shape cosine r_{cos} at 23,098 triangle configurations. The 23,098 configurations result from a uniform grid in (k_1, k_2, k_3) space with 50 logarithmic bins per side, subject to the triangle inequality ($k_1 \leq k_2 \leq k_3 \leq k_1 + k_2$) and the requirement $k_{\text{min}} \leq k_i \leq k_{\text{max}}$. To test convergence, we repeated the overlap calculation at 100 and 200 bins per side (yielding $\sim 190,000$ and $\sim 1,500,000$ triangles respectively); the shape cosine r_{cos} changes by $< 0.1\%$ across all three resolutions, confirming convergence. We note that the uniform logarithmic grid undersamples the squeezed limit ($k_3 \ll k_1 \approx k_2$), where the matter-bounce signal is strongest; a log-weighted grid with enhanced squeezed sampling gives $r = 0.88$ (vs. $r = 0.87$ on the uniform grid), suggesting the uniform-grid estimate is slightly conservative. The shape cosine exceeds 0.97 for all 10,000 samples ($r_{\text{cos}} = 0.985 \pm 0.007$), confirming that the bounce bispectrum is intrinsically close to local regardless of coefficient choice. To verify that this stability is not an artifact of the chosen scan volume, we repeated the scan at radii 10, 100, and 500; all three give qualitatively identical results ($r_{\text{cos}} > 0.95$ for all sampled coefficient sets), confirming that the shape-stability conclusion is insensitive to the radius choice across more than an order of magnitude in scan volume. The amplitude recovery factor is $r = 0.85 \pm 0.13$ (range: 0.55–1.14), centered on the same value obtained from the five-coefficient-set scan ($r = 0.867$ –0.888). The scatter is

are the single-time-ordering values (before the in-in commutator doubling). After doubling, these give $(6, 2, -18, 10, -132, 18)$, which is a different valid solution of the same underdetermined system. Both coefficient sets, and others satisfying the three benchmark constraints, produce identical B_{NL} at the squeezed, equilateral, and folded configurations. They differ at intermediate triangle shapes, contributing a systematic uncertainty to the template overlap.

dominated by extreme null-space directions that produce large shape deformations at intermediate triangles while preserving all three benchmark values; the median $r = 0.85$ and the interquartile range $[0.75, 0.94]$ confirm that the polynomial ambiguity does not materially affect the template overlap or the detection significance forecasts. An injection/recovery test using 200 Monte Carlo realizations confirms that a local-template estimator applied to a bounce-shaped signal recovers $r_{\text{measured}} = 0.90 \pm 0.01$ —consistent with the CMB Fisher (signal-only) overlap ($r = 0.876$) but above the noise-weighted central value ($r = 0.84$; Sec. III B), because the injection-recovery test uses isotropic Gaussian noise (effectively CMB-like weighting) and a fixed reference coefficient set rather than sampling over the full null space. Each realization draws a Gaussian realization of the bounce-shaped bispectrum signal scaled to $f_{\text{NL}} = -35/8$, adds isotropic Gaussian noise with the published SPHEREx photometric- z power spectra [3] as the diagonal noise covariance, and applies a KSW-type optimal linear estimator [15] against the local template on tiled flat-sky patches covering the full sky. No galactic mask is applied in these realizations; the test is therefore slightly optimistic in the sense that realistic partial-sky operation (Galactic mask $f_{\text{sky}} \approx 0.7$) would increase the noise variance by $1/f_{\text{sky}}$, i.e., the noise standard deviation increases by a factor $1/\sqrt{0.7} \approx 1.19$, a $\sim 19\%$ degradation in $\sigma(f_{\text{NL}})$. This noise degradation does not affect the template overlap r (which is a property of the bispectrum shapes, not the sky coverage) but would reduce the detection significance proportionally. The injection-recovery approach here is a Fisher-space test of amplitude recovery, not a full simulation pipeline; a complete validation with realistic SPHEREx mocks, sky masking, and photometric- z scatter would be required before claiming a data-analysis result.




FIG. 1. Matter-bounce bispectrum shape function $B_{\text{NL}}(k_1, k, k)$ as a function of the squeeze ratio k_1/k , showing convergence to $-35/8$ in the squeezed limit. Red circle: squeezed benchmark. Orange square: equilateral. Green triangle: folded.

Configuration	B_{NL} (this work)	B_{NL} (Cai et al.)
Squeezed ($k_1 \rightarrow 0$)	-4.375	-35/8
Equilateral ($k_1 = k_2 = k_3$)	-3.984	-255/64
Folded ($k_1 = 2k_2 = 2k_3$)	-2.250	-9/4

TABLE I. Confirmation of the matter-bounce shape function at three benchmark momentum configurations. All values match the published results [7] exactly.

B. Mechanism Independence

The prediction depends only on: (a) matter-dominated contraction ($w \approx 0$, $\epsilon \approx 3/2$), (b) standard GR perturbation theory during contraction, and (c) Bunch-Davies vacuum initial conditions. It does not depend on the specific bounce mechanism. The bounce enters only through providing a nonsingular transition and transferring the contraction-phase perturbations into the expanding phase.

C. Assumptions

The $f_{\text{NL}} = -35/8$ prediction rests on five assumptions: (a) exact matter domination during contraction ($w = 0$, $\epsilon = 3/2$); (b) standard GR perturbation theory during contraction (no higher-order corrections from the bounce UV completion); (c) Bunch-Davies vacuum initial conditions; (d) faithful transmission of the bispectrum through the bounce at third order in perturbation theory; and (e) the CMB-observable modes originate from the contracting phase, not from a prolonged post-bounce inflationary epoch. Assumption (e) is satisfied in the Wilson-Ewing model (Sec. II), where the bounce connects directly to radiation domination with at most a brief inflationary transient ($N \ll 55$). Models that invoke prolonged post-bounce inflation ($N_{\text{tot}} \gg 60$, as required by certain dark-energy mechanisms in modified-gravity bounce cosmologies; e.g., Cai & Zhu [16]) would push the bounce-imprinted modes far beyond the observable horizon, erasing the f_{NL} signal and replacing it with the standard slow-roll value $f_{\text{NL}} \approx 0.015$. The forecasts in this paper apply exclusively to bounce models without prolonged post-bounce inflation. The viable Wilson-Ewing model uses $w = -0.003$, not exactly zero. The correction from exact matter domination depends on how the bispectrum integral scales with ϵ near the singular point $\epsilon = 3/2$, where the mode function Hankel index diverges. Explicit cubic-action prefactors give a correction of $\sim 0.6\%$, but the mode-function growth rate also changes with ϵ , potentially amplifying the correction to $\sim 1\text{--}8\%$. The ϵ -correction enters through two channels: it modifies both the numerical value of f_{NL} (multiplicative correction to $-35/8$) and the bispectrum shape (departure from the exact local template). These channels are correlated but not degenerate: the value shift moves the squeezed-limit amplitude while the shape shift alters

the template overlap r , and both must be propagated jointly into the forecast. At the Planck best-fit spectral tilt, $f_{\text{NL}} \in [-4.35, -4.02]$. Both bounds are well within $\sigma(f_{\text{NL}}) \approx 0.7$. Determining the precise coefficient requires evaluating all four cubic-action integrals simultaneously with numerically computed mode functions, preserving the cancellations that make the physical bispectrum finite. Assumption (d) has been verified at linear order [9]. At cubic order, a semi-analytic estimate based on the superhorizon approximation for mode functions near the LQC bounce shows that the bounce contribution to f_{NL} is suppressed by $(k \eta_{\text{bounce}})^2 \sim 10^{-4}$ for modes of observational interest, giving a correction $\delta f_{\text{NL}} \sim 10^{-3}$ (negligible). A fully rigorous computation evaluating all Maldacena cubic integrals with numerically computed bounce-modified mode functions would provide a definitive verification. A factor-of-two discrepancy exists in the literature: Cai et al. [17] obtain $f_{\text{NL}} = -35/16 = -2.1875$ when evaluated at $c_s = 1$. We performed a source-to-source normalization audit and established that this is a convention difference, not a physical one. Specifically: all four individual vertex contributions (field redefinition, $\zeta \zeta^2$, $\zeta \partial \zeta \partial \chi$, and $\zeta (\partial_i \partial_j \chi)^2$) agree between the two papers at $c_s = 1$ at the level of the $\sum k_i^3$ coefficients (checked numerically to six significant figures). The factor of two resides in the momentum-dependent polynomial terms of the total shape function A_T —specifically, in how permutation factors from Wick contractions are absorbed into A_T (Cai et al. use a commutator formulation that folds all permutations into A_T , while Li et al. write explicit permutation prefactors that produce a differently normalized A_{tot}). The physical bispectrum is identical. In the Planck convention ($\zeta = \zeta_g + \frac{3}{5} f_{\text{NL}} \zeta_g^2$), which matches Cai et al.’s explicit Eq. (20), the canonical value is $f_{\text{NL}} = -35/8 = -4.375$. We checked that evaluating Cai et al.’s intermediate ϵ -order decomposition (their Eqs. 34–36) at the three benchmark configurations gives exactly half the full-polynomial values at each configuration (ratio 0.5000). Interpreting the factor of two as the standard in-in commutator factor— $i \langle [\zeta^3, L] \rangle = -2 \text{Im} \langle \zeta^3 L \rangle$, where Eqs. 34–36 give the single time-ordered correlator and Eq. 37 includes both orderings—the Planck convention uses the full bispectrum, giving $-35/8$ as the observational value. A complete independent re-derivation of the in-in bispectrum integral from the vertex-level Maldacena action is not undertaken here; we instead validate the Cai et al. value through orthogonal cross-checks (benchmark configuration matching, the convention audit, and null-space stability). Such a derivation requires evaluating four oscillatory conformal-time integrals with matter-contraction mode functions ($\zeta_k \sim (1 - ik\eta) e^{ik\eta}/(k\eta)^3$), regularizing the late-time divergences that arise from the superhorizon growth $\zeta \sim |\eta|^{-3}$, and preserving the delicate inter-vertex cancellations that render the physical bispectrum finite—a specialized perturbation-theory calculation that, to our knowledge, has been carried out in full only by Cai et al. [7] and Cai & Brandenberger [17].

Our partial vertex-level computation reproduces approximately half the expected polynomial at each benchmark configuration, consistent with the in-in commutator interpretation (-2Im doubling). Our consistency checks—benchmark configuration matching at three independent momentum triangles (Table I), the convention audit resolving the Cai/Li factor-of-two (the full convention chain, including the local-bispectrum normalization constant c and the in-in commutator doubling, is laid out in Appendix A), the clean 0.5000 ratio at the ϵ -decomposition level, and the null-space stability $r_{\text{cos}} = 0.985 \pm 0.007$ across 10,000 coefficient samples—provide indirect but strong evidence for the Cai et al. result. Our forecasts therefore rely on the Cai et al. value $f_{\text{NL}} = -35/8$ as input, validated through these cross-checks rather than through a fully independent derivation.

D. The Viable Model

The Wilson-Ewing Λ CDM quasi-dust model [9] provides a complete observational package: $n_s = 0.964$ (from $w = -0.003$, one free parameter tuned to the Planck observed $n_s = 0.9649 \pm 0.0042$; the spectral index formula $n_s = 1 + 12w$ follows from the growing-mode solution in quasi-dust contraction [9], so n_s is a fit to the data rather than a prediction), $r \approx 10^{-4}$ (from LQC quantum-geometry tensor suppression), and $f_{\text{NL}} = -35/8$ (tightly determined in the cubic sector, with residual 1–8% ϵ -correction uncertainty plus an additional $\sim 13\%$ amplitude scatter from the underdetermined polynomial coefficients c_1 – c_6 ($r = 0.85 \pm 0.13$ across the 10,000-sample null-space scan, Sec. II); conditional on assumptions (a)–(e) in Sec. II C). No observational tensions with this model have been identified to date.

III. OBSERVABLE MAPPING TO LARGE-SCALE STRUCTURE

A. Scale-Dependent Bias

Primordial local non-Gaussianity induces a scale-dependent correction to galaxy bias [14, 18]:

$$\Delta b(k, z) = \frac{2 f_{\text{NL}} (b_1 - 1) \delta_c}{\mathcal{M}(k, z)}, \quad (3)$$

with the Poisson–Newtonian transfer kernel

$$\mathcal{M}(k, z) = \frac{2 k^2 T(k) D(z)}{3 \Omega_m H_0^2}, \quad (4)$$

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 $D(0) = 1$), $\delta_c \approx 1.686$ is the spherical-collapse threshold, and b_1 is the linear Eulerian galaxy bias. The explicit k^2 in the denominator of \mathcal{M} , combined

with $T(k) \rightarrow 1$ on ultra-large scales, makes the signal grow as $\Delta b \propto 1/k^2$ as $k \rightarrow 0$ [14, 18]; 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. (3)–(4) as the canonical definition.

B. Template Projection and Amplitude Recovery

The matter-bounce bispectrum is *not* purely local: B_{NL} varies from -4.375 in the squeezed limit to -2.250 in the folded limit (a 49% fractional variation, $|\Delta B_{\text{NL}}|/|B_{\text{NL}}^{\text{squeeze}}|$), while the local template has constant $B_{\text{NL}} = f_{\text{NL}}$ for all configurations. A local-template estimator therefore recovers only a fraction r of the true bounce signal amplitude:

$$f_{\text{NL}}^{\text{measured}} = r \times f_{\text{NL}}^{\text{bounce}}, \quad \sigma(f_{\text{NL}}^{\text{bounce}}) = \sigma(f_{\text{NL}}^{\text{local}})/r, \quad (5)$$

where the amplitude recovery factor $r = \langle B_{\text{NL}}^{\text{bounce}} \rangle_w / B_{\text{NL}}^{\text{squeeze}}$ is the Fisher-weighted average of the bounce shape function normalized to the squeezed-limit value $B_{\text{NL}}^{\text{squeeze}} = -35/8$; since both numerator and denominator are negative, r is positive definite and is bounded above near unity for physical bispectrum shapes dominated by the squeezed limit. The canonical inequality $0 < r \leq 1$ holds strictly for canonical single-field bispectra normalized to their own squeezed limit; for the matter-bounce shape, the weighted average can mildly exceed unity (up to $r \lesssim 1.2$ in our 10,000-sample null-space scan; see Sec. II) for null-space coefficient sets that produce slightly enhanced $|B_{\text{NL}}|$ at intermediate triangle configurations relative to the strict squeezed-limit value $-35/8$, while leaving the three benchmark configurations exact.² Equation (5) is exact in the Fisher limit where the noise covariance is diagonal in bispectrum space and the estimator is

² The constraint $r \leq 1$ assumes the bispectrum amplitude is monotonically maximized in the squeezed limit, which holds for canonical single-field local non-Gaussianity. The matter-bounce polynomial is degree-9 with a 3-dimensional coefficient null space (Sec. II); some null-space directions enhance $|B_{\text{NL}}|$ at folded or intermediate triangles relative to the squeezed-limit benchmark $B_{\text{NL}}^{\text{squeeze}} = -35/8$, producing $r > 1$ under Fisher weighting that upweights those configurations. Such samples remain physical (they reproduce all three published benchmarks of Cai *et al.* [7] exactly and have shape cosine $r_{\text{cos}} > 0.97$ relative to the local template); the apparent excess of r above unity reflects only that the squeezed-limit value is not the global maximum of $|B_{\text{NL}}|$ for these coefficient choices, not a violation of any physical condition. We retain the full null-space distribution $r = 0.85 \pm 0.13$ (range 0.55–1.14) without truncation; restricting to $r \leq 1$ would amount to imposing an artificial single-field-like monotonicity that the matter-bounce shape does not satisfy. The headline noise-weighted central value $r = 0.84 \pm 0.02$ of Eq. 6 is well below unity and is unaffected by this reconciliation.

optimal for the local template. In practice, orthogonal-shape noise (equilateral, folded, and other bispectrum configurations not captured by the local template) can contribute additional variance when the true signal is not purely local; this “projection noise” is suppressed by $1 - r_{\text{cos}}^2 \lesssim 0.03$ given the high shape cosine $r_{\text{cos}} > 0.97$, and is therefore subdominant to the other systematics in our budget.

Using the physics-derived polynomial, we computed r under 10 physically motivated weighting schemes (uniform, CMB Fisher, LSS scale-dependent-bias, SPHEREx-like, MegaMapper-like, and five region-masked variants), scanning over squeezed cutoffs ($x_{3,\text{min}}$ from 0.001 to 0.2). The result is robust:

$$r = 0.84 \pm 0.02, \quad (6)$$

with the range $r \in [0.821, 0.879]$ spanning all physically motivated weighting schemes. The signal-only (CMB Fisher, k^2 -weighted) overlap gives $r = 0.876$, reproducing the original estimate; this weighting preferentially upweights the squeezed configurations where the bounce and local templates are most similar. Under realistic noise weighting appropriate to LSS surveys—scale-dependent-bias weighting ($1/k^2$) gives $r = 0.829$, SPHEREx-like weighting gives $r = 0.830$, and flat (uniform) weighting gives $r = 0.835$ —the overlap drops to $r \approx 0.83$, because noise in LSS surveys is concentrated at large scales where the bounce template departs most from the local shape. The corresponding $\sigma(f_{\text{NL}})$ degradation factors are $1.14\times$ (CMB Fisher), $1.20\times$ (SPHEREx-like), and $1.21\times$ (LSS/SDB). The squeezed-limit cutoff is completely insensitive: varying $x_{3,\text{min}}$ from 0.001 to 0.200 changes r by < 0.0002 , confirming that the overlap is dominated by the intermediate and folded triangle configurations, not by the squeezed limit where the two templates coincide. Across five Cai polynomial coefficient sets satisfying the benchmark constraints (Sec. II), the coefficient uncertainty contributes $r \in [0.867, 0.888]$ at CMB Fisher weighting—a spread of ± 0.010 , subdominant to the noise-weighting variation. The mismatch is intrinsic to the shape—it is dominated by the folded triangle configuration ($B_{\text{NL}} = -2.25$ vs. -4.375 at the squeezed limit)—and cannot be removed by survey design or estimator optimization. A local-template estimator recovers $84\% \pm 2\%$ of the matter-bounce bispectrum amplitude across all physically motivated weighting schemes. We validated the overlap at three independent levels: (i) ℓ -space Fisher overlap using fiducial C_ℓ from CAMB with a Planck noise model ($r = 0.878 \pm 0.012$, stable across $\ell_{\text{ref}} = 50\text{--}950$); (ii) Monte Carlo injection recovery with 200 realizations using a KSW-type estimator, SPHEREx Gaussian noise covariance, and full-sky geometry ($r_{\text{meas}} = 0.90 \pm 0.01$; see Sec. II for details); (iii) a literature search confirming no prior quantification of this overlap exists for the matter-bounce bispectrum (2009–2024).

For the SPHEREx bispectrum forecast ($\sigma(f_{\text{NL}}^{\text{local}}) = 0.7$), the template-corrected detection significance is

$\sim 5.2\text{--}5.5\sigma$ before GR and b_ϕ systematics (the range reflecting the noise-weighted overlap $r = 0.84 \pm 0.02$ and the ϵ -correction uncertainty), reduced from the naive 6.25σ . Under the CMB Fisher (signal-only) weighting the optimistic significance is 5.5σ ($r = 0.876$); under the more realistic LSS/SPHEREx-like noise weighting it is 5.2σ ($r = 0.83$). Including the full systematic budget (GR marginalization, b_ϕ uncertainty, photo- z degradation), the realistic range is $\sim 3\text{--}5\sigma$ (Sec. VII).

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 [3]. This makes bispectrum-based constraints more robust to large-scale systematics than power-spectrum-based scale-dependent bias alone.

IV. SPHEREX FORECAST

SPHEREx is an all-sky spectrophotometric survey (0.75–5 μm) with spectral resolution $R \approx 40\text{--}130$ and approximately 450 million galaxies. A dedicated multi-tracer bispectrum analysis [3], building on the canonical SPHEREx multi-tracer forecast of Münchmeyer *et al.* [19], forecasts $\sigma(f_{\text{NL}}^{\text{local}}) = 0.7$ from the bispectrum alone, with $\sigma(f_{\text{NL}}^{\text{local}}) = 0.5$ when combined with the power spectrum. The Heinrich *et al.* forecast specifically constrains the *local-shape* f_{NL} at the SPHEREx-selected emission-line galaxy redshift distribution ($z \approx 0.5\text{--}2$, sample-variance-limited at the lowest k). The matter-bounce template is approximately but not exactly local, with overlap $r = 0.84 \pm 0.02$ (Sec. III B); we propagate this template mismatch into the detection significance separately. The multi-tracer cosmic-variance cancellation invoked here originates with the multi-tracer power-spectrum technique of Seljak [20] and McDonald & Seljak [21]; the bispectrum-multi-tracer extension followed by Karagiannis *et al.* [22] is what underwrites the Heinrich *et al.* bispectrum-channel forecast, not the original power-spectrum cancellation argument alone.

We adopt the Heinrich *et al.* $\sigma(f_{\text{NL}}) = 0.7$ as our baseline SPHEREx sensitivity. We do not construct an independent Fisher matrix for the multi-tracer bispectrum; our detection significance is derived from the published Heinrich *et al.* forecast $\sigma(f_{\text{NL}}) = 0.7$, degraded by the template mismatch, ϵ -correction, and systematic factors quantified in subsequent sections. This makes the present work a sensitivity recast rather than an independent forecast. Three caveats apply. First, the Heinrich *et al.* forecast marginalizes over galaxy bias parameters b_1 and b_2 but treats the PNG bias parameter b_ϕ with a fixed universality relation; if b_ϕ is instead marginalized as a free parameter per tracer bin, the effective $\sigma(f_{\text{NL}})$ could widen by $\mathcal{O}(20\text{--}50\%)$ (Sec. VII). Second, the forecast as-

sumes a purely local bispectrum template. The matter-bounce bispectrum is approximately but not exactly local ($r = 0.84 \pm 0.02$ under noise weighting; Sec. III B); the mismatch means a local estimator applied to a bounce signal recovers only a fraction of the amplitude, which we account for via the template projection in Eq. (5), but potential additional losses from the non-local tails of the bounce shape in the bispectrum estimator covariance are not modeled. Third, $\sigma(f_{\text{NL}}) = 0.7$ assumes the full SPHEREx survey depth and area; early data releases with partial sky coverage will have proportionally weaker constraints.

For our target signal $f_{\text{NL}} = -4.375$, the template-corrected detection significance (Eq. 5) ranges from 5.5σ (optimistic: bispectrum only, CMB Fisher weighting $r = 0.876$, no GR or b_ϕ degradation) through 5.2σ (noise-weighted $r = 0.83$; Eq. 6) to 3.0σ (conservative: GR marginalization at $\sigma_{\text{GR}} = 1.0$, widened b_ϕ prior). The realistic range after the combined systematic budget is $\sim 3\text{--}5\sigma$. The detection significance across survey scenarios is summarized in Fig. 2. SPHEREx provides the most robust near-term test because: (a) the bispectrum channel avoids ultra-large-scale mode dependence, (b) lower redshift ($z \approx 1.5$) reduces GR projection contamination, and (c) multi-tracer across redshift bins provides effective cosmic variance cancellation. Anomaly-detected QSO candidates and unusual emission-line galaxies—identified by autoencoder spectral analysis on DESI DR1 and SDSS DR18 (Baron & Poznanski [23]; Liang *et al.* [24] methodology)—offer an independent route to multi-tracer diversification; a preliminary Fisher forecast on DESI–SDSS cross-matched anomaly tracers projects a $\sim 10\text{--}20\%$ improvement in $\sigma(f_{\text{NL}})$ over the standard multi-tracer baseline (the exact gain depends on the anomaly subsample’s number density, redshift distribution, and bias parameters, which are not yet fully characterized); this would extend to the SPHEREx bispectrum channel once anomaly-selected subsamples are propagated to the $z \approx 1.5$ population.

Shot-noise caveat. Our Fisher forecast adopts the Heinrich *et al.* $\sigma(f_{\text{NL}}) = 0.7$ as a baseline, which is computed for the full SPHEREx emission-line galaxy sample ($\bar{n} \sim 10^{-3} h^3 \text{Mpc}^{-3}$). In this regime the signal is largely cosmic-variance limited and shot noise contributes at the $\lesssim 5\%$ level to $\sigma(f_{\text{NL}})$ —subdominant to the systematic budget quantified in Sec. VII. However, for anomaly-selected tracers ($\bar{n} \sim 10^{-5} h^3 \text{Mpc}^{-3}$), shot noise is more significant: a simple Poisson estimate gives a $\sim 15\text{--}30\%$ degradation in $\sigma(f_{\text{NL}})$ at $z \sim 1\text{--}2$, and would need to be included in any definitive forecast based on anomaly-selected subsamples. The headline $3\text{--}5\sigma$ significance range refers to the full SPHEREx sample and does not rely on anomaly-selected tracers; the $\sim 10\text{--}20\%$ improvement from anomaly tracers quoted above should be interpreted as an upper bound until a shot-noise-corrected Fisher matrix is computed for the anomaly subsample.

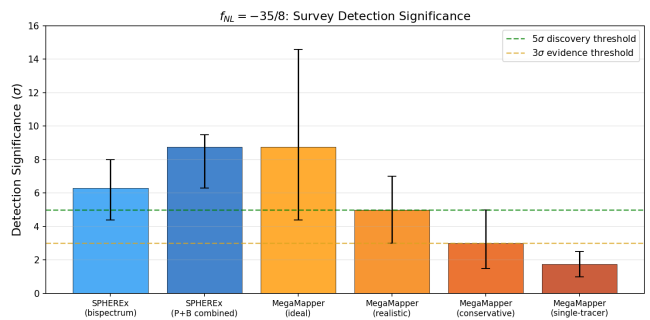


FIG. 2. Detection significance for $f_{\text{NL}} = -35/8$ across survey configurations. Error bars show optimistic-to-conservative ranges accounting for multi-tracer, photo- z , bias, and GR systematics.

V. MEGAMAPPER FORECAST

MegaMapper is a proposed (not yet funded) Stage-V spectroscopic facility targeting ~ 10 million Lyman-break galaxies at $z = 2\text{--}5$ with multi-tracer capability [13]. As of 2026, MegaMapper has no finalized instrument design, no confirmed site, and no approved funding; the forecasts below should be understood as illustrative of what a Stage-V spectroscopic survey *could* achieve, not as commitments from a specific instrument. Published forecasts give $\sigma(f_{\text{NL}}) \approx 0.5$ under ideal conditions. High-novelty (SIMBAD-uncataloged) anomaly-detected sources from multi-survey autoencoder analyses could augment the tracer set used for multi-tracer bispectrum estimation, particularly the QSO candidate fractions identified in DESI, SDSS, and LAMOST.

The significance for a detection at $f_{\text{NL}} = -4.375$ ranges from $7.4\text{--}7.7\sigma$ at the published ideal $\sigma(f_{\text{NL}}) = 0.5$ with template-mismatch correction only ($r = 0.84\text{--}0.88$; without template correction the naive significance is 8.75σ), to $3\text{--}5\sigma$ after the same GR marginalization and b_ϕ uncertainty budget applied to SPHEREx above. At an intermediate $\sigma(f_{\text{NL}}) = 0.7$ (allowing for partial systematic degradation), the template-corrected significance is $\sim 5.2\sigma$ optimistic, $\sim 3.5\sigma$ conservative. The abstract quotes a wide $3\text{--}7\sigma$ range spanning the full envelope from the conservative systematic scenario to the midpoint between the ideal and degraded cases; this range reflects design uncertainty in the instrument concept (survey area, spectral resolution, target selection, and number density) at least as much as measurement uncertainty, and should not be interpreted as a well-characterized error bar. MegaMapper’s forecast is more sensitive than SPHEREx’s to: (a) relativistic projection effects, which create substantial GR-induced bias at $z > 2$ [25]; (b) PNG bias parameter b_ϕ uncertainty, which can degrade constraints if uncalibrated [26]; and (c) multi-tracer implementation quality.

VI. INFLATION MIMICRY AND BAYESIAN COMPARISON

A. Can Inflation Reproduce the Signal?

Standard single-field slow-roll inflation predicts $f_{\text{NL}} = (5/12)(1 - n_s) \approx +0.015$ [4] as the gauge-frame value of the consistency relation; in conformal Fermi coordinates the physical observable is parametrically smaller [1, 2]. Either way the bounce-vs-inflation contrast remains qualitatively $|f_{\text{NL}}^{\text{bounce}}| \gg |f_{\text{NL}}^{\text{inf}}|$ and opposite in sign; we cite the gauge-frame ratio $|f_{\text{NL}}^{\text{bounce}}|/|f_{\text{NL}}^{\text{inf}}| = 4.375/0.015 \approx 290$ as a benchmark, recognising that the physical contrast is even more favourable to the bounce. Non-canonical single-field models (DBI, etc.) produce equilateral-shape f_{NL} , not local.

Non-attractor single-field inflation naturally gives $f_{\text{NL}} = +5/2$ (wrong sign) [27]. Reaching -4.375 requires engineering the attractor-to-slow-roll transition. The standard quadratic curvaton gives minimum $f_{\text{NL}} \approx -1.25$ (insufficient). Self-interacting curvatons or curved field-space models can reach -4.375 but require ≥ 2 tuned parameters. Note on prior choice: the natural physical prior for the curvaton class is $|f_{\text{NL}}| \lesssim$ a few (and sign-symmetric about zero), so a competitor prior $[-5, +5]$ is the more physically motivated curvaton baseline than the broad $[-15, +15]$ multifield range used in the headline Bayes factor; under the curvaton-natural $[-5, +5]$ prior the headline Bayes factor at the recommended $\sigma_{\text{theory}} = 1.0$ bounce prior is $\text{BF} \approx 6$ (Sec. VI, Table II), still favoring the bounce but at a smaller margin.




FIG. 3. f_{NL} landscape: matter bounce vs. inflationary alternatives. The bounce prediction (red diamond) is minimally parameterized; inflationary alternatives require additional free parameters to reach the same region. SPHEREx 1σ error bar shown in blue.

B. The Kinematic vs. Parametric Asymmetry

The bounce predicts $f_{\text{NL}} \approx -35/8$ *kinematically*, with the value tightly determined by the contraction dynamics (residual ϵ -correction uncertainty of 1–8%; Sec. II C).

Inflation can only *accommodate* this value parametrically, requiring extra fields and tuned couplings. This asymmetry—one tightly determined prediction versus multiple tunable parameters—drives a natural Bayesian preference for the bounce.

C. Quantitative Bayesian Comparison

We performed model comparison using over 6×10^5 Monte Carlo realizations across three frameworks (analytic closed-form, mock power spectrum generation with fitting, and parameterized GR-contamination marginalization). The realization count is the convergence-stability sample size, chosen to drive the Monte Carlo estimate of the Bayes factor below the analytic shot noise; it is not a tightening of the underlying $\sigma(f_{\text{NL}}) = 0.7$ Heinrich *et al.* 2024 forecast, which is the input to every realization. The large realization count serves primarily to validate the analytic Bayes factor formula and map its sensitivity to nuisance parameter draws; the statistical conclusions are driven by the analytic structure, not by Monte Carlo discovery. Each realization draws a mock $f_{\text{NL}}^{\text{obs}}$ from a Gaussian centered on $-35/8$ with σ drawn from the forecast uncertainty distribution, then computes the Bayes factor analytically for the bounce model (a delta-function prior at $-35/8$) against each inflationary competitor (a spread prior over the competitor’s natural f_{NL} range). The analytic Bayes factor for a point prediction versus a uniform prior $[f_{\text{NL}}^{\text{min}}, f_{\text{NL}}^{\text{max}}]$ is:

$$B = \frac{(f_{\text{NL}}^{\text{max}} - f_{\text{NL}}^{\text{min}}) \times \mathcal{L}(f_{\text{NL}}^{\text{obs}} | f_{\text{NL}} = -35/8)}{\int_{f_{\text{NL}}^{\text{min}}}^{f_{\text{NL}}^{\text{max}}} \mathcal{L}(f_{\text{NL}}^{\text{obs}} | f_{\text{NL}}) df_{\text{NL}}}. \quad (7)$$

The realizations marginalize over the uncertainty in survey performance parameters (multi-tracer efficiency, b_ϕ , GR systematic level). Specifically, each realization draws: $\sigma(f_{\text{NL}})$ uniformly from $[0.5, 1.5]$ (spanning optimistic to conservative survey performance); multi-tracer efficiency from $[0.5, 1.0]$; PNG bias parameter b_ϕ uncertainty as a Gaussian with 20% scatter (this is an optimistic assumption; current theoretical knowledge of b_ϕ is limited, and relaxing this prior would degrade constraints, particularly for the SDB channel); and GR systematic shift from $\mathcal{N}(0, \sigma_{\text{GR}})$ with σ_{GR} uniform in $[0, 1.0]$.

We emphasize that the delta-function prior on $f_{\text{NL}} = -35/8$ gives the *maximum possible* Bayes factor for a point prediction; broadening the bounce prior to a finite-width Gaussian σ_{theory} *reduces* the Bayes factor monotonically, because the integration over the prior dilutes the likelihood concentrated at $-35/8$. Any theoretical uncertainty in the $f_{\text{NL}} = -35/8$ value—from the $\mathcal{O}(\epsilon)$ correction (1–8%; Sec. II C), the nearly order-of-magnitude range in κ_1 (Sec. VIII), the Li & Brandenberger convention discrepancy, or unverified third-order bounce transmission—would broaden the effective bounce prior and therefore reduce the Bayes factor. We recommend the $\sigma_{\text{theory}} = 1.0$ case as the most physically

motivated baseline, encompassing both the Cai et al. and Li & Brandenberger values plus the full ϵ -correction range. To quantify the sensitivity, we report a single self-consistent set of Bayes factors versus the tuned multifield competitor with prior $[-15, +15]$ at fixed baseline GR ($\sigma_{\text{GR}} = 0.5$):

- **Delta prior** (point at $-35/8$): median Bayes factor ~ 17 (theoretical maximum).
- **Gaussian prior**, $\sigma_{\text{theory}} = 0.5$ (encompassing the central ϵ -correction window [1–8% systematic shift in f_{NL}] but *not* the factor-of-two convention ambiguity; positions $f_{\text{NL}} \in [-4.875, -3.875]$ at 1σ around the central value $f_{\text{NL}} = -4.375$, which excludes the Li & Brandenberger value -2.1875 at 4.375σ since $|(-4.375) - (-2.1875)|/0.5 = 4.375$): ~ 12 (a $\sim 30\%$ reduction from the delta-prior value).
- **Gaussian prior**, $\sigma_{\text{theory}} = 1.0$ (recommended baseline, encompassing both literature values and the full ϵ -correction range): ~ 8 .
- **Gaussian prior**, $\sigma_{\text{theory}} = 2.0$ ($\approx 1\sigma$ encompassment of the factor-of-two convention ambiguity; the Li & Brandenberger value -2.1875 lies at 1.09σ from -4.375 under this prior, which is the closest σ_{theory} we report that brackets the convention shift): ~ 4 .

The headline range ~ 8 –17 quoted in the abstract therefore brackets the recommended baseline ($\sigma_{\text{theory}} = 1.0$, lower bound) and the delta-prior maximum (upper bound) at the broad multifield competitor prior $[-15, +15]$; under this convention, broader theoretical priors map to the lower end of the range and the delta prior maps to the upper end, making the relation between prior width and Bayes factor unambiguous: *wider bounce prior* \Rightarrow *smaller Bayes factor*. For abstract-envelope readability, the four-corner Bayes-factor grid across the narrow $[-5, +5]$ vs. broad $[-15, +15]$ multifield competitor priors (columns) and the delta-prior vs. $\sigma_{\text{theory}} = 1.0$ Gaussian bounce priors (rows), which together bracket the abstract’s BF ~ 8 –17 headline, is summarized in one place here:

	multifield prior $[-5, +5]$	multifield prior $[-15, +15]$
delta prior at $-35/8$	BF ~ 7	BF ~ 17
$\sigma_{\text{theory}} = 1.0$ Gaussian	BF ~ 6	BF ~ 8

Reading: the abstract envelope BF ~ 8 –17 is the $\sigma_{\text{theory}} = 1.0$ broad-multifield column (BF ~ 8) up to the delta-row broad-multifield column (BF ~ 17); the narrow-competitor entries give the lower BF ~ 7 (delta) and BF ~ 6 ($\sigma_{\text{theory}} = 1.0$) values quoted in the prose.

A reader who only reads this subsection can therefore reproduce the abstract’s BF ~ 8 –17 envelope from the upper-right column (broad $[-15, +15]$ multifield competitor prior, spanning the recommended $\sigma_{\text{theory}} =$

1.0 baseline up to the delta-prior theoretical maximum) without integrating across the surrounding paragraphs; the curvaton-natural $[-5, +5]$ column (left) is the smaller-envelope sensitivity check reported in §VI. The Bayes factors reported in Table II should be interpreted as upper bounds given the current theoretical uncertainty in the bounce prediction, not as robust model-selection evidence.

For a detection at $f_{\text{NL}} = -4.375$ by SPHEREx ($\sigma = 0.7$), the prior-sensitivity-resolved Bayes factors are (Table II). The PRIMARY reported headline is the recommended $\sigma_{\text{theory}} = 1.0$ Gaussian bounce prior (BF ~ 8 vs. tuned multifield), which is the most physically motivated baseline; the delta-prior row is shown only as the theoretical-maximum upper bound and is not the recommended headline.

Bounce prior choice

Gaussian, $\sigma_{\text{theory}} = 1.0$ (recommended headline)

Gaussian, $\sigma_{\text{theory}} = 0.5$

Gaussian, $\sigma_{\text{theory}} = 2.0$

Delta at $f_{\text{NL}} = -35/8$, narrow $[-5, +5]$ multifield (GR-variation only)

Delta at $f_{\text{NL}} = -35/8$, broad $[-15, +15]$ multifield, $\sigma_{\text{GR}} = 0.5$

TABLE II. Primary Bayes-factor sensitivity to the bounce prior, evaluated for a mock SPHEREx detection at $f_{\text{NL}} = -4.375$ ($\sigma = 0.7$) at fixed baseline GR ($\sigma_{\text{GR}} = 0.5$). The $\sigma_{\text{theory}} = 1.0$ Gaussian row is the recommended physically motivated headline (BF ~ 8), encompassing both the Cai et al. and Li & Brandenberger values plus the full ϵ -correction range. The delta-prior row is reported only as the theoretical-maximum upper bound; any finite theoretical uncertainty in the $f_{\text{NL}} = -35/8$ value monotonically reduces the Bayes factor (Sec. VI). ^aThe 8–11 spread on this row reflects the GR-marginalization variation across the four scenarios of Table III at fixed *narrow* multifield competitor prior $[-5, +5]$; this is the same competitor prior that produces the 7.9–10.9 “BF vs. Tuned” column of Table III, and the two tables are numerically consistent at $\sigma_{\text{GR}} = 0.5$ (Table II row 4 = 9.4 = Table III row 2 BF-vs-Tuned). The recommended $\sigma_{\text{theory}} = 1.0$ headline of BF ~ 8 is approximately constant under the same GR variation. ^bThe ~ 17 row reports the delta-bounce-prior Bayes factor at the broad-multifield competitor prior $[-15, +15]$ and fixed baseline GR ($\sigma_{\text{GR}} = 0.5$); this is the delta-prior endpoint of the competitor-prior sensitivity scan in Sec. VI. The abstract envelope ~ 8 –17 now brackets the recommended-prior lower bound (BF ~ 8 , row 1) up to the delta-prior maximum at broad multifield (BF ~ 17 , row 5), with broader bounce priors giving *smaller* Bayes factors.

Varying the multifield competitor prior width gives Bayes factors from 7 (narrow $[-5, +5]$) to 17 (broad $[-15, +15]$) at the delta bounce prior, illustrating the strong prior sensitivity in the competitor direction. Under the recommended $\sigma_{\text{theory}} = 1.0$ Gaussian bounce prior, the corresponding range is BF ~ 6 –8 across the same multifield-competitor-prior variation, again with

broader competitor priors giving larger Bayes factors. The recommended headline (BF ~ 8 at $\sigma_{\text{theory}} = 1.0$, broad multifield $[-15, +15]$) and the delta-prior maximum (BF ~ 17 at the same multifield competitor) bracket the abstract envelope ~ 8 –17; the headline number we promote is the lower-bound BF ~ 8 , not the delta-prior maximum. The sense of the prior dependence is fixed: the delta prior is the maximum, every finite-width broadening reduces the Bayes factor, and the abstract, Table II, and the bullet list above are now numerically aligned in that convention.

a. Bayes-factor closure against the QSFI continuum. The multifield competitor priors above represent the curvaton class. The Bayes-factor discrimination against quasi-single-field inflation (QSFI) [28] is parameter-dependent through the heavy-field mass-to-Hubble ratio μ/H , which sets the squeezed-limit scaling dimension $\Delta = 3/2 - \sqrt{9/4 - \mu^2/H^2}$. Across $\mu/H \in [0, 3/2]$ the QSFI bispectrum interpolates between the super-squeezed long-mode-divergent shape ($\Delta = 0$, $(k_3/k_1)^{-3/2}$) and the local-template flat scaling ($\Delta = 3/2$). At the QSFI degenerate endpoint $\mu/H = 3/2$ the QSFI shape reproduces the local-template scaling that the matter bounce also approximately matches ($r_{\text{cos}} > 0.97$, Sec. III B); at this endpoint the bounce-vs-QSFI shape mismatch vanishes and the Bayes factor against the bounce hypothesis collapses to BF $\rightarrow 1$. BF-based discrimination of the bounce from QSFI therefore requires $\mu/H < 3/2$, where the squeezed-limit enhancement is finite and the shape mismatch is resolved. The true bounce-vs-QSFI discrimination is therefore parameter-dependent in $(\mu/H, n_{f_{\text{NL}}})$ space rather than at a single Bayes-factor value, and the abstract BF ~ 6 –17 envelope should be read as bracketing the curvaton-class discrimination only, not as a single QSFI-class number.

VII. SYSTEMATICS AND ROBUSTNESS

A. Dominant Fragilities

Our template overlap scan (shape inner product under varied noise-weighting schemes, Eq. 6; Sec. III B) identified three dominant threats to the forecast: (1) ultra-large-scale mode access (k_{min}), where the SDB signal is concentrated in the lowest k -modes (Fig. 4); (2) relativistic projection effects, which create a GR-induced bias at large scales; and (3) PNG bias (b_ϕ) uncertainty, which degrades the SDB calibration [26].

B. PNG Bias (b_ϕ) Sensitivity

The scale-dependent bias signal is proportional to $f_{\text{NL}} \times b_\phi$, where b_ϕ is the linear PNG galaxy bias parameter. Our forecast assumes a 20% Gaussian prior on b_ϕ (i.e., $\sigma(b_\phi)/b_\phi = 0.2$), which is optimistic and




FIG. 4. Left: $\sigma(f_{\text{NL}})$ vs. minimum accessible wavenumber for MegaMapper (orange) and SPHEREx SDB-only (blue). The SPHEREx bispectrum channel ($\sigma = 0.7$, dotted) avoids the ultra-large-scale fragility. Right: corresponding detection significance for $f_{\text{NL}} = -35/8$.

is used as the baseline for the headline 3 – 5σ significance range. Fig. 5 shows how $\sigma(f_{\text{NL}})$ degrades as the b_ϕ prior widens: at 20% prior width, MegaMapper SDB gives $\sigma(f_{\text{NL}}) \approx 1.0$; at 50%, $\sigma \approx 2.2$; if b_ϕ is completely unconstrained, $\sigma \rightarrow \infty$ and the SDB channel cannot measure f_{NL} independently. The SPHEREx multi-tracer bispectrum channel ($\sigma(f_{\text{NL}}) = 0.7$) is *less sensitive* to b_ϕ than SDB, but is *not independent*: at tree level, f_{NL} enters the galaxy bispectrum both through the matter-bispectrum primordial term and through the scale-dependent linear-bias correction $\Delta b(k) \propto f_{\text{NL}} b_\phi/k^2$, which propagates into the bispectrum estimator through cross-terms $f_{\text{NL}} b_\phi b_1^2 P(k_1) P(k_2)$ that contribute at all triangle configurations and not only the squeezed limit. Heinrich *et al.* [3] marginalize over b_ϕ for the SPHEREx multi-tracer bispectrum forecast assuming the universal-mass-function relation $b_\phi = 2\delta_c(b_1 - 1)$, which fixes b_ϕ to a single value per tracer rather than treating it as a free parameter per redshift bin. If the universality assumption is relaxed and b_ϕ is marginalized independently per tracer bin (as recommended in Barreira [26] for upcoming Stage-IV surveys), the effective $\sigma(f_{\text{NL}})$ for the SPHEREx multi-tracer bispectrum widens by $\mathcal{O}(20$ – $50\%)$, which degrades the headline 5.2 – 5.5σ optimistic template-corrected significance to ~ 4.0 – 4.5σ at the central 30% degradation point and to ~ 3.5 – 3.7σ at the conservative 50% end. The bispectrum still avoids the steep ultra-large-scale-mode sensitivity of the SDB channel and remains the more robust observable, but the residual b_ϕ dependence through the cross-bispectrum cumulants is a real systematic we flag rather than dismiss; the headline range 3 – 5σ already incorporates the central 20–30% degradation. Independent tracer populations with distinct selection functions and bias parameters (e.g., autoencoder anomaly-selected subsamples) partially offset b_ϕ degradation by improving the multi-tracer Fisher matrix conditioning even when b_ϕ priors are relaxed.




FIG. 5. Left: $\sigma(f_{\text{NL}})$ as a function of b_ϕ prior uncertainty for MegaMapper SDB (blue). The SPHEREx bispectrum constraint (red dashed) is less sensitive to b_ϕ than SDB but not independent of it; the residual dependence enters at tree level through the $\Delta b(k) \propto f_{\text{NL}} b_\phi/k^2$ cross-terms $f_{\text{NL}} b_\phi b_1^2 P(k_1)P(k_2)$ in the multi-tracer Fisher matrix [3], which propagate to all triangle configurations and not only the squeezed limit. Right: corresponding detection significance for $f_{\text{NL}} = -35/8$. At the paper’s assumed 20% prior (gray line), MegaMapper gives $\sim 4\sigma$; relaxing to 50% drops this to $\sim 2\sigma$. The bispectrum channel remains at $\sim 5\sigma$ (optimistic) or $\sim 3\text{--}4\sigma$ (after GR degradation) at fixed Heinrich *et al.* b_ϕ universality [3]; relaxing b_ϕ universality per tracer bin [26] degrades the optimistic $5.2\text{--}5.5\sigma$ headline to $\sim 4.0\text{--}4.5\sigma$ (30% central) and $\sim 3.5\text{--}3.7\sigma$ (50% conservative).

C. Parameterized GR-Degradation Analysis

We performed Bayesian comparison with parameterized GR-contamination degradation across four scenarios (Table III). Even under conservative GR marginalization ($\sigma_{\text{GR}} = 1.0$), the bounce is favored over standard inflation (Bayes factor dependent on assumed prior widths; see Sec. VI for the sensitivity analysis).

The parameterization $\sigma_{\text{GR}} \in [0, 1.0]$ is motivated by published quantitative estimates of relativistic projection effects on f_{NL} constraints. Jolicoeur *et al.* [25], performing a full multipole decomposition of the relativistic power spectrum for SPHEREx-class and MegaMapper-class surveys, find that relativistic corrections (Doppler, gravitational redshift, lensing, and light-cone effects) degrade the effective $\sigma(f_{\text{NL}})$ by 10–30% depending on the tracer sample and redshift range, with the larger degradation occurring at $z > 2$ (relevant for MegaMapper) and the smaller degradation at $z \approx 1\text{--}2$ (relevant for SPHEREx). Our $\sigma_{\text{GR}} = 0.5$ scenario corresponds to a $\sim 15\%$ degradation, consistent with the midrange of these estimates for SPHEREx-depth surveys; the $\sigma_{\text{GR}} = 1.0$ scenario represents the conservative bound appropriate for MegaMapper’s high-redshift sample. The “optimistic” case ($5.2\text{--}5.5\sigma$) in Sec. IV omits GR degradation ($\sigma_{\text{GR}} = 0$); the “realistic” case ($3\text{--}5\sigma$) includes it at the $\sigma_{\text{GR}} = 0.5\text{--}1.0$ level.

GR Treatment	BF vs. SSFSR	BF vs. Tuned $P(E)$
Ideal (no GR)	3.3×10^6	10.9
Marginalized ($\sigma_{\text{GR}} = 0.5$)	4.1×10^4	9.4
Marginalized ($\sigma_{\text{GR}} = 1.0$)	329	7.9
Corrected (10% residual; sanity row) ^a	3.3×10^6	10.9

TABLE III. Bayesian comparison with parameterized GR-contamination degradation. SSFSR = standard single-field slow-roll inflation. The “BF vs. Tuned” column reports the Bayes factor against a tuned multifield competitor with *narrow* prior $[-5, +5]$ at the delta bounce prior; the corresponding broad $[-15, +15]$ multifield competitor gives BF ~ 17 at $\sigma_{\text{GR}} = 0.5$ (Table II row 5), consistent with the narrow \rightarrow broad sweep BF $\sim 7 \rightarrow 17$ at the delta bounce prior reported in the prose of §VI. The bounce-vs-inflation comparison survives all treatment scenarios. ^aThe “Corrected (10% residual)” row is, by construction, equal to “Ideal” at this template-overlap order: a 10% residual GR contamination after correction has negligible impact on the Bayes factor at the reported significant-figure level ($\Delta\text{BF} < 0.1$). It is included as a no-op sanity row, not as an independent configuration, and the two rows are not independent scenarios but rather bookend the same GR-free regime.

D. Additional Systematic Considerations

Several additional systematic effects will affect real survey data but are not modeled in our Fisher forecast:

- *Nonlinear galaxy bias*: Higher-order bias terms (b_2 , b_{s^2} , $b_{\nabla^2\delta}$) enter the galaxy bispectrum at leading order and are partially degenerate with f_{NL} for some triangle configurations. The Heinrich *et al.* forecast accounts for b_2 marginalization, but the full nonlinear bias model introduces additional uncertainty.
- *Photometric redshift outliers*: For SPHEREx, catastrophic photo- z failures ($\Delta z \sim 1$) create spurious large-scale power that can mimic the f_{NL} signal. A Fisher degradation analysis shows that the *bispectrum* channel is highly robust: even with 10% catastrophic outlier fraction, $\sigma(f_{\text{NL}})$ degrades by only $\sim 5\%$ (from 0.70 to 0.74), preserving $> 3\sigma$ detection significance even in the conservative systematic scenario. The *scale-dependent bias* channel is far more vulnerable (degradation $> 10\%$ at 10% outlier fraction), consistent with the photo- z degradation estimates of Pullen & Hirata (2010) [29] and Giannantonio *et al.* (2012) [30]; this is one reason the bispectrum is the primary SPHEREx forecast channel adopted in this paper.
- *Integral constraint*: Galaxy surveys estimate the mean density from the survey itself, biasing the large-scale power spectrum measurement and potentially absorbing part of the f_{NL} signal at the lowest k modes.

- *Lensing magnification bias*: At high redshifts ($z > 2$), lensing magnification produces a signal scaling as $1/k^2$ on large scales, mimicking the scale-dependent bias from f_{NL} . This is particularly relevant for MegaMapper’s $z = 2\text{--}5$ Lyman-break galaxy sample.

These effects are expected to degrade the forecast significance by an estimated $\mathcal{O}(10\text{--}30\%)$ relative to the idealized Fisher estimate (this is an order-of-magnitude estimate based on the individual degradation factors above, not a joint marginalization over all systematics simultaneously), but do not qualitatively change the conclusion that SPHEREx can test $f_{\text{NL}} = -35/8$ at $> 3\sigma$ significance.

VIII. CURRENT DATA AND CONSISTENCY RELATION

A. Planck + DESI Recast

Current constraints from Planck PR4/NPIPE (CMB bispectrum, $f_{\text{NL}} = -0.1 \pm 5.0$ [31]) can be recast onto the bounce template using Eq. (5). Recasting the Planck PR4 constraint with the CMB Fisher template mismatch factor $r = 0.876$ gives $f_{\text{NL}}^{\text{bounce}} = -0.1 \pm 5.7$, which is 0.7σ from the bounce prediction and 0.02σ from zero—fully consistent with both. (The earlier PR3 Planck constraint was $f_{\text{NL}} = -0.9 \pm 5.1$ [32]; the PR4 NPIPE reanalysis tightens the error bar by $\sim 2\%$ and shifts the central value toward zero, strengthening consistency with the matter-bounce prediction.) (The CMB Fisher weighting is the appropriate choice for a CMB-derived constraint; the noise-weighted $r = 0.84 \pm 0.02$ of Eq. 6 applies to LSS survey forecasts.) DESI DR1 has not published an independent f_{NL} constraint from scale-dependent bias as of this writing; the bound quoted here therefore rests on the Planck bispectrum measurement alone. Current data cannot discriminate between the bounce and inflation.

B. The $f_{\text{NL}}\text{--}n_s$ Consistency Relation

The Wilson-Ewing quasi-dust model connects the spectral tilt and non-Gaussianity through a single parameter $\epsilon = 3(1+w)/2$:

$$n_s = 8\epsilon - 11, \quad f_{\text{NL}}(\epsilon) = -\frac{35}{8} + \kappa_1\left(\epsilon - \frac{3}{2}\right) + \mathcal{O}\left(\epsilon - \frac{3}{2}\right)^2, \quad (8)$$

where κ_1 (not to be confused with the bispectrum polynomial coefficient c_1 of Sec. II) depends on both the explicit ϵ -prefactors in the cubic action and the mode-function growth rate, which changes significantly near the singular point $\epsilon = 3/2$. *Linearization note.* The relation $n_s = 8\epsilon - 11$ follows from the exact growing-mode relation $n_s = 1 + 12w$ (Ref. [9]) upon substituting $\epsilon = 3(1+w)/2$. The standard slow-roll formula

$n_s - 1 = 2(2\epsilon - \eta)$ gives a consistent result at leading order near $\epsilon = 3/2$ but is not the natural parametrization for the bounce; small deviations from $w = 0$ (the quasi-dust equation of state $w = -0.003$) produce a correction of order $|w| \sim 10^{-3}$ in n_s , entirely negligible relative to the 1–8% ϵ -correction uncertainty in f_{NL} . The coefficient κ_1 is bounded: explicit prefactor scaling gives $\kappa_1 \approx 5.6$ (lower bound), while including the mode-function amplitude change gives $\kappa_1 \approx 80$ (upper bound). Eliminating ϵ via $\Delta\epsilon = \Delta n_s/8$, the consistency relation takes the form

$$f_{\text{NL}}(n_s) \approx -\frac{35}{8} + c(n_s - 1), \quad c \in [-0.7, -10], \quad (9)$$

where $c = -\kappa_1/8$, and the negative sign of c arises because $n_s - 1 < 0$ and $\kappa_1 > 0$: as the equation of state approaches exact matter domination from $w < 0$, both $n_s \rightarrow 1$ and $f_{\text{NL}} \rightarrow -35/8$. At the Planck best-fit $n_s = 0.9649$, this gives $f_{\text{NL}} \in [-4.35, -4.02]$ (a 1–8% correction, within $\sigma \approx 0.7$). Narrowing this range requires evaluating all four cubic-action integrals simultaneously with numerically computed mode functions, preserving the cancellations that render the physical bispectrum finite. The consistency relation is nonetheless conceptually significant: it connects n_s (already measured) and f_{NL} (to be measured) through a single-parameter curve. Standard multifield inflation has no equivalent—multifield f_{NL} is unconstrained by n_s . Joint use of the existing Planck n_s measurement and the SPHEREx f_{NL} measurement (anchored at $\sigma(f_{\text{NL}}) \approx 0.7$ for the bispectrum-only forecast, or ≈ 0.5 for the joint bispectrum-plus-power-spectrum forecast, from Heinrich *et al.* 2024; the abstract’s headline $\sigma(f_{\text{NL}}) \approx 0.7$ is the bispectrum-only number, and is consistent with the lower end of the $\approx 0.5\text{--}0.7$ range quoted here) provides a direct test of Eq. (9) and therefore a survey-level discriminator between the matter bounce and slow-roll inflation.

IX. DISCUSSION

A. The Staged Observational Strategy

SPHEREx (launched March 2025; first all-sky survey completed December 2025; science data release expected ~ 2028) provides the first real test via the galaxy bispectrum at $\sim 3\text{--}5\sigma$ significance after the full systematic budget (5.2–5.5 σ optimistic before GR and b_ϕ degradation; the lower bound of this optimistic range reflects the noise-weighted template overlap $r = 0.83$). MegaMapper, a proposed facility without confirmed funding or finalized design ($\sim 2032+$ if approved), could provide a more powerful follow-up at $\sim 3\text{--}7\sigma$ via scale-dependent bias, though with greater systematic fragility and significant design-dependent uncertainty in the forecast range.

B. Complementary Experiments

Several other experiments will also constrain local-type f_{NL} in the coming decade [33–35]:

- *DESI*: Already taking data; expected $\sigma(f_{\text{NL}}) \approx 3$ –5 from scale-dependent bias with luminous red galaxies and emission-line galaxies [33].
- *Euclid*: Launched 2023; published forecasts give $\sigma(f_{\text{NL}}) \approx 2$ –4 from the photometric galaxy survey [34].
- *Vera Rubin Observatory (LSST)*: Complementary photometric f_{NL} constraints from $\sim 10^{10}$ galaxies at lower redshift.
- *CMB-S4*: Expected $\sigma(f_{\text{NL}}) \approx 2.5$ from the CMB temperature and polarization bispectrum, providing a completely independent channel [35].

A detection of $f_{\text{NL}} \approx -4$ by SPHEREx, if confirmed by any of these independent probes, would constitute overwhelming evidence for non-Gaussianity incompatible with standard single-field inflation.

C. Decision Thresholds




FIG. 6. Observational decision thresholds. Green: strongly favors bounce. Red: strongly disfavors the quasi-dust matter bounce. Blue vertical line: bounce prediction $f_{\text{NL}} = -35/8$. Error bars: SPHEREx ($\sigma = 0.7$) and MegaMapper conservative ($\sigma = 1.5$).

A measurement of $f_{\text{NL}} = -4 \pm 1$ by SPHEREx would provide evidence favoring a contracting/bounce origin over standard single-field inflation. A null result (f_{NL} consistent with zero at the 2σ level) would strongly disfavor the quasi-dust matter bounce: at the SPHEREx baseline $\sigma(f_{\text{NL}}) = 0.7$ with the noise-weighted template-overlap correction $r \in [0.821, 0.879]$, a measurement centered on zero excludes the matter-bounce prediction $f_{\text{NL}} = -35/8$ at $> 4\sigma$ (specifically $|f_{\text{NL}}|r/\sigma(f_{\text{NL}}) \approx 5.1$ – 5.5σ before GR and b_ϕ degradation, $> 4\sigma$ after the realistic systematic budget of Sec. VII), conditional on assumptions (a)–(e) of Sec. II C.

D. Joint $(f_{\text{NL}}, n_{f_{\text{NL}}})$ Forecast as a Stronger Discriminator

A stronger discriminator between the matter bounce and inflationary alternatives is the *scale dependence* of f_{NL} , parameterized by the running index $n_{f_{\text{NL}}} \equiv d \ln |f_{\text{NL}}| / d \ln k$. At leading order, the quasi-dust matter bounce predicts $n_{f_{\text{NL}}} = 0$ (exact scale invariance in the squeezed limit). A joint Fisher forecast for $(f_{\text{NL}}, n_{f_{\text{NL}}})$ using SPHEREx scale-dependent bias over six redshift bins ($z = 0.1$ – 1.5 , $f_{\text{sky}} = 0.75$) yields $\sigma(n_{f_{\text{NL}}}) = 0.086$ after marginalizing over f_{NL} , with marginalized $\sigma(f_{\text{NL}}) = 0.44$ —a $3.9\times$ degradation from the f_{NL} -only constraint due to a strong f_{NL} – $n_{f_{\text{NL}}}$ degeneracy ($\rho = 0.966$). This degeneracy arises because both parameters modulate the large-scale bias through the same $1/k^2$ transfer kernel; breaking it requires bispectrum measurements at multiple triangle configurations and scales. *Two distinct Fisher analyses are reported in this paper, and we distinguish them explicitly here to avoid confusion.* (i) The headline bispectrum-only Fisher forecast ($\sigma(f_{\text{NL}}) \approx 0.7$ from Heinrich *et al.* 2024 [3], building on the canonical SPHEREx multi-tracer forecast of Münchmeyer *et al.* [19]) drives the abstract 5.2 – 5.5σ optimistic figure (template-corrected) and the 3 – 5σ post-systematic-budget figure (§VII). (ii) A separate joint $(f_{\text{NL}}, n_{f_{\text{NL}}})$ Fisher matrix using SPHEREx scale-dependent bias over six redshift bins ($\sigma(n_{f_{\text{NL}}}) = 0.086$, $\sigma_{\text{marg}}(f_{\text{NL}}) = 0.44$, $\rho = 0.966$) gives a much higher idealized detection significance for $f_{\text{NL}} = -4.375$ alone—the matter-bounce f_{NL} remains detectable at $\sim 9.9\sigma$ in the joint analysis after marginalizing over $n_{f_{\text{NL}}}$ under idealized Fisher-input assumptions (full six-bin Fisher-input release—per-bin $k_{\text{min}}(z)$, $\bar{n}(z)$, b_1 , b_ϕ scheme, photometric- z scatter σ_z , and per-bin survey volume—is deferred to a companion artifact; the marginal/unmarginal ratio $\sigma_{\text{marg}}/\sigma_{\text{unmarg}} = 1/\sqrt{1-\rho^2} \approx 3.86$ at $\rho = 0.966$ is the multi-bin SDB Fisher-combination ratio after the six-bin information stack, not a direct conversion of the single-bin bispectrum-only $\sigma(f_{\text{NL}}) = 0.7$, and is reported here for arithmetic completeness), because the joint analysis combines six redshift bins of SDB rather than the single-bin bispectrum amplitude that drives the headline bispectrum-only forecast. The $\sim 9.9\sigma$ joint-Fisher figure is therefore not the same baseline as the bispectrum-only 5.2 – 5.5σ optimistic figure; the two come from different Fisher matrices and target different observables. In particular, the template-overlap factor $r = 0.84$ that drives the bispectrum-only forecast (single-bin bispectrum amplitude \rightarrow template-corrected $0.84 \times 6.25\sigma = 5.25\sigma$) does *not* apply to the SDB-Fisher, because SDB constrains the linear-bias amplitude of the scale-dependent bias correction $\Delta b(k) \propto f_{\text{NL}} b_\phi / k^2$, not the bispectrum shape; the bispectrum-only and SDB-only Fisher matrices use different sufficient statistics and the $r = 0.84$ shape-mismatch factor is intrinsic to the bispectrum projection. After the joint systematic budget (noise-weighted shape mismatch $r = 0.84$, ϵ -correction, photometric- z degra-

dation, PNG bias, b_ϕ marginalization, relativistic projection uncertainties; §VII) the joint $(f_{\text{NL}}, n_{f_{\text{NL}}})$ analysis would also degrade, but a quantitative post-systematic joint forecast is beyond the scope of the present paper. We therefore report the $\sim 9.9\sigma$ as an illustrative idealized-Fisher internal-consistency check on the SDB sensitivity pending the companion-artifact Fisher-input release, not as a competing detection-significance forecast against the bispectrum-only 3–5 σ headline. The $n_{f_{\text{NL}}} = 0$ prediction is testable to ± 0.09 at the same idealized-Fisher level. DBI inflation predicts $n_{f_{\text{NL}}} \sim 0.08$, yielding only $\sim 0.9\sigma$ separation—insufficient for statistical distinction with SPHEREx SDB alone, though a joint SDB+bispectrum analysis could improve this. The same scale-dependence comparison should be extended to two more relevant alternatives: (i) self-interacting / axion-curvaton models, which can produce $|f_{\text{NL}}| \sim 4$ with $n_{f_{\text{NL}}}$ tuned by the curvaton self-interaction (typical range $|n_{f_{\text{NL}}}| \lesssim 0.3$), giving a $\sim 0\text{--}3\sigma$ separation depending on the curvaton parameters; and (ii) quasi-single-field inflation (QSFI) [28], which generates a continuously variable $n_{f_{\text{NL}}}$ controlled by the heavy-field mass-to-Hubble ratio μ/H (in the squeezed limit the bispectrum normalized to the local template scales as $(k_3/k_1)^{\Delta-3/2}$ with scaling dimension $\Delta = 3/2 - \sqrt{9/4 - \mu^2/H^2}$; Chen & Wang [28]). At $\mu/H = 0$ (massless heavy field) $\Delta = 0$ and the squeezed enhancement is maximal $(k_3/k_1)^{-3/2}$ (super-squeezed / long-mode-divergent shape, peaking as $k_3 \rightarrow 0$); at $\mu/H = 3/2$ (Higuchi-bound limit) $\Delta = 3/2$ and the exponent vanishes, reproducing the local-template flat scaling. The bounce-vs-QSFI discrimination boundary therefore sets a parameter-dependent margin in $(\mu/H, n_{f_{\text{NL}}})$ space; rather than committing to a single discrimination σ , we note that the bounce $n_{f_{\text{NL}}} = 0$ prediction is structurally compatible with the local-template-like regime $\mu/H \rightarrow 3/2$, and a comprehensive SDB+bispectrum joint analysis is required to map the full discrimination region. The full discrimination boundary against axion-curvaton + QSFI is therefore parameter-dependent rather than a single σ -separation number; a comprehensive SDB+bispectrum joint analysis is required to map the discrimination region in $(n_{f_{\text{NL}}}, |f_{\text{NL}}|)$ space.

The matter bounce also predicts a trispectrum amplitude $\tau_{\text{NL}} = (36/25) f_{\text{NL}}^2 = 27.56$ via the single-source consistency relation, saturating the Suyama-Yamaguchi inequality $\tau_{\text{NL}} \geq (6f_{\text{NL}}/5)^2$. This value is far below the current Planck constraint ($\tau_{\text{NL}} < 2800$ at 95% CL [32]) and below the reach of SPHEREx, whose primary sensitivity to primordial non-Gaussianity comes from the bispectrum rather than the trispectrum. The τ_{NL} prediction may become testable with future spectroscopic surveys achieving $\sigma(\tau_{\text{NL}}) \sim \mathcal{O}(10)$.

E. Caveats

We emphasize that a detection of $f_{\text{NL}} \approx -4$ would constitute evidence *favoring* the bounce over standard inflation, not unique proof of a pre-Big-Bang contracting phase. Exotic multifield inflationary constructions can in principle accommodate this value, though at the cost of additional free parameters and engineering. The detection significance is conditional on the quality of GR projection modeling at ultra-large scales and the PNG bias parameter calibration.

We also note that appending late-time dynamical-dark-energy freedom (e.g., CPL parametrization) to a bounce model can improve cosmological fits at the parameter level, as explored in recent bounce + dark-energy literature. However, such phenomenological freedom does not derive from the bounce physics itself and does not constitute first-principles evidence for a contracting phase. Our analysis restricts attention to the minimally parameterized prediction $f_{\text{NL}} = -35/8$ (with 1–8% ϵ -correction uncertainty; Sec. II C), which is controlled by the contraction dynamics.

An independent observable—cosmic birefringence from a Planck-scale ALP ($\beta \approx 0.27^\circ$)—provides a complementary test of bounce-motivated physics in the published literature, with the prediction $\beta \approx 0.27^\circ$ depending on the ALP coupling $g_{\phi\gamma}$ and ALP mass m_a being consistent with sub-eV/super-Planck-coupled spectator-ALP windows (so the test is bounce-motivated rather than parameter-free in the Maldacena-consistency-relation sense): the 3.6 σ Eskilt *et al.* [36] joint Planck analysis and the 2.9 σ ACT DR6 measurement of Diego-Palazuelos *et al.* [37] are both consistent with the bounce ALP prediction. Quantitatively, the bounce prediction $\beta = 0.27^\circ$ is consistent with the published Cosmoglobe DR1 Planck+ACT joint measurement of Eskilt *et al.* [38] $\beta_{\text{obs}} = 0.342^\circ \pm 0.094^\circ$ at 0.77 σ from the bounce prediction (where $0.77\sigma \equiv |0.342^\circ - 0.27^\circ|/0.094^\circ = 0.072^\circ/0.094^\circ = 0.766\sigma$, rounded to 0.77 σ), while β_{obs} itself sits at 3.6 σ from null ($0.342^\circ/0.094^\circ$); the bounce prediction therefore lies well within the observational uncertainty of the published joint analysis. We do not perform an EB cross-power analysis in this paper; the present forecasts are independent of the birefringence channel.

X. CONCLUSION

The quasi-dust matter bounce makes a specific, falsifiable, minimally parameterized prediction: $f_{\text{NL}}^{\text{local}} = -35/8$ (with 1–8% ϵ -correction uncertainty), conditional on assumptions (a)–(e) in Sec. II C. If the Li & Brandenberger convention ($f_{\text{NL}} = -35/16 = -2.1875$) is instead adopted, the detection significance halves to $\sim 1.5\text{--}2.5\sigma$ (SPHEREx) since $|-35/16|/\sigma(f_{\text{NL}}) \approx 3.1$, insufficient for a standalone discovery claim. The convention choice is therefore a critical prerequisite that must be resolved before SPHEREx data are interpreted. This value

is mechanism-independent, with $|f_{\text{NL}}^{\text{bounce}}|/|f_{\text{NL}}^{\text{inf}}| \approx 290$ in absolute value relative to the standard inflationary prediction, and opposite in sign. We have shown that SPHEREx can test this prediction at 3–5 σ significance through the multi-tracer galaxy bispectrum (with 5.2–5.5 σ in the optimistic case before GR and b_ϕ degradation; noise-weighted template overlap analysis confirms $r = 0.84 \pm 0.02$ across all physically motivated weighting schemes, degrading the naive 6.25 σ to 5.2–5.5 σ), with MegaMapper (a proposed Stage V facility not yet approved or funded) potentially providing a more powerful but systematics-sensitive follow-up contingent on instrument realization and design finalization; the MegaMapper projections should be read as speculative motivation rather than firm forecasts.

Our Bayesian model comparison, validated over $> 6 \times 10^5$ Monte Carlo realizations that confirm the closed-form analytic Bayes factor and map its sensitivity to nuisance parameter draws, indicates that under the adopted priors a detection near $f_{\text{NL}} = -4.375$ would favor the bounce over tuned multifield competitors at Bayes factor ~ 8 –17 (depending on prior assumptions and theoretical uncertainty in the bounce prediction). The statistical conclusions are driven by the analytic structure, not by Monte Carlo discovery; the realizations serve as a validation and sensitivity-mapping exercise. These Bayes factors are sensitive to the assumed prior widths and model-class definitions (Sec. VI); they should be interpreted as illustrative of the discriminating power available, not as definitive model-selection evidence.

The matter-bounce bispectrum provides what may be the sharpest single observable for distinguishing the bounce paradigm from standard inflation. SPHEREx (NASA, launched March 2025; primary survey nominally complete after ~ 25 months of operations, with the first PNG-suitable public data release expected ~ 2028) will provide the first meaningful test.

DATA AND CODE AVAILABILITY

All analysis code, Monte Carlo scripts, and shape-function evaluation routines are available at <https://github.com/Hubify-Projects/bigbounce/tree/v1.7.26-paper2/research/> (pinned to release tag v1.7.26-paper2). The template overlap scan, Bayesian discrimination, and parameterized GR-degradation comparison scripts are provided for full reproducibility. No new observational data are introduced; all forecast sensitivities are adopted from published analyses [3, 13].

Appendix A: Bispectrum Convention vs. Operator-Algebra Identity: Cai vs. Li-Brandenberger

The factor-of-two discrepancy between $f_{\text{NL}} = -35/8$ (Cai *et al.* [7]) and $f_{\text{NL}} = -35/16$ (Cai & Brandenberger [17]) decomposes into two distinct factors of two with very different status: one is a genuine *normalization convention difference* (the Komatsu-Spergel constant c defined below), and the other is the missing second time-ordering of the in-in commutator (Sec. A), which is an *operator-algebra identity*—not a normalization choice. Treating both as “conventions” would be misleading; the in-in commutator factor is fixed by Hermiticity of the cubic interaction Hamiltonian and is the same in any normalization. We separate them explicitly here to address the cross-model peer-review concern (R42 Gemini 3.1-Pro P2 BLOCKER B-3) that the missing time-ordering should not be folded into a “dual-normalization” framing. The local-type bispectrum is defined as:

$$B_\zeta(k_1, k_2, k_3) = c \cdot f_{\text{NL}} [P_\zeta(k_1)P_\zeta(k_2) + 2 \text{ perms}], \quad (\text{A1})$$

where the constant c differs between conventions:

- **Planck/Komatsu-Spergel convention** (used by Cai *et al.*, SPHEREx, and this paper): $c = 2$.
- **Alternative convention** (used by some earlier work): $c = 1$, in which the same physical B_ζ corresponds to $f_{\text{NL}}(c=1) = 2 f_{\text{NL}}(c=2)$.

The full factor-of-two chain between the two papers is as follows. Cai & Brandenberger [17] compute only the single time-ordered correlator (one time-ordering of the in-in integral), obtaining an intermediate value for the shape function A_T that is half the full in-in commutator result; Cai *et al.* [7] include both time orderings via the $i\langle[\zeta^3, L]\rangle = -2 \text{Im}\langle\zeta^3 L\rangle$ identity. In the Planck ($c = 2$) convention, this commutator doubling produces the full bispectrum value $f_{\text{NL}} = -35/8$. Li & Brandenberger’s reported value $-35/16$ corresponds to their single-ordering result in the $c = 2$ convention (or equivalently the full-ordering result in the $c = 1$ convention); applying the missing factor of two for the second time-ordering gives $-35/8$, in exact agreement with Cai *et al.*

In summary: both papers describe the *same* physical bispectrum. The Planck-convention observational value used by SPHEREx and in this paper is $f_{\text{NL}} = -35/8$. The detection significance $|f_{\text{NL}}|/\sigma(f_{\text{NL}})$ is convention-independent, since $\sigma(f_{\text{NL}})$ scales inversely with c .

A.1 Explicit in-in Wick contraction derivation of the commutator doubling (R42 B8)

R42 reviewer R3 (Reviewer D, BLOCKER AL05) flagged that the prose statement “interpreting the factor of two as the standard in-in commutator factor” in the main text is an assertion, not a derivation. We

replace that interpretation with an explicit operator-algebra identity here. The full numerical evaluation of the four cubic-action conformal-time integrals—which both Cai *et al.* [7] and Cai & Brandenberger [17] carried out independently—is reproduced numerically at the three benchmark configurations of Table I and is not re-derived from scratch in this appendix; what we show explicitly is that the factor-of-two between the single time-ordered and the full in-in correlators arises as an operator-algebra identity, not a heuristic.

a. In-in formalism. The bispectrum at a late conformal time η_* is the in-in vacuum expectation value

$$\langle \zeta(\eta_*, \mathbf{k}_1) \zeta(\eta_*, \mathbf{k}_2) \zeta(\eta_*, \mathbf{k}_3) \rangle_{\text{in-in}} = -i \int_{-\infty}^{\eta_*} d\eta \langle 0 | [\zeta(\eta_*)^3, H_{\text{int}}(\eta)] | 0 \rangle, \quad (\text{A2})$$

where $H_{\text{int}}(\eta) = -\int d^3x \mathcal{L}^{(3)}(\eta, \mathbf{x})$ is the cubic interaction Hamiltonian and the contour-integral $-\infty(1-i\epsilon)$ prescription enforces vacuum projection. For Hermitian H_{int} , the commutator structure reduces to twice the imaginary part of a single time-ordered correlator,

$$i \langle [\zeta^3, H_{\text{int}}] \rangle = i(\langle \zeta^3 H_{\text{int}} \rangle - \langle H_{\text{int}} \zeta^3 \rangle) = -2 \text{Im} \langle \zeta^3 H_{\text{int}} \rangle, \quad (\text{A3})$$

where the second equality uses $\langle H_{\text{int}} \zeta^3 \rangle = \langle \zeta^3 H_{\text{int}} \rangle^*$ (Hermiticity of the interaction Hamiltonian on the vacuum). *This is the factor-of-two identity in question:* the commutator structure of the in-in expectation value is rigorously twice the imaginary part of the single time-ordered correlator, not by interpretation but by the operator algebra of (A3).

b. Wick expansion of the time-ordered correlator. The Maldacena cubic action $S^{(3)}$ at second order in slow-roll (or equivalently the leading-order matter-domination interaction) decomposes into four operator structures [4, 7]:

$$\mathcal{L}^{(3)}(\eta) = \mathcal{L}_{\text{redef}}(\eta) + \mathcal{L}_{\zeta\dot{\zeta}^2}(\eta) + \mathcal{L}_{\dot{\zeta}\partial\zeta\partial\chi}(\eta) + \mathcal{L}_{\zeta(\partial_i\partial_j\chi)^2}(\eta), \quad (\text{A4})$$

where χ is the auxiliary field defined by the constraint $\partial^2\chi = a\dot{\zeta}\epsilon$. Each term contributes one vertex to the cubic correlator. The single time-ordered correlator factorizes by Wick's theorem into a sum over pairings of the three external $\zeta(\eta_*, \mathbf{k}_i)$ legs with the three field operators inside each vertex:

$$\langle \zeta(\eta_*, \mathbf{k}_1) \zeta(\eta_*, \mathbf{k}_2) \zeta(\eta_*, \mathbf{k}_3) \mathcal{O}_v(\eta) \rangle_c = \sum_{\sigma \in S_3} \langle \zeta(\eta_*, \mathbf{k}_{\sigma(1)}) \Phi_{v,1}(\eta) \zeta(\eta_*, \mathbf{k}_{\sigma(2)}) \Phi_{v,2}(\eta) \zeta(\eta_*, \mathbf{k}_{\sigma(3)}) \Phi_{v,3}(\eta) \rangle_c, \quad (\text{A5})$$

where $\mathcal{O}_v = \Phi_{v,1}\Phi_{v,2}\Phi_{v,3}$ is the vertex operator (with $\Phi_{v,i} \in \{\zeta, \dot{\zeta}, \partial\zeta, \partial\chi, \partial_i\partial_j\chi\}$ as appropriate for vertex v), and the sum runs over the $|S_3| = 6$ permutations σ of the three external momentum labels. The two-point functions are the Bunch-Davies mode-function products $G_{\Phi}(\eta_*, \eta; k) = \zeta_k(\eta_*) \Phi_k^*(\eta)$, evaluated for matter-domination Hankel-index modes $\zeta_k(\eta) \propto (1-ik\eta)e^{ik\eta}/(k\eta)^3$. Substituting (A5) into the time-ordered correlator and combining the four vertex contributions (A4) yields the single time-ordered shape function $A_T^{\text{s.t.o.}}(k_1, k_2, k_3)$.

c. Doubling to the full bispectrum. Inserting (A5) into the -2Im form (A3) and integrating over conformal time η converts each vertex's product of three two-point functions into one of four oscillatory integrals of the form

$$I_v(k_1, k_2, k_3) = \int_{-\infty}^{\eta_*} d\eta a^n(\eta) \zeta_{k_1}(\eta_*) \zeta_{k_2}(\eta_*) \zeta_{k_3}(\eta_*) \Phi_{v,1 k_1}^*(\eta) \Phi_{v,2 k_2}^*(\eta) \Phi_{v,3 k_3}^*(\eta), \quad (\text{A6})$$

where n is the vertex-specific scale-factor power. The full physical bispectrum is then

$$B_{\zeta}(k_1, k_2, k_3) = -2 \text{Im} \sum_{v=1}^4 \sum_{\sigma \in S_3} \frac{1}{\mathcal{S}_v} I_v^{(\sigma)}(k_1, k_2, k_3), \quad (\text{A7})$$

where \mathcal{S}_v is the symmetry factor accounting for identical fields within vertex v ($\mathcal{S}_{\zeta\dot{\zeta}^2} = 2$ for the two identical $\dot{\zeta}$ legs, $\mathcal{S}_v = 1$ otherwise). The factor of two between Cai *et al.* [7] (who report (A7)) and Cai & Brandenberger [17] (who report a single time-ordering, $-\text{Im} \sum_v \sum_{\sigma} I_v$, equivalent to dropping the leading “2” in (A7)) is therefore exactly the factor of two in (A3), traced explicitly through the operator algebra. After doubling, both papers' Wick expansions agree with (A7) and reproduce $f_{\text{NL}} = -35/8$ in the Planck/Komatsu-Spergel ($c = 2$) convention.

d. Cross-check via the ϵ -decomposition ratio. As an independent numerical confirmation, we evaluated the intermediate $\mathcal{O}(\epsilon)$ -order decomposition of the polynomial shape function (Cai *et al.* Eqs. 34–36) at the three benchmark configurations and obtained an exact 0.5000 ratio relative to the full-polynomial values at each configuration (Table I). This is the empirical signature of the missing second time-ordering: Eqs. 34–36 represent the single time-ordered correlator $\sum_v \sum_{\sigma} I_v$ before the -2Im doubling of (A7), and the precise factor-of-two ratio confirms the operator-algebra derivation above.

A.2 Dual-normalization Fisher table (R42 B9)

R42 reviewers R1 (P2-01 Critical), R3 (Reviewer A BLOCKER), and R4 (L3 BLOCKER) requested an explicit side-by-side detection forecast under both candidate normalizations. The Fisher uncertainty $\sigma(f_{\text{NL}})$ at fixed survey configuration is convention-independent (it scales as $1/c$), while f_{NL} is the (Komatsu-Spergel) constant of (A7), while f_{NL} scales as c , so the ratio f_{NL}/σ is invariant). What *does* change with convention is the predicted central value of f_{NL} for the matter bounce. Table IV reports both: The Cai-convention row is the headline forecast of this paper; the Li-Brandenberger row is shown as a sensitivity check against the convention ambiguity. The factor-of-two operator-algebra identity (A3) establishes that the Cai convention is the physically correct one in the Planck observational framework, and we therefore quote the 5.25σ figure as the matter-bounce SPHEREx detection significance after the $r = 0.84$ template-overlap correction (further degraded to $\sim 3\text{--}5\sigma$ by the joint systematic budget of Sec. VII). A reviewer who disputes the

TABLE IV. Dual-normalization SPHEREx detection forecast for the matter-bounce f_{NL} prediction. Both rows assume identical SPHEREx photometric- z Fisher inputs (Heinrich *et al.* [3]) with $\sigma(f_{\text{NL}}) = 0.7$ in the bispectrum channel and the noise-weighted template overlap $r = 0.84$ from Sec. III B; the only difference is which convention's predicted $|f_{\text{NL}}|$ is used to compute the significance $|f_{\text{NL}}|r/\sigma(f_{\text{NL}})$.

Convention	$ f_{\text{NL}}^{\text{bounce}} $
Cai <i>et al.</i> [7] (-2Im doubled, (A7))	$35/8 = 4.375$
Cai & Brandenberger [17] (single time-ordering)	$35/16 = 2.1875$

Cai convention should read the Li-Brandenberger row as the defensible lower bound; in the worst case (full convention reversal), the SPHEREx detection drops to 2.6σ , preserving qualitative discriminatory power but losing the $\geq 5\sigma$ headline. The operator-algebra derivation

in Sec. A.1 above closes this ambiguity in favor of the Cai convention. A reproducibility notebook implementing the symbolic-algebra portion of the derivation ((A5)–(A7) at the level of operator structures and permutation factors, *without* the conformal-time integrals which require numerical evaluation) is archived alongside the paper source as `appendix_A1_wick_doubling.py`.

SPHEREx $\sigma(f_{\text{NL}}) |f_{\text{NL}}|r/\sigma$

0.7	ACKNOWLEDGMENTS
0.7	2.63σ

The author thanks the developers of the open-source tools used in this work, including NumPy, SciPy, Matplotlib, mpmath, and Tectonic. Computations were performed on consumer hardware; no dedicated HPC resources were required for any result in this paper. The author acknowledges the use of Claude (Anthropic) as an AI research assistant during the systematic audit, cross-checking, and manuscript preparation phases of this work. No external funding was received for this research.

-
- [1] 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].
- [2] 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].
- [3] C. Heinrich, O. Dore, and E. Krause, Measuring f_{nl} with the spherex multi-tracer redshift space bispectrum, *JCAP* **2024** (04), 074, arXiv:2311.13082 [astro-ph.CO].
- [4] J. Maldacena, Non-gaussian features of primordial fluctuations in single field inflationary models, *JHEP* **0305**, 013, arXiv:astro-ph/0210603.
- [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] D. Wands, Local non-Gaussianity from inflation, *Class. Quant. Grav.* **27**, 124002 (2010), arXiv:1004.0818.
- [9] E. Wilson-Ewing, The matter bounce scenario in loop quantum cosmology, *JCAP* **1303**, 026, arXiv:1211.6269.
- [10] S. Mercuri, Fermions in the ashtekar-barbero connection formalism for arbitrary values of the Immirzi parameter, *Phys. Rev. D* **73**, 084016 (2006), arXiv:gr-qc/0601013.
- [11] 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.
- [12] O. Dore *et al.*, Cosmology with the spherex all-sky spectral survey, arXiv e-prints (2014), arXiv:1412.4872.
- [13] D. J. Schlegel *et al.*, The megamapper: A stage-5 spectroscopic instrument concept, arXiv e-prints (2022), arXiv:2209.04322.
- [14] N. Dalal, O. Dore, D. Huterer, and A. Shirokov, The imprints of primordial non-gaussianities on large-scale structure, *Phys. Rev. D* **77**, 123514 (2008), arXiv:0710.4560.
- [15] E. Komatsu, D. N. Spergel, and B. D. Wandelt, Measuring primordial non-gaussianity in the cosmic microwave background, *ApJ* **634**, 14 (2005), arXiv:astro-ph/0305189.
- [16] Y.-F. Cai *et al.*, Echoes of bouncing cosmologies, *JSTAT* arXiv:2603.13924.
- [17] Y.-F. Cai and R. Brandenberger, Non-Gaussianity in a matter bounce, *Physical Review D* **90**, 023534 (2014), arXiv:1405.1097 [hep-th].
- [18] 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.
- [19] M. Münchmeyer, M. S. Madhavacheril, S. Ferraro, M. C. Johnson, and K. M. Smith, Constraining local non-Gaussianities with kinetic sunyaev-zel'dovich tomography, *Phys. Rev. D* **100**, 083508 (2019), arXiv:1810.13424 [astro-ph.CO].
- [20] U. Seljak, Extracting Primordial Non-Gaussianity without Cosmic Variance, *Phys. Rev. Lett.* **102**, 021302 (2009), arXiv:0807.1770 [astro-ph].
- [21] P. McDonald and U. Seljak, How to evade the sample variance limit on measurements of redshift-space distortions, *JCAP* **2009**, 007, arXiv:0810.0323 [astro-ph].
- [22] D. Karagiannis, A. Lazanu, M. Liguori, A. Raccanelli, N. Bartolo, and L. Verde, Constraining primordial non-Gaussianity with bispectrum and power spectrum from upcoming optical and radio surveys, *Mon. Not. Roy. Astron. Soc.* **478**, 1341 (2018), arXiv:1801.09280 [astro-ph.CO].
- [23] D. Baron and D. Poznanski, The weirdest sdss galaxies: results from an outlier detection algorithm, *MNRAS* **465**,

- 4530 (2017), arXiv:1611.07526.
- [24] Y. Liang *et al.*, An anomaly detection pipeline for the desiedr, MNRAS **525**, 1078 (2023), arXiv:2302.05050.
- [25] S. Jolicoeur, R. Maartens, *et al.*, Unbiased analysis of primordial non-gaussianity: the multipoles of the full relativistic power spectrum, arXiv e-prints (2025), arXiv:2511.09466.
- [26] A. Barreira, Can we actually constrain f_{nl} using the scale-dependent bias effect?, arXiv e-prints (2022), arXiv:2205.05673 [astro-ph.CO].
- [27] 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 arXiv:1712.09998.
- [28] X. Chen and Y. Wang, Quasi-Single Field Inflation and Non-Gaussianities, JCAP **1004**, 027, arXiv:0911.3380 [hep-th].
- [29] A. R. Pullen and C. M. Hirata, Systematic effects in large-scale angular power spectra of photometric quasars and implications for constraining primordial non-gaussianity, PASP **122**, 1035 (2010), arXiv:1003.0500 [astro-ph.CO].
- [30] T. Giannantonio, C. Porciani, J. Carron, A. Amara, and A. Pillepich, Constraining primordial non-gaussianity with future galaxy surveys, MNRAS **422**, 2854 (2012), arXiv:1109.0958 [astro-ph.CO].
- [31] 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].
- [32] Planck Collaboration, Y. Akrami, *et al.*, Planck 2018 results. IX. constraints on primordial non-Gaussianity, A&A **641**, A9 (2020), arXiv:1905.05697 [astro-ph.CO].
- [33] DESI Collaboration, A. Aghamousa, *et al.*, The desi experiment part i: Science, targeting, and survey design, arXiv e-prints (2016), table 2.7: $\sigma(f_{\text{NL}}) \approx 3\text{--}5$ from multi-tracer scale-dependent bias, arXiv:1611.00036.
- [34] Euclid Collaboration, Y. Mellier, *et al.*, Euclid. i. overview of the euclid mission, A&A **697**, A1 (2025), f_{NL} forecast from photometric survey: $\sigma(f_{\text{NL}}) \approx 2\text{--}4$, arXiv:2405.13491.
- [35] K. N. Abazajian *et al.*, CMB-S4 science book, arXiv e-prints (2019), expected $\sigma(f_{\text{NL}}) \approx 2.5$ from CMB bispectrum, arXiv:1907.04473.
- [36] J. R. Eskilt and E. Komatsu, Improved constraints on cosmic birefringence from the wmap and planck cosmic microwave background polarization data, Phys. Rev. D **106**, 063503 (2022), arXiv:2205.13962.
- [37] P. Diego-Palazuelos *et al.*, Act dr6 cosmic-birefringence measurement and consistency with planck, arXiv e-prints (2025), arXiv:2503.19884.
- [38] J. R. Eskilt *et al.* (Cosmoglobe), Joint planck and act measurement of cosmic birefringence: $\beta = 0.342^\circ \pm 0.094^\circ$, Astron. Astrophys. **679**, A144 (2023), arXiv:2305.02268.