

Testing the Matter Bounce with Primordial Non-Gaussianity: A SPHEREx Sensitivity Recast with a MegaMapper Outlook

Houston Golden^{1, *}

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

(Dated: July 4, 2026)

Scope. This work is a sensitivity recast of a single externally published forecast, not an independent forecast: every quoted SPHEREx significance and Bayes factor rescales the imported Heinrich *et al.* [1] $\sigma(f_{\text{NL}}^{\text{local}}) \approx 0.7$ baseline by the explicit template-mismatch recast $(\hat{f}_{\text{NL}}^{\text{bounce}}, \sigma^{\text{bounce}}) = (\hat{f}_{\text{NL}}^{\text{local}}/r, \sigma^{\text{local}}/r)$ with $r = 0.84$ (Eq. 5) plus a heuristic additive-quadrature systematic budget, with no independent bispectrum Fisher matrix constructed here and cross-parameter correlations neglected in the recast. The headline ranges are therefore conditional sensitivity envelopes, not internally derived measurement precisions; because they are drawn from qualitatively different null procedures (signal-only CMB-Fisher weighting, LSS noise-weighting, and b_ϕ /GR-marginalized budgets), the individual σ endpoints within a range are *not directly comparable* to one another and should be read as a scoping envelope. *Load-bearing caveat* (\star): the entire forecast is conditional on assumption (d) — faithful third-order (cubic) bispectrum transmission through the bounce — which is verified only at linear order [2] and supported at cubic order solely by an order-of-magnitude superhorizon-scaling estimate ($\delta f_{\text{NL}} \sim 10^{-3}$; not a derived bound); this is the single weakest link, and a full cubic in-in computation across an explicit bounce is the #1 follow-up required to convert the prediction into a derived result. With that framing fixed: 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), in the scalar-only matter-bounce class defined by assumptions (a)–(f) of Sec. IIC — in particular (e) no prolonged post-bounce inflation, (f) negligible fermion-sourced torsion during contraction, and (d) faithful third-order bispectrum transmission through the bounce (verified at linear order [2] and supported at cubic order by a superhorizon-scaling estimate). The bounce-vs-inflation discrimination is dual-pronged: in the Planck/local-template gauge convention used by SPHEREx and BOSS, the slow-roll value $f_{\text{NL}}^{\text{inf}} \approx 0.015$ at $n_s = 0.9649$ [3, 4] gives a contrast $|f_{\text{NL}}^{\text{bounce}}|/|f_{\text{NL}}^{\text{inf}}| \approx 290$; in the conformal-Fermi physical-observer frame the squeezed-limit consistency relation [5, 6] sends $f_{\text{NL}}^{\text{local}} \rightarrow 0$ for single-field slow-roll, so a non-zero physical-frame local detection (modulo CFC gradient/projection corrections) would disfavor the single-field slow-roll attractor. The forecast estimators measure the gauge-frame f_{NL} directly; the CFC statement is a complementary theoretical discriminator, not the on-sky observable.

We forecast tests of this prediction with SPHEREx (launched March 2025, primary survey through ~ 2027 , first PNG-suitable release expected ~ 2028) and the proposed MegaMapper via scale-dependent bias and the galaxy bispectrum. A known factor-of-two discrepancy exists in the literature: Cai *et al.* [7] report $f_{\text{NL}}^{\text{local}} = -35/8$ while Li *et al.* [8] report $-35/16$ (at $c_s = 1$), both in the identical $f_{\text{NL}} = 10A/(3 \sum k_i^3)$ normalization and the identical squeezed limit. We audited both source calculations and find their total shape-function polynomials agree coefficient-for-coefficient at $c_s = 1$; the factor of two lies in the local-limit reduction / permutation bookkeeping of that shared shape function and is *not* resolved by our checks: it is neither the Komatsu–Spergel local-template constant (c) convention (both papers already fix the same c) nor a demonstrable dropped in-in time-ordering (a dropped ordering would halve Li’s entire polynomial, which it does not). Settling it requires an independent numerical re-derivation of the full four-vertex in-in integrals, which we do not perform. We therefore adopt the Cai *et al.* value $f_{\text{NL}}^{\text{local}} = -35/8$ as the headline input and carry the Li *et al.* value $-35/16$ as a genuine alternative-literature robustness branch (final paragraph), *not* as a definitively closed convention question. We quantify the template mismatch between the matter-bounce and local templates: a local estimator recovers 83%–88% of the bounce signal across noise-weighting schemes ($r \in [0.829, 0.876]$; CMB Fisher signal-only $r = 0.876$, 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). Only the noise-weighted $r \approx 0.83$ enters the SPHEREx significance; r is applied as a shape-weighted degradation to the Heinrich *et al.* baseline rather than recomputed as an independent cross-Fisher matrix, making this a sensitivity recast rather than an independent forecast.

The Heinrich *et al.* multi-tracer galaxy bispectrum forecast under the local-template normalization $B_\zeta^{\text{local}}(k_1, k_2, k_3) = (6f_{\text{NL}}^{\text{local}}/5)[P_\zeta(k_1)P_\zeta(k_2) + 2 \text{ perms}]$ achieves $\sigma(f_{\text{NL}}^{\text{local}}) \approx 0.7$ [1]. After template-mismatch correction we obtain bispectrum-only $5.2\text{--}5.5\sigma$ at $f_{\text{NL}} = -35/8$ (CMB Fisher to LSS noise-weighted endpoints), reducing to a realistic $\sim 2.6\text{--}5.5\sigma$ after the systematic budget (mismatch, ϵ -correction, polynomial-null-space scatter ± 0.13 in r at fixed sampling measure with 16th-percentile $r_{16\text{th}} = 0.70$ as a distributional robustness bound separate from the conservative

floor which uses the noise-weighted central $r = 0.84$ of Table IV, photometric- z degradation, PNG-bias b_ϕ marginalization, and relativistic projection); these systematics are combined additively in quadrature, assuming the individual contributions are uncorrelated, so the realistic 2.6–5.5 σ range is a scoping sensitivity envelope under an additive-quadrature heuristic systematic budget, *not* a joint-covariance forecasted measurement precision (the true joint constraint could be weaker, or marginally stronger, than this envelope); a full bispectrum joint Fisher over the systematic nuisances is not performed here. Moreover, the imported Heinrich *et al.* multi-tracer covariance was constructed for a purely local template; any additional variance from the non-local tails of the bounce shape in the full multi-tracer estimator covariance is not modeled and is absorbed into this systematic envelope (a full re-derivation of the Heinrich Fisher at the bounce fiducial, which would convert this into a derived bound, is out of scope and flagged for follow-up). Separately, and as an independently computed result kept distinct from this bispectrum-only headline, a joint scale-dependent-bias (SDB) Fisher matrix for $(f_{\text{NL}}, n_{f_{\text{NL}}})$ (`c8_fnl_running_fisher.json`, Planck 2018, CAMB 1.6.6; SDB signal of §III, forecast reported in §IX) shows that the $f_{\text{NL}}-n_{f_{\text{NL}}}$ running degeneracy — and, with the per-sample linear biases co-marginalized, the b_1-f_{NL} degeneracy — *loosen* the SDB f_{NL} constraint by 2.0–4.6 \times ($\sigma_{\text{unmarg}}(f_{\text{NL}}) = 1.53 \rightarrow \sigma_{\text{marg}}(f_{\text{NL}}) = 3.08$ fixed-bias \rightarrow 7.06 bias-marginalized; anti-correlation $\rho \approx -0.87$ to -0.97), with the running the dominant direction. This computed SDB result is a subordinate cross-check on the scale dependence, not a competitor to the bispectrum-only headline, and the bispectrum template-overlap factor $r = 0.84$ does not apply to it. Because the input amplitude itself carries the unresolved Cai–Li factor-of-two, we headline the forecast as an *amplitude-conditional range* spanning both published values, $f_{\text{NL}}^{\text{local}} \in [-35/16, -35/8]$ (Li–Cai): the bispectrum-only significance is 2.6–5.5 σ (Li lower endpoint to Cai upper endpoint) and the realistic post-systematic-budget significance is 1.3–5.5 σ , whose lower bound reflects the conservative Li $-35/16$ input and whose upper bound reflects the optimistic Cai $-35/8$ input. We do *not* headline the optimistic Cai-only 5.2–5.5 σ / 2.6–5.5 σ ranges alone while the factor-of-two is open; those Cai-input values are the upper edge of the reported range (Cai adopted as the fiducial input, Li carried as a genuine downward robustness branch, final paragraph). Both endpoints rest on the single imported Heinrich *et al.* $\sigma(f_{\text{NL}}^{\text{local}}) \approx 0.7$ baseline [1] recast for template mismatch, and are a sensitivity recast rather than an independent cross-Fisher forecast. MegaMapper, proposed but not yet funded, could reach $\sigma(f_{\text{NL}}) \approx 0.5$ ideally, projecting an illustrative 3–7 σ envelope that reflects design uncertainty as much as measurement uncertainty.

A closed-form Bayesian comparison validated across three independent 10^5 -realization Monte Carlo ensembles (§VI, Table II) finds that a SPHEREx detection near $f_{\text{NL}} = -4.375$ favors the bounce over tuned multifield competitors at Bayes factor $\text{BF} \approx 9$ (recommended $\sigma_{\text{theory}} = 1.0$ Gaussian bounce prior, broad multifield $[-15, +15]$ competitor) up to $\text{BF} \approx 14$ at the delta-prior theoretical maximum, evaluated under the noise-weighted $r \approx 0.84$ template-mismatch bookkeeping consistent with the significance headline (the $r \rightarrow 1$ bookkeeping endpoint gives $\text{BF} \approx 10$ –17 and is reported in Table II); Table II reports the $r \rightarrow 1$ endpoint values while the abstract headline applies the noise-weighted $r \approx 0.84$ rebooking $\sigma_{\text{eff}} = \sigma(f_{\text{NL}}^{\text{local}})/r$ to those entries; a curvaton-natural $[-5, +5]$ competitor narrows this to $\text{BF} \approx 4$ –7. These Bayes factors should be read as illustrative of the discriminating power available given the current theoretical uncertainty in the bounce prediction — not as definitive model-selection evidence — and as $\gg 1$ against the parameter-free single-field slow-roll prediction. A SPHEREx null would disfavor the quasi-dust matter bounce benchmark at the same post-systematic-budget level as a detection (~ 1.3 –5.5 σ across the Li–Cai amplitude range; the exclusion arithmetic is symmetric).

The Li *et al.* [8] value $f_{\text{NL}} = -35/16$ is a genuine alternative literature value for the same quantity, in the same normalization and squeezed limit, differing from the Cai *et al.* [7] value $-35/8$ by an unresolved factor of two (Appendix A). We adopt the Cai *et al.* value throughout and assess robustness to the discrepancy as a real stress-test branch: a reader adopting $-35/16$ would see every significance halve — to ~ 2.6 –2.75 σ optimistic and ~ 1.3 –2.75 σ post-systematic-budget — preserving qualitative discriminatory power but losing the $\geq 5\sigma$ headline. Because the factor of two is not definitively resolved here, this branch is a genuine downward robustness bound on the forecast, not merely a hypothetical (see Appendix A.2 and Sec. II C).

I. INTRODUCTION

The inflationary paradigm provides a remarkably successful framework for generating the observed spec-

trum 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 [3] (gauge-frame value; the conformal-Fermi-frame equivalent differs by $\mathcal{O}(\text{slow-roll})$ corrections [5, 6] and is not the bounce-discriminating quantity here, since the bounce-

* houston@hubify.com

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 [9, 10].

A distinctive prediction of the matter bounce is a large, negative, minimally parameterized local-type non-Gaussianity $f_{\text{NL}} = -35/8 = -4.375$ [7, 11] at zeroth order in the ϵ -expansion (the first-order coefficient κ_ϵ has an order-of-magnitude range 5.6–80, Sec. VIII, so the value is tightly determined at leading order but the first correction carries substantial theoretical uncertainty). 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. IIC), 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 robust across the bounce class without prolonged post-bounce inflation (conditional in particular on assumption (d), faithful cubic-order transmission through the bounce, verified only at linear order; Sec. IIC): it depends only on the contracting-phase dynamics, not on the specific UV completion that produces the bounce [2]; it is conditional on the assumptions about the bounce transition listed in Sec. IIC, in particular assumption (e) which restricts the prediction to the Wilson-Ewing class (no prolonged post-bounce inflation). The term “mechanism-independent” as it appears in earlier matter-bounce literature refers to UV-completion independence within this restricted bounce class *and is conditional on faithful cubic-order bispectrum transmission through the bounce* (assumption (d), Sec. IIC, verified only at linear order in Ref. [2]); it is not genuine model independence across the full bounce-cosmology landscape. Bounce models that invoke prolonged post-bounce inflation (e.g., as required by some dark-energy-from-bounce constructions) erase the f_{NL} signal and replace it with the standard slow-roll value, and bounce-transition dynamics that violate faithful third-order transfer would re-introduce mechanism dependence at the order being tested. 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 [12]; Freidel *et al.* [13]), rendering the Barbero-Immirzi parameter invisible in all scalar observables. This decoupling holds in the *scalar-only sector with no fermion matter present*;

if fermion matter is reinstated, the Hehl-Datta-Mercuri four-fermion contact term sourced by $(\bar{\psi}\gamma^5\gamma^a\psi)^2$ activates torsion and breaks the Holst topological-invariance argument, so the Barbero-Immirzi parameter γ_{BI} re-enters scalar observables through the dim-6 four-fermion channel and a possible ΔN_{eff} contribution. The matter-bounce $f_{\text{NL}} = -35/8$ prediction in this work is robust *within the scalar-only matter-bounce class* (see Assumption (f) in Sec. IIC: fermion energy density during contraction and bounce is negligible, so the dim-6 four-fermion operator does not source torsion or reactivate γ_{BI} in the contracting-phase cubic action); the broader “UV-completion-independent” framing should therefore be read as UV-completion independence *within the scalar-only contracting phase*, not as independence across the full ECH operator space. Bounce models with significant fermion sectors during contraction would require an explicit bound on $(\bar{\psi}\gamma^5\gamma^a\psi)^2$ before $f_{\text{NL}} = -35/8$ can be quoted in that broader class; the present forecasts do not apply to such models without that additional input.

The next generation of galaxy surveys—SPHEREx [14] and MegaMapper [15] (a proposed Stage V spectroscopic facility, not yet approved or funded)—will constrain local-type f_{NL} at precision substantially beyond current Planck bounds through the scale-dependent bias effect [16] and the galaxy bispectrum [1]. 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. The length of this manuscript reflects the breadth of that scope — a source-level audit of the $f_{\text{NL}} = -35/8$ prediction, an explicit template-mismatch and null-space quantification, a fully itemized systematic budget, and a closed-form Bayesian comparison cross-validated against three Monte Carlo ensembles — rather than the size of any single incremental result; readers seeking only the headline sensitivity may consult the abstract, Table IV, and Sec. IV.

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 [3] 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 $P \equiv P(k_1, k_2, k_3)$ is a degree-9 homogeneous polynomial in the wavenumbers (throughout this paper,

three-argument P denotes this polynomial only; power spectra always carry a field subscript and a single argument, $P_\zeta(k)$ or $P_\Phi(k)$). The configuration-dependent nonlinearity amplitude, whose squeezed limit is the local nonlinearity parameter, is:

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

where $k \equiv k_2 \approx k_3$ denotes the hard-mode scale and $k_1 \ll k$ is the squeezed (long-wavelength) mode. Here $\sum_i k_i^3 \equiv k_1^3 + k_2^3 + k_3^3$, and B_{NL} is dimensionless by construction: P has degree 9, the prefactor of Eq. (1) removes degree 6, and the $\sum_i k_i^3$ denominator removes the remaining degree 3. Note that B_{NL} retains its full dependence on the coefficients (c_1, \dots, c_6) through P via A_T ; no cancellation of P occurs between Eqs. (1) and (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 $\sum k_i^9$, $\sum_{i \neq j} k_i^7 k_j^2$, $\sum_{i \neq j} k_i^6 k_j^3$, $\sum_{i \neq j} k_i^5 k_j^4$, $\sum_{i \neq j \neq l} k_i^5 k_j^2 k_l^2$, $\sum_{i \neq j \neq l} k_i^4 k_j^3 k_l^2$ with six monomial coefficients (c_1, \dots, c_6) . All orbit sums run over *ordered* index tuples: $\sum_{i \neq j}$ contains six ordered terms, and $\sum_{i \neq j \neq l}$ (all three indices pairwise distinct) contains the six ordered permutations — the convention implemented in the committed evaluation scripts (`null_space_analysis.py`). These six monomials are the *Cai-physics-restricted subset* of the fully S_3 -symmetric degree-9 orbit space that is generated by the matter-bounce cubic-action vertex structure: they are exactly the six S_3 -orbits that appear with non-zero coefficient in the explicit vertex-level derivation of Ref. [7] (their Eq. 37, partitions $(9, 0, 0)$, $(7, 2, 0)$, $(6, 3, 0)$, $(5, 4, 0)$, $(5, 2, 2)$, $(4, 3, 2)$), arising from Wick-contracting the four cubic operators $\mathcal{L}_{\text{redef}}$, $\mathcal{L}_{\zeta\dot{\zeta}^2}$, $\mathcal{L}_{\zeta\partial\zeta\partial\chi}$, $\mathcal{L}_{\zeta(\partial_i\partial_j\chi)^2}$ against the matter-domination Hankel-index mode functions $\zeta_k \propto (1 - ik\eta) e^{ik\eta}/(k\eta)^3$ (App. A). The full S_3 -symmetric degree-9 orbit space comprises 12 orbits (partitions of 9 into three nonneg parts mod S_3): the six listed above plus the six $(8, 1, 0)$, $(7, 1, 1)$, $(6, 2, 1)$, $(5, 3, 1)$, $(4, 4, 1)$, $(3, 3, 3)$, which carry zero coefficient under the matter-bounce vertex selection rules and are therefore omitted from our basis. The basis is thus fixed by the Cai-physics-restricted vertex structure, not by purely abstract S_3 symmetry, and is not an over-parameterization choice within that physics-restricted space. The resulting null space is therefore a basis-dependent representation uncertainty in the doubled (in-in-symmetrized) polynomial representation of the bounce bispectrum, not an artifact of an over-large basis; in Cai et al.’s single-time-ordering form (six coefficients, three benchmarks, but constrained by the explicit vertex-level derivation) the redundancy is fixed by the derivation, while in our doubled representation the same three benchmarks under-constrain the same six symmetric coefficients and propagate as a ± 0.13 scatter in r . **Important scope of the underdetermination claim:**

the six-monomial expansion above is *this paper’s* symmetrization choice for the degree-9 polynomial P , not Cai et al.’s; [7] expresses P in single-time-ordering form prior to the in-in commutator doubling and does not encounter underdetermination at the level of their derivation. The three-constraint vs. six-coefficient mismatch arises specifically when we recompile the doubled polynomial into our symmetrized monomial basis to enable numerical evaluation across triangle configurations; in that internal basis the system is underdetermined and 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 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 (each constraint row is the B_{NL} prefactor times the six symmetric monomials evaluated at the unit-scale benchmark triangles $(k_1, k_2, k_3) = (1, 1, 1)$, $(2, 1, 1)$, and $(10^{-4}, 1, 1)$, with no per-column rescaling). The three benchmark kinematics — equilateral, folded, squeezed — produce well-separated rows; we emphasize that $\sigma_3/\sigma_1 \approx 0.3$ is an empirical property of these benchmark rows in this normalization, far from any rank-deficiency tolerance, and we make no claim of a theoretical lower bound on the conditioning tied to the kinematic separation. The rank is therefore exactly 3 (full row rank) and the null space is exactly 3-dimensional, as required for a

¹ The coefficients printed in Eq. (37) of [7]— $(3, 1, -9, 5, -66, 9)$ —are expressed in that paper’s own monomial normalization, which absorbs Wick-permutation factors differently from the symmetrized basis used here; they are not directly transplantable into our basis (direct evaluation of those coefficients, or their in-doubled values, in our basis does not satisfy the three benchmark constraints; artifact `c9i_epsilon_ratio_check.json`). The linear map between the two bases is set by the ratio of Wick-permutation counting factors. For example, the $(7, 2, 0)$ partition has trivial little group (all three exponents distinct, one of them zero), so the full S_3 orbit of ordered monomials has size $|S_3|/|\text{stab}(7, 2, 0)| = 6/1 = 6$. Cai et al.’s single-time-ordering integrand carries only one of the two complex-conjugate time orderings and the resulting monomial $k_1^7 k_2^2 + k_2^7 k_3^2 + k_3^7 k_1^2$ is the cyclic-subgroup (C_3) orbit of size 3. After in-in commutator doubling (-2Im), the doubled integrand carries the full S_3 orbit. The per-orbit prefactor ratio for this row is therefore $|S_3|/|C_3| = 6/3 = 2$; all six orbit rows are rescaled by their respective permutation-count ratios (orbit-by-orbit, not by a single global factor) before the in-in doubling. The resulting transformation matrix is full-rank (rank 6, verified numerically via `c9i_epsilon_ratio_check.json`) but the direct transplant fails because the per-orbit Wick-factor ratios are orbit-dependent, not a single global rescaling. Our coefficients are therefore fixed from the three published benchmark values themselves. All coefficient sets satisfying those three constraints produce identical B_{NL} at the squeezed, equilateral, and folded configurations but differ at intermediate triangle shapes, contributing the null-space systematic uncertainty to the template overlap quantified in this section.

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 shape cosine is the normalized Euclidean inner product over the triangle grid between the bounce shape $S_{\text{bounce}} = B_{\text{NL}} \times S_{\text{local}}$ (here $S_{\text{local}} \propto 1/k_1^3 k_2^3 + 2$ perms is the dimensionless local-template shape function and S_{templ} is the corresponding unit-normalized local template, $S_{\text{templ}} \propto S_{\text{local}}$) and $r_{\text{cos}} = \sum S_{\text{bounce}} S_{\text{templ}} / \sqrt{\sum S_{\text{bounce}}^2 \sum S_{\text{templ}}^2}$, computed with uniform (unweighted) measure on the dimensionless triangle-ratio grid; the local-estimator Fisher weight $w \propto S_{\text{local}}^2$ enters only the amplitude recovery r , not r_{cos} (implementation in `null.space.analysis.py`, Data and Code Availability). The radius and uniform sampling measure are conventional choices, and the uniform Euclidean measure in this monomial basis is not invariant under linear reparametrizations of the monomials or rotations within the null space — an equally valid basis choice would yield a different ± 0.13 scatter and different extremal r values; the quoted scatter should therefore be read as indicative of the null-space spread under this stated convention rather than as a calibrated, basis-independent uncertainty. The basis-independent robust statement is the shape-cosine stability ($r_{\text{cos}} > 0.95$ across scan radii 10–500, below).

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 $\sim N^3$ growth expected for a three-dimensional (k_1, k_2, k_3) grid under the ordering and triangle-inequality cuts, i.e. $\sim 2^3 = 8\times$ for $50 \rightarrow 100$ and $\sim 4^3 = 64\times$ for $50 \rightarrow 200$ bins per side, not the N^2 of a fixed-shape-ratio grid); the shape cosine r_{cos} changes by $< 0.1\%$ across all three resolutions, confirming convergence. Because both templates are scale-free (the bounce B_{NL} is a degree-0 ratio of homogeneous polynomials and the local template is scale-invariant), the overlap depends only on the triangle shape ratios, not on the absolute $(k_{\text{min}}, k_{\text{max}})$ values or on cosmological parameters; cosmology enters only the ℓ -space Fisher validation channel (fiducial CAMB C_ℓ with a Planck noise model; Sec. III B).

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 en-

hanced 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 at radius 50 ($r_{\text{cos}} = 0.985 \pm 0.007$), 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 $r_{\text{cos}} > 0.97$ bound applies specifically to the radius-50 scan used for the headline distribution; the $r_{\text{cos}} > 0.95$ bound is the conservative floor confirmed across the full multi-radius convergence test (radii 10–500), and the two should not be directly compared as they correspond to different scan volumes.

The amplitude recovery factor is $r = 0.85 \pm 0.13$ (range: 0.55–1.14); the five-coefficient-set scan value ($r = 0.867$ –0.888, CMB-Fisher weighting) lies within the interquartile range [0.75, 0.94] of this null-space distribution, offset from the median 0.85 by ~ 0.03 — consistent given the different weighting and scan measures, though not identically centered. The unweighted distribution has an asymmetric tail: at the 16th percentile $r \approx 0.70$, giving a conservative lower endpoint $4.375 \times 0.70 / 0.7 \approx 4.4\sigma$ (pre-systematic), while the 84th percentile $r \approx 0.99$ gives $\approx 6.2\sigma$ (the [0.75, 0.94] interquartile endpoints correspond to 4.7σ and 5.9σ , consistent with the per-sample 16th–84th band below); these tails are not propagated into the headline significance (which uses the noise-weighted central value $r = 0.84 \pm 0.02$ of Eq. 6), but they set the conservative floor under extreme null-space coefficient variation.² The scatter is 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. Propagating the full scatter onto the detection significance, the 16th–84th percentile range of $|f_{\text{NL}}| r / \sigma(f_{\text{NL}})$ across the 10,000 coefficient samples is 4.4 – 6.2σ (median 5.3σ ; per-sample propagation artifact released with the paper’s code — see Data and Code Availability). These percentiles are computed at the pre-systematic baseline ($\sigma(f_{\text{NL}}) = 0.7$) and should be compared against the

² The $r = 0.85 \pm 0.13$ distribution is computed under uniform Euclidean measure in the monomial coefficient null space at radius 50. The headline noise-weighted $r = 0.84 \pm 0.02$ (Eq. 6) uses CMB/LSS-motivated noise weighting and is not a percentile of this unweighted distribution; the two quantities probe different aspects of the template overlap (null-space coefficient freedom vs. noise-weighting scheme variation) and their central values are consistent to within ± 0.01 .

pre-systematic $5.2\text{--}5.5\sigma$ noise-weighting endpoints at the central r , not against the post-budget 3σ floor of §IV: the 16th-percentile 4.4σ draw sits $\sim 0.8\sigma$ below the pre-systematic $5.2\text{--}5.5\sigma$ band but remains well above the post-systematic floor. Pushed through the conservative GR-marginalization budget ($\sigma_{\text{GR}} = 1.0$ in quadrature with $\sigma(f_{\text{NL}}) = 0.7$), the same 16th-percentile null-space draw maps to $4.4\sigma \times (0.7/\sqrt{0.7^2 + 1.0^2}) \approx 2.5\sigma$, which is below the 3σ GR-only floor and consistent with the $\sim 2.6\text{--}2.7\sigma$ all-combined endpoint of §IV; the headline $5.2\text{--}5.5\sigma$ quotes the two noise-weighting endpoints at the central r , not percentiles of this distribution.

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 [1] as the diagonal noise covariance, and applies a KSW-type optimal linear estimator [17] 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 $\sigma(f_{\text{NL}})$. We stress that this $1/\sqrt{f_{\text{sky}}}$ statement is a CMB-estimator heuristic appropriate only to this two-dimensional flat-sky test; it does not transfer to a three-dimensional galaxy-bispectrum analysis with photometric redshifts, where masking enters through the survey window and mode coupling, and we do not use this scaling in any quantitative forecast. 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 using a two-dimensional flat-sky CMB-style estimator — not the three-dimensional galaxy-bispectrum estimator a SPHEREx analysis would employ — and 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 $B_{\text{NL}}(k_1, k, k)$ (dimensionless, Eq. (2)) as a function of the dimensionless 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 ($k_1=2k_2=2k_3$) ^a	-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. ^aThe folded row sits on the degenerate boundary $k_1 = k_2 + k_3$ and is evaluated as the limit of the sequence $k_1 = 2k$, $k_2 = k_3 = k$.

B. UV-Completion Independence (Conditional on Faithful Cubic-Order Transfer)

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. The prediction is therefore *UV-completion-independent within the Wilson-Ewing class*: any UV completion that produces a non-singular matter-dominated contraction satisfying (a)–(c) and faithfully transfers the cubic-order bispectrum across the bounce reproduces the same $-35/8$ value. We emphasize “UV-completion-independent” rather than the stronger “mechanism-independent” phrasing: faithful third-order (cubic) transmission of the bispectrum through the bounce is itself an assumption (assumption (d), Sec. II C) verified only at linear order [2], and the specific bounce-transition dynamics govern whether (d) holds. The bounce enters only through (i) providing a nonsingular transition and (ii) transferring the contraction-phase cubic correlator into the expanding phase; failures of (ii) at third order would re-introduce mechanism dependence at the order being tested.

C. Assumptions

The $f_{\text{NL}} = -35/8$ prediction rests on six 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; (e) the CMB-observable modes originate from the contracting phase, not from a prolonged post-bounce inflationary epoch; and (f) negligible fermion energy density during the contracting phase and bounce transition (so the Hehl-Datta-Mercuri four-fermion contact term $\langle \bar{\psi}\gamma^5\gamma^a\psi \rangle^2$ does not activate torsion or reactivate the Barbero-Immirzi parameter in the scalar cubic action; the prediction is therefore exact within the scalar-only Einstein-Cartan-Holst class of models, and conditional on fermion contributions during contraction being suppressed). Assumption (f) is the closure of the ECH-decoupling caveat noted by Hehl-Datta-Mercuri: a scalar-only model satisfies it trivially, while bounce models with significant fermion sectors during contraction would require an explicit bound on $\langle \bar{\psi}\gamma^5\gamma^a\psi \rangle^2$ before $f_{\text{NL}} = -35/8$ can be quoted in that broader class. A rigorous order-of-magnitude bound on the threshold $\rho_{\text{fermion}}/\rho_{\text{scalar}}$ requires evaluating the four-fermion contact term $\langle \bar{\psi}\gamma^5\gamma^a\psi \rangle^2$ inside the contracted Einstein-Cartan-Holst cubic action and is not undertaken here. We therefore treat assumption (f) as an externally imposed constraint: the prediction $f_{\text{NL}} = -35/8$ applies to the scalar-only sector, and application to models with substantial fermion populations during contraction would require a per-model bound on this operator before the same numerical value can be quoted. 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., Zhu & Cai [18]) 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 $\epsilon = 3/2$. The matter-contraction Hankel index ν is finite at $\epsilon = 3/2$ ($\nu = 3/2 + \mathcal{O}(\epsilon - 3/2)$, Wilson-Ewing dust contraction); the sensitivity of cubic-order quantities to ϵ near this point is driven not by a divergent Hankel index but by the explicit $A_T \propto 1/\epsilon^3$ prefactor channel and by the $|\eta|^{-\nu}$ mode-function amplitude growth. Explicit cubic-action prefactors give a correction of $\sim 0.6\%$, but the mode-function growth rate also changes with ϵ , potentially am-

plifying 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 [2]. At cubic order, a semi-analytic order-of-magnitude estimate based on the superhorizon approximation for mode functions near the LQC bounce suggests 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 if the superhorizon scaling holds; this is a scaling estimate, not a derived bound). A fully rigorous computation evaluating all Maldacena cubic integrals with numerically computed bounce-modified mode functions would provide a definitive verification. Assumption (d) is the weakest link of the present derivation: it is verified at linear order [2] and supported at cubic order only by the order-of-magnitude superhorizon scaling argument above, so every “robust across the bounce class” statement in this paper is conditional on assumption (d), and only the full numerical computation just described would upgrade it to a derived bound. A genuine factor-of-two discrepancy exists in the literature: Li *et al.* [8] obtain $f_{\text{NL}} = -35/16 = -2.1875$ when evaluated at $c_s = 1$ (their Eq. 5.1, $f_{\text{NL}}^{\text{local}} = -165/16 + 65/8c_s^2$), against Cai *et al.*'s [7] $f_{\text{NL}} = -35/8 = -4.375$ (their Eqs. 38–39). We performed a source-to-source audit of both published calculations and report its outcome honestly, including its limits. Both papers use the *identical* local-template normalization $f_{\text{NL}} = 10 A/(3\sum k_i^3)$ (Cai Eq. 21 = Li Eq. 4.20) with $\Phi = \frac{3}{5}\zeta$ (Cai Eq. 20), and both quote the *same* squeezed local limit $k_1 \ll k_2 = k_3$; so the factor of two is neither the Komatsu-Spergel local-template constant (c) convention (both papers already fix the same c ; see Appendix A) nor a difference of limit or normalization. Moreover, the two papers' *total shape-function polynomials agree coefficient-for-coefficient* at $c_s = 1$: Cai's Eq. (37) prefactor $3/256$ with monomial coefficients $(1, -9, +5, -66, +9)$ matches Li's Eq. (4.19) at $c_s = 1$ term for term (Li's $(10 - 9c_s^2) \rightarrow 1$, etc.). A dropped in-in time-ordering would halve the *entire* polynomial, not just the local-limit amplitude; because Li's full polynomial is not halved relative to Cai's, the discrepancy *cannot* be attributed to Li omitting the second time-ordering of the -2Im commutator. The residual factor of two therefore resides in the local-limit reduction / permutation book-keeping by which the dominant local amplitude is extracted from that shared shape function, and *its origin*

is not settled by the checks in this paper. The generic in-in commutator identity $i\langle[\zeta^3, L]\rangle = -2\text{Im}\langle\zeta^3 L\rangle$ (verified symbolically in Appendix A) is a true operator-algebra statement, but it does *not* establish that either published value omitted an ordering, and we no longer claim it resolves the Cai/Li factor of two. A complete independent re-derivation of the in-in bispectrum integral from the vertex-level Maldacena action—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 superhorizon divergences $\zeta \sim |\eta|^{-3}$, and preserving the inter-vertex cancellations that render the physical bispectrum finite—would be required to determine which local-limit value is correct; this specialized calculation has to our knowledge been carried out in full only by Cai et al. [7] and Li *et al.* [8], and we do not undertake it here. We therefore *adopt* the Cai et al. value $f_{\text{NL}} = -35/8$ as the headline input—benchmark-configuration matching at three independent momentum triangles (Table I) and null-space stability ($r_{\text{cos}} = 0.985 \pm 0.007$ across 10,000 coefficient samples) confirm we reproduce the Cai et al. shape function—while *explicitly disclosing* that the Cai/Li factor of two is an unresolved literature discrepancy and carrying $-35/16$ as a genuine downward robustness branch (Appendix A, Table V), under which every significance halves.

D. The Viable Model

The Wilson-Ewing Λ CDM quasi-dust model [2] 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 [2], so n_s is a fit to the data rather than a prediction), a tensor-to-scalar ratio $r_t \approx 10^{-4}$ (from LQC quantum-geometry tensor suppression; we write r_t for the tensor-to-scalar ratio throughout to avoid collision with the template-overlap amplitude recovery factor r of Secs. II–III B), 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)–(f) in Sec. II C). We are not aware of observational tensions with this model within current uncertainties.

III. OBSERVABLE MAPPING TO LARGE-SCALE STRUCTURE

A. Scale-Dependent Bias

Primordial local non-Gaussianity induces a scale-dependent correction to galaxy bias [16, 19]:

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

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 $D(0) = 1$), wavenumbers k are comoving and quoted in $h \text{ Mpc}^{-1}$ throughout (we work in units $c = 1$ and likewise express H_0 in $h \text{ Mpc}^{-1}$, so that \mathcal{M} is dimensionless and Δb carries no residual units), $\delta_c \approx 1.686$ is the spherical-collapse threshold, and b_1 is the linear Eulerian galaxy bias. Since $\Delta b \propto 1/\mathcal{M}$ and $\mathcal{M} \propto k^2$ on ultra-large scales (where $T(k) \rightarrow 1$), the signal grows as $\Delta b \propto 1/k^2$ as $k \rightarrow 0$ [16, 19]; 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:

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

Only the noise-weighted central value $r = 0.84 \pm 0.02$ (Eq. 6) enters σ_{eff} and every headline significance in this paper; the broader null-space distribution $r = 0.85 \pm 0.13$ (range 0.55–1.14; Sec. II) is a basis-measure stress band that is never propagated into $\sigma_{\text{eff}} = \sigma(f_{\text{NL}}^{\text{local}})/r$, so no $r > 1$ sample can shrink the effective uncertainty below the local-template baseline. Here 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 $r \lesssim 1.14$ 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.³ Equation (5) is exact in the Fisher limit where the noise covariance is diagonal in bispectrum space and the estimator is 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 the shape cosine r_{cos} (computed in the unweighted shape metric of Sec. II, not in the survey Fisher metric of the estimator: the suppression is a heuristic shape-similarity indicator, not an estimator-mismatch variance bound under SPHEREx weighting): using the mean $r_{\text{cos}} \approx 0.985$ from the 10,000-sample scan, $1 - r_{\text{cos}}^2 \approx 0.03$; using only the conservative lower bound $r_{\text{cos}} > 0.97$, $1 - r_{\text{cos}}^2 < 0.06$. Either bound confirms the projection noise is subdominant to the other systematics in our budget.

Using the physics-derived polynomial, we computed r under 10 physically motivated weighting schemes (uniform, $w = 1$; CMB Fisher, $w \propto k^2$; LSS scale-dependent-bias, $w \propto 1/k^2$; SPHEREx-like and MegaMapper-like survey-noise variants; and five region-masked variants, where w is the Fisher weight entering the weighted average $r = \langle B_{\text{NL}} \rangle_w / D_{\text{NL}}^{\text{squeeze}}$ defined below Eq. 5), scanning over squeezed cutoffs ($x_{3,\text{min}}$ from 0.001 to 0.2, where $x_3 \equiv k_3/k_1$ is the squeezed-limit ratio with $x_3 \rightarrow 0$ corresponding to the squeezed limit $k_3 \ll k_1 \approx k_2$). *Index-labeling note:* the overlap scan parametrizes triangles with k_1 held fixed as a hard reference scale and k_3 taken as the long (squeezed) mode (so $x_3 = k_3/k_1 \rightarrow 0$ is the squeezed limit, with $k_3 \leq k_2 \leq k_1$ in the scan grid); this interchanges the index roles relative to the benchmark convention of Sec. II, where k_1 denotes the long mode

and $k_2 \approx k_3$ the hard modes. The relabeling is purely a parametrization choice — B_{NL} is permutation-symmetric in (k_1, k_2, k_3) , so the squeezed-cutoff result below is unchanged by which index carries the long mode. The result is robust:

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

with the range $r \in [0.829, 0.876]$ spanning all physically motivated weighting schemes (the three noise-weighted values 0.829 [scale-dependent-bias, $1/k^2$], 0.830 [SPHEREx-like], 0.835 [flat/uniform], together with the signal-only CMB-Fisher value 0.876 — four values total; the per-realization spread from the ℓ -space Fisher-overlap artifact (Data and Code Availability) is wider, [0.856, 0.895], but is dominated by Monte-Carlo noise of the ℓ -space Fisher integrand rather than weighting-scheme variation). The ± 0.02 of Eq. (6) characterizes the spread about the central 0.84 dominated by the noise-weighted schemes; the signal-only CMB-Fisher endpoint 0.876 lies just outside that band and is therefore always quoted separately as the optimistic endpoint rather than absorbed into the noise-weighted uncertainty. 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 relative to the signal-only CMB-Fisher weighting — which preferentially upweights the squeezed configurations where the bounce and local templates coincide — the LSS noise-weighting shifts relative weight onto the intermediate and folded configurations where their integrated mismatch is largest, so the overlap drops below the signal-only CMB-Fisher endpoint; this $r_{\text{LSS}} < r_{\text{CMB}}$ ordering emerges from the full Fisher-weighted shape integration (`null_space_analysis.py`, Data and Code Availability). 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. (This cutoff test is not in tension with the ~ 0.01 shift under the log-weighted squeezed-enhanced grid of Sec. II: the cutoff test removes the most extreme squeezed triangles at fixed sampling measure — configurations where the two templates agree and which therefore carry almost no mismatch weight — whereas the grid-reweighting test changes the sampling density across *all* configurations, including the intermediate and folded regions that dominate the mismatch; the two procedures probe different sensitivities and their different magnitudes are expected.) Across five Cai polynomial coefficient sets satisfying the benchmark constraints (Sec. II),

³ 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.

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 if one uses a local-template estimator; a shape-matched estimator optimized for the bounce bispectrum could recover $r \rightarrow 1$ at the cost of losing the local-template interpretability that connects to standard f_{NL} constraints. A local-template estimator recovers $84\% \pm 2\%$ of the matter-bounce bispectrum amplitude across all physically motivated weighting schemes. We cross-checked the overlap at three independent levels; the ℓ -space Fisher result has limited commensurability with the 3D LSS bispectrum and should be read as a consistency check rather than an independent validation: (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). The full bispectrum-shape coefficient map and the per-configuration overlap values for all six monomial coefficient sets satisfying the Cai et al. benchmarks are archived as `phase3_bispectrum_shape_overlap.json` (released with the paper’s code; see Data and Code Availability), enabling independent reproduction of the $r = 0.84 \pm 0.02$ noise-weighted central value.

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 $\sim 2.6\text{--}5.5\sigma$ (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 [1]. 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 ($0.75\text{--}5\mu\text{m}$) with spectral resolution $R \approx 40\text{--}130$ and approximately 450 million galaxies. DBI inflation is intentionally not propagated through the local-template SPHEREx forecast in this section: it has vanishing squeezed-limit local amplitude and an equilateral shape, so a local-template estimator does not discriminate against it; the relevant bispectrum-shape comparison against DBI is deferred to §VI (where the multifield-tuned competitor envelope is drawn) and the DBI scoping discussion in §VI (the parenthetical at the end of the joint $(f_{\text{NL}}, n_{f_{\text{NL}}})$ subsection). A dedicated multi-tracer bispectrum analysis [1], building on the canonical SPHEREx multi-tracer forecast lineage of Doré *et al.* [14] (the foundational SPHEREx galaxy-survey forecast paper; Münchmeyer *et al.* [20] is the CMB kinetic-Sunyaev-Zel’dovich tomography companion forecast for the same parameter target), 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 Heinrich *et al.* Fisher forecast is constructed at the local-template fiducial $f_{\text{NL}} = 0$; applying the resulting $\sigma(f_{\text{NL}}) \approx 0.7$ at the bounce-fiducial $f_{\text{NL}} = -4.375$ relies on the leading-order linearization that the Fisher matrix is approximately invariant under fiducial shifts of order the parameter uncertainty (a standard but non-trivial Fisher-forecast assumption). As an order-of-magnitude check: the leading non-Gaussian correction to the bispectrum covariance from a non-zero fiducial f_{NL} enters at second order in f_{NL} through the connected six-point function. The dimensionless fractional covariance correction at fixed k then scales as

$$\frac{\delta C}{C_{\text{Gauss}}} \sim \frac{f_{\text{NL}}^2 \Delta_\zeta^2(k)}{N_{\text{modes}}(k)}, \quad (7)$$

where $\Delta_\zeta^2(k) \equiv k^3 P_\zeta(k)/(2\pi^2) \approx 2.1 \times 10^{-9}$ is the dimensionless curvature power spectrum and $N_{\text{modes}}(k) \sim V_{\text{survey}} k^2 \delta k/(2\pi^2)$ is the dimensionless mode count in a thin k -shell. At SPHEREx scales $k \sim 0.01\text{--}0.1 h \text{Mpc}^{-1}$ with $f_{\text{NL}} = -4.375$, $f_{\text{NL}}^2 \Delta_\zeta^2 \sim 4 \times 10^{-8}$, so $\delta C/C \ll 10^{-3}$ even before dividing by the mode count. The propagated fractional shift in $\sigma(f_{\text{NL}})$ follows from $\delta\sigma/\sigma \sim \frac{1}{2} \delta C/C$ and is therefore $\lesssim 5 \times 10^{-4}$, well below the percent level. Eq. (7) is a heuristic primordial-field scaling check in the curvature perturbation ζ and is *not* a covariance-level derivation in the galaxy field: it does not include linear-bias, quadratic-bias, shot-noise, or the six-point galaxy-bispectrum covariance terms that enter the full multi-tracer covariance of Ref. [1], and the quoted $\lesssim 5 \times 10^{-4}$

figure should be read as a heuristic dimensional check rather than a derived bound. A full re-derivation of the Heinrich Fisher matrix at the bounce fiducial would convert this scaling check into a derived bound, but requires the full galaxy-field survey covariance machinery of Ref. [1]. The multi-tracer cosmic-variance cancellation invoked here originates with the multi-tracer power-spectrum technique of Seljak [21] and McDonald & Seljak [22]; the bispectrum-multi-tracer extension followed by Karagiannis *et al.* [23] 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 assumes 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 $r = 0.84$; Eq. 6) to 3.0σ (the conservative floor reported here is defined as: $\sigma_{\text{GR}} = 1.0$ added in quadrature to the baseline $\sigma(f_{\text{NL}}) = 0.7$, without the additional widened b_ϕ prior; explicitly, $4.375 \times 0.84 / \sqrt{0.7^2 + 1.0^2} \approx 3.00\sigma$; using the slightly lower noise-weighted endpoint $r = 0.829$ gives $\approx 2.97\sigma \approx 3.0\sigma$, so the floor is robust to the noise-weighting choice across the $r = 0.829\text{--}0.84$ noise-weighted range). Adding the widened b_ϕ prior of §VII (which moves the combined-sample baseline $\sigma(f_{\text{NL}})$ from 0.7 to 0.9–1.0; the $\sigma(f_{\text{NL}}) = 0.7$ Heinrich *et al.* baseline is the full multi-tracer combined-sample forecast, not a per-redshift-bin uncertainty) brings the all-combined endpoint to $\sim 2.6\text{--}2.7\sigma$; this is the honest cumulative-systematics endpoint and is reported here as such, distinct from the 3.0σ GR-only floor of the headline range. The realistic range after the combined systematic bud-

get is $\sim 2.6\text{--}5.5\sigma$, with the lower endpoint reflecting the all-combined cumulative budget. 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 [24]; Liang *et al.* [25] 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; the quoted gain is an upper bound pending the shot-noise-corrected Fisher analysis described in the shot-noise caveat below); 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: the naive Poisson amplitude scaling gives $\sigma(f_{\text{NL}})$ inflated by $\sqrt{11} \approx 3.3\times$ relative to the shot-noise-free Fisher limit at $\bar{n} \sim 10^{-5}$ (scaling as $\sigma(f_{\text{NL}})_{\text{shot}}/\sigma(f_{\text{NL}})_{\text{CV}} \sim \sqrt{1 + 1/(\bar{n}P_0)}$ where $P_0 \sim 10^4 h^{-3} \text{Mpc}^3$; for $\bar{n} \sim 10^{-5}$, $1/(\bar{n}P_0) \sim 10$), while the bispectrum-estimator effective degradation at the squeezed-limit modes that dominate f_{NL} sensitivity is the more relevant $\sim 15\text{--}30\%$ at $z \sim 1\text{--}2$, because the squeezed limit preferentially downweights the highest- k modes where shot noise is most severe, and would need to be included in any definitive forecast based on anomaly-selected subsamples. For the 10% catastrophic photo- z outlier fraction mentioned above, the 5% degradation in the bispectrum $\sigma(f_{\text{NL}})$ follows from a first-order Fisher correction: catastrophic outliers scatter modes across the $1/k^2$ bias kernel at rate $f_{\text{cat}} = 0.1$, diluting the effective PNG signal by $\sim f_{\text{cat}}^2/(1 + f_{\text{cat}})^2 \approx 0.008$, well below the 5% level; the dominant effect is a net smearing of the cross-bin Fisher matrix at the $\sim 5\%$ level consistent with published photo- z degradation estimates [26]. The headline $2.6\text{--}5.5\sigma$ 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, plotted at the template-corrected values used in the text. SPHEREx bars (left to right): the naive uncorrected $|f_{\text{NL}}|/\sigma(f_{\text{NL}}) = 4.375/0.70 = 6.25\sigma$ (hatched gray; shown only for reference, *not used in any headline*); the template-corrected optimistic bispectrum forecast $5.2\text{--}5.5\sigma$ ($|f_{\text{NL}}|/\sigma$ with noise-weighted $r = 0.84$, §IV); the realistic post-systematic-budget envelope $2.6\text{--}5.5\sigma$ (§VII); and the all-combined conservative endpoint $2.6\text{--}2.7\sigma$ (effective σ_{eff} ranges from $\sqrt{0.9^2 + 1.0^2} = 1.35$ at the central b_ϕ -widened scenario to $\sqrt{1.0^2 + 1.0^2} = 1.41$ at the conservative endpoint; §IV). MegaMapper bars: template-corrected ideal $7.4\text{--}7.7\sigma$ (§V); the illustrative $3\text{--}7\sigma$ design-uncertainty envelope; conservative; and single-tracer. Error bars span the optimistic endpoint (published ideal $\sigma(f_{\text{NL}})$ with template-overlap correction only) to the conservative endpoint (full §VII budget: $r = 0.84$ overlap, ϵ -correction, photometric- z degradation, PNG bias, b_ϕ marginalization, and GR-projection marginalization (σ_{GR} , Table III)).

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 [15]. 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 $2.6\text{--}5.5\sigma$ after the same GR marginalization and b_ϕ uncertainty budget applied to SPHEREx above (for illustration only; this budget is not independently calibrated to MegaMapper’s higher-redshift $z = 2\text{--}5$ sensitivities, where GR projection effects are expected to be substantially larger [27]). At an intermediate $\sigma(f_{\text{NL}}) = 0.7$ (allowing for partial systematic degradation), the template-corrected significance is $\sim 5.2\sigma$ optimistic, $\sim 3.2\sigma$ conservative (combining the noise-

weighted template overlap $r = 0.84$, $\sigma(f_{\text{NL}}) = 0.7$, and a 30% b_ϕ prior widening that moves the combined-sample baseline $\sigma(f_{\text{NL}})$ to ≈ 0.9 : $4.375 \times 0.84/\sqrt{0.7^2 + 0.9^2} \approx 3.2\sigma$; with $\sigma_{\text{GR}} = 1.0$ and no b_ϕ widening the floor is $\sim 3.0\sigma$; this is a separate illustrative stress test using Table IV combination rules, not independently calibrated to MegaMapper’s higher- z systematics). 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 $z > 2$ [27]; (b) PNG bias parameter b_ϕ uncertainty, which can degrade constraints if uncalibrated [28]; 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 $f_{\text{NL}} = (5/12)(1 - n_s) \approx +0.015$ [3] as the gauge-frame value of the consistency relation; in conformal Fermi coordinates the physical observable is parametrically smaller [5, 6]. 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) [29]. 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 $\text{BF} \approx 4$ (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 three independent Monte Carlo ensembles of 10^5 realizations each (3×10^5 aggregate) across three frameworks (analytic closed-form, mock power spectrum generation with fitting, and parameterized GR-contamination marginalization). The three ensembles use *framework-specific priors* (delta-vs-flat $[-15, +15]$ for the analytic closed-form; lognormal σ -base for the mock-validation; σ_{GR} Gaussian for the GR-contamination marginalization) and are not three frames of a single ensemble—they are three cross-checks of the analytic Bayes-factor formula. The four-corner Bayes-factor prior grid reported in this section (curvaton-natural $[-5, +5]$ vs. broad multifield $[-15, +15]$ competitor priors \times delta vs. $\sigma_{\text{theory}} = 1.0$ Gaussian bounce priors) is computed analytically via the closed-form Bayes-factor integration (the analytic formula given later in this section); the Monte Carlo ensembles span only a subset of these corners (the delta-vs-flat $[-15, +15]$ corner and the GR-marginalization variation), and the curvaton-natural $[-5, +5]$ column plus the $\sigma_{\text{theory}} = 1.0$ Gaussian row are derived analytically from the same closed-form integration without a dedicated Monte Carlo scan. 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. The closed-form values quoted in this section and in Tables II and III are deterministic integrals carrying no Monte Carlo uncertainty; the released recompute script additionally reports realization-marginalized mean/median/geometric-mean Bayes factors over 2×10^5 draws per σ_{eff} as a convergence and stability check (e.g., for the delta-prior/narrow-competitor cell at $\sigma_{\text{eff}} = 0.7$, closed-form 7.0 vs. realization-marginalized median 6.6), with realization marginalization always lowering the central value; the spreads quoted in this paper come from prior and scenario variation, not from Monte Carlo noise. 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:

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

a. Closed-form derivation of the Bayes factor. For a Gaussian likelihood $\mathcal{L}(f_{\text{NL}}^{\text{obs}} | f_{\text{NL}}) = \mathcal{N}(f_{\text{NL}}^{\text{obs}}; f_{\text{NL}}, \sigma_{\text{eff}}^2)$ with $f_{\text{NL}}^{\text{obs}} = -35/8$ (mock detection at the bounce prediction) and a uniform competitor prior on $W \equiv f_{\text{NL}}^{\text{max}} - f_{\text{NL}}^{\text{min}}$, the denominator of Eq. (8) evaluates to an error-function difference:

$$\begin{aligned} B &= \frac{W \cdot \mathcal{N}(-35/8; -35/8, \sigma_{\text{eff}}^2)}{\int_{f_{\text{NL}}^{\text{min}}}^{f_{\text{NL}}^{\text{max}}} \mathcal{N}(-35/8; f_{\text{NL}}, \sigma_{\text{eff}}^2) df_{\text{NL}}} \\ &= \frac{W / (\sqrt{2\pi} \sigma_{\text{eff}})}{\Phi\left(\frac{f_{\text{NL}}^{\text{max}} + 35/8}{\sigma_{\text{eff}}}\right) - \Phi\left(\frac{f_{\text{NL}}^{\text{min}} + 35/8}{\sigma_{\text{eff}}}\right)}, \end{aligned} \quad (9)$$

where $\Phi(z) \equiv \frac{1}{2}[1 + \text{erf}(z/\sqrt{2})]$ is the standard normal CDF. This is the *exact* closed-form result; all tabled Bayes factors are computed from Eq. (9) without approximation.

In the large- W limit where the competitor prior is broad relative to σ_{eff} and the bounce prediction $-35/8$ is well inside the prior interval, both CDF arguments are large in magnitude and the denominator approaches unity; the Bayes factor reduces to the Gaussian-peak approximation:

$$B \approx \frac{W}{\sqrt{2\pi} \sigma_{\text{eff}}} \quad (W \gg \sigma_{\text{eff}}). \quad (10)$$

b. Numerical self-consistency check. For the four tabled values at $\sigma_{\text{eff}} = 0.7$ (the $r \rightarrow 1$ no-rebooking endpoint):

- Delta prior, broad $[-15, +15]$ ($W = 30$, $\sigma_{\text{eff}} = 0.7$): exact $B = 17.10$; approx $30/(\sqrt{2\pi} \times 0.7) = 17.07$ — error 0.18% (sub-percent; both CDF tails are $< 10^{-7}$).
- Delta prior, $\sigma_{\text{eff}} = 0.833$ ($r = 0.84$ rebooked): exact $B = 14.36$; approx $30/(\sqrt{2\pi} \times 0.833) = 14.34$ — error 0.14%.
- Gaussian bounce prior $\sigma_{\text{theory}} = 1.0$, broad $[-15, +15]$, $\sigma_{\text{eff}} = 0.7$: exact $B = 9.80$; approx $= 17.10/\mathcal{C}$ where \mathcal{C} is the prior-convolution correction (not the simple $W/(\sqrt{2\pi}\sigma_{\text{eff}})$ formula, which applies only to the delta-prior case). Equation (10) thus applies to the *delta-prior* corner only; the Gaussian-bounce-prior corners require the full Eq. (9) with the bounce likelihood replaced by the prior-convolved marginal.
- **Gaussian-bounce prior** ($\sigma_{\text{theory}} = 1.0$), narrow $[-5, +5]$ competitor ($W = 10$, $\sigma_{\text{eff}} = 0.7$): exact $B = 4.01$ from Eq. (9) with the bounce likelihood replaced by the prior-convolved marginal (Gaussian-bounce-prior result; the closed-form CDF integration over the $\sigma_{\text{theory}} = 1.0$ Gaussian is required here). *Note:* Eq. (10) is the **delta-prior** large- W approximation ($B \approx W/(\sqrt{2\pi}\sigma_{\text{eff}})$) and is *not applicable* to the Gaussian-bounce-prior case. As a self-consistency cross-check only: if one incorrectly applied the delta-prior formula Eq. (10) to this narrow-competitor cell, it would give $10/(\sqrt{2\pi} \times 0.7) = 5.69$ versus the exact Gaussian-bounce-prior $B = 4.01$ — a 42% discrepancy arising from two compounding sources: (i) the prior mismatch (delta vs. Gaussian bounce prior), and (ii) non-negligible CDF tail terms with $W = 10$ and $\sigma_{\text{eff}} = 0.7$ (the competitor-prior denominator is $\Phi((5 + 35/8)/0.7) - \Phi((-5 + 35/8)/0.7) = \Phi(13.4) - \Phi(-0.893) \approx 1 - 0.186 = 0.814$, so the finite lower tail ≈ 0.186 — the upper tail is negligible — reduces the denominator by $\approx 18\%$; for the delta-prior narrow case this *raises* B from the large- W approximation 5.69 to the exact $5.69/0.814 \approx 7.0$, while for the Gaussian-bounce case the reduction below 5.69 to 4.01 is dominated by the prior-width broadening). The narrow-prior $[-5, +5]$ column therefore uses Eq. (9) exclusively for both the delta-prior row ($B \approx 7.0$, exact CDF) and the Gaussian-bounce-prior row ($B = 4.01$, exact prior-convolved CDF); Eq. (10) is *not valid* for any narrow-competitor entry. All tabled narrow-prior values (4.7–7.0 and 4.01) are exact CDF evaluations under their respective prior specifications.

To summarize the equation-to-prior assignment unambiguously: Eq. (9) (the closed-form CDF integration) is the primary formula for *all* tabled entries. Eq. (10) ($B \approx W/(\sqrt{2\pi}\sigma_{\text{eff}})$) is a large- W approximation applicable *only* within the delta-prior row, and only for the broad

$[-15, +15]$ competitor where the 0.18% sub-percent approximation error confirms the large- W condition is met; for the narrow $[-5, +5]$ competitor at $W = 10$ and $\sigma_{\text{eff}} = 0.7$, Eq. (10) gives only 5.69 (the finite lower CDF tail raises the exact result to $B \approx 7.0$ relative to the large- W approximation, because the tail reduces the competitor-prior denominator, increasing the Bayes factor), so the narrow delta-prior entry $B \approx 7.0$ is an exact CDF result from Eq. (9), not from Eq. (10). Eq. (9) is *required* (not optional) for every Gaussian-bounce-prior entry: the Gaussian-bounce-prior narrow-competitor result $B = 4.01$ comes from Eq. (9) with the bounce likelihood replaced by the prior-convolved marginal, and applying Eq. (10) to that cell is a category error that inflates the result by 42% (giving the spurious 5.69). Readers comparing rows across Table II should therefore note that the delta-prior row and the Gaussian-bounce-prior row differ in both prior specification *and* required equation: the Gaussian row uses Eq. (9) exclusively, and within the delta row only the broad $[-15, +15]$ entry permits the Eq. (10) approximation as a cross-check.

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 (0.6–8%; Sec. II C), the nearly order-of-magnitude range in κ_ϵ (Sec. VIII), or unverified third-order bounce transmission—would broaden the effective bounce prior and therefore reduce the Bayes factor. (The Li *et al.* alternative-literature value $-35/16$ is *not* folded into this physical-uncertainty budget—it is a separate unresolved factor-of-two literature discrepancy (App. A), tracked as a distinct robustness branch—but we report where it falls under each prior.) We recommend the $\sigma_{\text{theory}} = 1.0$ case as the most physically motivated baseline: it covers the full ϵ -correction range comfortably (the 0.6–8% shift is $\lesssim 0.35$ in f_{NL} , well inside 1σ), while the Li *et al.* robustness-branch value -2.1875 lies at 2.19σ under this prior — bracketed within its $\sim 2\sigma$ tail but *not* within 1σ (full 1σ encompassment of the stress-test value requires $\sigma_{\text{theory}} \gtrsim 2.0$, the case reported below). 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 [0.6–8% systematic shift in f_{NL}] but *not* the Li *et al.* robustness-branch value; positions $f_{\text{NL}} \in [-4.875, -3.875]$ at 1σ around the central value $f_{\text{NL}} = -4.375$, which excludes the Li *et al.* value -2.1875 at 4.375σ since $|(-4.375) - (-2.1875)|/0.5 = 4.375$): ~ 14 (scipy.stats.norm: BF=13.91 at broad $[-15, +15]$).
- **Gaussian prior**, $\sigma_{\text{theory}} = 1.0$ (recommended baseline: covers the full ϵ -correction range within 1σ ; the Li *et al.* value -2.1875 sits at 2.19σ , i.e. inside the prior’s $\sim 2\sigma$ tail but not within 1σ): ~ 10 (BF=9.80, broad $[-15, +15]$ prior).
- **Gaussian prior**, $\sigma_{\text{theory}} = 2.0$ ($\approx 1\sigma$ encompassment of the Li *et al.* robustness-branch value; the Li *et al.* value -2.1875 lies at 1.09σ from -4.375 under this prior, which is the closest σ_{theory} we report that brackets the branch shift; quoted as a robustness check, not as a physically motivated prior width): ~ 6 (scipy.stats.norm: BF=5.65 broad).

The recommended-to-theoretical-maximum range ~ 10 – 17 — quoted in the abstract as the $r \rightarrow 1$ bookkeeping endpoint, with the abstract headline reading $\text{BF} \approx 9$ – 14 under the noise-weighted $r \approx 0.84$ bounce-amplitude bookkeeping (see the template-mismatch bookkeeping paragraph below) — 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*. In addition to the discrete grid, a continuous marginalization over the prior-width hyperparameter $\sigma_{\text{theory}} \sim \mathcal{U}[0.5, 2.0]$ (uniform hyperprior spanning the reported grid, same closed-form conventions) gives $\text{BF} = 8.8$ at the broad $[-15, +15]$ competitor and $\text{BF} = 3.6$ at the narrow $[-5, +5]$ competitor — between the $\sigma_{\text{theory}} = 0.5$ and 2.0 endpoints and close to the recommended $\sigma_{\text{theory}} = 1.0$ values (9.80 broad, 4.01 narrow); the marginal value sits slightly below the recommended-prior value because the uniform hyperprior weights the wide-prior end equally. The bounce-over-multifield ranking ($\text{BF} > 1$) holds across the entire hyperprior support, so it is stable under continuous prior-width variation and is not an artifact of the discrete grid (three-point grid converged to < 0.03 in log-evidence; computation released as `c91_sigma_theory_continuous_marginalization.py`, see Data and Code Availability).

For abstract-envelope readability, the four-corner Bayes-factor grid is summarized in one place below. Columns: narrow $[-5, +5]$ vs. broad $[-15, +15]$ multifield competitor priors. Rows: delta-prior vs. $\sigma_{\text{theory}} =$

1.0 Gaussian bounce priors. Together they bracket the abstract’s $r \rightarrow 1$ endpoint envelope $\text{BF} \sim 10$ – 17 (the abstract headline reads $\text{BF} \approx 9$ – 14 under the noise-weighted $r \approx 0.84$ bounce-amplitude bookkeeping). *Reading:* the abstract’s $r \rightarrow 1$ endpoint envelope $\text{BF} \sim 10$ – 17 is the $\sigma_{\text{theory}} = 1.0$ broad-multifield column ($\text{BF} \sim 10$) up to the delta-row broad-multifield column ($\text{BF} \sim 17$); the curvaton-natural narrow-competitor column gives the lower-envelope $\text{BF} \sim 4$ ($\sigma_{\text{theory}} = 1.0$) and $\text{BF} \sim 7$ (delta).

The narrow-competitor entries give the lower $\text{BF} \sim 7$ (delta) and $\text{BF} \sim 4$ ($\sigma_{\text{theory}} = 1.0$) values quoted in the prose.

A reader who only reads this subsection can therefore reproduce the abstract’s $r \rightarrow 1$ endpoint envelope $\text{BF} \sim 10$ – 17 (rebooked to the headline $\text{BF} \approx 9$ – 14 under the noise-weighted $r \approx 0.84$ bookkeeping) 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.

c. Worked example: reproducing the abstract $\text{BF} \approx 9$. Starting from the recommended $\sigma_{\text{theory}} = 1.0$ Gaussian bounce prior and the broad multifield competitor prior $[-15, +15]$ (Eq. 8), the abstract’s headline $\text{BF} \approx 9$ follows in three steps. (1) Set the bounce likelihood width to $\sigma_{\text{eff}} = \sigma(f_{\text{NL}}^{\text{local}})/r = 0.7/0.84 \approx 0.833$ (noise-weighted $r = 0.84$ rebooking). (2) Evaluate the analytic Bayes-factor formula with Gaussian bounce prior $\mathcal{N}(-35/8, \sigma_{\text{theory}}^2 = 1.0)$, uniform competitor prior on $[-15, +15]$, and mock detection $\hat{f}_{\text{NL}} = -35/8$: $\text{BF} = \mathcal{L}(\hat{f}_{\text{NL}}|\text{bounce prior})/\mathcal{L}(\hat{f}_{\text{NL}}|\text{multifield prior})$. Using `scipy.stats.norm` with $\sigma_{\text{eff}} = 0.833$ reproduces $\text{BF} \approx 9.2$ (the abstract rounds to ≈ 9). (3) At the $r \rightarrow 1$ no-rebooking endpoint ($\sigma_{\text{eff}} = 0.7$) the same formula gives $\text{BF} \approx 9.8$ (≈ 10 in the abstract’s $r \rightarrow 1$ Table II entry); the modest difference (9.2 vs 9.8) reflects the σ_{eff} inflation from 0.7 to 0.833 under the $r = 0.84$ rebooking. Both values reflect the competitor-prior sensitivity emphasized in Sec. VI: because $\text{BF} \propto W$ in the broad-competitor regime (Eq. 10), widening $[-15, +15]$ to $[-20, +20]$ ($W : 30 \rightarrow 40$) raises the Bayes factor by the factor $\approx 4/3$ (e.g. $9.8 \rightarrow 13.1$ at the $r \rightarrow 1$ endpoint), while the curvaton-natural $[-5, +5]$ competitor reduces the headline to $\text{BF} \approx 4$.

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 ($\text{BF} \sim 10$ vs. tuned multifield, broad $[-15, +15]$ column), which is the most physically motivated baseline; the delta-prior row is shown only as the theoretical-maximum upper bound

	multifield prior $[-5, +5]$	multifield prior $[-15, +15]$
delta prior at $-35/8$	BF ~ 7	BF ~ 17
$\sigma_{\text{theory}} = 1.0$ Gaussian	BF ~ 4	BF ~ 10

and is not the recommended headline.

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 ~ 4 –10 across the same multifield-competitor-prior variation (scipy.stats.norm recompute: 4.01 narrow, 9.80 broad), again with broader competitor priors giving larger Bayes factors. The recommended headline (BF ~ 10 at $\sigma_{\text{theory}} = 1.0$, broad multifield $[-15, +15]$) and the delta-prior maximum (BF ~ 17 at the same multifield competitor) bracket the $r \rightarrow 1$ endpoint envelope ~ 10 –17 (abstract headline ≈ 9 –14 after the noise-weighted $r \approx 0.84$ rebooking); the headline number we promote is the lower bound (BF ~ 10 at $r \rightarrow 1$, ≈ 9 rebooked), 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.

d. Template-mismatch bookkeeping of the Bayes factors. The mock detection and the $\sigma = 0.7$ likelihood width above place both the detection and the bounce prediction at $f_{\text{NL}} = -35/8$ with the survey’s local-template uncertainty applied directly — the $r \rightarrow 1$ (no-template-mismatch) bookkeeping endpoint. The rebooking to noise-weighted $r \approx 0.84$ proceeds by replacing the effective uncertainty with $\sigma_{\text{eff}} = \sigma(f_{\text{NL}}^{\text{local}})/r \approx 0.7/0.84 \approx 0.83$ in Eq. (8), or equivalently evaluating the likelihood at the noise-weighted prediction $r \times f_{\text{NL}}^{\text{bounce}} \approx -3.675$ rather than at $f_{\text{NL}}^{\text{bounce}} = -35/8$; both rescalings yield the same modest shift in the Bayes-factor four-corner grid reported below. Under the Eq. (5) bookkeeping in bounce-amplitude space, $\sigma_{\text{eff}} = \sigma(f_{\text{NL}}^{\text{local}})/r \approx 0.83$ at the noise-weighted $r = 0.84$, and the closed-form Eq. (8) four-corner grid reduces modestly: 17.1 \rightarrow 14.4 (delta, broad), 9.8 \rightarrow 9.2 ($\sigma_{\text{theory}} = 1.0$, broad), 7.0 \rightarrow 6.2 (delta, narrow), 4.0 \rightarrow 4.0 ($\sigma_{\text{theory}} = 1.0$, narrow); the $r \rightarrow 1$ envelope BF ~ 10 –17 correspondingly reads ~ 9 –14 in strict bounce-amplitude bookkeeping — the bookkeeping the abstract now adopts for its headline, with the $r \rightarrow 1$ values retained in Table II. The alternative fully measured-space bookkeeping (prediction and detection at $r f_{\text{NL}}^{\text{bounce}} \approx -3.68$ with $\sigma = 0.7$ and unrescaled competitor priors, since the multifield competitor is itself local-shaped) gives a similar modest reduction (7.0 \rightarrow 5.9 at the delta/narrow corner; the broad-competitor cells are unchanged). No qualitative conclusion changes under either consistent bookkeeping; the rescaling computation is released with the paper’s code (see Data and Code Availability).

e. Bayes-factor closure against the QSFI continuum. The multifield competitor priors above represent the cur-

vaton class. The Bayes-factor discrimination against quasi-single-field inflation (QSFI) [30] 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 local-template scaling ($\Delta = 0$ at $\mu/H = 0$, the massless-isocurvature / curvaton-like limit, where the shape-to-local ratio $(k_3/k_1)^\Delta$ is flat) and an intermediate shape whose squeezed amplitude is suppressed by $(k_3/k_1)^{3/2}$ relative to local ($\Delta = 3/2$ at $\mu/H = 3/2$). At the QSFI degenerate endpoint $\mu/H \rightarrow 0$ the QSFI shape reproduces the local-template scaling that the matter bounce also approximately matches ($r_{\text{cos}} > 0.97$, Sec. III B); near this endpoint shape-based discrimination weakens substantially, and the residual Bayes factor is set by the amplitude prior-predictive distributions and nuisance modeling rather than by shape mismatch — shape degeneracy alone does not force BF $\rightarrow 1$. BF-based *shape* discrimination of the bounce from QSFI therefore requires μ/H appreciably above 0, where the squeezed-limit suppression relative to the local template 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 envelope (BF ≈ 9 –14 headline; ~ 10 –17 at the $r \rightarrow 1$ endpoint) should be read as bracketing the curvaton-class discrimination only, not as a single QSFI-class number.

VII. SYSTEMATICS AND ROBUSTNESS

Scope of the systematic budget (stated up front). Every systematic contribution in this section, and the consolidated budget of Table IV, is combined by addition in quadrature ($\sigma_{\text{eff}} = \sqrt{\sigma_{\text{base}}^2 + \sum_i \sigma_i^2}$). This is a transparent scoping *heuristic*, not a joint multi-tracer marginalized Fisher: correlations between nuisances (notably the b_1 – b_ϕ – f_{NL} large-scale degeneracy) can tighten or loosen the combined budget depending on the sign of the covariance, so the realistic ~ 2.6 – 5.5σ envelope is a sensitivity recast rather than a self-consistent joint-covariance forecast. For the strongest of these degeneracies a joint Fisher has now been computed (c8_fnl_running_fisher.json; Planck 2018 cosmology, CAMB 1.6.6; SPHEREx public-product inputs, Doré+2014-validated; §III). In the SDB power-spectrum channel the f_{NL} -only constraint $\sigma(f_{\text{NL}}) = 1.53$ (2.86σ detection at $f_{\text{NL}} = -4.375$, biases fixed) degrades to $\sigma_{\text{marg}}(f_{\text{NL}}) = 3.08$ once the running is marginalized (a $2.0\times$ degradation, anti-correlation $\rho = -0.87$, detection 1.42σ), and to 7.06 once the per-sample linear biases are additionally co-marginalized (a $4.6\times$ degrada-

Bounce prior choice	BF vs. tuned multifield [−15, +15], $r \rightarrow 1$ endpoint ($r=0.84$ headline in parens)	BF vs. SSFSR
Gaussian, $\sigma_{\text{theory}} = 1.0$ (recommended headline)	~ 10 ($r=0.84$: 9.2)	$\gg 1$
Gaussian, $\sigma_{\text{theory}} = 0.5$	~ 14	$\gg 1$
Gaussian, $\sigma_{\text{theory}} = 2.0$	~ 6	$\gg 1$
Delta at $f_{\text{NL}} = -35/8$, narrow [−5, +5] multifield (GR-variation only)	4.7–7.0 ^a	$> 10^2$
Delta at $f_{\text{NL}} = -35/8$, broad [−15, +15] multifield (GR fixed)	$\sim 17^b$ ($r=0.84$: 14.4)	$> 10^5$

All BF entries are the $r \rightarrow 1$ (no-template-mismatch, $\sigma_{\text{eff}} = 0.7$) endpoint. The promoted abstract headline BF ≈ 9 –14 applies the noise-weighted $r = 0.84$ rebooking $\sigma_{\text{eff}} = \sigma(f_{\text{NL}}^{\text{local}})/r = 0.83$ to these entries (Eq. 5): the recommended row 1 maps 9.8 \rightarrow 9.2 and the delta/broad row maps 17 \rightarrow 14.4, bracketing the 9–14 headline. The 0.55–1.14 null-space r band (Table IV) is a distributional stress band and never enters σ_{eff} .

TABLE II. Primary Bayes-factor sensitivity to the bounce prior, evaluated for a mock SPHEREx detection at $f_{\text{NL}} = -4.375$ ($\sigma = 0.7$) with the GR-contamination amplitude fixed (no GR marginalization), except row 4, which reports the GR-marginalization spread. “BF vs. SSFSR” evaluates standard single-field slow-roll as a point hypothesis at its parameter-free prediction $f_{\text{NL}}^{\text{inf}} \approx 0.015$ (effectively $f_{\text{NL}} = 0$ at survey precision); the tabulated entries are exact closed-form Eq. (8) evaluations under the stated assumptions, but because their magnitudes are strongly prior-width and point-hypothesis dependent, they should be *interpreted* only at the order-of-magnitude level. The $\sigma_{\text{theory}} = 1.0$ Gaussian row is the recommended physically motivated headline (BF ~ 10 at broad multifield [−15, +15], BF ~ 4 at narrow [−5, +5]): it covers the full ϵ -correction range within 1σ , while the Li *et al.* robustness-branch value -2.1875 lies at 2.19σ under this prior (inside the $\sim 2\sigma$ tail, not within 1σ ; see Sec. VI). 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 4.7–7.0 spread on this row reflects the GR-marginalization variation across the scenarios of Table III at fixed *narrow* multifield competitor prior [−5, +5]; this is the same competitor prior that produces the 4.7–7.0 “BF vs. Tuned” column of Table III, and the two tables are numerically identical by construction (both from the closed-form Eq. (8) recompute, `c9g_bf_table_recompute.py`). The recommended $\sigma_{\text{theory}} = 1.0$ headline of BF ~ 10 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] with the GR-contamination amplitude fixed (no GR marginalization; $\sigma_{\text{eff}} = 0.7$); GR marginalization at $\sigma_{\text{GR}} = 0.5$ reduces it to ~ 14 . This is the delta-prior endpoint of the competitor-prior sensitivity scan in Sec. VI. The $r \rightarrow 1$ endpoint envelope ~ 10 –17 quoted in the abstract brackets the recommended-prior lower bound (BF ~ 10 , row 1) up to the delta-prior maximum at broad multifield (BF ~ 17 , row 5), with broader bounce priors giving *smaller* Bayes factors; the abstract headline BF ≈ 9 –14 applies the noise-weighted $r \approx 0.84$ bounce-amplitude bookkeeping to these endpoints (see the template-mismatch bookkeeping paragraph).

tion, $\rho = -0.97$, detection 0.62σ). This resolves the open sign in the caveat above: the b_1 – f_{NL} – $n_{f_{\text{NL}}}$ correlations *loosen* rather than tighten the constraint, with the running $n_{f_{\text{NL}}}$ the dominant degradation direction — so the additive-quadrature budget is refined, not contradicted, by an explicit joint covariance in the channel where this degeneracy is largest. This SDB degradation does not propagate to the bispectrum-only headline of §IV, which accesses additional triangle configurations that break the f_{NL} – $n_{f_{\text{NL}}}$ degeneracy; a full joint bispectrum covariance is the one element still combined heuristically here.

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 [28].

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 is used as the baseline for the headline 2.6–5.5 σ significance range. This 20% prior is motivated by the expected precision of the SPHEREx multi-tracer sample under the universal-mass-function (UMF) relation $b_\phi = 2\delta_c(b_1 - 1)$, which anchors b_ϕ to the linear bias b_1 with a theoretical scatter of $\mathcal{O}(15$ –20%) relative to N-body calibrations of the UMF relation [28]; Heinrich *et al.* [1] adopt UMF universality to fix b_ϕ per tracer population, corresponding approximately to this 20% level. When the UMF universality assumption is relaxed (per-tracer-bin free marginalization), the effective $\sigma(f_{\text{NL}})$ degrades by $\mathcal{O}(20$ –50%), which is why the headline range extends to the b_ϕ -30% and 50% scenarios in Table IV. Fig. 5 shows how $\sigma(f_{\text{NL}})$ degrades as the b_ϕ prior widens: at 20% prior width,




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) is less sensitive to the ultra-large-scale fragility. Right: corresponding detection significance for $f_{\text{NL}} = -35/8$.

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. (Here P denotes the late-time *matter* power spectrum, not the primordial P_ζ ; the primordial normalization enters Δb through the transfer kernel $\mathcal{M}(k, z)$ of Eq. (4), and the spherical-collapse threshold δ_c is carried inside $b_\phi = 2\delta_c(b_1 - 1)$, so it is not double-counted between Δb and the cross-terms.) Heinrich *et al.* [1] 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 [28] for upcoming Stage-IV surveys), the effective $\sigma(f_{\text{NL}})$ for the SPHEREx multi-tracer bispectrum widens by $\mathcal{O}(20\text{--}50\%)$, which degrades the headline 5.2–5.5 σ optimistic template-corrected significance to $\sim 4.0\text{--}4.2\sigma$ at the central 30% degradation point and to $\sim 3.5\text{--}3.7\sigma$ 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 2.6–5.5 σ already incorporates the central 20–30% degradation. Independent tracer popula-

tions 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.

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). In addition to the discrete $\sigma_{\text{GR}} \in \{0, 0.5, 1.0\}$ grid, a continuous marginal-likelihood integration over $\sigma_{\text{GR}} \sim \mathcal{U}[0, 1]$ (uniform prior, same likelihood conventions as the Table III recompute) gives $\text{BF} = 6.0$ vs. the tuned narrow competitor — inside the discrete-grid spread 4.7–7.0, with the three-point grid converged to < 0.01 in log-evidence on this column — and $\text{BF} \approx 8.6 \times 10^3$ vs. SSFSR, which sits between the $\sigma_{\text{GR}} = 0.5$ and $\sigma_{\text{GR}} = 1.0$ cells because the exponentially σ -sensitive SSFSR evidence is dominated by the conservative end of the prior (computation released as `c9k_gr_continuous_marginalization.py`; see Data and Code Availability). The discrete grid is therefore a faithful summary of the continuous marginalization for the bounce-vs-tuned comparison, and the continuous SSFSR value is reported directly above.

The parameterization $\sigma_{\text{GR}} \in [0, 1.0]$ is an internal stress-test amplitude grid, motivated by — but not a direct translation of — published demonstrations that relativistic projection effects materially contaminate f_{NL} inference. Addis *et al.* [27], analyzing the multipoles of the full relativistic power spectrum, show that ne-




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 [1], 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; the SPHEREx multi-tracer bispectrum combined-sample baseline is $\sigma(f_{\text{NL}}) = 0.7$ at this point, not a per-redshift-bin uncertainty), MegaMapper gives $\sim 4\sigma$; relaxing the per-tracer-bin b_ϕ prior to 30% widens the combined-sample SPHEREx $\sigma(f_{\text{NL}})$ to ≈ 0.9 – 1.0 , and relaxing to 50% drops MegaMapper to $\sim 2\sigma$. The bispectrum channel remains at $\sim 5\sigma$ (optimistic) or ~ 3 – 4σ (after GR degradation) at fixed Heinrich *et al.* b_ϕ universality [1]; relaxing b_ϕ universality per tracer bin [28] degrades the optimistic 5.2–5.5 σ headline to ~ 4.0 – 4.2σ (30% central) and ~ 3.5 – 3.7σ (50% conservative).

glecting relativistic effects (Doppler, gravitational redshift, lensing, and light-cone effects) produces best-fit f_{NL} biases of order $\sim 3\sigma$ for a Euclid-like H α survey and $\sim 20\sigma$ for a MegaMapper-like Lyman-break sample, with a 15–20% improvement available from bright/faint multi-tracer splitting. They do not report a $\sigma(f_{\text{NL}})$ degradation for an SPHEREx-class survey; our σ_{GR} values are therefore stress-test amplitudes chosen to span a plausible contamination scale, not figures calibrated to their Fisher results. Under the additive-quadrature model of Table III ($\sigma_{\text{eff}} = \sqrt{\sigma_{\text{base}}^2 + \sigma_{\text{GR}}^2}$, calibrated to the conservative contamination scale of Addis *et al.* rather than to a direct Fisher translation), the $\sigma_{\text{GR}} = 0.5$ scenario corresponds to a $\sim 23\%$ inflation of the effective uncertainty ($\sigma_{\text{eff}}/0.7 = 0.860/0.700 = 1.23$); the $\sigma_{\text{GR}} = 1.0$ scenario is the conservative bound, appropriate in spirit for the high-redshift samples where Addis *et al.* find the largest relativistic biases. The “optimistic” case (5.2–5.5 σ) in Sec. IV omits GR degradation ($\sigma_{\text{GR}} = 0$); the “realistic” case (2.6–5.5 σ) includes it at the $\sigma_{\text{GR}} = 0.5$ – 1.0 level.

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 simplified Fisher degradation estimate indicates 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 the 2.6–5.5 σ post-budget envelope of Sec. VII ($\gtrsim 2.5\sigma$ in the most conservative all-combined 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) [31] and Giannantonio *et al.* (2012) [26]; 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-

GR Treatment	BF vs. SSFSR	BF vs. Tuned	$P(\text{BF} > 3)$ vs. SSFSR
Ideal (no GR)	3.5×10^8	7.0	99.9%
Marginalized ($\sigma_{\text{GR}} = 0.5$)	4.5×10^5	6.1	99.0%
Marginalized ($\sigma_{\text{GR}} = 1.0$)	6.4×10^2	4.7	93.1%
<i>Corrected (10% residual; = Ideal, verification only)^a</i>	3.5×10^8	7.0	99.9%

TABLE III. Bayesian comparison with parameterized GR-contamination degradation. SSFSR = standard single-field slow-roll inflation. All Bayes factors are closed-form Eq. (8) evaluations at the central mock detection $\hat{f}_{\text{NL}} = -35/8$ with the GR-contamination nuisance added in quadrature to the SPHEREx baseline, $\sigma_{\text{eff}} = \sqrt{0.7^2 + \sigma_{\text{GR}}^2}$; $P(\text{BF} > 3)$ is the fraction of 2×10^5 mock-detection realizations $\hat{f}_{\text{NL}} \sim \mathcal{N}(-35/8, \sigma_{\text{eff}})$ exceeding $\text{BF} = 3$ against SSFSR (script and output released with the paper’s code, `c9g_bf_table_recompute.py`; see Data and Code Availability; the binomial Monte-Carlo standard errors on the $P(\text{BF} > 3)$ column are < 0.1 percentage points at 2×10^5 realizations, below the quoted precision). 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 no-GR cell (7.0) coincides with the narrow→broad sweep $\text{BF} \sim 7 \rightarrow 17$ at the delta bounce prior reported in the prose of §VI, and the corresponding broad $[-15, +15]$ competitor gives $\text{BF} \sim 14$ at $\sigma_{\text{GR}} = 0.5$. The bounce-vs-inflation comparison survives all treatment scenarios. ^aThe “Corrected (10% residual)” row is a bookkeeping verification row, not an independent configuration: it implements the limit in which GR contamination is corrected and the post-correction residual is neglected at this template-overlap order, so it equals “Ideal” by construction. For calibration, a literal residual nuisance of $\sigma_{\text{GR}} = 0.05$ (10% of the $\sigma_{\text{GR}} = 0.5$ scenario amplitude surviving correction) gives $\sigma_{\text{eff}} = 0.702$ and $\text{BF vs. Tuned} = 6.99$ ($\Delta\text{BF} = 0.01$, below the reported precision), while the exponentially σ -sensitive BF-vs-SSFSR column would read 3.1×10^8 rather than 3.5×10^8 ; the row is therefore reported in the strict zero-residual limit. The table thus contains three independent GR scenarios plus one verification row, and rows 1 and 4 bookend the same GR-free regime. As in Table II, the “BF vs. Tuned” entries use the $r \rightarrow 1$ template bookkeeping ($\sigma_{\text{eff}} = 0.7$ baseline before GR); the noise-weighted $r = 0.84$ rebooking of Table II reduces the delta/narrow no-GR value $7.0 \rightarrow 6.2$.

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 the realistic $2.6\text{--}5.5\sigma$ level ($> 2.5\sigma$ even at the conservative all-combined endpoint). For one-place reconciliation of the significance windows quoted in this paper: the realistic SPHEREx-only window is $2.6\text{--}5.5\sigma$ (post-systematic-budget); the conservative endpoint of the b_ϕ 50%-prior bispectrum degradation is $\sim 3.5\text{--}3.7\sigma$ (§VII, within the realistic window); and the wider $3\text{--}7\sigma$ envelope quoted for the survey program combines that conservative SPHEREx endpoint with the MegaMapper optimistic ceiling ($\sim 7\sigma$, a proposed-facility projection, not a SPHEREx forecast).

E. Consolidated Systematic Budget

Table IV consolidates all systematic contributions to the SPHEREx bispectrum detection significance in one place. Each row identifies the source, its numerical value or range, whether it degrades the numerator (signal amplitude), the denominator (effective $\sigma(f_{\text{NL}})$), or both, the combination rule, and the resulting cumulative σ_{eff} at $f_{\text{NL}} = -35/8$. The significance at each row is

$$|f_{\text{NL}}| \times r / \sigma_{\text{eff}} \text{ with } |f_{\text{NL}}| = 4.375.$$

VIII. CURRENT DATA AND CONSISTENCY RELATION

A. Planck + DESI Recast

Current constraints from Planck PR4/NPIPE (CMB bispectrum, $f_{\text{NL}} = -0.1 \pm 5.0$ [32]) 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 0.75σ from the bounce prediction ($| -4.375 + 0.1 | / 5.71$; recasting the central value as well, $-0.1 / 0.876 = -0.11$, leaves the tension at 0.75σ) and 0.02σ from zero—fully consistent with both. (The earlier PR3 Planck constraint was $f_{\text{NL}} = -0.9 \pm 5.1$ [33]; 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 LRG and QSO analyses report $f_{\text{NL}}^{\text{loc}}$ bounds at $\sigma \approx 9\text{--}10$: $f_{\text{NL}}^{\text{loc}} = -3.6_{-9.1}^{+9.0}$ from the combined LRG+QSO sample [34]; $f_{\text{NL}}^{\text{loc}} = -3.3 \pm 9.2$ from the QSO assembly-bias analysis [35] (these are distinct analyses), consistent with both bounce and inflation at current precision. Recasting via $r = 0.84$ gives $\sigma(f_{\text{NL}}^{\text{bounce}}) \approx \sigma_{\text{loc}} / r \approx 11$, far too weak to discriminate. The bound quoted here therefore remains consistent with Planck alone, and current LSS data cannot discriminate

Systematic source	Value / range	Acts on	Combination rule	σ (detection)
Naive uncorrected (ref. only)	$ f_{\text{NL}} /\sigma = 4.375/0.7$	—	not used in headline	6.25σ (no template correction; see caption)
Template-corrected baseline	$r = 0.84 \pm 0.02$ [0.829–0.876]	numer. ($\times r$)	multiplicative	5.2–5.5 σ headline
ϵ -correction	0.6–8% in f_{NL}	numer. ($\pm\delta f_{\text{NL}}$)	add. quadrature	$\lesssim 0.4\sigma$ effect
Null-space r scatter	$r = 0.85 \pm 0.13$; 16th pctile $r = 0.70$	numer. ($\times r$)	distributional	floor $\approx 4.4\sigma$
b_ϕ (20%, baseline)	no degradation	—	—	5.2–5.5 σ optimistic
b_ϕ (30%, central)	$\sigma(f_{\text{NL}}) \rightarrow 0.9$	denom. (repl.)	Fisher widening / baseline replacement	$\sim 4.1\sigma$
b_ϕ (50%, conserv.)	$\sigma(f_{\text{NL}}) \rightarrow 1.0$	denom. (repl.)	Fisher widening / baseline replacement	~ 3.5 – 3.7σ
GR ($\sigma_{\text{GR}} = 0$)	no GR	—	—	5.2–5.5 σ
GR ($\sigma_{\text{GR}} = 0.5$)	$\sigma_{\text{eff}} = \frac{\sigma_{\text{eff}}}{\sqrt{0.7^2 + 0.5^2}} = 0.86$	= denom. (\oplus)	add. quadrature	$\sim 4.3\sigma$
GR ($\sigma_{\text{GR}} = 1.0$)	$\sigma_{\text{eff}} = \frac{\sigma_{\text{eff}}}{\sqrt{0.7^2 + 1.0^2}} = 1.22$	= denom. (\oplus)	add. quadrature	$\sim 3.0\sigma$
All combined 30% + GR 1.0)	$(b_\phi \sigma_{\text{eff}}) = \frac{\sigma_{\text{eff}}}{\sqrt{0.9^2 + 1.0^2}} = 1.35$	= denom. (\oplus)	add. quadrature	$\sim 2.7\sigma$
All combined 50% + GR 1.0)	$(b_\phi \sigma_{\text{eff}}) = \frac{\sigma_{\text{eff}}}{\sqrt{1.0^2 + 1.0^2}} = 1.41$	= denom. (\oplus)	add. quadrature	$\sim 2.6\sigma$ conserv.

TABLE IV. Consolidated systematic budget for the SPHEREx bispectrum detection of $f_{\text{NL}} = -35/8 = -4.375$. **Row types:** the first row (naive uncorrected) is shown for reference only and is *not* used in any headline; its 6.25σ figure is not directly comparable to the template-corrected 5.2 – 5.5σ headline because it uses a distinct null procedure (no template-mismatch correction applied). The template-corrected baseline row ($r = 0.84$) is the starting point for all subsequent rows; the ϵ -correction and null-space rows are *distributional* contributions (not cumulative denominators); the b_ϕ and GR rows are *cumulative* additions to the effective denominator σ_{eff} ; the all-combined rows stack both contributions simultaneously. “Acts on” refers to the signal numerator ($|f_{\text{NL}}| \times r$) or the effective denominator σ_{eff} . Combination rule “ \oplus ” denotes addition in quadrature ($\sigma_{\text{eff}} = \sqrt{\sigma_{\text{base}}^2 + \sigma_{\text{syst}}^2}$), a transparent scoping choice described in Sec. II C (additive-quadrature rather than joint-marginalized Fisher). Significance is $|f_{\text{NL}}| \times r / \sigma_{\text{eff}}$ with $r = 0.84$ (noise-weighted central value). The b_ϕ -30% row uses $\sigma(f_{\text{NL}}) \rightarrow 0.9$ (30% widening over the baseline 0.7; the effective denominator increases to 0.9 as the b_ϕ marginalization replaces the baseline). The all-combined rows add GR degradation on top of the b_ϕ -widened denominator using the additive-quadrature systematic budget defined in Sec. VII. The ~ 2.6 – 5.5σ post-systematic-budget envelope quoted in the abstract spans the all-combined conservative floor (bottom row) to the no-GR optimistic ceiling (template-corrected baseline row with b_ϕ -20% assumption). The effective denominator σ_{eff} is defined in Sec. IV (Eq. 5) as the quadrature combination of $\sigma(f_{\text{NL}}^{\text{local}})$ and any additional systematic nuisances; all σ_{eff} values in this table follow that convention.

between the bounce and inflation.

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

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

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

where κ_ϵ (not to be confused with the bispectrum polynomial coefficient c_1 of Sec. II) depends on both the explicit $A_T \propto 1/\epsilon^3$ prefactor in the cubic action and the $|\eta|^{-\nu}$ mode-function amplitude (with $\nu = 3/2$ at $\epsilon = 3/2$ and finite by the Wilson-Ewing dust scaling), both of which change with ϵ near $\epsilon = 3/2$. *Linearization note.* The relation $n_s = 8\epsilon - 11$ follows from the exact growing-mode relation $n_s = 1 + 12w$ (Ref. [2]) 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 0.6–8% ϵ -correction uncertainty in f_{NL} . The coefficient κ_ϵ is bounded by two physically distinct contributions to the ϵ -derivative of the Cai et al. [7] cubic-action prefactor $A_T \propto 1/\epsilon^3$ at $\epsilon = 3/2$ (the explicit-prefactor channel) and to the Hankel-index dependence of the matter-contraction mode functions $\zeta_k \propto |\eta|^{-\nu(\epsilon)}$ with $\nu = 3/2 + \mathcal{O}(\epsilon - 3/2)$ (the mode-function channel; cf. Wilson-Ewing [2] for the linear-order scaling and Cai *et al.* [7] App. B for the cubic-vertex sensitivity). Explicit prefactor scaling alone gives $\kappa_\epsilon \approx 5.6$ (lower bound, holding mode functions fixed at $\epsilon = 3/2$); the mode-function amplitude channel ($|\eta|^{-\nu}$ growth with ν finite at $3/2$) adds a multiplicative correction which we bound schematically at $\approx 14\times$ the prefactor-only value, giving $\kappa_\epsilon \approx 80$ as the upper endpoint; the upper endpoint is a schematic scaling bound, not a derived coefficient. The order-of-magnitude span [5.6, 80] therefore reflects the well-known sensitivity of cubic-order quantities to the combined $A_T \propto 1/\epsilon^3$ prefactor and the $|\eta|^{-\nu}$ mode-function amplitude channels near exact matter domination, not a free parameter; pinning κ_ϵ to better than this range requires the full four-vertex numerical evaluation discussed in Sec. II C. Even the conservative endpoint shifts f_{NL} by only $\kappa_\epsilon |\Delta\epsilon| \approx 80 \times 0.0045 \approx 0.36$ — well inside the recommended $\sigma_{\text{theory}} = 1.0$ bounce prior — so no Bayes-factor conclusion in Sec. VI hinges on the precise value of this endpoint. Eliminating ϵ via $\Delta\epsilon = \Delta n_s/8$, the consistency relation takes the form

$$f_{\text{NL}}(n_s) \approx -\frac{35}{8} - c'(n_s - 1), \quad c' \equiv \kappa_\epsilon/8 \in [0.7, 10], \quad (12)$$

where $c' \equiv \kappa_\epsilon/8 > 0$ follows directly from $\Delta\epsilon = (n_s - 1)/8$ and the leading-order coefficient $-\kappa_\epsilon$ of the ϵ -form above; as the equation of state approaches exact matter domination from $w < 0$ (so $\epsilon \rightarrow 3/2^-$ and $n_s \rightarrow 1^-$), both deviations vanish and $f_{\text{NL}} \rightarrow -35/8$. For $\epsilon < 3/2$ (the quasi-dust regime $w < 0$), $(\epsilon - 3/2) < 0$ and the correction $-\kappa_\epsilon(\epsilon - 3/2) > 0$, so f_{NL} moves *less negative* than $-35/8$ — consistent with the Planck- n_s range quoted below. At the Planck best-fit $n_s = 0.9649$, this gives $f_{\text{NL}} \in [-4.35, -4.02]$ (a 0.6–8% 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; that dedicated four-vertex in-in computation is outside the scope of this forecast paper, and its impact on the present results is bounded by the $\kappa_\epsilon |\Delta\epsilon| \approx 0.36$ shift quoted above, inside the recommended $\sigma_{\text{theory}} = 1.0$ bounce prior. 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 ≈ 0.5 – 0.7 range quoted here) provides a direct test of Eq. (12) 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 ~ 2.6 – 5.5σ significance after the full systematic budget (5.2 – 5.5σ optimistic before GR and b_ϕ degradation; the lower bound of this optimistic range reflects the noise-weighted template overlap $r = 0.84$). MegaMapper, a proposed facility without confirmed funding or finalized design ($\sim 2032+$ if approved), could provide a more powerful follow-up at ~ 3 – 7σ 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 [36–38]:

- *DESI*: Already taking data; expected $\sigma(f_{\text{NL}}) \approx 3$ – 5 from scale-dependent bias with luminous red galaxies and emission-line galaxies [36]. (The DESI, Euclid, and LSST figures quoted in this list are primarily scale-dependent-bias / power-spectrum forecasts, a different channel from the SPHEREx multi-tracer *bispectrum* baseline of this paper; they are complementary probes, not like-for-like sensitivity comparisons.)
- *Euclid*: Launched 2023; published forecasts give $\sigma(f_{\text{NL}}) \approx 2$ – 4 from the photometric galaxy survey [37].
- *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 [38].

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 on the measured local-template f_{NL} (horizontal axis). Shaded regions, left to right: dark green — measurement near the $-35/8$ prediction, strongly favors the bounce; light green — negative f_{NL} of bounce-compatible magnitude, supports the bounce; dark red (legend label “bounce excluded”) — measurement consistent with zero, disfavoring the quasi-dust matter bounce while remaining consistent with standard single-field inflation; light red — positive f_{NL} , disfavoring the bounce and instead supporting exotic multifield inflation. Blue vertical line: bounce prediction $f_{\text{NL}} = -35/8$; black dotted line: standard inflation $f_{\text{NL}} \approx 0$. Error bars: SPHEREx ($\sigma = 0.7$) and MegaMapper conservative ($\sigma = 1.5$, illustrative: between the $\sigma \approx 0.5$ ideal and the b_ϕ -50% degraded $\sigma \approx 2.2$ of Sec. VII, not an independently calibrated forecast), both centered on the bounce prediction.

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 $r \in [0.829, 0.876]$, a measurement centered on zero excludes the matter-bounce prediction $f_{\text{NL}} = -35/8$ at the same significance as a detection would carry (specifically $|f_{\text{NL}}| r / \sigma(f_{\text{NL}}) \approx 5.1\text{--}5.5\sigma$ before GR and b_ϕ degradation, degrading to the $\sim 2.6\text{--}5.5\sigma$ post-budget envelope after the realistic systematic budget of Sec. VII, since the exclusion arithmetic is symmetric to the detection arithmetic), conditional on assumptions (a)–(f) of Sec. II C.

D. The Joint $(f_{\text{NL}}, n_{f_{\text{NL}}})$ Scale-Dependent-Bias Fisher Forecast

Motivation, methodology, and scope. The forecast reported in this subsection is a separately computed scale-dependent-bias (SDB) Fisher matrix for $(f_{\text{NL}}, n_{f_{\text{NL}}})$, methodologically distinct from the imported Heinrich *et al.* bispectrum forecast of §IV: it is a *new galaxy-covariance Fisher computation* (c8_fnl_running_fisher.json, Planck 2018, CAMB 1.6.6), not a recast of an external $\sigma(f_{\text{NL}})$, and it constrains the SDB linear-bias amplitude $\Delta b(k) \propto f_{\text{NL}} b_\phi / k^2$ over six SPHEREx low-redshift bins rather than the bispectrum shape. Its motivation is that the *running* of f_{NL} is a sharper bounce-vs-inflation dis-

criminator than the amplitude alone; its headline result is that marginalizing over $n_{f_{\text{NL}}}$ degrades the SDB f_{NL} constraint by $2.0\text{--}4.6\times$ ($\sigma_{\text{unmarg}}(f_{\text{NL}}) = 1.53 \rightarrow \sigma_{\text{marg}}(f_{\text{NL}}) = 3.08$ fixed-bias $\rightarrow 7.06$ bias-marginalized; $\sigma(n_{f_{\text{NL}}}) = 0.295/0.596$; anti-correlation $\rho \approx -0.87$ to -0.97). It is reported here as a *subordinate cross-check* on the scale dependence, not as the paper’s headline forecast, which remains the bispectrum-only $5.2\text{--}5.5\sigma$ of §IV.

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: the dimensionless shape function B_{NL} is a degree-9 polynomial divided by a degree-9 normalization — degree-zero in the overall momentum scale — so the squeezed-limit amplitude carries no running at leading order). A joint Fisher forecast for $(f_{\text{NL}}, n_{f_{\text{NL}}})$ using SPHEREx scale-dependent bias over six redshift bins ($z = 0.1\text{--}1.5$, $f_{\text{sky}} = 0.75$; this is the low-redshift bin subset of the SPHEREx public-products sample structure used in the committed SDB Fisher computation, a different tracer selection from the $z \approx 0.5\text{--}2$ emission-line sample underlying the Heinrich *et al.* bispectrum forecast of Sec. IV — the two channels quote different redshift ranges by construction) yields $\sigma(n_{f_{\text{NL}}}) = 0.295$ after marginalizing over f_{NL} (0.596 with the per-sample linear biases additionally marginalized as free nuisances), with marginalized $\sigma(f_{\text{NL}}) = 3.08$ (7.06 bias-marginalized)—a $2.0\times$ ($4.6\times$) degradation from the f_{NL} -only constraint $\sigma_{\text{unmarg}}(f_{\text{NL}}) = 1.53$ due to a strong anti-correlated $f_{\text{NL}}\text{--}n_{f_{\text{NL}}}$ degeneracy ($\rho = -0.87$ fixed-bias, -0.969 bias-marginalized; here ρ is defined on the reduced post-marginalization 2×2 Fisher sub-covariance over $(f_{\text{NL}}, n_{f_{\text{NL}}})$ after the per-sample linear biases have been profiled out). With biases co-marginalized the $n_{f_{\text{NL}}}$ -fixed baseline is itself $\sigma(f_{\text{NL}}) = 1.75$ rather than 1.53 , and the two-parameter identity $\sigma_{\text{marg}} = \sigma_{\text{unmarg}} / \sqrt{1 - \rho^2}$ applied to that baseline reproduces the quoted value ($1.75 / \sqrt{1 - 0.969^2} = 7.06$); the $4.6\times$ factor above is referenced to the fixed-bias 1.53 for comparability with the fixed-bias column. 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 [1], building on the canonical SPHEREx multi-tracer galaxy-survey forecast lineage of Doré *et al.* [14]) drives the abstract $5.2\text{--}5.5\sigma$ optimistic figure (template-corrected) and the $2.6\text{--}5.5\sigma$ post-systematic-budget figure (§VII). (ii) A separate joint $(f_{\text{NL}}, n_{f_{\text{NL}}})$ Fisher matrix using SPHEREx scale-dependent bias over the same six redshift bins ($\sigma(n_{f_{\text{NL}}}) = 0.295/0.596$, $\sigma_{\text{marg}}(f_{\text{NL}}) = 3.08/7.06$, fixed-bias/bias-marginalized). In the SDB-only joint analysis the matter-bounce $f_{\text{NL}} = -4.375$

is detectable at only 1.4σ (fixed-bias) to 0.6σ (bias-marginalized) after marginalizing over $n_{f_{\text{NL}}}$: the SDB joint channel is a *subordinate* cross-check on the running, not a competitor to the bispectrum-only headline. The joint-SDB figures are therefore not the same baseline as the bispectrum-only $5.2\text{--}5.5\sigma$ 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 degradation, 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. The SDB-only $\sigma_{\text{unmarg}}(f_{\text{NL}}) = 1.53$ is consistent in scaling with, though $\approx 2.2\times$ weaker than, the bispectrum-only $\sigma(f_{\text{NL}}) = 0.7$ baseline and with the published SPHEREx SDB forecast lineage (Doré+2014, Heinrich+2024, Münchmeyer+2019), as expected for the six-bin subset of the full eleven-sample stack. The $n_{f_{\text{NL}}} = 0$ prediction is testable to ± 0.30 (fixed-bias) to ± 0.60 (bias-marginalized) at the same idealized-Fisher level. (DBI inflation is distinguished from local-type non-Gaussianity by its vanishing squeezed-limit amplitude and equilateral shape, not by its $n_{f_{\text{NL}}}$ running, so a meaningful DBI-vs-bounce SDB comparison would require the bispectrum-shape channel rather than the local $n_{f_{\text{NL}}}$ channel; we do not include DBI in the SDB discriminator list for this reason.) 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) [30], 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 $(k_3/k_1)^\Delta$ with scaling dimension $\Delta = 3/2 - \sqrt{9/4 - \mu^2/H^2}$; Chen & Wang [30]). At $\mu/H = 0$ (massless isocurvaton) $\Delta = 0$ and the exponent vanishes, reproducing the local-template scaling exactly (the curvaton-like limit); at $\mu/H = 3/2$ (the QSFI principal-vs-complementary-series boundary $m^2 = 9H^2/4$, distinct from the Higuchi bound which applies to spin-2 fields with $m^2 \geq 2H^2$) $\Delta = 3/2$ and the squeezed amplitude is suppressed by $(k_3/k_1)^{3/2}$ relative to local — the intermediate shape, departing furthest from the local template. 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 $\mu/H \rightarrow 0$, and a comprehensive SDB+bispectrum joint analysis is required to map the full discrimination region.

Channel hierarchy and sub-labeling note. This subsection (Sec. IX.D in the standard *Physical Review D* sectioning of the submitted manuscript) reports the joint scale-dependent-bias Fisher analysis as a *subordinate discriminator* within the Discussion; it is not a stand-alone SDB forecast section. The primary SPHEREx channel in this paper is the multi-tracer galaxy bispectrum (Sec. IV, $\sigma(f_{\text{NL}}) \approx 0.7$); the SDB channel (Sec. III, $\sigma_{\text{unmarg}}(f_{\text{NL}}) = 1.53$, $\sigma_{\text{marg}}(f_{\text{NL}}) = 3.08\text{--}7.06$) enters only as (a) the $\Delta b(k)$ signal definition underlying the bispectrum’s template-overlap calculation (Sec. III B) and (b) the joint $(f_{\text{NL}}, n_{f_{\text{NL}}})$ running discriminator reported here. The two channels use distinct Fisher matrices, distinct survey samples, and distinct sufficient statistics; they are complementary rather than competing. Readers seeking the headline SPHEREx forecast should consult Sec. IV; the joint SDB discussion in this subsection provides an additional model-discrimination diagnostic conditional on the bispectrum detection.

No trispectrum prediction is derived in this paper. The Suyama-Yamaguchi (SY) inequality $\tau_{\text{NL}} \geq (6f_{\text{NL}}/5)^2$ is governed by single- vs multi-source structure of the curvature perturbation, not by overlap of the local bispectrum template against the matter-bounce shape: for a structurally single-source bounce, a local single-source analogy would suggest $\tau_{\text{NL}} \approx (36/25)f_{\text{NL}}^2$ up to configuration-form-factor corrections, independent of the template-overlap $r \in [0.829, 0.876]$ reported in Sec. III B. A local single-source analogy with $|f_{\text{NL}}| = 35/8$ would correspond to an order-of-magnitude $\tau_{\text{NL}} \sim (36/25)f_{\text{NL}}^2 \approx 27.56$, far below the current Planck constraint ($\tau_{\text{NL}} < 2800$ at 95% CL [33]) and below SPHEREx sensitivity, but this is not used in any forecast in this paper; a quantitative τ_{NL} prediction for the matter bounce requires a direct trispectrum computation, which we do not undertake here. The τ_{NL} channel may become testable with future spectroscopic surveys achieving $\sigma(\tau_{\text{NL}}) \sim \mathcal{O}(10)$.

E. Caveats

a. Response to recurring referee concerns (signpost). Because this paper is a sensitivity *recast* of a single externally published baseline, several concerns recur predictably; each is already addressed in the body and is collected here so it need not be re-derived on a first reading. (i) “*The 2.6–5.5 σ and BF $\approx 9\text{--}14$ are not an independent multi-tracer Fisher forecast.*” Correct, and stated as such: the abstract *Scope* sentence and §III B/§IV declare that every significance and Bayes factor rescales the single imported Heinrich *et al.* [1] $\sigma(f_{\text{NL}}^{\text{local}}) \approx 0.7$ baseline through the explicit map $(\hat{f}^{\text{bounce}}, \hat{\sigma}^{\text{bounce}}) =$

($\hat{f}^{\text{local}}/r, \sigma^{\text{local}}/r$), $r = 0.84$ (Eq. 5); no independent bispectrum Fisher is constructed here. (ii) “Cubic-order transmission through the bounce is not computed, only linear.” Correct, and elevated to the *load-bearing caveat* (\star) in the abstract, in the Conclusion, and in the paragraph immediately below: assumption (d) is verified only at linear order [2], supported at cubic order solely by a superhorizon-scaling estimate, and a full cubic in-in computation is the stated #1 follow-up. (iii) “The Cai/Li factor of two rests on a symbolic identity, not a from-scratch four-vertex in-in re-derivation.” Correct, and we have accordingly *stopped claiming* the symbolic identity resolves it. As stated in the abstract and Appendix A: the two papers use the same normalization and the same squeezed limit and their shape-function polynomials agree coefficient-for-coefficient at $c_s = 1$, so the factor of two is neither the Komatsu–Spergel c -convention nor a demonstrable dropped time-ordering; it is an *unresolved literature discrepancy* in the local-limit reduction of a shared shape function. We adopt the Cai *et al.* value $-35/8$ as the headline input and carry the Li *et al.* value $-35/16$ as a genuine downward robustness branch (every significance halves), pending an independent from-scratch four-vertex in-in re-derivation, which we do not perform here. (iv) “A local-template Fisher is applied to a non-local shape through a scalar r ; the estimator covariance is not modeled.” Stated in the abstract and §III B: cross-parameter correlations are neglected in the recast, and any additional variance from the non-local tails of the bounce shape in the imported multi-tracer covariance is absorbed into the systematic envelope, with a full Heinrich-Fisher re-derivation at the bounce fiducial flagged for follow-up. (v) “The systematic budget is additive-in-quadrature, so 2.6σ is not necessarily conservative.” Stated at the head of §VII: the budget is an explicit heuristic ($\sigma_{\text{eff}} = \sqrt{\sigma_{\text{base}}^2 + \sum_i \sigma_i^2}$), not a joint-covariance forecast, and the realistic range is labeled a scoping envelope throughout; the one degeneracy that could dominate ($b_1\text{-}f_{\text{NL}}\text{-}n_{f_{\text{NL}}}$) has been computed explicitly (§VII, `c8.fnl_running_fisher.json`) and *loosens* rather than tightens the constraint. (vi) “BF $\approx 9\text{--}14$ is prior-dependent and should be illustrative.” Stated in the abstract and Conclusion: these Bayes factors are explicitly “illustrative of the discriminating power available . . . not . . . definitive model-selection evidence,” and their sensitivity to prior width is both mapped (§VI) and stated. None of these are unmodeled gaps; they are the deliberate, labeled scope of a sensitivity recast, and the follow-up computations that would convert each conditional statement into a derived bound are named in the text. *Single-source limitation (explicit)*. Because every significance and Bayes factor in this paper is a post-hoc rescaling of a single published Fisher forecast (Heinrich *et al.* [1]) recast to the bounce fiducial, the quoted σ endpoints and Bayes factors inherit that forecast’s survey specification, tracer assumptions, and covariance modeling, and are therefore *not statistically independent con-*

firmations of one another or of the bounce sensitivity — they are one baseline propagated through several book-keeping conventions. The concrete step required to break this single-source dependence is an independent bounce-fiducial multi-tracer bispectrum Fisher re-run (or a second, independently constructed forecast) that does not reuse the Heinrich *et al.* covariance; we identify this as the necessary follow-up and do not claim it here.

b. Leading theoretical uncertainty: cubic-order transmission through the bounce. The entire $f_{\text{NL}} = -35/8$ prediction depends on assumption (d) of Sec. II C: faithful third-order (bispectrum) transfer of the curvature perturbation through the nonsingular bounce. This is verified only at *linear* order in Wilson-Ewing [2]; at cubic order it is supported here only by a superhorizon-scaling estimate, not by an explicit calculation of the cubic transfer across the bounce surface. We stress that this is a *core theoretical uncertainty of the prediction itself*, not merely a forecast caveat: a bounce transition that does not faithfully transmit the third-order curvature correlator would shift, suppress, or reshape f_{NL} at the order being tested, and could in principle move it outside the recommended $\sigma_{\text{theory}} = 1.0$ bounce prior. A full cubic-order in-in computation of the curvature bispectrum across an explicit nonsingular bounce model is not performed in this work and is flagged as the most important follow-up required to upgrade the prediction from “conditional on assumption (d)” to “derived.” Until that computation exists, every significance and Bayes factor in this paper should be read as conditional on faithful cubic-order transmission.

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.

c. Auxiliary consistency check: cosmic birefringence. (This paragraph is an independent auxiliary check; the headline f_{NL} forecasts of this paper are independent of this channel.) Bounce-motivated physics can accommodate a spectator axion-like particle (ALP) coupling that generically produces cosmic birefringence; however, any specific numerical prediction for the rotation angle β de-

pends sensitively on the ALP coupling $g_{\phi\gamma}$ and mass m_a , which are free parameters in the spectator sector. We therefore note only that current CMB polarization measurements show a non-zero rotation signal: the Eskilt & Komatsu [39] joint WMAP+Planck measurement gives $\beta_{\text{obs}} = 0.342^\circ \pm 0.094^\circ$ (3.6σ from null); the ACT DR6 measurement of Diego-Palazuelos *et al.* [40] gives 2.9σ from null; and the Cosmoglobe DR1 II reanalysis [41] reports $\beta = 0.35^\circ \pm 0.70^\circ$, consistent with all three. Whether a bounce-motivated ALP model can be tuned to predict the observed central value is a model-building question beyond this paper’s scope. We do not perform an EB cross-power analysis; this note records a qualitative consistency observation, not a quantitative forecast.

X. CONCLUSION

The quasi-dust matter bounce makes a specific, falsifiable, minimally parameterized prediction: $f_{\text{NL}}^{\text{local}} = -35/8$ (with 0.6–8% ϵ -correction uncertainty), conditional on assumptions (a)–(f) in Sec. II C. *Load-bearing caveat* (\star): the single weakest link is assumption (d), faithful third-order (cubic) bispectrum transmission through the bounce, verified only at linear order [2] and supported at cubic order solely by an order-of-magnitude superhorizon-scaling estimate ($\delta f_{\text{NL}} \sim 10^{-3}$; a scaling argument, not a derived bound). Every “robust across the bounce class” statement in this paper is conditional on (d), and a full cubic in-in computation across an explicit bounce — our stated #1 follow-up — is required to convert the forecast from conditional to derived. As a robustness remark: a reader adopting the Li *et al.* alternative published value $f_{\text{NL}} = -35/16$ — which differs from the Cai *et al.* value $-35/8$ by an unresolved factor-of-two literature discrepancy that we adopt Cai’s side of but do not definitively settle (Appendix A) — would see the pre-systematic raw ratio drop to $|-35/16|/\sigma(f_{\text{NL}}) \approx 3.1$, propagating to a post-budget ~ 1.3 – 2.75σ headline. Because this discrepancy is not resolved here, this is a genuine downward robustness bound on the forecast, not merely a hypothetical. This value is robust across the Wilson-Ewing bounce class (Sec. II C, assumption e: no prolonged post-bounce inflation). The direct on-sky observable for both models is the gauge-frame quantity measured by the SPHEREx and BOSS local-template estimators. In the gauge frame, the Maldacena consistency relation gives $f_{\text{NL}}^{\text{inf, gauge}} = -(5/12)(n_s - 1) \approx 0.015$ for slow-roll single-field inflation; the bounce-vs-inflation gauge-frame amplitude ratio is therefore $|f_{\text{NL}}^{\text{bounce, gauge}}|/|f_{\text{NL}}^{\text{inf, gauge}}| \approx 290$, with the bounce prediction opposite in sign. *This gauge-frame contrast is the comparison relevant to the survey forecast and the only apples-to-apples observable-to-observable statement.* As a separate, purely theoretical point about the structure of the consistency relation: in the conformal-Fermi physical-observer frame, the squeezed-limit consistency relation [5, 6] implies

that $f_{\text{NL}}^{\text{inf, phys}} \rightarrow 0$ at leading order in slow-roll (with $\mathcal{O}(\text{slow-roll})$ residuals), so a non-zero physical-frame local detection would disfavor the single-field slow-roll attractor on theoretical grounds; this is a complementary theoretical discriminator, not the observable that SPHEREx measures. We retain the gauge-frame $\sim 290\times$ ratio as *the* bounce-vs-inflation discriminator for the survey forecast, and confine the physical-frame statement to its proper role as a theoretical consistency-relation point rather than carrying it into the headline gauge-frame comparison. We have shown that SPHEREx can test this prediction at a realistic post-budget 1.3 – 5.5σ significance through the multi-tracer galaxy bispectrum, where this range is amplitude-conditional and spans both published input values $f_{\text{NL}}^{\text{local}} \in [-35/16, -35/8]$ (Li-Cai): the lower bound reflects the conservative Li $-35/16$ input and the upper bound the optimistic Cai $-35/8$ input, so we do not foreground the Cai-only 2.6 – 5.5σ while the factor-of-two is open — this realistic range being a scoping sensitivity envelope under the additive-quadrature heuristic systematic budget of §VII, not a joint-covariance forecasted measurement precision, with no full bispectrum joint Fisher over the systematic nuisances performed here — (the Cai-input case reaches 5.2 – 5.5σ optimistically 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 6.25σ to 5.2 – 5.5σ), 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 3×10^5 Monte Carlo realizations across three framework-specific ensembles 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 $\text{BF} \approx 9$ – 14 under the noise-weighted $r \approx 0.84$ template-mismatch bookkeeping (~ 10 – 17 is the $r \rightarrow 1$ bookkeeping endpoint; both depend 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/main/research/> and archived at Zenodo (DOI inserted at submission). The joint $(f_{\text{NL}}, n_{f_{\text{NL}}})$ scale-dependent-bias Fisher computation referenced in §VII (c8_fnl_running_fisher.py, with its SPHEREx public-product survey inputs and output matrices) is included in the same repository, as is the GR-degradation Bayes-factor recompute of Table III (c9g_bf_table_recompute.py with its JSON output). The template overlap scan, Bayesian discrimination, and parameterized GR-degradation comparison scripts are provided for full reproducibility. Named artifacts referenced in the text are archived in the same repository and Zenodo release: the per-sample null-space significance propagation (c9h_nullspace_significance_propagation.json), the ϵ -ratio basis check (c9i_epsilon_ratio_check.json), the ℓ -space Fisher-overlap output (phase3_fisher_overlap.json), the null-space scan (null_space_analysis.py), the template-mismatch Bayes-factor bookkeeping check (c9j_bf_template_rescale.py with its JSON output), and the continuous GR-marginalization convergence check (c9k_gr_continuous_marginalization.py with its JSON output), and the continuous bounce-prior-width marginalization check (c9l_sigma_theory_continuous_marginalization.py with its JSON output), and the full bispectrum-shape coefficient map and per-configuration overlap values (phase3_bispectrum_shape_overlap.json, generated by null_space_analysis.py). No new observational data are introduced; all forecast sensitivities are adopted from published analyses [1, 15].

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

The factor-of-two discrepancy between $f_{\text{NL}} = -35/8$ (Cai et al. [7]) and $f_{\text{NL}} = -35/16$ (Li et al. [8], evaluated at $c_s = 1$) is a *genuine, unresolved literature discrepancy*, and we present it as such rather than as a resolved convention question. Two facts constrain what it can be. First, it is *not* the Komatsu-Spergel local-template normalization constant c defined below: both papers write the same c and the same amplitude normalization $f_{\text{NL}} = 10A/(3 \sum k_i^3)$ (Cai Eq. 21 = Li Eq. 4.20), so no c -rescaling separates them. Second, it is *not* a demonstrable dropped in-in time-ordering: the two papers' total shape-function polynomials agree coefficient-for-coefficient at $c_s = 1$ (Cai Eq. 37 = Li Eq. 4.19 with

$(10 - 9c_s^2) \rightarrow 1$), whereas a missing -2Im time-ordering would halve the *entire* polynomial, not merely the extracted local amplitude. Indeed Li et al.'s own in-in integral (their Eq. 4.13) is written explicitly as $-2 \times 2 \text{Im} \int$, i.e. the full in-in expression carrying both time-orderings, so no ordering is dropped and the discrepancy cannot be a single-ordering artifact. The residual factor of two therefore lies in the local-limit reduction / permutation bookkeeping that maps the shared shape function to the dominant local amplitude, and settling which value is correct requires a from-scratch numerical evaluation of the four-vertex in-in integrals that we do not perform. The generic in-in commutator identity of Sec. A is a true operator-algebra statement, but — contrary to an earlier version of this appendix — it does *not* establish that Li et al. omitted a time-ordering, and we do not use it to close the discrepancy. We adopt the Cai et al. value $-35/8$ as the headline input and carry $-35/16$ as a real downward robustness branch (Sec. A 0 d). The local-type bispectrum is defined, in terms of the matter-era Bardeen potential $\Phi = \frac{3}{5}\zeta$ (the field in which the Komatsu-Spergel constant is conventionally written), as:

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

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_{Φ} corresponds to $f_{\text{NL}}(c=1) = 2 f_{\text{NL}}(c=2)$.

In terms of the curvature perturbation the identical $c = 2$ convention reads $B_{\zeta} = \frac{6}{5} f_{\text{NL}} [P_{\zeta}(k_1)P_{\zeta}(k_2) + 2 \text{ perms}]$ — the local-template normalization of the Heinrich et al. forecast quoted in the abstract: with $P_{\zeta} = (5/3)^2 P_{\Phi}$ and $B_{\zeta} = (5/3)^3 B_{\Phi}$, the Φ -field $c = 2$ form maps exactly onto the ζ -field 6/5 form with the *same* f_{NL} ($B_{\zeta} = (5/3)^3 \cdot 2 f_{\text{NL}} P_{\Phi}^2 [+ \text{ perms}] = (5/3)^3 (3/5)^4 \cdot 2 f_{\text{NL}} P_{\zeta}^2 [+ \text{ perms}] = \frac{6}{5} f_{\text{NL}} P_{\zeta}^2 [+ \text{ perms}]$). The f_{NL} constrained by the SPHEREx estimator and the $f_{\text{NL}} = -35/8$ predicted by the matter bounce are therefore the same quantity, with no residual conversion factor between the appendix and survey normalizations.

We caution against the tempting narrative that Li et al. [8] simply computed a single time-ordering and that doubling their result recovers $-35/8$. That narrative is *not supported* by Li et al.'s published shape function: their total A_{tot} (Eq. 4.19) reduces at $c_s = 1$ to the same degree-9 polynomial with the same coefficients as Cai et al.'s A_T (Eq. 37), so Li's calculation is *not* a factor of two below Cai's at the level of the shape function. A dropped -2Im time-ordering would suppress the entire shape polynomial by two; it does not. The factor of two appears only in the scalar local amplitude extracted in the squeezed limit ($-35/8$ vs $-35/16$), i.e. in the reduction/permutation bookkeeping applied to a shared shape function, and we do not have an independent derivation

that fixes which reduction is correct. We therefore treat the discrepancy as unresolved.

In summary: both papers evaluate the same four-vertex in-in shape function in the same normalization and the same squeezed limit, yet report local amplitudes differing by exactly two, and this paper does not resolve which is correct. We *adopt* the Cai *et al.* observational value $f_{\text{NL}} = -35/8$ used by SPHEREx-oriented forecasts, and carry $-35/16$ as a genuine downward robustness branch (Sec. A 0 d). The detection significance $|f_{\text{NL}}|/\sigma(f_{\text{NL}})$ is normalization-independent in the narrow, provable sense that both f_{NL} and the Heinrich *et al.* $\sigma(f_{\text{NL}}) = 0.7$ are expressed in the same ζ -field 6/5 template normalization (the explicit Φ -to- ζ mapping above), so numerator and denominator share one convention and no rescaling enters; more generally both scale as $1/c$ under a change of the local-template constant, so the ratio is invariant under a consistent change of c . This c -invariance is *separate* from the unresolved Cai/Li factor of two, which is *not* a c -rescaling and does change

$|f_{\text{NL}}|$ (and hence the significance) by two between the two branches.

A.1 Explicit in-in Wick contraction derivation of the commutator doubling

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 it 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 Li *et al.* [8] 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 op-

erator 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 [3, 7]:

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

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:

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

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 ($\mathcal{S}_{\zeta\zeta^2} = 2$ for the two identical ζ legs, $\mathcal{S}_v = 1$ otherwise). Equation (A7) is the standard full in-in bispectrum, and both Cai *et al.* [7] and Li *et al.* [8] carry out this full four-vertex evaluation with the -2Im commutator structure; consistent with this, their published total shape-function polynomials agree at $c_s = 1$. We emphasize that the identity (A3) does *not* by itself imply that either paper omitted a time-ordering: it fixes the relation between a single time-ordered correlator and the full commutator, but does not adjudicate the residual factor-of-two difference in the two papers' extracted local amplitudes ($-35/8$ vs $-35/16$), which survives even though both use the full -2Im structure and the same shape polynomial. That residual factor of two is the unresolved literature discrepancy discussed in Appendix A, not a consequence of (A3).

d. Status of the ϵ -decomposition factor of two. Within Cai *et al.*'s own calculation, Eqs. 34–36 represent the single time-ordered correlator $\sum_v \sum_{\sigma} I_v$ before the -2Im doubling of (A7), so their ratio to Cai's full result is exactly $1/2$ as an operator-algebra identity — the symbolic verification above, not an empirical measurement. This is an *internal* step of Cai *et al.* and does *not* map onto the Cai-vs-Li discrepancy: Li *et al.* also report the doubled full result at $c_s = 1$ with the same shape polynomial, so their $-35/16$ is not Cai's undoubled intermediate. We note for transparency that the printed coefficients of Cai *et al.* Eq. (37), $(3, 1, -9, 5, -66, 9)$, are expressed in that paper's own monomial normalization and are *not* directly transplantable into this paper's symmetrized basis: direct evaluation of those coefficients (or their doubles) in our basis does not reproduce the benchmark values (artifact `c9i_epsilon_ratio_check.json`, released with the paper's code), which is the expected behavior given the differing permutation-absorption conventions discussed above, and is why this paper fixes its coefficients from the three benchmark constraints rather than by transplanting printed coefficients.

A.2 Time-ordering sensitivity Fisher table

Table V provides an explicit side-by-side detection forecast under the adopted Cai *et al.* value and the Li *et al.* alternative-literature value. The detection significance at fixed survey configuration is convention-independent (under a consistent change of the Komatsu–Spergel constant c of (A7) at fixed physical bispectrum, both f_{NL} and $\sigma(f_{\text{NL}})$ scale as $1/c$, so the ratio $|f_{\text{NL}}|/\sigma(f_{\text{NL}})$ is invariant). What differs between the two rows of Table V is the predicted central value of $|f_{\text{NL}}|$ for the matter bounce, and that difference is *not* a c -rescaling of one physical bispectrum: it is the unresolved Cai/Li factor-of-two literature discrepancy (Appendix A). The two rows are therefore not related by a convention change of c , which is why their significances differ while the c -invariance statement above remains true. Because we do not resolve which local-limit value is correct, the Li row is a genuine downward robustness bound, not merely a hypothetical. Table V reports both: The Cai row is the adopted headline forecast of this paper; the Li *et al.* row is a genuine downward robustness branch corresponding to the alternative published local amplitude. We quote the 5.25σ figure as the matter-bounce SPHEREx detection significance after the $r = 0.84$ template-overlap correction (further degraded to ~ 2.6 – 5.5σ by the joint systematic budget of Sec. VII), while stating explicitly that if the Li *et al.* local-limit reduction is the correct one, every significance halves — the Li row drops the SPHEREx detection to 2.6σ , preserving qualitative discriminatory power but losing the $\geq 5\sigma$ headline. We do *not* claim the operator-algebra identity of Sec. A.1 resolves this ambiguity in favor of Cai; it does not (both papers use the full commutator and agree at the shape-function level), and the correct local amplitude can be settled only by an independent from-scratch evaluation of the four-vertex in-in integrals, which is left to future work. A reproducibility notebook implementing the symbolic-algebra portion of the derivation ((A5)–(A7) at the level of operator structures and permutation factors, *without* the

TABLE V. Sensitivity of the SPHEREx detection significance to single- vs full-time-ordering in the matter-bounce f_{NL} derivation. Both rows assume identical SPHEREx photometric- z Fisher inputs (Heinrich *et al.* [1]) 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 $|f_{\text{NL}}|$ value is used to compute the significance $|f_{\text{NL}}| r / \sigma(f_{\text{NL}})$. The Cai row is the adopted headline value; the Li row is the alternative published value at $c_s = 1$ and is shown as a genuine downward robustness branch reflecting the unresolved Cai/Li factor-of-two literature discrepancy (Appendix A).

time-ordering branch	$ f_{\text{NL}}^{\text{bounce}} $	SPHEREx $\sigma(f_{\text{NL}})$	$ f_{\text{NL}} r / \sigma$
Cai <i>et al.</i> [7] (adopted headline value)	$35/8 = 4.375$	0.7	5.25σ
Li <i>et al.</i> [8] (alternative published value at $c_s = 1$; robustness branch)	$35/16 = 2.1875$	0.7	2.63σ

conformal-time integrals which require numerical evaluation) is archived alongside the paper source as `appendix_A1_wick_doubling.py`.

ACKNOWLEDGMENTS

The author thanks the developers of the open-source tools used in this work, including NumPy, SciPy, Mat-

plotlib, mpmath, and Tectonic. Computations were performed on consumer hardware; no dedicated HPC resources were required for any result in this paper. AI-assisted software tooling (Anthropic Claude) was used for code development, consistency checking, and manuscript preparation; all scientific claims, derivations, and editorial decisions are the author’s responsibility. No external funding was received for this research.

-
- [1] C. Heinrich, O. Dore, and E. Krause, Measuring f_{nl} with the spherex multi-tracer redshift space bispectrum, *Phys. Rev. D* **109**, 123511 (2024), [arXiv:2311.13082 \[astro-ph.CO\]](#).
 - [2] E. Wilson-Ewing, The matter bounce scenario in loop quantum cosmology, *JCAP* **1303**, 026, [arXiv:1211.6269](#).
 - [3] J. Maldacena, Non-gaussian features of primordial fluctuations in single field inflationary models, *JHEP* **0305**, 013, [arXiv:astro-ph/0210603](#).
 - [4] Planck Collaboration, N. Aghanim, *et al.*, Planck 2018 results. VI. Cosmological parameters, *Astron. Astrophys.* **641**, A6 (2020), table 2 (TT,TE,EE+lowE+lensing): $n_s = 0.9649 \pm 0.0042$, [arXiv:1807.06209](#).
 - [5] 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 \[astro-ph.CO\]](#).
 - [6] T. Tanaka and Y. Urakawa, Dominance of Gauge Artifact in the Consistency Relation for the Primordial Bispectrum, *JCAP* **1105**, 014, [arXiv:1103.1251 \[astro-ph.CO\]](#).
 - [7] Y.-F. Cai, W. Xue, R. Brandenberger, and X. Zhang, Non-gaussianity in a matter bounce, *JCAP* **0905**, 011, [arXiv:0903.0631 \[astro-ph.CO\]](#).
 - [8] Y.-B. Li, J. Quintin, D.-G. Wang, and Y.-F. Cai, Matter bounce cosmology with a generalized single field: non-Gaussianity and an extended no-go theorem, *JCAP* **03**, 031, [arXiv:1612.02036 \[astro-ph.CO\]](#).
 - [9] D. Wands, Duality invariance of cosmological perturbation spectra, *Phys. Rev. D* **60**, 023507 (1999), [arXiv:gr-qc/9809062](#).
 - [10] 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](#).
 - [11] D. Wands, Local non-Gaussianity from inflation, *Class. Quant. Grav.* **27**, 124002 (2010), [arXiv:1004.0818](#).
 - [12] 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](#).
 - [13] L. Freidel, D. Minic, and T. Takeuchi, Quantum gravity, torsion, parity violation and all that, *Phys. Rev. D* **72**, 104002 (2005), [arXiv:hep-th/0507253](#).
 - [14] O. Dore *et al.*, Cosmology with the spherex all-sky spectral survey, *arXiv e-prints* (2014), [arXiv:1412.4872](#).
 - [15] D. J. Schlegel *et al.*, The MegaMapper: A stage-5 spectroscopic instrument concept for the study of inflation and dark energy, *arXiv e-prints* (2022), [arXiv:2209.04322](#).
 - [16] N. Dalal, O. Dore, D. Huterer, and A. Shirokov, The imprints of primordial non-gaussianities on large-scale structure: scale-dependent bias and abundance of virialized objects, *Phys. Rev. D* **77**, 123514 (2008), [arXiv:0710.4560](#).
 - [17] 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](#).
 - [18] M. Zhu and Y.-F. Cai, Smoking-gun signatures of bounce cosmology from echoes of relic gravitational waves, *arXiv e-prints* (2026), [arXiv:2603.13924](#).
 - [19] 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](#).
 - [20] 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 \[astro-ph.CO\]](#).
 - [21] U. Seljak, Extracting Primordial Non-Gaussianity without Cosmic Variance, *Phys. Rev. Lett.* **102**, 021302 (2009), [arXiv:0807.1770 \[astro-ph\]](#).

- [22] P. McDonald and U. Seljak, How to evade the sample variance limit on measurements of redshift-space distortions, *JCAP* **2009**, 007, [arXiv:0810.0323 \[astro-ph\]](#).
- [23] 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 \[astro-ph.CO\]](#).
- [24] D. Baron and D. Poznanski, The weirdest SDSS galaxies: results from an outlier detection algorithm, *MNRAS* **465**, 4530 (2017), [arXiv:1611.07526](#).
- [25] Y. Liang *et al.*, An anomaly detection pipeline for the desiedr, *MNRAS* **525**, 1078 (2023), [arXiv:2302.05050](#).
- [26] 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 \[astro-ph.CO\]](#).
- [27] C. Addis, S. L. Guedezounme, J. Hammond, C. Clarkson, F. Montano, S. Camera, S. Jolicoeur, and R. Maartens, Unbiased analysis of primordial non-gaussianity: the multipoles of the full relativistic power spectrum, *arXiv e-prints* (2025), [arXiv:2511.09466](#).
- [28] A. Barreira, Can we actually constrain f_{nl} using the scale-dependent bias effect?, *arXiv e-prints* (2022), [arXiv:2205.05673 \[astro-ph.CO\]](#).
- [29] Y.-F. Cai, X. Chen, M. H. Namjoo, M. Sasaki, D.-G. Wang, and Z. Wang, Revisiting non-gaussianity from non-attractor inflation models, *JCAP* **2018** (05), 012, [arXiv:1712.09998 \[astro-ph.CO\]](#).
- [30] X. Chen and Y. Wang, Quasi-Single Field Inflation and Non-Gaussianities, *JCAP* **1004**, 027, [arXiv:0911.3380 \[hep-th\]](#).
- [31] 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 \[astro-ph.CO\]](#).
- [32] 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 \[astro-ph.CO\]](#).
- [33] Planck Collaboration, Y. Akrami, *et al.*, Planck 2018 results. IX. constraints on primordial non-Gaussianity, *A&A* **641**, A9 (2020), [arXiv:1905.05697 \[astro-ph.CO\]](#).
- [34] E. Chaussidon *et al.*, Constraining Primordial Non-Gaussianity with DESI 2024 LRGs and QSOs, *J. Cosmol. Astropart. Phys.* (2024), dESI DR1 LRG combined PNG constraint; $f_{\text{NL}}^{\text{loc}} = -3.6_{-9.1}^{+9.0}$; $\sigma \approx 9-10$, [arXiv:2411.17623](#).
- [35] A. Fondi *et al.*, Assembly bias and local Primordial non-Gaussianity from DESI DR1 Quasars, *arXiv e-prints* (2025), dESI QSO assembly-bias PNG analysis; $f_{\text{NL}}^{\text{loc}} = -3.3 \pm 9.2$; distinct from Chaussidon *et al.* LRG+QSO combined result, [arXiv:2602.12357](#).
- [36] DESI Collaboration, A. Aghamousa, *et al.*, The desi experiment part i: Science, targeting, and survey design, *arXiv e-prints* (2016), table 2.7: $\sigma(f_{\text{NL}}) \approx 3-5$ from multi-tracer scale-dependent bias, [arXiv:1611.00036](#).
- [37] Euclid Collaboration, Y. Mellier, *et al.*, Euclid. i. overview of the euclid mission, *A&A* **697**, A1 (2025), f_{NL} forecast from photometric survey: $\sigma(f_{\text{NL}}) \approx 2-4$, [arXiv:2405.13491](#).
- [38] 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](#).
- [39] 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](#).
- [40] P. Diego-Palazuelos and E. Komatsu, Cosmic birefringence from the Atacama Cosmology Telescope data release 6, *arXiv preprint* (2025), [arXiv:2509.13654 \[astro-ph.CO\]](#).
- [41] J. R. Eskilt *et al.* (Cosmoglobe), Cosmoglobe DR1 results. II. constraints on isotropic cosmic birefringence from reprocessed WMAP and Planck LFI data, *Astron. Astrophys.* **679**, A144 (2023), reports $\beta = 0.35^\circ \pm 0.70^\circ$ from WMAP+Planck LFI alone (no ACT), [arXiv:2305.02268](#).