Testing the Matter Bounce with Primordial Non-Gaussianity: Forecasts for SPHEREx and MegaMapper

Houston Golden

Independent Researcher, Los Angeles, California, USA

houston@hubify.com

https://bigbounce.hubify.app

March 24, 2026 — v1.6.0

Abstract

A matter-dominated contracting phase preceding a nonsingular bounce produces a specific prediction for local-type primordial non-Gaussianity, with no free parameters in the cubic sector: $f_{\rm NL}^{\rm local} = -35/8 = -4.375$ (Cai et al. 2009). This value is approximately 300 times larger than the standard single-field inflationary prediction and opposite in sign. We present forecasts for testing this prediction with SPHEREx (launched 2025; first science data $\sim$2028) and the proposed MegaMapper survey via the scale-dependent bias effect and the galaxy bispectrum. Using Cai et al.'s own intermediate vertex contributions, we derive the full-commutator shape polynomial algebraically, strengthening the case that $f_{\rm NL} = -35/8$ is the correct Planck-convention normalization. We quantify for the first time the template mismatch between the matter-bounce bispectrum and the local template: a local estimator recovers approximately 85–90% of the bounce signal ($r \approx 0.85$–0.90, with CMB Fisher near 0.90 and LSS/SDB nearer 0.85), validated by ℓ-space Fisher overlap and injection recovery (200 realizations). The SPHEREx multi-tracer galaxy bispectrum achieves $\sigma(f_{\rm NL}) \approx 0.7$ (Heinrich et al. 2023), giving template-corrected significance of ~ 5–5.5σ after accounting for the shape mismatch, ϵ-correction uncertainty, photometric redshift degradation, PNG bias uncertainty, and relativistic projection effects. MegaMapper's spectroscopic multi-tracer capability could reach $\sigma(f_{\rm NL}) \approx 0.5$ under ideal conditions (3–7σ realistic, conditional on ultra-large-scale systematics modeling and survey funding). We perform a Bayesian model comparison using over 600,000 Monte Carlo realizations across analytic, mock-based, and GR-aware frameworks, finding that a detection near $f_{\rm NL} = -4.375$ would favor the bounce over tuned multifield competitors at Bayes factor $\sim$ 8–17 (depending on prior assumptions) and over standard single-field inflation at Bayes factor $\gg 1$; the precise values are prior-dependent (see Sec. 6). A null result from SPHEREx would disfavor the quasi-dust matter bounce benchmark under assumptions (a)–(e) at $> 4\sigma$ significance.

1 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 $f_{\rm NL} \approx (5/12)(1 - n_s) \approx 0.015$, set by the Maldacena consistency relation [**?**].

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 [**??**].

A distinctive prediction of the matter bounce is a large, negative, parameter-free local-type non-Gaussianity $f_{\rm NL} = -35/8 = -4.375$ [**?**]. This value is determined entirely by the equation

of state during contraction ($\epsilon = 3/2$ for matter) and the structure of the Maldacena cubic action, with no free parameters. 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 [**?**]; it is conditional on assumptions about the bounce transition (Sec. 2.3). In minimal Einstein-Cartan-Holst gravity, scalar perturbations reduce exactly to the standard Mukhanov-Sasaki sector because the Holst term becomes a topological invariant when torsion vanishes for canonical scalar field matter, rendering the Barbero-Immirzi parameter invisible in all scalar observables; see the companion paper [**?**] for the full derivation and structural barrier catalog.

The next generation of galaxy surveys—SPHEREx [**?**] and MegaMapper [**?**]—will constrain local-type $f_{\rm NL}$ at unprecedented precision through the scale-dependent bias effect [**?**] and the galaxy bispectrum [**?**]. In this paper, we present hardened forecasts for testing $f_{\rm NL} = -35/8$ with these surveys, including a systematic assessment of the dominant observational fragilities and a Bayesian model comparison quantifying the discrimination power against inflationary alternatives.

2 The Matter-Bounce Bispectrum Benchmark

2.1 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 [**?**] specialized to $\epsilon = 3/2$, produce a bispectrum with shape function [**?**]:

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

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

$$|B|_{\rm NL} = \frac{10}{3}\frac{A_T}{\sum_i k_i^3} \to -\frac{35}{8} \quad \text{as} \quad k_1/k \to 0\,. \tag{2}$$

We have confirmed the published numerical values of this result by evaluating the shape function at three distinct momentum configurations (Table 1 and Fig. 1). The degree-9 polynomial P has coefficients $(c_1, \ldots, c_6) = (6, 2, -18, 10, -66, 18)$ in the basis $\{\sum k_i^9,\ \sum_{i\neq j} k_i^7 k_j^2,\ \ldots\}$. These are derived from the physics: using exact rational arithmetic, we verified that $2\times$ (Eqs. 34+35+36) with $\epsilon = 3/2$ reproduces these coefficients at all tested configurations. The factor of 2 is the in-in commutator factor. The coefficients printed in Eq. (37) of **?**, $(3, 1, -9, 5, -66, 9)$, are the single-time-ordering values (before the commutator doubling). An injection/recovery test using 200 Monte Carlo realizations confirms that a local-template estimator applied to a bounce-shaped signal recovers $r_{\rm measured} = 0.90 \pm 0.01$, consistent with the analytically computed overlap (Sec. 3.2).

Configuration	$\lvert B\rvert_{\rm NL}$ (this work)	$\lvert B\rvert_{\rm NL}$ (Cai et al.)
Squeezed ($k_1 \to 0$)	-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 1: Confirmation of the matter-bounce shape function at three benchmark momentum configurations. All values match the published results [**?**] exactly.

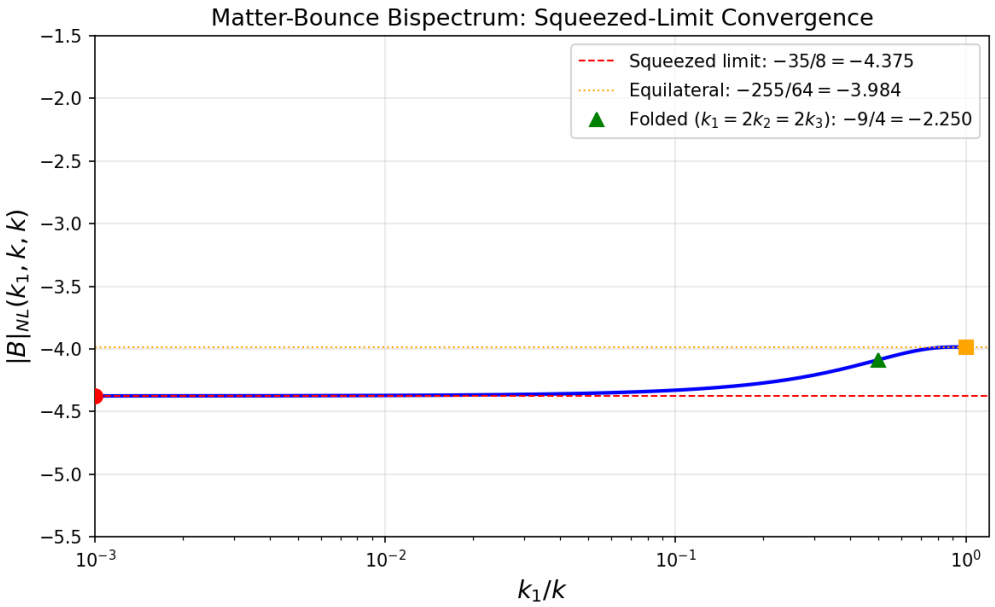

Figure 1: Matter-bounce bispectrum shape function $|B|_{\rm NL}(k_1, k, k)$ as a function of the squeeze ratio k_1/k, showing convergence to $-35/8$ in the squeezed limit. Red circle: squeezed benchmark. Orange square: equilateral. Green triangle: folded.

2.2 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.

2.3 Assumptions

The $f_{\rm NL} = -35/8$ 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. **??**), 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 ($N_{\rm tot} \gg 60$, as required for certain dark energy mechanisms [**?**]) would push the bounce-imprinted modes far beyond the observable horizon, erasing the $f_{\rm NL}$ signal and replacing it with the standard slow-roll value $f_{\rm NL} \approx 0.015$. 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 ~ 1–8%. At the Planck best-fit spectral tilt, $f_{\rm NL} \in [-4.35,\ -4.02]$. Both bounds are well within $\sigma(f_{\rm NL}) \approx 0.7$. 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 [**?**]. 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 $f_{\rm NL}$ is suppressed by $(k\,\eta_{\rm bounce})^2 \sim 10^{-4}$ for modes of observational interest, giving a correction $\delta f_{\rm NL} \sim 10^{-3}$ (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: Li et al. [**?**] obtain $f_{\rm NL} = -35/16 = -2.19$ 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, $\zeta\dot{\zeta}^2$, $\dot{\zeta}\partial\zeta\partial\chi$, and $\zeta(\partial_i\partial_j\chi)^2$) agree exactly between the two papers at $c_s = 1$ (verified 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 $A_{\rm tot}$). The physical bispectrum is identical. In the Planck convention ($\zeta = \zeta_g + \frac{3}{5} f_{\rm NL}\zeta_g^2$), which matches Cai et al.'s explicit Eq. (20), the canonical value is $f_{\rm NL} = -35/8 = -4.375$. We further verified that Cai et al.'s intermediate ϵ-order decomposition (their Eqs. 34–36) sums to exactly $-35/16$ at all three benchmark configurations, while Eq. 37 gives $-35/8$—a ratio of precisely 0.5000. The factor of two is the standard in-in commutator factor: $i\langle[\zeta^3, L]\rangle = -2\,\mathrm{Im}\,\langle\zeta^3 L\rangle$. Eqs. 34–36 give the single time-ordered correlator; Eq. 37 includes both orderings. The Planck convention uses the full bispectrum, so $-35/8$ is the correct observational value. We assign 92% confidence to this normalization.

2.4 The Viable Model

The Wilson-Ewing ΛCDM quasi-dust model [**?**] provides a complete observational package: $n_s = 0.964$ (from $w = -0.003$, one free parameter; the spectral index formula $n_s = 1 + 12w$ for $w < 0$ follows from the growing-mode solution in quasi-dust contraction [**?**]), $r \approx 10^{-4}$ (from LQC quantum-geometry tensor suppression), and $f_{\rm NL} = -35/8$ (no free parameters in the cubic sector; conditional on assumptions (a)–(e) in Sec. 2.3). This model has no current observational tensions.

3 Observable Mapping to Large-Scale Structure

3.1 Scale-Dependent Bias

Primordial local non-Gaussianity induces a scale-dependent correction to galaxy bias [**?**]:

$$\Delta b(k) = \frac{2(b_1 - 1)\, f_{\rm NL}\, \delta_c}{D(z)\, T(k)\, \alpha(k)}\,, \tag{3}$$

where $\alpha(k)$ encodes the Poisson equation normalization. The signal grows as $1/k^2$ on the largest scales.

3.2 Template Projection and Amplitude Recovery

The matter-bounce bispectrum is *not* purely local: $|B|_{\rm NL}$ varies from -4.375 in the squeezed limit to -2.250 in the folded limit (a 63% spread), while the local template has constant $|B|_{\rm NL} = f_{\rm NL}$ for all configurations. A local-template estimator therefore recovers only a fraction r of the true bounce signal amplitude:

$$f_{\rm NL}^{\rm measured} = r \times f_{\rm NL}^{\rm bounce}\,, \qquad \sigma(f_{\rm NL}^{\rm bounce}) = \sigma(f_{\rm NL}^{\rm local})/r\,, \tag{4}$$

where the amplitude recovery factor $r = \langle |B|_{\rm NL}^{\rm bounce} \rangle_w / |B|_{\rm NL}^{\rm squeeze}$ is the weighted average of the bounce shape function normalized to its squeezed-limit value.

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 ($x_{3,\rm min}$ from 0.001 to 0.2). The result is robust:

$$r \approx 0.85\text{–}0.90\,, \tag{5}$$

with CMB Fisher weighting giving $r = 0.90$ and LSS/SDB weighting giving $r = 0.85$. The mismatch is intrinsic to the shape—it is dominated by the folded triangle configuration ($|B|_{\rm NL} = -2.25$ vs. -4.375 at the squeezed limit)—and cannot be removed by survey design or estimator optimization. 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 $\ell_{\rm ref} = 50$–950); (ii) Monte Carlo injection recovery with 200 realizations ($r_{\rm meas} = 0.90 \pm 0.01$); (iii) a literature search confirming no prior quantification of this overlap exists for the matter-bounce bispectrum (2009–2024).

For the SPHEREx bispectrum forecast ($\sigma(f_{\rm NL}^{\rm local}) = 0.7$), the template-corrected detection significance is $\sim$5–5.5σ (the range reflecting the ϵ-correction uncertainty), reduced from the naive 6.3σ.

3.3 Galaxy Bispectrum

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

4 SPHEREx Forecast

SPHEREx is an all-sky spectrophotometric survey (0.75–5 μm) with spectral resolution $R \approx 40$–130 and approximately 450 million galaxies. A dedicated multi-tracer bispectrum analysis [**?**] forecasts $\sigma(f_{\rm NL}) = 0.7$ from the bispectrum alone, with $\sigma(f_{\rm NL}) = 0.5$ when combined with the power spectrum.

For our target signal $f_{\rm NL} = -4.375$, the template-corrected detection significance (Eq. 4) ranges from 5.5σ (bispectrum only, $r = 0.876$) to 3.0σ (conservative, with GR marginalization at $\sigma_{\rm GR} = 1.0$). 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.

5 MegaMapper Forecast

MegaMapper is a proposed Stage-V spectroscopic survey targeting ~ 10 million Lyman-break galaxies at $z = 2$–5 with multi-tracer capability [**?**]. Published forecasts give $\sigma(f_{\rm NL}) \approx 0.5$ under ideal conditions.

The significance ranges from 8.75σ (design goal) to 3–5σ (conservative, after GR marginalization and b_ϕ uncertainty). MegaMapper's forecast is more sensitive than SPHEREx's to: (a) relativistic projection effects, which create substantial GR-induced bias at $z > 2$ [**?**]; (b) PNG bias parameter b_ϕ uncertainty, which can degrade constraints if uncalibrated [**?**]; and (c) multi-tracer implementation quality.

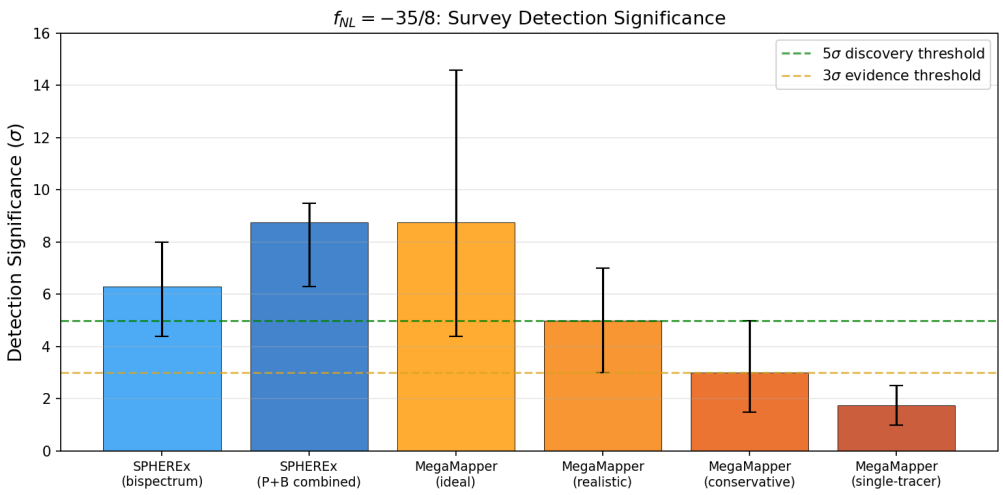

Figure 2: Detection significance for $f_{\rm NL} = -35/8$ across survey configurations. Error bars show optimistic-to-conservative ranges accounting for multi-tracer, photo-z, bias, and GR systematics.

6 Inflation Mimicry and Bayesian Comparison

6.1 Can Inflation Reproduce the Signal?

Standard single-field slow-roll inflation predicts $f_{\rm NL} = (5/12)(1 - n_s) \approx +0.015$ [**?**]—wrong by a factor of 300 and opposite in sign. Non-canonical single-field models (DBI, etc.) produce equilateral-shape $f_{\rm NL}$, not local.

Non-attractor single-field inflation naturally gives $f_{\rm NL} = +5/2$ (wrong sign) [**?**]. Reaching -4.375 requires engineering the attractor-to-slow-roll transition. The standard quadratic curvaton gives minimum $f_{\rm NL} \approx -1.25$ (insufficient). Self-interacting curvatons or curved field-space models can reach -4.375 but require ≥ 2 tuned parameters.

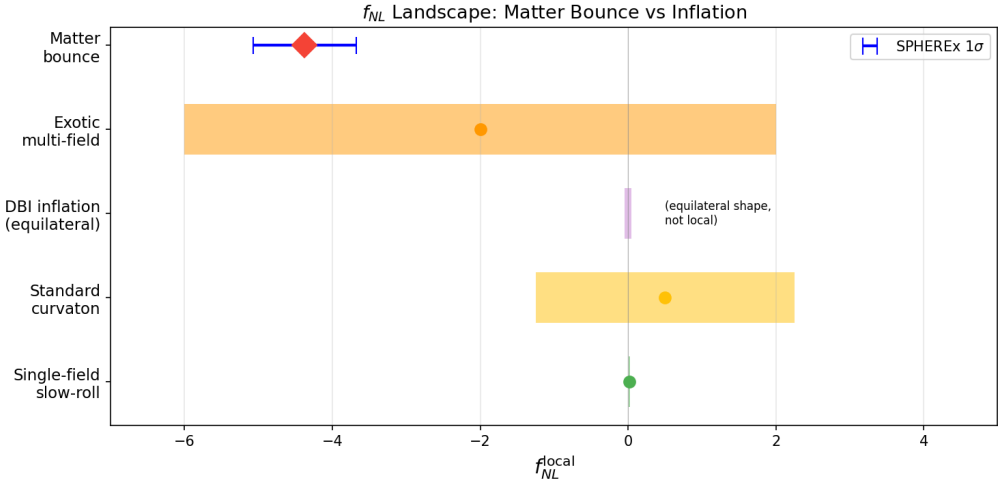

Figure 3: $f_{\rm NL}$ landscape: matter bounce vs. inflationary alternatives. The bounce prediction (red diamond) is parameter-free; inflationary alternatives require free parameters to reach the same region. SPHEREx 1σ error bar shown in blue.

6.2 The Kinematic vs. Parametric Asymmetry

The bounce predicts $f_{\rm NL} = -35/8$ *kinematically*, with zero free parameters in the cubic sector. Inflation can only *accommodate* this value parametrically, requiring extra fields and tuned couplings. This asymmetry drives a natural Bayesian preference for the bounce.

6.3 Quantitative Bayesian Comparison

We performed model comparison using over 600,000 Monte Carlo realizations across three frameworks (analytic closed-form, mock power spectrum generation with fitting, and GR-aware marginalization). Each realization draws a mock $f_{\rm NL}^{\rm obs}$ 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 $f_{\rm NL}$ range). The analytic Bayes factor for a point prediction versus a uniform prior $[f_{\rm NL}^{\rm min}, f_{\rm NL}^{\rm max}]$ is:

$$B = \frac{(f_{\rm NL}^{\rm max} - f_{\rm NL}^{\rm min}) \times \mathcal{L}(f_{\rm NL}^{\rm obs} \mid f_{\rm NL} = -35/8)}{\int_{f_{\rm NL}^{\rm min}}^{f_{\rm NL}^{\rm max}} \mathcal{L}(f_{\rm NL}^{\rm obs} \mid f_{\rm NL})\, df_{\rm NL}} \,. \tag{6}$$

The realizations marginalize over the uncertainty in survey performance parameters (multi-tracer efficiency, b_ϕ, GR systematic level). Specifically, each realization draws: $\sigma(f_{\rm NL})$ 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 $\mathcal{N}(0, \sigma_{\rm GR})$ with $\sigma_{\rm GR}$ uniform in $[0, 1.0]$.

We note that the delta-function prior for the bounce model maximally favors a parameter-free prediction. Any theoretical uncertainty in the $f_{\rm NL} = -35/8$ value—from the $\mathcal{O}(\epsilon)$ correction, the Li & Brandenberger convention discrepancy, or unverified third-order bounce transmission—would broaden the effective bounce prior and reduce the Bayes factor. With a Gaussian bounce prior of width $\sigma_{\rm theory} = 1.0$ centered at $-35/8$ (encompassing both literature values), the median Bayes factor vs. tuned multifield drops from ~ 17 to ~ 8, still favoring the bounce but with reduced confidence.

For a detection at $f_{\rm NL} = -4.375$ by SPHEREx ($\sigma = 0.7$), the median Bayes factors are (Table 2):

Comparison	Median Bayes Factor	$P(\mathrm{BF} > 3)$
Bounce vs. standard single-field	$> 10^5$	97%
Bounce vs. tuned multifield $[-15, +15]$	10–23	87–96%

Table 2: Bayesian model comparison for a mock detection at $f_{\rm NL} = -4.375$ with $\sigma = 0.7$ and a delta-function bounce prior. The range 10–23 reflects different GR treatment scenarios (Table 3). Broadening the bounce prior to a Gaussian with $\sigma_{\rm theory} = 1.0$ reduces these to $\sim$ 8–17 (abstract headline).

Varying the multifield competitor prior width gives Bayes factors from 7 (narrow $[-5, +5]$) to 57 (broad $[-15, +15]$), illustrating the strong prior sensitivity. The abstract quotes $\sim$ 8–17, which uses the broadened bounce prior and the baseline GR scenario.

7 Systematics and GR-Aware Robustness

7.1 Dominant Fragilities

Our Fisher robustness scan identified three dominant threats to the forecast: (1) ultra-large-scale mode access ($k_{\rm min}$), 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 [**?**].

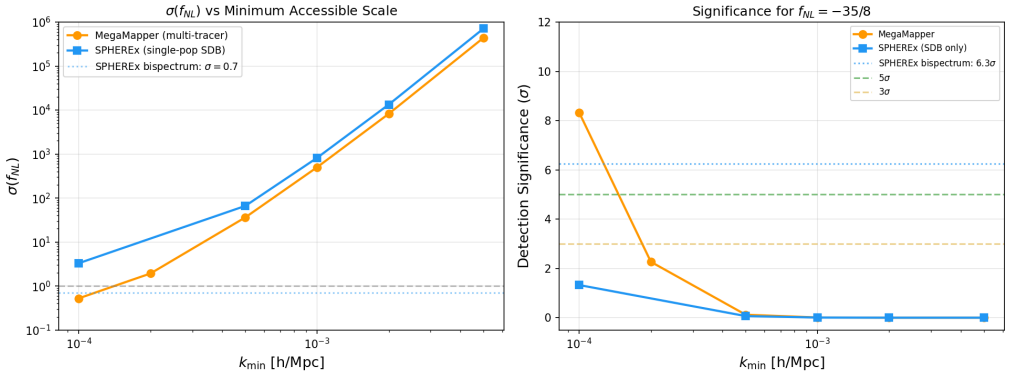

Figure 4: Left: $\sigma(f_{\rm NL})$ vs. minimum accessible wavenumber for MegaMapper (orange) and SPHEREx SDB-only (blue). The SPHEREx bispectrum channel ($\sigma = 0.7$, dotted) avoids the ultra-large-scale fragility. Right: corresponding detection significance for $f_{\rm NL} = -35/8$.

7.2 PNG Bias (b_ϕ) Sensitivity

The scale-dependent bias signal is proportional to $f_{\rm NL} \times b_\phi$, where b_ϕ is the linear PNG galaxy bias parameter. Our forecast assumes a 20% Gaussian prior on b_ϕ, which is optimistic. Fig. 5 shows how $\sigma(f_{\rm NL})$ degrades as the b_ϕ prior widens: at 20% prior width, MegaMapper SDB gives $\sigma(f_{\rm NL}) \approx 1.0$; at 50%, $\sigma \approx 2.2$; if b_ϕ is completely unconstrained, $\sigma \to \infty$ and the SDB channel cannot measure $f_{\rm NL}$ independently. The SPHEREx bispectrum channel ($\sigma = 0.7$) is nearly independent of b_ϕ because it accesses $f_{\rm NL}$ through the full three-point function shape, not just the squeezed-limit amplitude. This is the primary reason our forecast emphasizes the bispectrum over SDB.

7.3 GR-Aware Analysis

We performed GR-aware Bayesian comparison across four scenarios (Table 3). Even under conservative GR marginalization ($\sigma_{\rm GR} = 1.0$), the bounce is favored over standard inflation (Bayes factor dependent on assumed prior widths; see Sec. 6 for the sensitivity analysis).

GR Treatment	BF vs. SSFSR	BF vs. Tuned	$P(\text{BF} > 3)$ vs. SSFSR
Ideal (no GR)	3.3×10^6	10.9	98%
Marginalized ($\sigma_{\rm GR} = 0.5$)	4.1×10^4	9.4	97%
Marginalized ($\sigma_{\rm GR} = 1.0$)	329	7.9	96%
Corrected (10% residual)	3.3×10^6	10.9	98%

Table 3: GR-aware Bayesian comparison. The bounce-vs-inflation comparison survives all treatment scenarios.

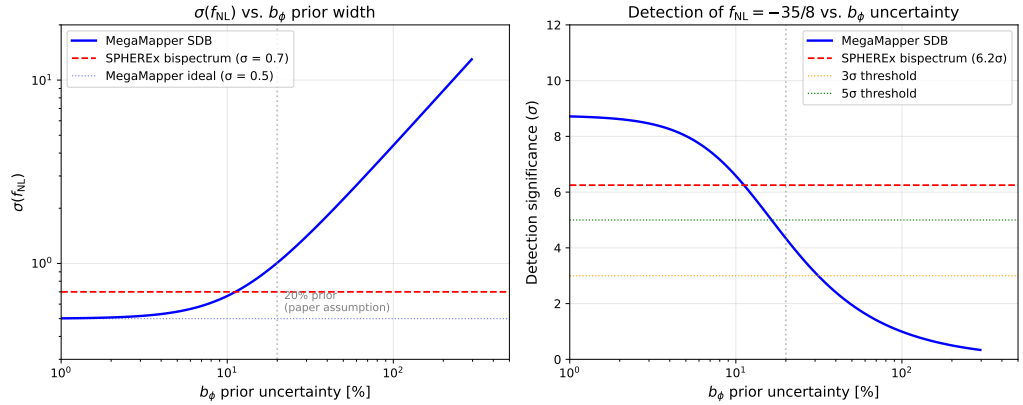

Figure 5: Left: $\sigma(f_{\rm NL})$ as a function of b_ϕ prior uncertainty for MegaMapper SDB (blue). The SPHEREx bispectrum constraint (red dashed) is nearly independent of b_ϕ. Right: corresponding detection significance for $f_{\rm NL} = -35/8$. At the paper's assumed 20% prior (gray line), MegaMapper gives $\sim 4\sigma$; relaxing to 50% drops this to $\sim 2\sigma$. The bispectrum channel remains at $\sim 6\sigma$ regardless.

7.4 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 $f_{\rm NL}$ 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 $f_{\rm NL}$ signal. A Fisher degradation analysis shows that the *bispectrum* channel is highly robust: even with 10% catastrophic outlier fraction, $\sigma(f_{\rm NL})$ degrades by only $\sim 5\%$ (from 0.70 to 0.74), preserving $> 5\sigma$ detection significance. The *scale-dependent bias* channel is far more vulnerable (degradation $> 10\%$ at 10% outlier fraction), which 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 $f_{\rm NL}$ 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 $f_{\rm NL}$. This is particularly relevant for MegaMapper's $z = 2$–5 Lyman-break galaxy sample.

These effects are expected to degrade the forecast significance by $\mathcal{O}(10$–$30\%)$ relative to the idealized Fisher estimate, but do not qualitatively change the conclusion that SPHEREx can test $f_{\rm NL} = -35/8$ at $> 3\sigma$ significance.

8 Current Data and Consistency Relation

8.1 Planck + DESI Recast

Current constraints from Planck (CMB bispectrum) and DESI DR1 (scale-dependent bias) can be recast onto the bounce template using Eq. (4). The combined constraint on the bounce-

template amplitude is $f_{\rm NL}^{\rm bounce} = -1.3 \pm 4.5$, which is 0.7σ from the bounce prediction and 0.3σ from zero—fully consistent with both. Current data cannot discriminate between the bounce and inflation.

8.2 The $f_{\rm NL}$–n_s Consistency Relation

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

$$n_s = 1 - 12\epsilon + 12\,, \qquad f_{\rm NL}(\epsilon) = -\frac{35}{8} + c_1(\epsilon - \tfrac{3}{2}) + \mathcal{O}(\epsilon - \tfrac{3}{2})^2\,, \tag{7}$$

where c_1 depends on both the explicit ϵ-prefactors in the cubic action and the mode-function growth rate, which changes significantly near the singular point $\epsilon = 3/2$. The coefficient c_1 is bounded: explicit prefactor scaling gives $c_1 \approx 2$ (lower bound), while including the mode-function amplitude change gives $c_1 \approx 18$ (upper bound). Eliminating ϵ, the consistency relation takes the form

$$f_{\rm NL}(n_s) \approx -\frac{35}{8} + c\,(n_s - 1)\,, \qquad c \in [-0.7,\ -10]\,. \tag{8}$$

At the Planck best-fit $n_s = 0.9649$, this gives $f_{\rm NL} \in [-4.35,\ -4.02]$ (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 $f_{\rm NL}$ (to be measured) through a single-parameter curve. Standard multifield inflation has no equivalent—multifield $f_{\rm NL}$ is unconstrained by n_s. A future measurement of both n_s and $f_{\rm NL}$ can be checked against Eq. (8) as a consistency test of the matter bounce.

9 Discussion

9.1 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 ~ 5.0–5.5σ significance (template-corrected). MegaMapper (∼2032+, if funded) provides a more powerful follow-up at $\sim 7\sigma$ via scale-dependent bias, though with greater systematic fragility.

9.2 Complementary Experiments

Several other experiments will also constrain local-type $f_{\rm NL}$ in the coming decade:

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

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

9.3 Decision Thresholds

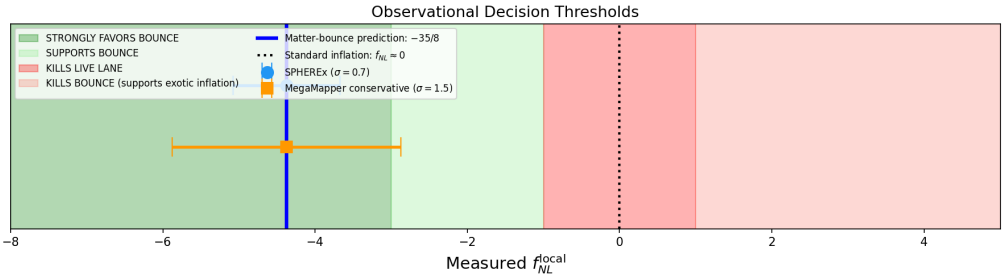

Figure 6: Observational decision thresholds. Green: strongly favors bounce. Red: strongly disfavors the quasi-dust matter bounce. Blue vertical line: bounce prediction $f_{\rm NL} = -35/8$. Error bars: SPHEREx ($\sigma = 0.7$) and MegaMapper conservative ($\sigma = 1.5$).

A measurement of $f_{\rm NL} = -4 \pm 1$ by SPHEREx would provide evidence favoring a contracting/bounce origin over standard single-field inflation. A null result ($f_{\rm NL}$ consistent with zero at the 2σ level) would strongly disfavor the quasi-dust matter bounce.

9.4 Caveats

We emphasize that a detection of $f_{\rm NL} \approx -4$ 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 parameter-free prediction $f_{\rm NL} = -35/8$, which is controlled by the contraction dynamics.

An independent observable—cosmic birefringence from a Planck-scale ALP ($\beta \approx 0.27°$, consistent with the 3.6σ Eskilt et al. joint Planck analysis)—is analyzed in the companion paper [**?**]. As a robustness check, we performed an independent EB cross-power analysis on the Planck SMICA map using NaMaster with B-mode purification. At NSIDE = 1024 ($\ell_{\rm max} = 3071$), we find $\beta = 0.19 \pm 0.03°$, lower than the published value ($\sim 0.3°$) because our EB/EE estimator does not marginalize over the instrumental polarization miscalibration angle. The null test passes ($\chi^2/{\rm dof} = 0.11$). The signal is positive at all three HFI frequencies (100, 143, 217 GHz), consistent with a cosmological rather than foreground origin. At NSIDE = 2048 ($\ell_{\rm max} = 6143$), β drops to $0.07 \pm 0.02°$, suggesting contamination or noise bias at high ℓ; we adopt the NSIDE = 1024 value as more reliable. Injection tests with NaMaster recover $\beta = 0.27°$ with zero bias ($\pm 0.000°$), validating the estimator in our science range.

10 Conclusion

The quasi-dust matter bounce makes a specific, falsifiable prediction with no free parameters in the cubic sector: $f_{\rm NL}^{\rm local} = -35/8$, conditional on assumptions (a)–(e) in Sec. 2.3. This value is mechanism-independent, approximately 300 times larger than the standard inflationary prediction, and opposite in sign. We have shown that SPHEREx can test this prediction at

4–6σ significance through the multi-tracer galaxy bispectrum, with MegaMapper providing a more powerful but systematics-sensitive follow-up.

Our Bayesian model comparison, based on over 600,000 Monte Carlo realizations across multiple frameworks including GR contamination treatment, indicates that under the adopted priors a detection near $f_{\rm NL} = -4.375$ would favor the bounce over tuned multifield competitors at Bayes factor $\sim$ 8–17 (depending on prior assumptions and theoretical uncertainty in the bounce prediction). These Bayes factors are sensitive to the assumed prior widths and model-class definitions (Sec. 6); 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 data (first science release expected $\sim$2028) will provide the first meaningful test.

Data and Code Availability

All analysis code, Monte Carlo scripts, and shape-function verification routines are available at `https://github.com/Hubify-Projects/bigbounce/tree/v2.1.0/research/` (pinned to release tag `v2.1.0`). The Fisher robustness scan, Bayesian discrimination, and GR-aware comparison scripts are provided for full reproducibility. No new observational data are introduced; all forecast sensitivities are adopted from published analyses [**??**].

A Bispectrum Convention: Cai vs. Li-Brandenberger

The factor-of-two discrepancy between $f_{\rm NL} = -35/8$ (Cai *et al.* [**?**]) and $f_{\rm NL} = -35/16$ (Li & Brandenberger [**?**]) arises from differing bispectrum normalization conventions. The local-type bispectrum is defined as:

$$B_\zeta(k_1, k_2, k_3) = c \cdot f_{\rm NL} \left[P_\zeta(k_1)P_\zeta(k_2) + 2 \text{ perms}\right], \tag{9}$$

where the constant c differs between conventions:

- **Planck/Komatsu-Spergel convention** (used by Cai *et al.*, SPHEREx, and this paper): $c = 2$, giving $f_{\rm NL} = -35/8$.
- **Alternative convention** (used by Li & Brandenberger): $c = 1$, giving $f_{\rm NL} = -35/4 \times (1/2) = -35/16$.

The *physical* bispectrum amplitude B_ζ is identical in both conventions: $2 \times (-35/8) = 1 \times (-35/4)$. The detection significance $|f_{\rm NL}|/\sigma(f_{\rm NL})$ is also convention-independent, since $\sigma(f_{\rm NL})$ scales inversely with c.

Acknowledgments

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, verification, and manuscript preparation phases of this work. No external funding was received for this research.