A removability criterion for radiative corrections to the Starobinsky attractor

in weakly nonlocal gravity

arXiv:2608.17215

Daris Samart

Khon Kaen Particle Physics and Cosmology Theory group (KKPaCT)
Khon Kaen University

SHEAP 2026, Modified Gravity
Naresuan University, 13–20 August 2026

Three things to take away ✓

1. Gravity can be given an ultraviolet completion in which the classical inflationary solution is untouched. Every effect on observables is then a genuine quantum effect.
2. Whether a correction to the inflaton potential moves n_s is decided by its algebraic form, not by its size. One derivative test settles it.
3. Applied to the gravitational one-loop flow, the test says the flow is invisible. Quantum-gravitational running cannot explain the ACT shift.

Everything else in this hour is the reasoning behind these three lines.

Outline

The observational problem

What inflation gives us to measure

Single-field slow-roll inflation predicts a nearly scale-invariant spectrum of curvature perturbations. Two numbers carry almost all the information.

The observables

$$n_s - 1 = \frac{d \ln \mathcal{P}_\zeta}{d \ln k} \qquad \text{(scalar tilt)}, \qquad r = \frac{\mathcal{P}_T}{\mathcal{P}_\zeta} \qquad \text{(tensor-to-scalar ratio)}.$$

A third, the amplitude $A_s = 2.1 \times 10^{-9}$, is fitted rather than predicted. Remember this. It will do a lot of work later.

Add one operator to Einstein gravity,

$$S = \frac{M_{\rm P}^2}{2} \int d^4x \sqrt{-g} \left(R + \frac{R^2}{6M_S^2} \right).$$

The R^2 term propagates an extra scalar, the scalaron. A conformal transformation $g_{\mu\nu}^E = f'(R) g_{\mu\nu}^J$ with $\phi = \sqrt{3/2}\, M_{\rm P} \ln f'$ makes it canonical,

$$V_{\rm St}(\phi) = \frac{3}{4} M_{\rm P}^2 M_S^2 \left(1-u\right)^2, \qquad u \equiv e^{-a\phi}, \qquad a \equiv \sqrt{\tfrac{2}{3}}\,\frac{1}{M_{\rm P}}.$$

A plateau at large ϕ. The flatness is not tuned. It follows from the exponential in the frame transformation.

Why it is called an attractor ✓

On the plateau the slow-roll parameters lose all memory of the parameters,

$$\epsilon_V \to \frac{3}{4N^2}, \qquad \eta_V \to -\frac{1}{N}, \qquad n_s \to 1 - \frac{2}{N}, \qquad r \to \frac{12}{N^2}.$$

Numbers at $N_* = 55$

$$M_S = 1.28 \times 10^{-5} M_{\rm P} = 3.1 \times 10^{13}\ \mathrm{GeV}, \quad n_s = 0.96498, \quad r = 3.498 \times 10^{-3}.$$

One input, A_s, fixes M_S. Everything else is output. There is no shape parameter to adjust.

ACT DR6 pulls the other way ✓

ACT DR6 prefers a higher spectral index than the attractor value,

$$\Delta n_s \simeq +9.3 \times 10^{-3} \quad \text{at } N_* = 55.$$

The model has no dial to turn. So the community turned to corrections.

- R^3 and higher curvature terms
- marginal deformations $R^2 \to R^{2(1-\alpha)}\mu^{2\alpha}$
- logarithms of the field from matter loops
- polynomial corrections in the canonical field

Every one of them works. Percent-level changes move n_s into the preferred region.

The puzzle, and the two questions

If percent-level corrections move the prediction so easily, then quantum gravity, which certainly corrects the potential, should matter too.

Question 1, practical

Can the renormalization-group flow of the *gravitational* sector supply the ACT displacement?

Question 2, structural

Given an arbitrary δV, what property decides whether it moves n_s at all?

The literature answers the second implicitly: bigger correction, bigger effect. We will see that this is the wrong variable, and that fixing it answers the first.

But first: we need a theory to compute in

To ask what quantum gravity does to inflation we need a quantum theory of gravity that

1. is well defined in the ultraviolet,
2. has no ghosts,
3. and does not already ruin the classical prediction before we even start.

The next part builds such a theory from scratch. If you already know infinite-derivative gravity, treat it as a refresher.

Why gravity needs an ultraviolet completion

Einstein gravity as an effective field theory

Expanding $g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}/M_{\mathrm{P}}$, the graviton propagator behaves as

$$\Pi(k) \sim \frac{1}{k^2}, \qquad \text{vertices} \sim \frac{k^2}{M_{\mathrm{P}}}.$$

Superficial degree of divergence grows with the number of loops. New counterterms are needed at every order, each with its own free coefficient.

Consequence

Einstein gravity is predictive only below M_{P}. Inflation at $H \sim 10^{13}\,\mathrm{GeV}$ is close enough that the question of the completion is not academic.

Stelle's theorem, and its price ✓

Add all curvature-squared operators,

$$S = \int d^4x\sqrt{-g}\left(\frac{M_{\mathrm{P}}^2}{2}R + \alpha_C C^2 + \alpha_R R^2 + \alpha_{\mathrm{GB}}\mathcal{G}_{\mathrm{GB}}\right).$$

Now the propagator falls as $1/k^4$ and the theory is renormalizable. Decompose it into spin projections,

$$\Pi \sim \frac{1}{k^2} - \frac{1}{k^2 + m_2^2} + \cdots$$

The price

The improved ultraviolet behaviour comes from a second pole with the wrong-sign residue: a massive spin-2 ghost. Unitarity, or stability, is lost.

The idea: leave the polynomials ✓

Yes. An entire function with no zeros. In one complex variable these are exactly

$$e^{H(z)}, \qquad H \text{ entire.}$$

So replace the polynomial form factor by an exponential of an entire function of the d'Alembertian,

$$\text{polynomial in } \Box \quad \longrightarrow \quad e^{H(\Box/\Lambda_*^2)}.$$

Two immediate consequences

- e^H never vanishes, so the propagator gains no new poles and no new states
- e^H can grow arbitrarily fast in the ultraviolet, so loop integrals converge

This is the origin of the phrase *weakly nonlocal*, or *infinite-derivative*, gravity.

A terminology hazard, worth ten seconds

Two very different constructions carry the word “nonlocal”.

Ultraviolet nonlocality	**Infrared nonlocality**
form factor $e^{H(\Box)}$, entire, zero free	operators built on $\Box^{-1}$
super-renormalizable	not super-renormalizable
modifies the deep ultraviolet	modifies the late universe
this talk	Deser–Woodard type models

The distinction will matter again at the end: an infrared-nonlocal deformation is not a function of the instantaneous potential, so it falls outside our classification by construction.

Ghost and pole control, more carefully

Write the dressed inverse propagator schematically as

$$\Pi^{-1}(k) = P(k^2)\, e^{H(k^2/\Lambda_*^2)}.$$

- **Poles** come from zeros of Π^{-1}. Since $e^H \neq 0$ everywhere, the poles are exactly the zeros of P, that is those of the local theory.
- **Residues** at those poles are unchanged up to the finite factor e^H, so no wrong-sign state is introduced.
- **Convergence** is controlled by the growth of H along the relevant directions in the complex plane.

The trade that is left

Choosing H to be both fast growing and stable on a given background is non-trivial. Zeros appearing at loop level, and background-induced states, are where the current literature lives.

Super-renormalizability and one-loop exactness ✓

Let γ be the asymptotic polynomial degree of the form factor in D dimensions.

Power counting

For $\gamma > D/2$ the superficial degree of divergence stops growing with the loop order. Only a finite number of diagrams diverge, and with $N \geq 3$ they all sit at one loop.

Why this is a gift for us

The local divergent running is one-loop exact. There is no higher-loop local running to worry about. Any statement we make about the one-loop flow is a statement about the whole local flow.

Lineage: Krasnikov 1987; Kuz'min 1989; Tomboulis 1997; Modesto and Rachwał 2014.

Nonlocal Starobinsky: a short history

- **Briscese, Marcianò, Modesto, Saridakis 2013.** Inflation in super-renormalizable gravity.
- **Briscese, Modesto, Tsujikawa 2014.** A super-renormalizable or finite completion of the Starobinsky theory.
- **Koshelev, Modesto, Rachwał, Starobinsky 2016.** Exact R^2 inflation occurs inside a nonlocal ultraviolet-complete gravity.
- **Edholm 2017.** Constraints on the nonlocality scale from the tensor-to-scalar ratio.
- **Koshelev, Kumar, Starobinsky 2023.** Generalized nonlocal R^2-like inflation, and its non-Gaussianities.

The recurring difficulty

Earlier proposals struggled to be renormalizable *and* stable at the same time, and often the nonlocal operators disturbed the classical solution that one wanted to keep.

Where the field stands

Modesto and Orlando, 2026

A nonlocal completion of a general $f(R)$ theory in the Einstein frame that shares with $f(R)$ the same background solutions and the same equations of motion for perturbations, at linear and nonlinear level.

Consequences, in their words:

- classical cosmological observables are not affected by the nonlocal operators needed for the quantum completion
- the model overcomes the incompatibility of renormalizability and stability of previous proposals
- the theory is at least super-renormalizable, and can be finite

This is the completion we will work in. Let us look at how the construction achieves it.

The completion we use

The construction ✓

Take any local action $S_{\rm loc}$ and add a term quadratic in its own equations of motion,

$$S_{\rm NL} = S_{\rm loc} + \frac{1}{2}\int d^4x\sqrt{-g}\, E_i \left(e^{H(\Delta/\Lambda_*^2)} - 1\right)^i{}_k \left(\Delta^{-1}\right)^{kj} E_j,$$

$$E_i = \frac{\delta S_{\rm loc}}{\delta \Phi^i}, \qquad \Delta_{ij} = \frac{\delta E_i}{\delta \Phi^j}, \qquad H \text{ entire, so } e^H \text{ has no zeros.}$$

Read the structure: the added piece is built entirely from the local equations of motion E_i and vanishes quadratically when they are satisfied.

Property one: the background survives exactly ✓

Vary the added term. Because it is quadratic in E_i, its first variation is linear in E_i.

$$\frac{\delta S_{\mathrm{NL}}}{\delta \Phi} = E + \mathcal{O}(E) \quad \Longrightarrow \quad E_i = 0 \;\; \text{still solves the completed theory.}$$

What this buys us

The Starobinsky background, the slow-roll trajectory, and every classical prediction are inherited exactly. Not approximately, not up to $\mathcal{O}(\Lambda_*^{-2})$.

This is the pedagogical heart of the setup. Any deviation in the observables must therefore be a quantum effect, with nowhere else to hide.

Property two: the spectrum survives too

On shell the second variation is simply dressed,

$$\delta^2 S_{\mathrm{NL}}\Big|_{E=0} = \delta\Phi\,\Delta\,e^{H(\Delta/\Lambda_*^2)}\,\delta\Phi.$$

Since e^H is zero free, the tree-level poles remain those of the local theory,

$$k^2 = 0 \qquad \text{and} \qquad k^2 = -M_S^2.$$

In the two spin channels,

$$\Pi_2^{-1} \propto k^2\left[1 - \frac{4\alpha_C}{M_{\mathrm{P}}^2}k^2\right], \qquad \Pi_0^{-1} \propto k^2\left[1 + \frac{12\alpha_R}{M_{\mathrm{P}}^2}k^2\right], \qquad \alpha_R^{\mathrm{tree}} = \frac{M_{\mathrm{P}}^2}{12M_S^2}.$$

Graviton plus scalaron. Nothing else.

The stage is now clean

What we have

- a super-renormalizable, ghost-free quantum theory of gravity
- whose classical inflationary solution is exactly the Starobinsky one
- and whose local divergent running terminates at one loop

So the question sharpens

Everything classical is fixed. The only thing that can move n_s is the running of the couplings. Does it?

Renormalization-group improvement

What renormalization-group improvement means

The one-loop effective potential contains large logarithms,

$$V_{\text{1loop}} \sim V + \#\ln\frac{\mathcal{O}}{\mu^2} + \cdots$$

Improvement means: choose μ to kill the large logarithm at each point, and let the couplings run instead.

$$V \;\longrightarrow\; V\big(\text{couplings}(\mu)\big), \qquad \mu = \mu(\phi).$$

The scheme dependence is physical here

The choice of $\mu(\phi)$ is not a relabelling. Different prescriptions resum different logarithms and, as we will see, give different answers.

The ingredients

- Six local couplings run,

$$\left\{M_{\mathrm{P}}^2,\ \rho_\Lambda,\ M_S^2,\ \xi,\ \alpha_C,\ \alpha_R\right\}$$

- Divergences from the Seeley–DeWitt coefficient for $\Delta = -\Box + E$,

$$a_2 = \frac{1}{180}\left(R_{\mu\nu\rho\sigma}^2 - R_{\mu\nu}^2\right) + \frac{1}{2}\left(E - \frac{R}{6}\right)^2 + \cdots$$

- Poles converted to beta functions by

$$\Gamma \supset + \int \sqrt{-g}\,\omega O, \quad \Gamma_{\mathrm{div}} = +\frac{X}{2\bar{\epsilon}} \int \sqrt{-g}\, O \implies \mu\frac{d\omega}{d\mu} = -X.$$

Watch the signs. For the potential the pair is $(-V, 1)$ because $\mathcal{L} \supset -V$, so $\beta_V = +X_V$.

The scale choice

The gravitational sector responds to the background curvature. On the slow-roll solution that is the Hubble rate, so

$$\mu = H(\phi).$$

An apparently innocent alternative is $\mu^2 = |V''|$. It gives

$$\ell_B \simeq -\frac{a\phi}{2},$$

a logarithm growing linearly in the field. We will see at the end that this is not a scheme ambiguity but a different physical prescription, and that it produces an unphysical local maximum.

One exact identity ✓

The technical foundation of everything that follows

$$\ell = \ln \frac{1-u}{1-u_0} = \frac{1}{2} \ln \frac{V_{\mathrm{St}}(\phi)}{V_{\mathrm{St}}(\phi_0)}$$

The second equality is exact, not a plateau expansion.

In words: with the curvature prescription, the renormalization-group time is one half of the logarithm of the tree potential. Keep this line in mind. The entire result is a consequence of it.

The deformation the flow produces ✓

$$V_{\mathrm{RG}} = V_{\mathrm{St}}\,[1+\Delta]\,,$$

$$\Delta = (\gamma_M + b_G)\,\ell - b_G\,\frac{a\phi\,u}{1-u}\,\ell + \frac{\rho_\Lambda + B_\Lambda \ell}{V_{\mathrm{St}}} + \frac{\Delta V_{\mathrm{1loop}}}{V_{\mathrm{St}}}$$

Four running parameters, four structures. The Coleman–Weinberg piece is $\mathcal{O}(10^{-16})$ for $\xi = 0$, so it is dropped.

They collapse under one derivative

$$\frac{V_{\mathrm{St}}}{V'_{\mathrm{St}}}\,\Delta' = \frac{K}{2}, \qquad K \equiv \gamma_M + b_G + \hat{B}_\Lambda - 2\hat{\rho}_\Lambda.$$

Four couplings enter the observables through one constant. Why that particular combination is the right question is the subject of the next part.

Removable and non-removable deformations

Step back: what can observations actually see? ✓

In the slow-roll N-formalism with $\epsilon = -d\ln H/dN$,

$$n_s - 1 = -2\epsilon - \frac{d\ln\epsilon}{dN}, \qquad r = 16\epsilon.$$

Both are functionals of the single function $\epsilon(N)$. Now run the logic backwards.

- $d\ln H/dN = -\epsilon$ determines H up to one multiplicative constant
- $d\phi/dN = M_{\mathrm{P}}\sqrt{2\epsilon}$ determines ϕ up to one additive constant

Two integration constants. Two directions the data cannot resolve.

The equivalence theorem ✓

Exact statement

Two single-field slow-roll models predict identical n_s and r, with $\epsilon(N)$ matching over a finite interval, **if and only if** their potentials differ by a constant amplitude rescaling combined with a constant inflaton shift.

Why exactly these two are invisible:

- the **amplitude** is absorbed by A_s, which we fit and do not predict
- the **field origin** is meaningless, since N is counted from the end of inflation

Analogy. Changing the volume of a melody and transposing its key do not change the tune, because the ear reads intervals. Here n_s and r are the intervals.

Definitions, stated once and used throughout ✓

Write any correction to the potential multiplicatively,

$$V(\phi) = V_{\mathrm{St}}(\phi)\,[1 + \Delta(\phi)]\,.$$

Δ is the deformation. It is a *function* of ϕ, not a number.

Removable

Δ is **removable** if an amplitude rescaling together with a constant inflaton shift undoes it, so the leading predictions are unchanged.

Non-removable

Δ is **non-removable** otherwise. It changes the shape of $\epsilon(N)$ and therefore moves n_s.

Neither definition mentions the size of Δ.

The criterion ✓

Demand that the induced displacement of the e-fold mapping be uniform along the trajectory. That single requirement gives

$$\boxed{\mathcal{R}[\Delta] \equiv \frac{V_{\rm St}}{V'_{\rm St}}\,\Delta' = \text{const} \iff \Delta = c_1 + c_2 \ln \frac{V_{\rm St}}{A_*}}$$

Equivalently, on the plateau,

$$V \longrightarrow A_* \left(\frac{V}{A_*}\right)^{1+c_2}.$$

How to read it

A removable deformation raises the potential to a slightly different power. It rescales the amplitude, which A_s absorbs, and rescales the coefficient of $e^{-a\phi}$, which a field redefinition absorbs.

How to use it, in three steps ✓

On the Starobinsky plateau the criterion collapses. With $u = e^{-a\phi}$,

$$\frac{V_{\rm St}}{V'_{\rm St}} = \frac{1-u}{2au}, \qquad \Delta' = -au\,\frac{d\Delta}{du} \qquad \Longrightarrow \qquad \boxed{\mathcal{R}[\Delta] = -\frac{1-u}{2}\,\frac{d\Delta}{du}}$$

The factors of a and u cancel completely.

The recipe

1. Write Δ as a function of u, using $a\phi = \ln(1/u)$ and $R = 3M_S^2(1-u)/u$ where needed.
2. Take one derivative with respect to u.
3. Multiply by $-(1-u)/2$ and ask whether the answer still depends on u.

No slow-roll integrals. No spectra. One derivative.

Worked example one: an extra R^2

For any addition to $f(R)$,

$$\Delta = -\frac{6M_S^2}{R^2}\,\delta f(R), \qquad R = 3M_S^2\,\frac{1-u}{u}.$$

Take $\delta f = c_2 R^2$. Then Δ is a constant, so $\Delta' = 0$ and

$$\mathcal{R} = 0 = \mathrm{const}.$$

Removable, identically

And we know why: shifting the coefficient of R^2 just renormalizes M_S, which A_s reabsorbs. The criterion reproduces a fact we already knew. That is the sanity check.

Worked example two: an extra R^3 ✓

Now take $\delta f = c_3 R^3$. Using the same lemma,

$$\mathcal{R}\left[\Delta_{(n)}\right] = -\frac{(n-2)c_n}{2}\frac{(1-u)^{n-2}}{u^{n-1}} \simeq -\frac{(n-2)c_n}{2}\left(\frac{4N}{3}\right)^{n-1}.$$

For $n = 3$ this is $\mathcal{O}(N)$, and it grows along the trajectory.

The general statement

Only $n = 2$ is removable. Every higher power violates the criterion at a rate that increases with n.

Notice what did the work: not the coefficient c_n, but the u-dependence.

Worked example three: the one that fools you

Take $\Delta = c\,u$. It looks nothing like $\ln V_{\mathrm{St}}$, so it should fail.

$$\frac{d\Delta}{du} = c \qquad \Longrightarrow \qquad \mathcal{R} = -\frac{1-u}{2}\,c = -\frac{c}{2} + \mathcal{O}(u).$$

Constant. It passes.

Why

On the plateau $\ln V_{\mathrm{St}} = 2\ln(1-u) \simeq -2u$, so

$$c\,u \simeq -\frac{c}{2}\ln V_{\mathrm{St}}.$$

It is a member of the removable class in disguise.

Moral: inspecting Δ by eye is not enough. Take the derivative.

The complete invisible class

Exact invisibility and the criterion define *different* classes, meeting only in the constant. Together they span

$$\Delta_{\mathrm{inv}} = c_1 + c_2 \ln \frac{V_{\mathrm{St}}}{A_*} + c_3 \frac{V'_{\mathrm{St}}}{V_{\mathrm{St}}}.$$

A geometric identification

This is the tangent space of the E-model family $V = A\left(1 - e^{-a(\phi-\phi_0)}\right)^p$ at $p = 2$, with directions ∂_A, ∂_p, ∂_{ϕ_0}.

Evaluating the criterion on it,

$$\mathcal{R}\left[\Delta_{\mathrm{inv}}\right] = c_2 + c_3\,\mathcal{G}, \qquad \mathcal{G} = \frac{V''_{\mathrm{St}}}{V'_{\mathrm{St}}} - \frac{V'_{\mathrm{St}}}{V_{\mathrm{St}}} = -\frac{a}{1-u} = \mathcal{O}(1).$$

Read at $\mathcal{O}(u) = \mathcal{O}(1/N)$, the criterion is necessary *and* sufficient.

Results

How badly does "size matters" fail? ✓

Exact slow-roll solutions at $N_* = 55$.

Change applied	Size	Δn_s
Amplitude $A \to 1.15A$	15%	**exactly zero**
Field shift $\phi_0 + 0.6\, M_{\mathrm{P}}$	very large	**exactly zero**
Exponent $p : 2 \to 2.4$	20%	-1.8×10^{-4}
An R^3 term	0.0067%	$+9.7 \times 10^{-3}$

The first two are visible by eye in a plot of V and move nothing. The last lies on top of the tree curve and reaches the ACT central value. Four orders of magnitude in size, opposite conclusions.

The deformation from the flow, plotted

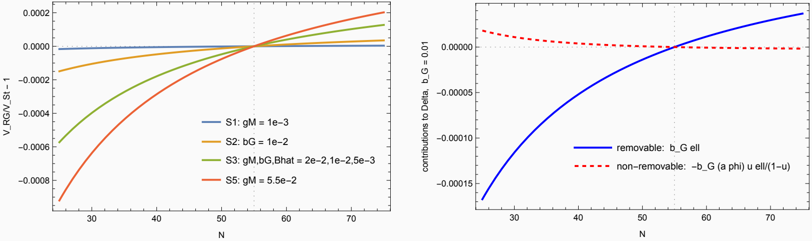

Left: $V_{\rm RG}/V_{\rm St} - 1$ across the observable window. Right: the b_G channel split into its removable part and the non-removable exponent shift.

Where the flow lands ✓

Resumming $dV/dt = KV$ with $t = \ell$ gives $V = V(0)e^{K\ell}$, and the exact identity turns the exponential into a power of the tree potential,

$$V_{\mathrm{RG}} \propto \left(1 - e^{-a\phi}\right)^{p} \left[1 + \mathcal{O}(u \ln u)\right], \qquad p = 2 + K.$$

The result

The one-loop flow moves the model *along* the E-model family. It never leaves it at leading order. And the predictions of that family are exponent independent.

One-loop exactness closes the question: no higher-loop local running reopens it.

The truncation constant is not a free parameter

The scalaron mass is the R^2 coefficient in a different basis, so

$$\gamma_M = b_G - \frac{B_R}{\alpha_R}, \qquad \omega_{0,R} = \alpha_R^{\mathrm{tree}} = \frac{M_{\mathrm{P}}^2}{12 M_S^2} = \frac{N_*^2}{288\pi^2 A_s} = 5.1 \times 10^8.$$

$$\Longrightarrow \qquad \gamma_M = b_G \left[1 + \mathcal{O}(10^{-5})\right], \qquad K = 2b_G + \hat{B}_\Lambda - 2\hat{\rho}_\Lambda.$$

Standard Model matter

Only nonminimally coupled scalars feed B_R, so for the Standard Model the Higgs alone,

$$\Delta n_s = 1.2 \times 10^{-14} \left(\xi_H - \tfrac{1}{6}\right)^2.$$

Specifying the matter content therefore predicts K rather than parametrizing it.

What survives: the residuals

The leading shift cancels, so what remains is the difference between two members of the family,

$$\Delta n_s^{\mathrm{shape}} \xrightarrow[\kappa \ll 1]{} \frac{K}{N^2}\left[1.458 - \frac{3}{4}\ln N\right], \qquad \left.\frac{\Delta r}{r}\right|^{\mathrm{shape}} = \frac{3K}{4N}\left[\ln\frac{4N}{3} - 1\right].$$

- one power of $1/N$ below the naive estimate
- the coefficient of $\ln N$ is $-3/4$ in *every* removable channel
- reproduces the exact numerical solution to better than 5% over $45 \leq N \leq 65$

Residuals, channel by channel

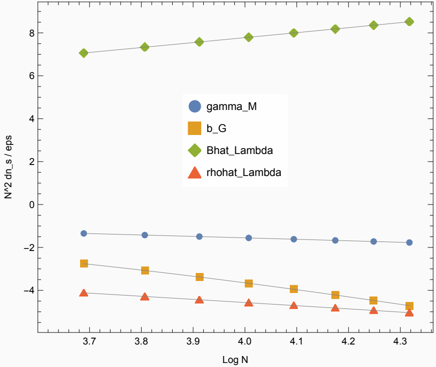

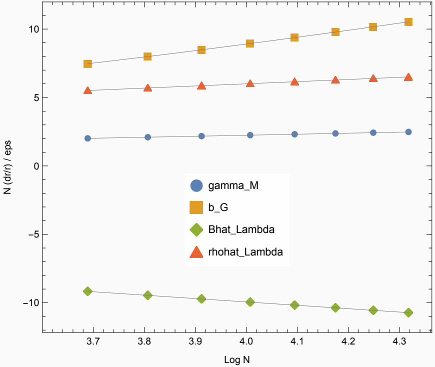

The γ_M channel is purely removable and matches the analytic coefficients $(1.458, -3/4)$ to better than one percent.

The one channel that is not protected

Removability protects the *shape*. It says nothing about N_*, which the post-inflationary history sets.

The e-fold channel

Contributes at the 22% level relative to the shape channel, with an opposite sign.

This is the single route by which gravitational running reaches the observables without being removable. It is also the least suppressed physical channel, and a natural target for future work.

Bounding the effect from above ✓

Matching the ultraviolet and infrared descriptions gives a window for the nonlocality scale,

$$1.2 \times 10^{14}\ \mathrm{GeV} \lesssim \Lambda_* \lesssim 1.8 \times 10^{16}\ \mathrm{GeV},$$

and, for any perturbative completion satisfying the naturalness condition,

$$\boxed{\Delta n_s < 3.7 \times 10^{-4}}$$

Compare

ACT wants $+9.3 \times 10^{-3}$. That is a factor of 25 larger. The conclusion is structural and does not rest on the explicit matching estimates.

Where the model sits

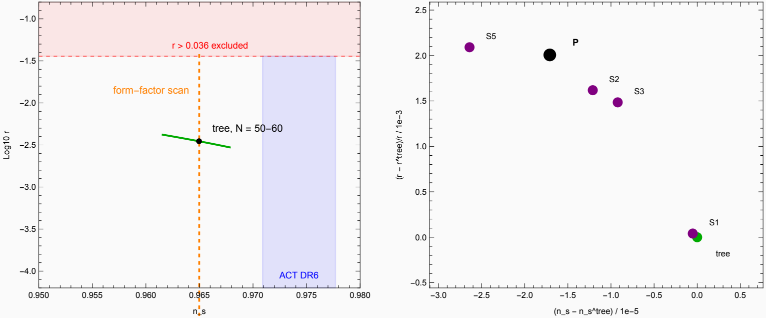

Tree trajectory and benchmark points in the (n_s, r) plane. The observational bands are indicative.

Discussion

Why matter is different from gravity ✓

The whole distinction is which logarithm appears.

	logarithm	criterion
gravitational flow at $\mu = H$	$\ln V_{\mathrm{St}}$	constant
matter loops	$\ln \phi$, $a\phi$	grows with N

A logarithm of the *potential* is removable by construction, by the exact identity. A logarithm of the *field* is not, because ϕ and $\ln V_{\mathrm{St}}$ are not proportional on the plateau. This is why matter-induced corrections shift these models and purely gravitational ones do not.

Every deformation in one table

✓

$\Delta(u)$	$\mathcal{R}[\Delta] = -\frac{1-u}{2}\, d\Delta/du$	scaling	verdict
generated by the flow			
$K\ell$	$K/2$	exactly constant	removable
$b_G\gamma_M\ell^2$	$b_G\gamma_M\ell$	$\mathcal{O}(u)$	no
$\hat{B}_\Lambda\ell/(1-u)^2$	$-\hat{B}_\Lambda(2\ell-1)/2(1-u)^2$	const $+\mathcal{O}(u)$	no
$\hat{\rho}_\Lambda/(1-u)^2$	$-\hat{\rho}_\Lambda/(1-u)^2$	const $+\mathcal{O}(u)$	no
calibration			
$c\,u$	$-c(1-u)/2$	const $+\mathcal{O}(u)$	removable
$V_{\mathrm{St}}(\phi+\delta)$	$-(1-s)(1-su)/(1-u)^2$	const $+\mathcal{O}(u)$	removable
proposed in the literature			
$-c(1-u)/u$ $[R^3]$	$-c(1-u)/2u^2$	$\mathcal{O}(N^2)$	no
$2\alpha\ln\frac{1-u}{u}$ [marginal]	α/u	$\mathcal{O}(N)$	no
$\delta\ln(\phi/M_{\mathrm{P}})$	$\delta(1-u)/2au\phi$	$\mathcal{O}(N/\ln N)$	no
$c\,(a\phi)^2$	$c(1-u)a\phi/u$	$\mathcal{O}(N\ln N)$	no
$\Box^{-1}R$	not a function of V_{St}	—	outside

Every proposal that raises n_s fails, not because it is large but because its $\mathcal{R}$ grows with N.

What "constant" means operationally

Variation of $\mathcal{R}$ across the observable window, $N = 40$ to $N = 55$.

deformation	variation of $\mathcal{R}$
$c\,u$	0.4%
field shift	0.4%
$\delta \ln(\phi/M_{\mathrm{P}})$	20.6%
marginal	25.7%
$c\,(a\phi)^2$	31.0%
R^3	45.0%

A separation by a factor of fifty or more. The test is not marginal, and it needs no fitting, no likelihood, and no spectra. We do not read this as an argument against those proposals. It locates the origin of their effect in the algebraic structure rather than in the magnitude.

Could the flow do it with a larger coupling?

Ask what amplitude a linear extrapolation would need to reach $\Delta n_s = +9.3 \times 10^{-3}$.

- literature entries: $|\varepsilon_{\mathrm{ACT}}| \ll 1$
- flow-generated entries: $|\varepsilon_{\mathrm{ACT}}| \gtrsim 1$

What the second line means

At those amplitudes the "correction" exceeds the tree potential by one to three orders of magnitude. The number records an impossibility, not a value.

Evaluated exactly there, the shift falls three to fourteen times *short*. The nonlinearity works against the displacement, not for it.

The scale choice was physical after all

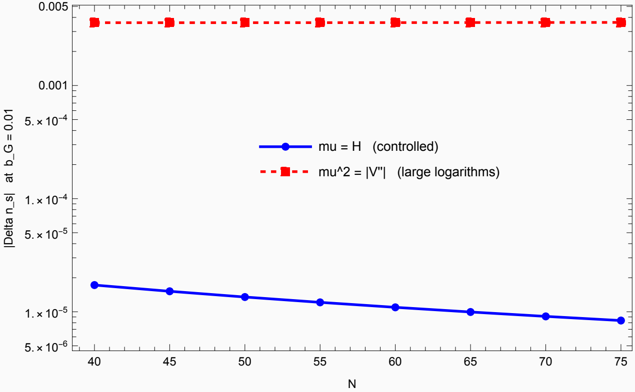

$\mu = H$ gives $\Delta n_s = -1.2 \times 10^{-5}$. The choice $\mu^2 = |V''|$ gives -3.6×10^{-3} and an unphysical local maximum.

Prescription, not scheme

- $\mu = H$ tracks the curvature that the gravitational sector actually responds to on the slow-roll solution
- $\mu^2 = |V''|$ resums logarithms of an exponentially varying field-space quantity

The point

The curvature prescription forces the deformation to be removable, and it does so through the exact identity rather than by construction.

Related work resumming a comoving-scale flow, or using functional renormalization-group improvements, or quadratic gravity with large matter multiplicity, is consistent with this. Those mechanisms rest on contributions of exactly the non-removable type the criterion identifies.

What we assumed

1. The ghost-free condition on the form factor is tuned for a specific background. On an inflationary background these theories generically carry additional complex-mass states, outside the present truncation.
2. The truncation is purely gravitational. The contrast with matter is the objective, not an oversight.
3. The improvement resums leading logarithms only, so non-removable residual coefficients are evaluated numerically.
4. Matching fixes the parametric dependence on Λ_* but leaves two dimensionless parameters undetermined without a concrete form factor.

Putting it together

The whole argument in six steps ✓

1. The completion adds a term quadratic in the equations of motion, so the classical Starobinsky solution and its spectrum are inherited exactly. *Every effect on observables is quantum.*
2. Improving at $\mu = H(\phi)$ makes the flow time equal to $\frac{1}{2}\ln V_{\rm St}$, by an exact identity.
3. Observables are functionals of $\epsilon(N)$ alone, and reconstruction leaves two integration constants. Those are the directions data cannot see.
4. Demanding a uniform displacement of the e-fold map gives $\mathcal{R}[\Delta] = (V_{\rm St}/V'_{\rm St})\Delta' = \mathrm{const}$, whose solution space is the tangent space of the E-model family at $p = 2$.
5. The gravitational flow satisfies the criterion. It moves the model along that family with $p = 2 + K$, and the family is exponent independent.
6. What survives is $\mathcal{O}(\ln N/N^2)$, bounded above by 3.7×10^{-4}. A factor of 25 short of ACT.

What is new here

1. The invisible class is characterized completely, and identified geometrically as the tangent space of the E-model family at $p = 2$. That upgrades the criterion from necessary to necessary and sufficient.
2. Removability, not magnitude, is the decisive property. This inverts the assumption the surrounding literature runs on.
3. A no-go: local one-loop running of the gravitational sector cannot supply the ACT displacement in this class of completions.
4. The nonlinearity works against the displacement. Evaluated exactly at the required amplitude, the shift falls three to fourteen times short.
5. The truncation constant is predicted, not fitted, because the measured amplitude fixes $\alpha_R = 5.1 \times 10^8$.

Why you might care, whichever side you work on

If you build inflationary models

A fast test for any new proposal, applied *before* the full numerics. It also tells you in advance which directions are futile.

If you work on quantum gravity

A concrete observational handle on super-renormalizable completions, and the result that gravitational running is invisible in the two-point spectra.

If you analyse CMB data

A statement of which part of theory space ACT can discriminate, and which part it cannot in principle. The (n_s, r) plane becomes a diagnostic, not a fitting space.

And the general lesson transfers: wherever observables are functionals of one function, the integration constants of its reconstruction are always the invisible directions.

Glossary

Glossary I: the gravity side

Term	Meaning	Form
Form factor	entire function of $\Box$, no zeros, so no new poles	$e^{H(\Box/\Lambda_*^2)}$
Weakly nonlocal	ultraviolet nonlocality. Not $\Box^{-1}$ models	H entire
Super-renormalizable	finitely many divergent diagrams, all at one loop	$\gamma > D/2$, $N \geq 3$
The completion	adds a term quadratic in the equations of motion	$S_{\rm loc} + \frac{1}{2}E\,(e^{H} - 1)\Delta^{-1}E$
Scalaron	the scalar carried by R^2, the inflaton here	mass $M_S = 3.1 \times 10^{13}$ GeV

One picture to keep: the form factor is a factor that never vanishes, so it can make loops converge without adding states.

Glossary II: the criterion

Term	Meaning	Form
Deformation	fractional change of the potential, a function of ϕ	$V = V_{\mathrm{St}}(1+\Delta)$
Removable	undone by amplitude rescaling plus a field shift	$\Delta = c_1 + c_2 \ln(V_{\mathrm{St}}/A_*)$
Non-removable	changes the shape of $\epsilon(N)$, so it moves n_s	$\Delta n_s = \mathcal{O}(\mathcal{R}/N)$
The criterion	one derivative decides which of the two it is	$\mathcal{R} = \dfrac{V_{\mathrm{St}}}{V'_{\mathrm{St}}}\Delta'$
Invisible class	all three directions data cannot resolve	$c_1 + c_2 \ln V_{\mathrm{St}} + c_3 \dfrac{V'_{\mathrm{St}}}{V_{\mathrm{St}}}$
E-model family	where those directions live, exponent independent	$V = A(1 - e^{-a(\phi-\phi_0)})^p$

One picture to keep: volume and key do not change a melody. Intervals do.

Glossary III: the numbers worth remembering

Symbol	What it is	Value or scaling
ℓ	the flow time, an exact identity	$\ell = \frac{1}{2}\ln\big(V_{\mathrm{St}}/V_{\mathrm{St}}(\phi_0)\big)$
K	the one constant four couplings collapse into	$p = 2 + K,\ \ \mathcal{R} = K/2$
α_R	R^2 coefficient, fixed by A_s	5.1×10^8
e-fold channel	the only unprotected route	22%, opposite sign
Residual	what survives after the cancellation	$\mathcal{O}(\ln N/N^2)$
Matching bound	the most the flow can ever do	$< 3.7 \times 10^{-4}$, ACT wants 9.3×10^{-3}

One picture to keep: a 15% change that does nothing, and a 0.0067% change that does everything.

Conclusions

What the criterion does

One derivative decides

$$\mathcal{R}[\Delta] = \frac{V_{\mathrm{St}}}{V'_{\mathrm{St}}}\Delta'$$

constant $\Rightarrow$ removable, no shift at leading order, residual $\mathcal{O}(\ln N/N^2)$

varying $\Rightarrow$ non-removable, $\Delta n_s = \mathcal{O}(\mathcal{R}/N)$

No cosmological observables need be recomputed. The invisible class is completely characterized, as the tangent space of the E-model family at $p = 2$.

What we learn about quantum gravity and inflation ✓

- In a completion that leaves the classical solution exact, the gravitational one-loop flow at $\mu = H$ is a removable direction.
- It moves the model along the E-model family with $p = 2 + K$, and that family is exponent independent.
- Residuals are $\mathcal{O}(\ln N/N^2)$. The e-fold channel is the only unprotected route, at 22% and opposite sign.
- Matching bounds the shift below 3.7×10^{-4}, a factor of 25 short of ACT.

The no-go

If the tension between Starobinsky and ACT is confirmed, quantum-gravitational running within this class is unlikely to resolve it. A real deviation would point to tree-level structures or to the matter sector.

Open directions, if you are looking for a project

- **The squeezed limit.** A removable deformation is an amplitude rescaling plus a field shift, so the single-field consistency relation suggests the same protection extends to the bispectrum. Proving it would lift the criterion from two-point to three-point functions.
- **Fully coupled matter.** The interplay of removable and non-removable corrections once matter is included properly.
- **The reheating channel.** The least suppressed physical channel here, and the one that sets N_*.
- **Background-induced states.** A complete de Sitter treatment of the dressed propagator.

Removability, not magnitude.

Backup: pointers into the literature

Backup: heat-kernel conventions

The literature form is usually quoted for $D = -(\nabla^2 + E')$. The substitution $E' = -E$ reverses the sign of the cross term,

$$\left.\frac{180E'^2 + 60RE' + 5R^2}{360}\right|_{E'=-E} = \frac{180E^2 - 60RE + 5R^2}{360} = \frac{1}{2}\left(E - \frac{R}{6}\right)^2 .$$

Omitting that term, or keeping it without reversing its sign, destroys the $(\xi - 1/6)$ structure and with it the conformal fixed point.

Backup: the operator basis and the sign reversal

The pairs (ω, O) entering $\Gamma \supset + \int \sqrt{-g}\, \omega O$,

$$\left(\tfrac{1}{2} M_{\mathrm{P}}^2, R\right), \ \left(\alpha_C, C^2\right), \ \left(\alpha_R, R^2\right), \ \left(\alpha_{\mathrm{GB}}, \mathcal{G}_{\mathrm{GB}}\right), \ \left(-\rho_\Lambda, 1\right), \ \left(-V, 1\right), \ \left(-\tfrac{1}{2}\xi, R\phi^2\right).$$

The final three carry an explicit minus sign and therefore reverse $\mu\, d\omega/d\mu = -X$. This is the origin of $\beta_V = +X_V$, confirmed against the Coleman–Weinberg potential.

Backup: why a field shift never matters, for any δ

Let $V_\delta(\phi) \equiv V_{\rm St}(\phi + \delta)$ and put $\tilde{\phi} = \phi + \delta$. Because δ is a constant, the kinetic term is untouched, $(\partial\phi)^2 = (\partial\tilde{\phi})^2$, so

$$S\left[\phi; V_\delta\right] = S\left[\tilde{\phi}; V_{\rm St}\right].$$

It is the same theory with a relabelled variable, not a new model. Everything follows at once,

$$\epsilon(\phi) = \epsilon_{\rm St}(\tilde{\phi}), \qquad \phi_{\rm end} \to \phi_{\rm end} - \delta, \qquad N(\phi) = N_{\rm St}(\tilde{\phi}).$$

The one line that does the work

The observables are read at fixed N, not at fixed ϕ. Since N is counted from the end of inflation, the shift cancels between the endpoint and the evaluation point. Therefore $\epsilon(N)$ is identical and so are n_s and r, exactly, for any δ.

The anchor is the end of inflation, not the origin of field space.

Backup: the demonstration, and the trap

Exact solver at $N_* = 55$.

δ/M_P	ϕ_end	ϕ_*	$\phi_* + \delta$	n_s
-20	20.9402	25.3529	5.3529	0.964977222226
-1	1.9402	6.3529	5.3529	0.964977222226
0	0.9402	5.3529	5.3529	0.964977222226
$+1$	-0.0598	4.3529	5.3529	0.964977222226
$+20$	-19.0598	-14.6471	5.3529	0.964977222226
$+50$	-49.0598	-44.6471	5.3529	0.964977222226

$|\Delta n_s| \leq 1.1 \times 10^{-16}$ and $|\Delta r/r| \leq 5.7 \times 10^{-15}$ throughout. Machine precision.

The trap

Compare the two potentials at a *fixed field value* instead, and $\delta = 0.6\, M_\mathrm{P}$ appears to give $\Delta n_s = +1.4 \times 10^{-2}$, larger than the ACT shift. The entire effect is an artefact of the wrong anchor.

Backup: the field-shift direction at finite amplitude

A finite shift $V_{\rm St}(\phi+\delta)$ is exactly invisible for every δ, verified over $-50 \le \delta/M_{\rm P} \le +20$ to machine precision. The criterion evaluated on it,

$$\mathcal{R}_{\rm shift} = -\frac{(1-s)(1-su)}{(1-u)^2}, \qquad s = e^{-a\delta} \quad \xrightarrow[\delta\to 0]{} \quad \delta\,\mathcal{G}.$$

- $\delta > 0$: variation saturates near 1%, any size passes
- $\delta < 0$: degrades near $|\delta| \simeq 3.9\,M_{\rm P}$, and fails at $|\delta| = \phi_*$, where the pivot reaches the minimum

The deformation itself is already 50% there. The criterion degrades only where the shift stops being a small correction.

Backup: is this consistent with effective field theory?

The worry

If the size of a correction does not decide the effect, has the effective-field-theory expansion been abandoned?

No. The claim is compressed. What we actually have is

$$\Delta n_s = \mathcal{O}\left(\frac{\mathcal{R}[\Delta]}{N}\right), \qquad \mathcal{R}[\Delta] = \frac{V_{\mathrm{St}}}{V'_{\mathrm{St}}}\Delta',$$

which is strictly linear in the coupling. Double the coupling, double the shift. Verified over several decades. The algebraic form fixes the coefficient, not the linearity.

Backup: where the enhancement comes from

Per unit coupling, at $N_* = 55$,

channel	Δn_s per unit amplitude
the flow, $K\ell$	-5.1×10^{-4}
R^3	-1.4×10^{2}

A ratio of 2.7×10^5 at equal coupling. Its origin is kinematic,

$$\frac{V'_{\mathrm{St}}}{V_{\mathrm{St}}} = \frac{2au}{1-u} = 0.021 \implies \frac{V_{\mathrm{St}}}{V'_{\mathrm{St}}} = 47.8 \simeq \mathcal{O}(N).$$

Reading

The plateau is flat, so $V_{\mathrm{St}}/V'_{\mathrm{St}}$ is large. A deformation can be tiny as a *function* and still have a large $\mathcal{R}$, because the observables read its derivative structure. For R^3 at $c = 6.7 \times 10^{-5}$ the potential moves by half a percent, $|\Delta| = 5 \times 10^{-3}$, yet $\mathcal{R} = 0.21$.

Backup: removable directions are redundant directions

In an effective field theory, operators removable by a field redefinition do not affect the S-matrix, at any value of their coefficients. Nobody calls that a violation. It is how one builds an operator basis. Here the map $V \mapsto (n_s, r)$ through $\epsilon(N)$ has two redundancies.

- **Amplitude.** A_s is a fitted input, not a prediction. Rescaling V is reabsorbed, exactly as refitting α_s leaves ratios alone.
- **Field origin.** $\phi \to \phi + \delta$ is a field redefinition, and N is counted from the end of inflation.

And the protection is not infinite

$$\Delta n_s^{\mathrm{shape}} = \frac{K}{N^2}\left[1.458 - \frac{3}{4}\ln N\right] \neq 0.$$

Suppression by two powers of $1/N$, proportional to K. If it were exact to all orders, *that* would be worrying.

Backup: two general lemmas

Powers of the curvature

Using $\Delta = -6M_S^2\delta f(R)/R^2$ with $R = 3M_S^2(1-u)/u$,

$$\mathcal{R}\left[c_n R^n\right] = -\frac{(n-2)c_n}{2}\,\frac{(1-u)^{n-2}}{u^{n-1}} \simeq -\frac{(n-2)c_n}{2}\left(\frac{4N}{3}\right)^{n-1}.$$

The prefactor vanishes at $n = 2$. That single zero is why R^2 is special.

Powers of the canonical field

$$\mathcal{R}\left[c\,(a\phi)^k\right] = \frac{ck(1-u)(a\phi)^{k-1}}{2u}.$$

At $k = 1$ this returns the field-linear part of the marginal deformation, at $k = 2$ the polynomial entry.

Backup: a subtlety in the table

The two vacuum-energy channels are listed as not removable, yet their $\mathcal{R}$ varies by only 1.8% and 0.9% across the window, close to the removable entries. No contradiction. Their $\mathcal{R}$ has the form

$$\mathcal{R} = \text{const} + \mathcal{O}(u).$$

The constant part does nothing. The $\mathcal{O}(u)$ part produces $\Delta n_s = \mathcal{O}(u/N) = \mathcal{O}(1/N^2)$, which is exactly the residual order. That is why they sit in the upper block and give shifts of the same size as the flow.

The honest statement

The criterion separates by the scaling with N, not by a binary constant or not. Three behaviours exist: exactly constant, constant up to $\mathcal{O}(u)$, and growing as a power of N. The table compresses them into two columns.

Backup: numerical procedure

- Exact slow-roll solver. ϕ_{end} from $\epsilon_V = 1$, then ϕ_* from the e-fold integral, iterated on the pivot.
- Python reference implementation, `sympy` for the symbolic work and `scipy` `brentq` and `quad` for the numerics.
- Mathematica notebooks reproduce it and self-test against reference values exported to JSON.
- Every equation checked symbolically. Every quoted number recomputed independently.

Backup: the no-ghost window