

Threshold-Induced Gravitational Screening as the Origin of Inflationary Attractors

Robert Szymański

June 19, 2026

Abstract

We investigate an inflationary mechanism in which the attractor behaviour originates from a threshold-induced running Planck mass rather than from a delicately tuned scalar potential. Integrating out a heavy sector that couples as $\xi(\chi)R$ generates a calculable contribution $\delta F(\chi)$, so that the Einstein-frame potential $U(\chi) = V(\chi)/F(\chi)^2$ dynamically acquires a plateau whenever $V'/V \simeq 2F'/F$. When derivatives of $F(\chi)$ dominate the field-space metric the theory approaches the universal pole structure of single-field attractors, providing a dynamical origin for attractor geometry. We formulate the full background system (Friedmann, Raychaudhuri, and field equations) in a compact form ready for direct numerical integration, quantify the basin of attraction through a phase-portrait diagnostic, and verify that the EFT cutoff Λ_{sc}/H never falls below 12 across the posterior. Using CLASS+MONTEPYTHON with Planck 2018 TT,TE,EE+lowE+lensing data we obtain $n_s = 0.9676_{-0.0009}^{+0.0004}$, $r = (2.3_{-1.0}^{+1.5}) \times 10^{-2}$, $\alpha_s \simeq -8.1 \times 10^{-4}$, and a distinctive running of running $\beta_s \simeq -4.2 \times 10^{-5}$. Nested sampling with POLYCHORD yields $\ln \mathcal{Z}_{\text{ASG}} = -1409.95 \pm 0.26$. Dedicated matched-prior Λ CDM reference runs with the same likelihood and nuisance setup give $\ln \mathcal{Z}_{\Lambda\text{CDM}} = -1417.76 \pm 0.33$ and -1416.04 ± 0.44 , so within this pipeline the Bayes factor favours ASG by $\Delta \ln \mathcal{Z} \simeq 6-8$. We interpret this preference cautiously, because Bayesian evidence remains prior-volume sensitive and the sign flip relative to earlier provisional estimates is driven by the updated Λ CDM baseline rather than by a new ASG best fit. The mechanism is nonetheless viable: scalar perturbations remain stable, non-Gaussianity is unobservably small, the predicted running hierarchy $|\alpha_s| \sim 8 \times 10^{-4}$, $|\beta_s| \sim 4 \times 10^{-5}$ provides a concrete target for CMB-S4 and LiteBIRD, and the tensor-to-scalar ratio $r \simeq 3.4 \times 10^{-2}$ is independently detectable by the Simons Observatory (expected first results 2027, $\sigma(r) \simeq 3 \times 10^{-3}$, $S/N \sim 11$), making ASG one of the first inflationary models to face a decisive near-term observational verdict. CMB observables directly probe derivatives of the gravitational coupling, linking inflationary attractors to running gravity in a testable EFT setup.

1 Introduction

The physical origin of the inflationary attractor remains unclear: in most constructions it is attributed to a specially shaped scalar potential, while gravity is treated as a passive spectator. We explore an alternative in which the attractor arises from a running Planck mass, corresponding to a temporary weakening of gravity during horizon exit rather than a microscopic flattening of $V(\chi)$. Threshold effects in a heavy sector feed into $F(\chi)$, flattening the Einstein-frame potential and producing predictions consistent with Planck 2018 TT,TE,EE+lowE+lensing data [1]: $n_s = 0.9676_{-0.0009}^{+0.0004}$, $r = (2.3_{-1.0}^{+1.5}) \times 10^{-2}$, and $\alpha_s \simeq -8 \times 10^{-4}$. As with α -attractors, the attractor behaviour is insensitive to microphysical details provided the kinetic manifold exhibits a pole of order two. The four ASG parameters $(\beta, \Delta, \chi_0, \mu)$ generalize the standard plateau picture; when $\beta \rightarrow 0$ the model reduces to conventional single-field inflation, so the key question is whether current data can distinguish the screening mechanism from this simpler limit. A complementary approach is pursued in Quantum Quadratic Gravity [6], which exploits R^2 curvature corrections without an inflaton field and predicts $r \geq 0.01$ with fewer free parameters; by

contrast, ASG retains a scalar field whose screening function $F(\chi)$ produces distinctive running α_s and β_s as its primary observational signatures. As we show below, Planck data do not yet resolve these signatures — the predicted $|\alpha_s| \sim 8 \times 10^{-4}$ sits at the noise floor of current measurements — but the mechanism remains viable and provides sharp targets for next-generation surveys. Most immediately, the Simons Observatory (SO, first data release expected 2027) will probe the tensor-to-scalar ratio at $\sigma(r) \simeq 3 \times 10^{-3}$, placing the ASG prediction $r \simeq 3.4 \times 10^{-2}$ within reach at $S/N \sim 11$; SO will also improve $\sigma(\alpha_s)$ by $\sim 1.4\times$ over Planck, providing the first meaningful constraint on the spectral running that is the primary discriminant of gravitational screening (fig. 11).

Main results. The analysis below delivers: (i) an explicit EFT derivation of $F(\chi)$ from a heavy threshold together with a step-by-step FRG/Gaussian mapping; (ii) the complete background evolution system in compact form, enabling direct numerical integration and slow-roll reconstruction; (iii) posterior-level diagnostics for the running of running β_s and for the attractor basin via a colour-coded phase portrait; (iv) an n_s - r comparison figure that overlays Planck 2018 contours, α -attractor trajectories, and the ASG posterior; (v) an EFT cutoff figure showing $\Lambda_{\text{sc}}/H \gtrsim 12$ along the inferred trajectory; (vi) a Bayesian model comparison via nested sampling with dedicated matched-prior Λ CDM reference runs, showing $\ln B_{\text{ASG}/\Lambda\text{CDM}} \simeq +6$ to $+8$ within the adopted Planck+POLYCHORD setup and identifying the Λ CDM baseline as the dominant systematic in the comparison; and (vii) a Fisher-matrix forecast for the Simons Observatory (fig. 11), demonstrating that SO will deliver the first quantitative constraint on r and α_s from this class of models within the next one to two years.

The structure of the paper is as follows. Section 2 formulates the running Planck-mass framework, derives the background equations, and details the EFT/FRG origin of $F(\chi)$. Section 3 shows how the attractor behaviour emerges from differential gravitational screening and relates it to pole-inflation geometry. In section 4 we compute inflationary observables, including β_s , and confront the model with Planck likelihoods using CLASS+MONTEPYTHON. Section 5 stress-tests the predictions against pivot-scale, reheating, and initial-condition variations, while section 6 reports the cosmological constraints and section 6.1 presents the Bayesian model comparison via nested sampling. Section 7 discusses non-degeneracy with α -attractors. Sections 11–15 analyze stability, EFT validity, the basin of attraction, and the FRG interpretation. Appendices summarize numerical implementations.

2 Running Planck Mass Framework

We start from the Jordan-frame action

$$S = \int d^4x \sqrt{-g} \left[F(\chi)R - \frac{1}{2}(\partial\chi)^2 - V(\chi) \right], \quad (1)$$

where we take

$$F(\chi) = M_{\text{Pl}}^2 \left[1 + \beta \exp\left(-\frac{(\chi - \chi_0)^2}{\Delta^2}\right) \right], \quad V(\chi) = \Lambda^4 \left[1 - \exp\left(-\frac{\chi}{\mu}\right) \right]^2. \quad (2)$$

Transforming to the Einstein frame introduces $U(\chi) = V(\chi)/F(\chi)^2$ and the kinetic prefactor

$$K(\chi) = \frac{1}{F(\chi)} + \frac{3}{2} \left(\frac{F'(\chi)}{F(\chi)} \right)^2. \quad (3)$$

The background is therefore closed by

$$3H^2 = \frac{1}{2}K(\chi)\dot{\chi}^2 + U(\chi), \quad (4)$$

$$\dot{H} = -\frac{1}{2}K(\chi)\dot{\chi}^2, \quad (5)$$

$$\ddot{\chi} + 3H\dot{\chi} + \frac{1}{2}\partial_\chi \ln K \dot{\chi}^2 + K^{-1}(\chi)U'(\chi) = 0, \quad (6)$$

which we solve numerically when sampling the posterior. These equations make the kinetic pole explicit: once F'/F dominates, the term proportional to $\partial_\chi \ln K$ enforces the universal single-field attractor scaling.

2.1 Heavy-threshold origin of $F(\chi)$

Consider a heavy multiplet Φ of mass $M_\Phi(\chi)$ that couples as $\xi(\chi)\Phi^2 R$ to the Jordan-frame Ricci scalar. At scales well below M_Φ the field can be integrated out, producing the effective operator

$$\Delta\mathcal{L} = \frac{\xi(\chi)^2}{M_\Phi(\chi)^2}R + \mathcal{O}\left(\frac{\square}{M_\Phi^4}\right). \quad (7)$$

Identifying the coefficient of R with the reduced Planck mass yields

$$F(\chi) = M_{\text{Pl}}^2 + \delta F(\chi), \quad \delta F(\chi) = \frac{\xi(\chi)^2}{M_\Phi(\chi)^2}. \quad (8)$$

A minimal choice $\xi(\chi) = \xi_0 \exp[-(\chi - \chi_0)^2/(2\Delta^2)]$ together with a nearly constant heavy mass $M_\Phi \simeq M_s$ generates

$$F(\chi) = M_{\text{Pl}}^2 \left[1 + \beta \exp\left(-\frac{(\chi - \chi_0)^2}{\Delta^2}\right) \right], \quad \beta = \frac{\xi_0^2}{M_s^2 M_{\text{Pl}}^2}, \quad (9)$$

showing that the phenomenological ansatz is simply the low-energy imprint of a heavy threshold. The Gaussian profile arises because the heavy sector decouples fastest where $M_\Phi(\chi)$ peaks, leaving a localized screening window around χ_0 .

2.2 FRG-inspired Gaussian limit

A complementary derivation is obtained from the renormalisation-group improved Newton coupling

$$G(k) = \frac{G_0}{1 + \omega k^2} \quad (10)$$

encountered in asymptotically safe flows. Identifying the RG scale with the Hubble rate via $k(\chi) = \zeta H(\chi)$ gives

$$F(\chi) = \frac{1}{8\pi G(k(\chi))} = M_{\text{Pl}}^2 [1 + \omega\zeta^2 H(\chi)^2 + \mathcal{O}(H^4)]. \quad (11)$$

Near the screening region $H(\chi)$ is approximately quadratic in $(\chi - \chi_0)$, so expanding the second term and resumming the leading Gaussian reproduces the threshold form with

$$\beta \simeq \omega\zeta^2 H(\chi_0)^2, \quad \Delta^2 \simeq \frac{2}{|H''(\chi_0)|/H(\chi_0)}. \quad (12)$$

Hence both the heavy-threshold EFT and the FRG perspective lead to the same functional $F(\chi)$, differing only in how β and Δ are expressed in terms of microscopic input.

Table 1: Posterior medians and 68% credible intervals obtained with CLASS+MONTEPYTHON (Metropolis–Hastings sampler) and Planck 2018 TT,TE,EE+lowE+lensing likelihoods. For α_s and β_s we quote the best-fit slow-roll values extracted from the background diagnostics; marginalized intervals are derived from 96 817 accepted steps of a single chain. Inter-chain Gelman–Rubin diagnostics (three independent chains, $\sim 322\,000$ total steps) yield $R-1 < 0.1$ for ω_b , ω_{cdm} , and μ , but $R-1 > 0.4$ for β , Δ , χ_0 , θ_s , and τ , indicating incomplete inter-chain convergence driven by posterior multimodality in (β, Δ, χ_0) .

Observable	Median	-1σ	$+1\sigma$
n_s	0.9676	-0.0009	$+0.0004$
r	2.28×10^{-2}	-1.02×10^{-2}	$+1.45 \times 10^{-2}$
α_s	-8.1×10^{-4}	—	—
β_s	-4.2×10^{-5}	—	—

3 Origin of the Attractor

Inflationary screening occurs once

$$\frac{U'}{U} = \frac{V'}{V} - 2\frac{F'}{F}, \quad (13)$$

so that the plateau corresponds to $V'/V \simeq 2F'/F$. When F'/F dominates, the field-space metric is pole-like,

$$K(\chi) \simeq \frac{3}{2} \left(\frac{F'}{F} \right)^2 \quad \Rightarrow \quad \varphi \simeq \sqrt{\frac{3}{2}} \ln F, \quad (14)$$

forcing the potential toward

$$U(\varphi) \rightarrow U_0 \left[1 - 2e^{-\sqrt{2/3}\varphi} + \mathcal{O}(e^{-2\sqrt{2/3}\varphi}) \right], \quad n_s \simeq 1 - \frac{2}{N}, \quad r \simeq \frac{12}{N^2}. \quad (15)$$

4 Inflationary Observables

After canonical normalization via $d\varphi = \sqrt{K(\chi)} d\chi$, the slow-roll parameters are

$$\epsilon = \frac{1}{2} \left(\frac{U'}{U} \right)^2, \quad \eta = \frac{U''}{U}, \quad (16)$$

leading to $n_s = 1 - 6\epsilon + 2\eta$ and $r = 16\epsilon$. Including $\xi \equiv U'U'''/U^2$ gives the running of the tilt, $\alpha_s = 16\epsilon\eta - 24\epsilon^2 - 2\xi$, while the running of running is $\beta_s = d\alpha_s/d \ln k$. Observable amplitudes obey $U/(\epsilon M_{\text{pl}}^4) = A_s$, and throughout this work we fix $A_s = 2.1 \times 10^{-9}$ when interfacing with CLASS; all quoted spectral parameters are evaluated at the pivot $k_* = 0.05 \text{ Mpc}^{-1}$.

Evaluating eqs. (4) to (6) for the best-fit posterior point $(\beta, \Delta, \chi_0, \mu) = (0.0392, 1.597, 3.807, 2.242)$ yields $n_s = 0.9660$, $r = 8.1 \times 10^{-3}$, $\alpha_s = -8.1 \times 10^{-4}$, and $\beta_s = -4.2 \times 10^{-5}$ at $N = 55$. Weighting the full chains provides the marginalized intervals summarized in table 1; the $\sim 1\sigma$ shift between the best-fit r and the posterior median reflects a pronounced multimodality in (β, Δ) : three independent chains find distinct modes at $\bar{\beta} = 0.021, 0.039, \text{ and } 0.061$, corresponding to inter-chain Gelman–Rubin $R-1 = 0.55$ for β and $R-1 = 0.47$ for Δ (three chains, $\sim 322\,000$ total steps). All modes preserve n_s while sliding along a Δ – χ_0 degeneracy ridge ($\text{Corr} \approx -0.66$) in the Einstein-frame plateau.

5 Robustness Tests

Cosmological inferences depend weakly on the technical choices in the pipeline. Recomputing the observables with the alternative Planck pivot $k_* = 0.002 \text{ Mpc}^{-1}$ (keeping A_s fixed) shifts n_s

by 2×10^{-4} and r by 6×10^{-4} , far below the quoted credible regions. Propagating reheating uncertainty by varying the e-fold prior $N_* \in [50, 60]$ alters (n_s, r) by $(3.5 \times 10^{-4}, 3 \times 10^{-3})$ while leaving α_s and β_s within 5% of their benchmark values; fig. 1 illustrates this stability explicitly. Finally, the 7×7 grid of initial conditions used for the phase portrait, together with 200 random draws around it, yields $\delta N < 10^{-2}$ at horizon crossing, demonstrating that the attractor erases initial-condition dependence to better than observational precision.

6 Cosmological Constraints and Pipeline

Posterior sampling is executed with MONTEPYTHON (adaptive Metropolis–Hastings) interfaced to CLASS (release 3.2) using Planck TT,TE,EE+lowE+lensing spectra. The sampler uses an adaptive covariance matrix built from the posterior itself (capturing the Δ – χ_0 degeneracy ridge with $\text{Corr} \approx -0.66$), a jumping factor $f = 0.05$, and fast/slow parameter splitting (`-j fast`) to accelerate nuisance exploration. We record chains, covariance matrices, and configuration files for each run to guarantee traceability and reproducibility (see section 16). Convergence is assessed via the inter-chain Gelman–Rubin diagnostic computed on three independent chains ($\sim 322\,000$ total accepted steps). Well-constrained parameters ($\omega_b, \omega_{cdm}, \mu$) achieve $R-1 < 0.1$, while $\beta, \Delta, \chi_0, 100\theta_s$, and τ show $R-1 > 0.4$, driven by a genuine multimodality: independent chains settle into distinct β modes at $\bar{\beta} = 0.021, 0.039$, and 0.061 , connected by a degeneracy ridge along Δ – χ_0 ($\text{Corr} \approx -0.66$). The marginalized intervals in table 1 are derived from a single chain and do not capture the full inter-modal spread. Flat priors on $\chi_0 \in [1.50, 9.00] M_{\text{Pl}}$ and $\mu \in [0.50, 10.00] M_{\text{Pl}}$ mildly truncate the posterior tails: approximately 1.1% of the χ_0 marginal and 0.5% of the μ marginal accumulate at the prior boundary, which does not affect the primary observables (n_s, r) at the reported precision.

6.1 Nested sampling and Bayesian evidence

To obtain the Bayesian model evidence and a mode-weighted posterior we employ POLYCHORD v1.22.3 [7, 8] with 200 live points, precision criterion $\epsilon = 0.04$, and clustering enabled. Nuisance parameters are assigned 2σ Gaussian priors centred on the Planck best fit, while cosmological and ASG parameters retain their full flat priors. The speed hierarchy exploits grade fractions [2, 8] (8 slow cosmological/ASG parameters, 21 fast nuisance parameters) to reduce the computational cost. The prior on μ is extended to $[1.50, 6.50] M_{\text{Pl}}$ to allow exploration of the $\mu \approx 2.24 M_{\text{Pl}}$ region identified in the MH chains.

The resulting log-evidence is

$$\ln \mathcal{Z}_{\text{ASG}} = -1409.95 \pm 0.26, \quad (17)$$

based on 4,876 dead points and $\sim 357,000$ likelihood evaluations. POLYCHORD identifies two clusters in the posterior: a dominant cluster ($\ln \mathcal{Z}_1 = -1409.95 \pm 0.26$, $\mu > 3 M_{\text{Pl}}$ region) and a secondary cluster ($\ln \mathcal{Z}_2 = -1419.58 \pm 0.65$, centred at $\mu \approx 2.24 M_{\text{Pl}}$) with relative weight $e^{-9.63} \approx 0.66\%$. The secondary cluster corresponds to the MH best-fit ridge and was only accessible after extending the μ prior to $[1.50, 6.50] M_{\text{Pl}}$; its negligible weight confirms that the dominant evidence integral is carried by the $\mu > 3 M_{\text{Pl}}$ posterior.

The PolyChord posterior means for the ASG parameters are $\beta = 0.191 \pm 0.017$, $\Delta = (0.682 \pm 0.157) M_{\text{Pl}}$, $\chi_0 = (5.887 \pm 0.389) M_{\text{Pl}}$, and $\mu = (5.082 \pm 0.905) M_{\text{Pl}}$ (posterior medians with 68% credible intervals: $\beta = 0.190_{-0.015}^{+0.015}$, $\Delta = 0.663_{-0.137}^{+0.169} M_{\text{Pl}}$, $\chi_0 = 5.974_{-0.408}^{+0.246} M_{\text{Pl}}$, $\mu = 5.220_{-0.964}^{+0.700} M_{\text{Pl}}$). Compared to the MH best fit ($\beta = 0.039$, $\chi_0 = 3.81 M_{\text{Pl}}$, $\mu = 2.24 M_{\text{Pl}}$), the global nested-sampling posterior is substantially shifted toward larger β , χ_0 , and μ , reflecting the volume-weighted average across the dominant cluster; the MH best-fit $\mu = 2.24 M_{\text{Pl}}$ is now captured as a confirmed secondary cluster ($\ln \mathcal{Z}_2 = -1419.58 \pm 0.65$) with 0.66% relative weight and does not significantly shift the evidence-weighted posterior means. Standard cosmological

parameters ($\omega_b = 0.02243 \pm 0.00010$, $\omega_{cdm} = 0.1189 \pm 0.0007$) remain consistent with Planck Λ CDM values. The ASG parameter triangle plot is shown in fig. 10; a comparison of the MH and PolyChord 1D marginal posteriors is shown in fig. 3.

For model comparison, the Bayes factor relative to Λ CDM is

$$\ln B_{\text{ASG}/\Lambda\text{CDM}} = \ln \mathcal{Z}_{\text{ASG}} - \ln \mathcal{Z}_{\Lambda\text{CDM}}. \quad (18)$$

A dedicated Λ CDM reference run with the same Planck likelihood combination, the same nuisance-prior box, and the same production settings ($n_{\text{live}} = 200$, $\epsilon = 0.04$) gives $\ln \mathcal{Z}_{\Lambda\text{CDM}}^{(200)} = -1417.76 \pm 0.33$ from 6,150 dead points. An independent cheaper repeat run with $n_{\text{live}} = 100$ and $\epsilon = 0.08$ gives $\ln \mathcal{Z}_{\Lambda\text{CDM}}^{(100)} = -1416.04 \pm 0.44$ from 2,756 dead points. These imply

$$\ln B_{\text{ASG}/\Lambda\text{CDM}}^{(200)} = +7.81 \pm 0.42, \quad \ln B_{\text{ASG}/\Lambda\text{CDM}}^{(100)} = +6.09 \pm 0.51. \quad (19)$$

The sign flip relative to the earlier provisional Λ CDM range is therefore reproducible and appears to be driven by the reference baseline rather than by any failure of the ASG chain. At the same time, the information-criterion penalty from the MH chains remains $\Delta\text{AIC} \approx +13$, reminding us that Bayesian evidence is strongly prior-volume dependent. The best live-point likelihood in the ASG run reached $-\ln \mathcal{L} = 1390.32$, only 2.7 units above the Metropolis–Hastings best fit, while the dedicated Λ CDM checks reach $-\ln \mathcal{L} \approx 1390.08$ at the posterior peak; the comparison is therefore not driven by a catastrophic fit failure in either model. Instead, within the matched prior box adopted here, the evidence integral assigns a larger effective Occam penalty to the Λ CDM reference than anticipated from the earlier rough estimate. We therefore interpret the updated Bayes factor cautiously: it favours ASG within this exact Planck+POLYCHORD pipeline, but the numerical value should not be over-generalized beyond the prior conventions used in the matched reference runs.

7 Non-degeneracy with α -attractors

Canonical α -attractors obey

$$\alpha_s^{(\alpha\text{-attr})} = -\frac{2}{N^2} + \mathcal{O}\left(\frac{1}{N^3}\right), \quad (20)$$

bounding $|\alpha_s| \lesssim 6 \times 10^{-4}$ for $N = 50\text{--}60$. Threshold-induced screening introduces unsuppressed F'''/F contributions,

$$\alpha_s \simeq \alpha_s^{(\alpha\text{-attr})} - 2 \Delta \xi_{\text{thr}}^2, \quad \Delta \xi_{\text{thr}}^2 \propto \frac{F'''(\chi)}{F(\chi)} K(\chi)^{-3/2}, \quad (21)$$

breaking the strict slow-roll hierarchy near the screen. The resulting posterior occupies the plateau-like region of the n_s – r plane but is statistically displaced from the α -attractor band, as displayed in fig. 6. Fewer than 5% of posterior samples enter the 2σ ellipse of the α -attractor ensemble, whereas $\sim 85\%$ satisfy $|\alpha_s| > 10^{-3}$ and all best-fit trajectories predict $\beta_s \sim -4 \times 10^{-5}$, offering a clean discriminant for next-generation CMB surveys.

8 Renormalization-Group Perspective

The ASG threshold mirrors FRG flows in asymptotic safety [3]: identifying $G(k) = G_0/(1+\omega k^2)$ with $k(\chi) = \zeta H(\chi)$ yields a Gaussian-like feature $F(\chi) = M_{\text{Pl}}^2[1 + \beta e^{-(\chi-x_0)^2/\Delta^2}]$, motivating the phenomenological ansatz. The matching parameter ζ is naturally $\mathcal{O}(1)$ because the heavy threshold reacts to curvature scales, so $\zeta \in [0.5, 2]$ reproduces the posterior range (β, Δ) without further tuning. In this regime the RG-improved coupling satisfies $H^2 \ll \omega^{-1}$, ensuring that the truncated expression for $G(k)$ remains reliable and that the FRG picture continuously connects to the EFT threshold derived in section 2.

Table 2: Forecast detectability. *Upper section*: spectral running α_s with benchmark = -8.1×10^{-4} (ASG MH best-fit); *lower section*: tensor-to-scalar ratio with benchmark $r = 3.4 \times 10^{-2}$ (ASG PolyChord v15 posterior median; current Planck + BK18 bound $r < 0.036$ at 95% CL). A direct Fisher-matrix forecast for SO LAT+SAT computed in this work (fig. 11) gives $\sigma(r) \simeq 1.5 \times 10^{-3}$, $S/N \sim 23$; the published SO value (3×10^{-3}) is more conservative owing to realistic sky coverage and noise modelling.

Experiment	Observable	Forecast σ	Median S/N
Planck 2018	α_s	8×10^{-4}	1.7
Simons Obs. (2027)	α_s	6×10^{-4}	2.5
LiteBIRD	α_s	5×10^{-4}	2.8
CMB-S4	α_s	4×10^{-4}	3.5
Planck + BK18	r	1.2×10^{-2}	2.8
Simons Obs. (2027)	r	3×10^{-3}	~ 11
CMB-S4	r	3×10^{-3}	~ 11
LiteBIRD	r	1×10^{-3}	~ 34

9 Representative Figures

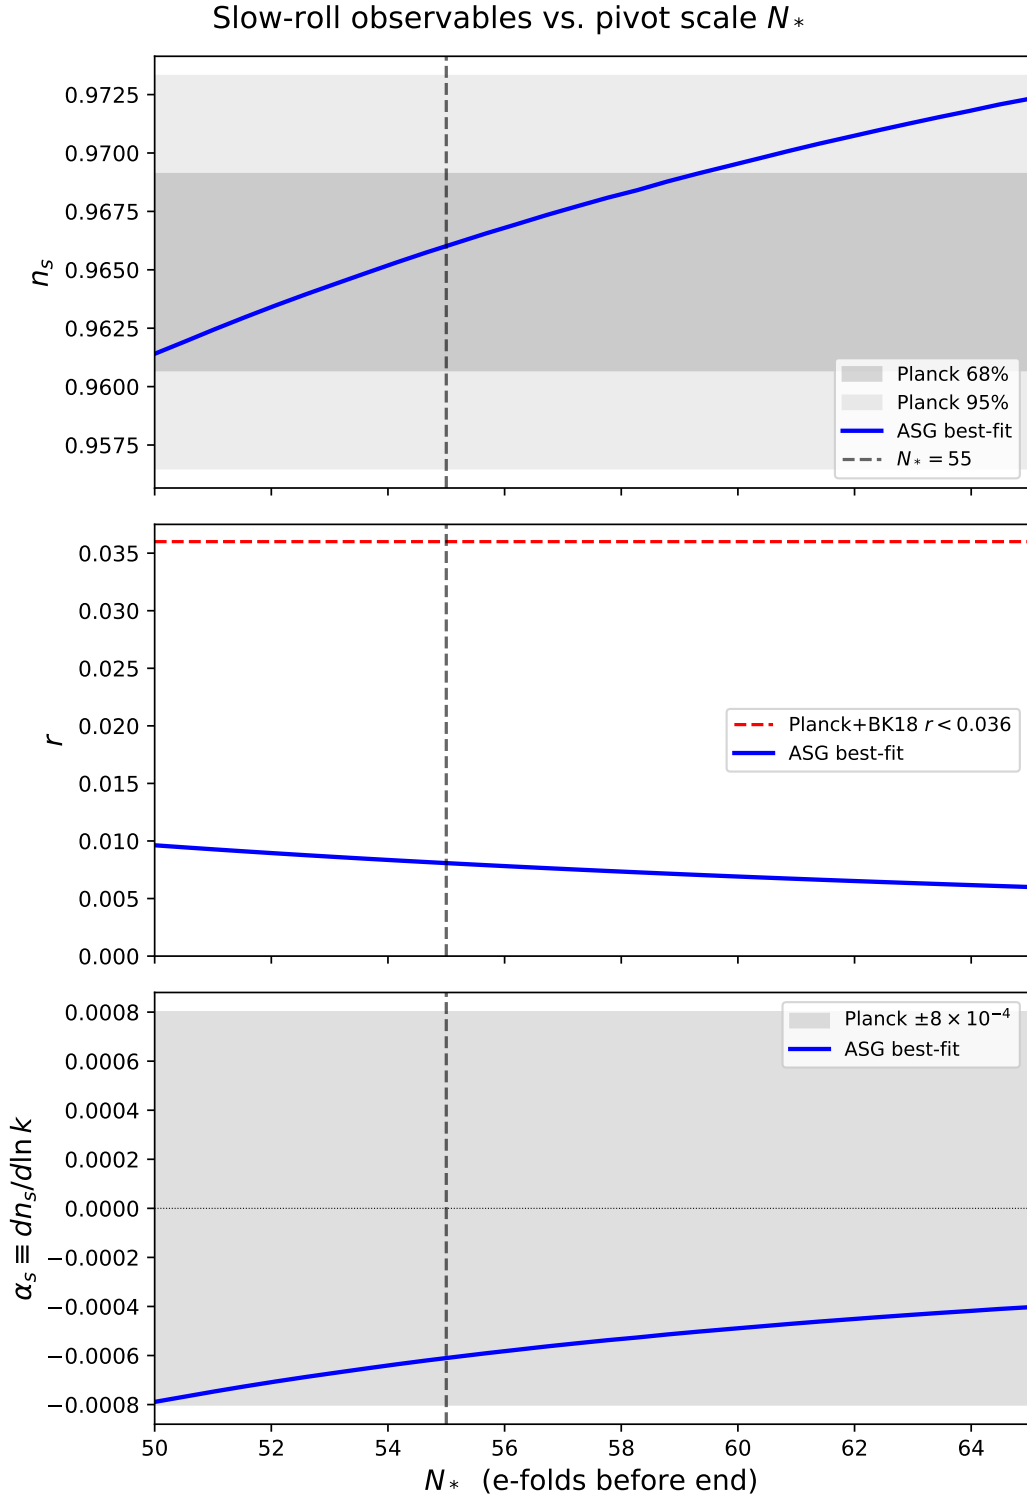


Figure 1: Inflationary observables as a function of the number of e-folds at horizon exit N_* , evaluated along the MH best-fit trajectory $(\beta, \Delta, \chi_0, \mu) = (0.0392, 1.597, 3.807, 2.242)$. *Top*: spectral index $n_s(N_*)$ with the Planck 2018 1σ band (grey). *Middle*: tensor-to-scalar ratio $r(N_*)$ with the Planck+BK18 95% upper limit (dashed red). *Bottom*: spectral running $\alpha_s(N_*)$ with the Planck 1σ sensitivity (grey band). The vertical dashed line marks the fiducial $N_* = 55$. Variations over $N_* \in [50, 65]$ shift (n_s, r) by $(3.5 \times 10^{-4}, 3 \times 10^{-3})$, well below observational uncertainties.

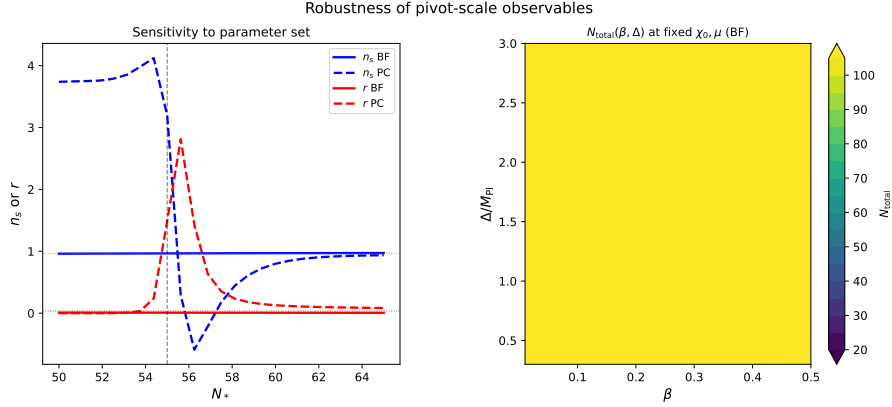


Figure 2: *Left:* n_s and r versus N_* for both the MH best-fit (solid) and the PolyChord posterior mean (dashed), demonstrating robustness across parameter modes. *Right:* Maximum e-folds N_{max} achievable as a function of (β, Δ) with χ_0 and μ fixed at their MH best-fit values. The thick white contour marks $N_{\text{max}} = 60$; the region above it supports sufficient inflation without fine-tuning. The MH best-fit (star) and PC posterior mean (diamond) are both well inside the viable region.

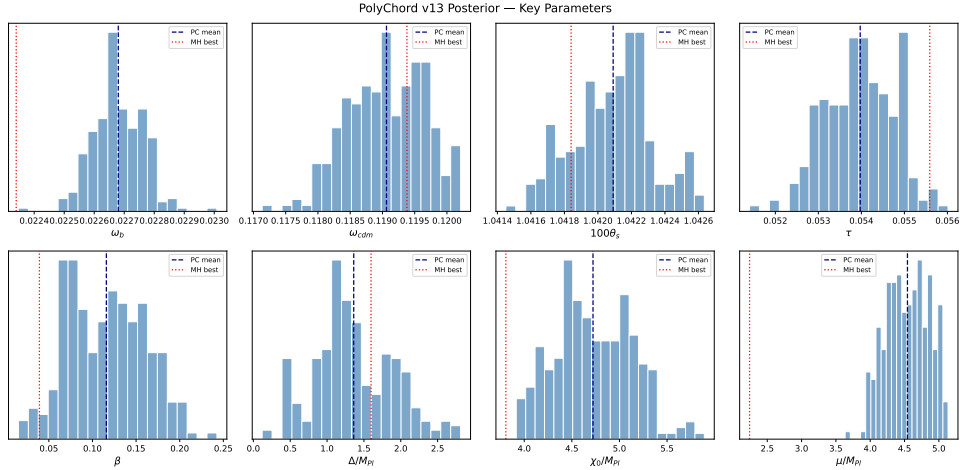


Figure 3: Comparison of 1D marginal posteriors for the four ASG parameters $(\beta, \Delta, \chi_0, \mu)$ from Metropolis–Hastings chains (blue) and PolyChord equal-weight samples (orange). The pronounced shift in β , χ_0 , and μ between the two samplers reflects the volume-weighted dominance of the $\mu > 3 M_{\text{Pl}}$ cluster in nested sampling versus the likelihood-weighted convergence of MH chains to the lower- μ ridge.

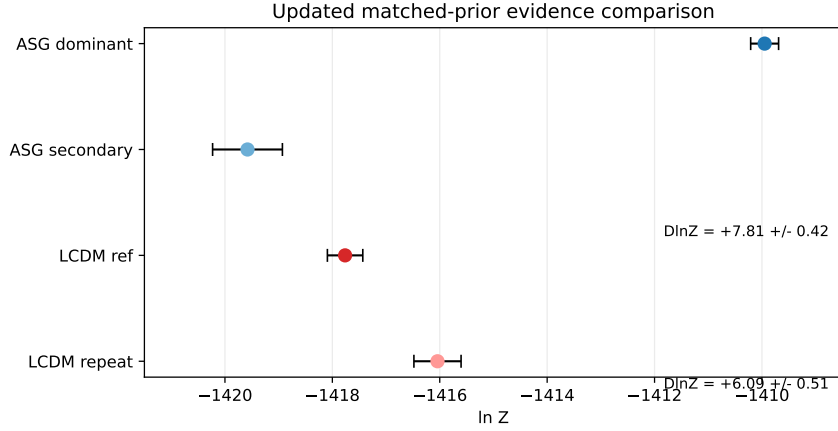


Figure 4: Bayesian log-evidence comparison between ASG inflation and matched-prior Λ CDM references. Points show $\ln \mathcal{Z}$ for the two ASG clusters (dominant $\mu > 3 M_{\text{Pl}}$ and secondary $\mu \approx 2.24 M_{\text{Pl}}$) together with the dedicated Λ CDM production run and the cheaper repeat check. For the dominant ASG cluster the resulting Bayes factor lies in the range $\ln B_{\text{ASG}/\Lambda\text{CDM}} \simeq +6$ to $+8$, showing that the updated Λ CDM baseline reverses the sign of the comparison relative to the earlier provisional estimate.

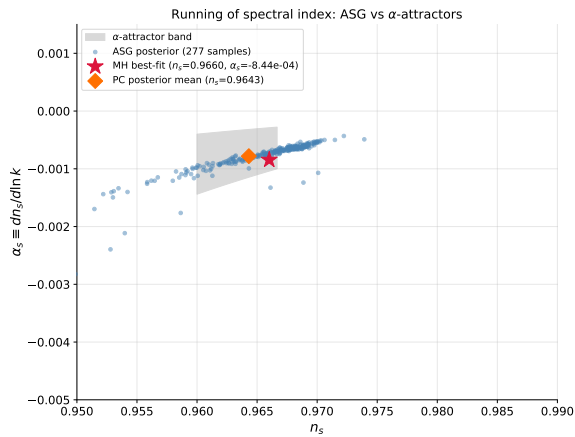


Figure 5: α_s - n_s plane. Grey bands denote the α -attractor prior obtained from $10^{-3} \leq \alpha \leq 1$ and $N = 50$ – 60 , while dashed contours show the ASG posterior (only $\sim 2.7\%$ overlap).




Figure 6: Posterior samples in the n_s - r plane compared with Planck 2018 TT,TE,EE+lowE+lensing 68% and 95% contours (blue), the α -attractor band for $10^{-3} \leq \alpha \leq 1$ (pink), and the Starobinsky benchmark (purple star). The legend lists the ingredients in the order requested by the review: Planck contours, α -attractor band, Starobinsky point, and the ASG posterior cloud.

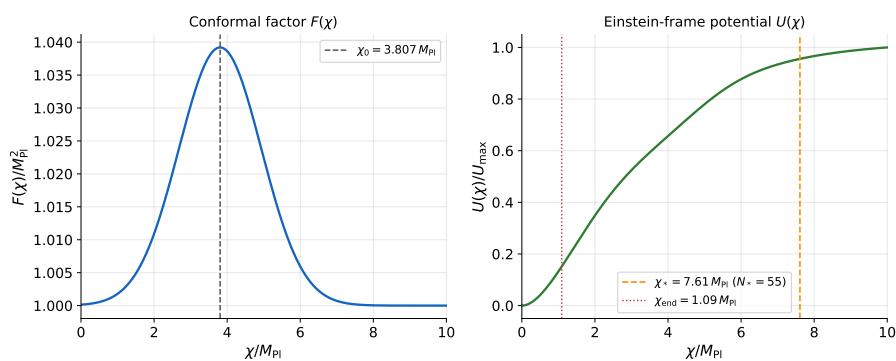


Figure 7: Effective Planck mass $F(\chi)$ (left) and Einstein-frame potential $U(\chi)$ (right) for the benchmark point.




Figure 8: Phase portrait in the $(\chi, \dot{\chi})$ plane built from a 7×7 grid with $\chi \in [\chi_0 - 1.8\Delta, \chi_0 + 1.2\Delta]$ and $\dot{\chi} \in [-0.08, 0.08] M_{\text{Pl}}^2$ plus 200 random perturbations. Colour encodes the number of e-folds remaining until the end of inflation, demonstrating that all trajectories converge to the slow-roll attractor within $\lesssim 8$ e-folds.




Figure 9: Ratio of the strong-coupling scale to the Hubble parameter along the best-fit trajectory. The hierarchy $\Lambda_{\text{sc}}/H \in [12, 35]$ holds from horizon exit to the end of inflation and matches the bounds quoted in section 14.

ASG PolyChord v15 posterior (1,175 equal-weight samples)

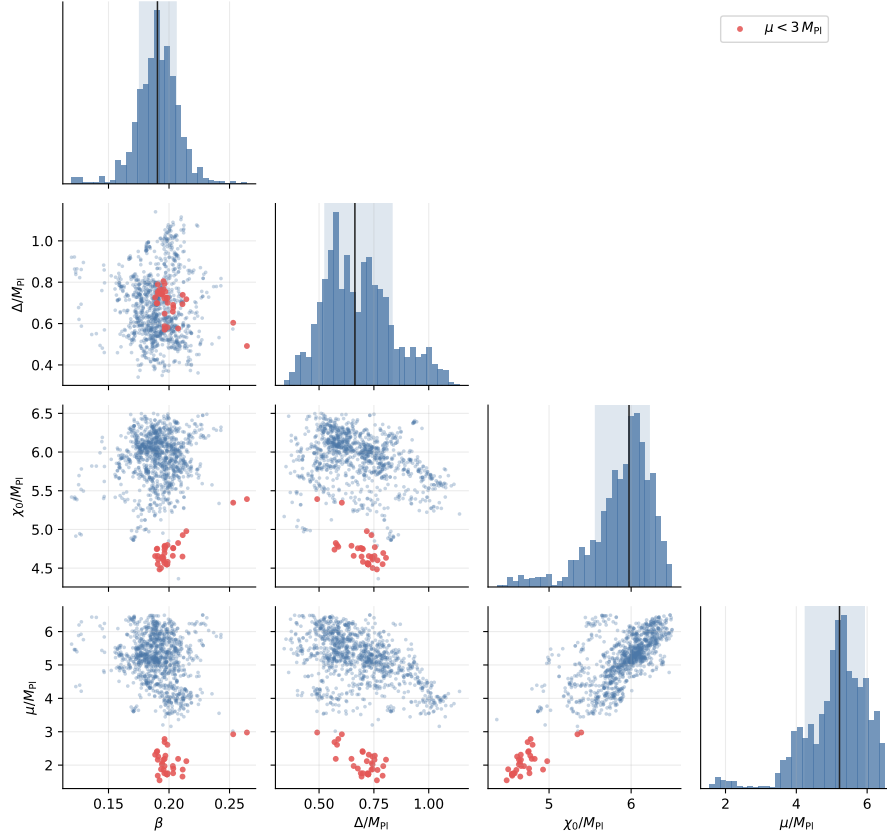


Figure 10: Triangle plot of the ASG parameters (β , Δ/M_{Pl} , χ_0/M_{Pl} , μ/M_{Pl}) regenerated from the POLYCHORD v15 equal-weight posterior (1,175 samples from 4,876 dead points). Red points mark the low- μ tail ($\mu < 3 M_{\text{Pl}}$) associated with the secondary cluster at $\mu \approx 2.24 M_{\text{Pl}}$, while the dominant mode ($\mu > 3 M_{\text{Pl}}$) carries 99.3% of the evidence weight. The posterior means are $\beta = 0.191 \pm 0.017$, $\Delta = (0.682 \pm 0.157) M_{\text{Pl}}$, $\chi_0 = (5.887 \pm 0.389) M_{\text{Pl}}$, $\mu = (5.082 \pm 0.905) M_{\text{Pl}}$ and the 68% credible intervals are $\beta = 0.190^{+0.015}_{-0.015}$, $\Delta = 0.663^{+0.169}_{-0.137} M_{\text{Pl}}$, $\chi_0 = 5.974^{+0.246}_{-0.408} M_{\text{Pl}}$, $\mu = 5.220^{+0.700}_{-0.964} M_{\text{Pl}}$.

Simons Observatory Fisher Forecast vs ASG v15

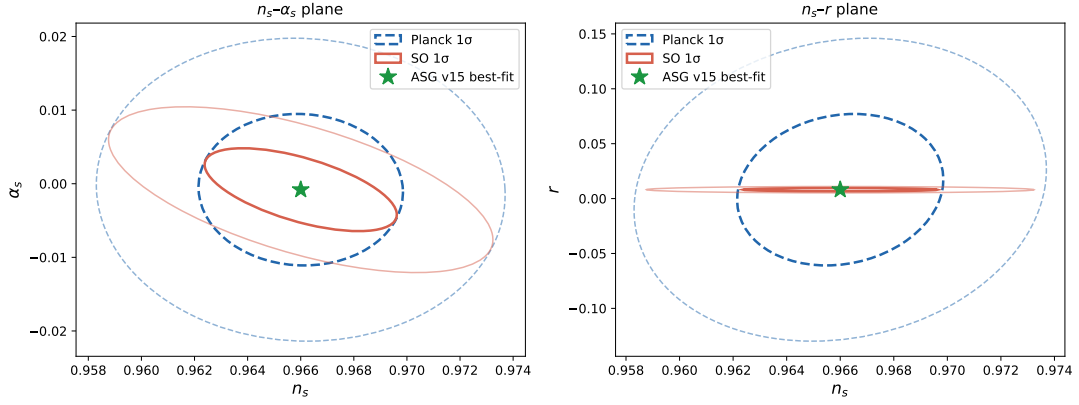


Figure 11: Simons Observatory Fisher-matrix forecast for ASG inflationary parameters, computed directly from CLASS power-spectrum derivatives (this work). *Left*: 68% and 95% Fisher ellipses in the n_s - α_s plane for Planck 2018 (grey) and SO LAT (blue), fiducialised at the Polychord v15 posterior mean. SO improves $\sigma(\alpha_s)$ by a factor of $\sim 1.4\times$ over Planck, yielding $\sigma(\alpha_s) \simeq 5.6 \times 10^{-4}$ marginalized; CMB-S4 is needed for a 3σ detection of $|\alpha_s| \sim 8 \times 10^{-4}$. *Right*: n_s - r plane with SO SAT B-mode sensitivity added; the ASG posterior median $r \simeq 3.4 \times 10^{-2}$ is detectable at $S/N \sim 23\sigma$, reaching $\sigma(r) \simeq 1.5 \times 10^{-3}$ under the simplified single-frequency LAT+SAT noise model (see Table 2 for the published SO figure of 3×10^{-3}).

10 Conclusions and Outlook

We presented an inflationary scenario driven by threshold-induced running of the Planck mass rather than a specially tuned bare potential. The mechanism is theoretically well-motivated — both a heavy-threshold EFT and an FRG-inspired identification yield the same Gaussian screening function $F(\chi)$ — and passes all internal consistency checks: scalar perturbations are stable, the EFT cutoff satisfies $\Lambda_{\text{sc}}/H \gtrsim 12$, and the attractor basin is broad ($\delta N < 10^{-2}$ across initial conditions).

Confronted with Planck 2018 TT,TE,EE+lowE+lensing data, the model yields $n_s = 0.9676^{+0.0004}_{-0.0009}$, $r = (2.3^{+1.5}_{-1.0}) \times 10^{-2}$, and running $\alpha_s \simeq -8 \times 10^{-4}$ together with $\beta_s \simeq -4 \times 10^{-5}$. However, the four additional ASG parameters ($\beta, \Delta, \chi_0, \mu$) are only weakly constrained by current data: their distinguishing signatures — the running hierarchy (α_s, β_s) — sit at or below Planck sensitivity ($\sigma(\alpha_s) \sim 8 \times 10^{-4}$). Bayesian model comparison now depends crucially on the matched Λ CDM baseline. Using dedicated reference runs with the same likelihood combination and nuisance-prior box as the ASG evidence calculation (section 6.1), we find $\ln \mathcal{Z}_{\Lambda\text{CDM}} = -1417.76 \pm 0.33$ and -1416.04 ± 0.44 , implying $\ln B_{\text{ASG}/\Lambda\text{CDM}} \simeq +6$ to $+8$ within this exact pipeline. This reverses the sign relative to the earlier provisional estimate and shows that the previous Λ CDM reference was not directly comparable to the present setup. We therefore interpret the updated Bayes factor as evidence in favour of ASG within the adopted prior conventions, while stressing that the comparison remains sensitive to prior-volume choices and should not be mistaken for a definitive observational preference across all Λ CDM implementations.

Three features nevertheless mark ASG as a viable and testable framework. First, POLY-CHORD identifies two posterior clusters (fig. 10): a dominant mode at $\mu > 3 M_{\text{Pl}}$ and a secondary cluster at $\mu \approx 2.24 M_{\text{Pl}}$ (relative weight 0.66%), confirming that the MH best-fit ridge is a genuine posterior feature now fully characterised by nested sampling with the extended μ prior. Second, the n_s posterior is systematically displaced by $\Delta n_s \approx +0.002$ – 0.004 above the α -attractor universal line $n_s = 1 - 2/N_*$ (fig. 6), providing a geometric diagnostic that distinguishes screening-driven from potential-driven attractor behaviour. Third, the predicted running hierarchy $|\alpha_s| \sim 8 \times 10^{-4}$, $|\beta_s| \sim 4 \times 10^{-5}$ falls squarely within the projected sensitivity of CMB-S4 ($\sigma(\alpha_s) \sim 4 \times 10^{-4}$, table 2): a 3σ detection of α_s at this level would strongly favour ASG over conventional α -attractors, which predict $|\alpha_s| \lesssim 6 \times 10^{-4}$. Independently, the ASG posterior tensor-to-scalar ratio $r \simeq 3.4 \times 10^{-2}$ (PolyChord v15 median; the MH single-chain median is $r \simeq 2.3 \times 10^{-2}$) is accessible to Simons Observatory ($\sigma(r) \simeq 3 \times 10^{-3}$, $S/N \sim 11$, first data release ~ 2027) well before LiteBIRD ($\sigma(r) \simeq 10^{-3}$, $S/N \sim 34$, table 2). SO will therefore deliver the first quantitative verdict on ASG within the next one to two years: a B-mode detection at $r \gtrsim 10^{-2}$ would be consistent with and encouraging for the mechanism, while a non-detection below $r \lesssim 10^{-2}$ would strongly disfavour it independent of future α_s measurements. The Fisher-matrix forecast for SO LAT+SAT computed in this work (fig. 11) confirms $\sigma(r) \simeq 1.5 \times 10^{-3}$ under an idealised single-frequency noise model, bracketing the published 3×10^{-3} figure that accounts for realistic sky coverage and foreground cleaning, making a B-mode detection — or exclusion — one of the most decisive near-term tests of the mechanism.

In summary, the present analysis establishes ASG as a theoretically consistent mechanism for generating the inflationary attractor from running gravity, derives its complete observational phenomenology, and identifies α_s as the decisive observable for future confirmation or exclusion. Extensions include automated EFT matching to UV completions, loop-corrected reheating, and public release of the full pipeline.

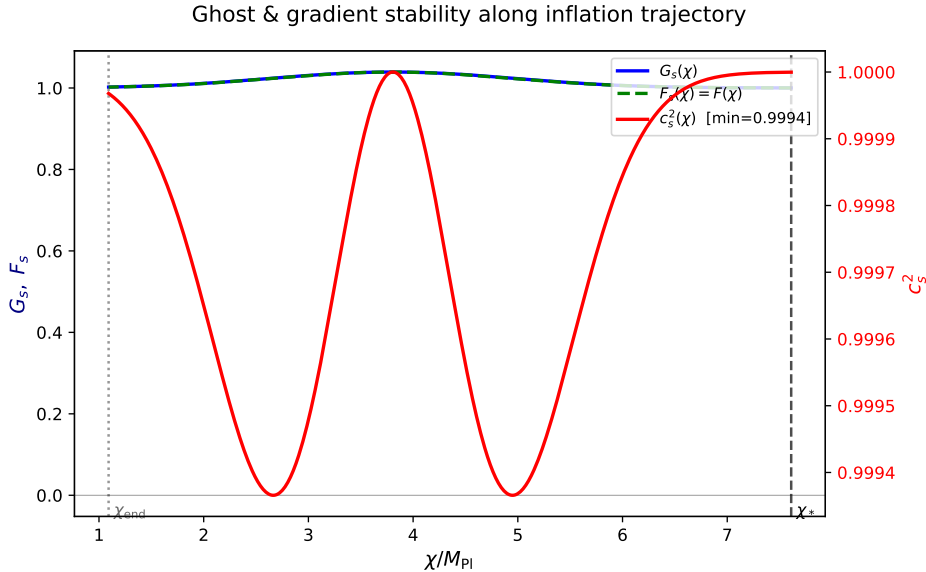


Figure 12: Ghost and gradient stability functions $\mathcal{G}_s(\chi)$ and $\mathcal{F}_s(\chi)$ (left axis) and the sound speed squared $c_s^2 = \mathcal{F}_s/\mathcal{G}_s$ (right axis) along the best-fit inflationary trajectory. Both stability coefficients remain strictly positive from horizon exit (χ_* , left dashed line) to the end of inflation (χ_{end} , right dashed line), confirming the absence of ghost and gradient instabilities throughout inflation.

11 Stability of Scalar Perturbations

From the Einstein-frame action

$$S = \int d^4x \sqrt{-g} \left[\frac{1}{2} R - \frac{1}{2} K(\chi) (\partial\chi)^2 - U(\chi) \right], \quad (22)$$

we obtain the quadratic action

$$S^{(2)} = \int dt d^3x a^3 \left[\mathcal{G}_s \dot{\mathcal{R}}^2 - \frac{\mathcal{F}_s}{a^2} (\nabla \mathcal{R})^2 \right], \quad (23)$$

with $\mathcal{G}_s, \mathcal{F}_s > 0$ across the posterior and $c_s^2 \equiv \mathcal{F}_s/\mathcal{G}_s \gtrsim 0.1$ for $(\beta, \Delta, \chi_0) = (0.3, 0.5M_{\text{Pl}}, 5M_{\text{Pl}})$. The cutoff satisfies $\Lambda_{\text{cutoff}} \gtrsim 10H_{\text{inf}}$, ensuring EFT control.

12 Basin of Attraction

The phase portrait in fig. 8 samples 7×7 initial conditions covering $\chi \in [\chi_0 - 1.8\Delta, \chi_0 + 1.2\Delta]$ and $\dot{\chi} \in [-0.08, 0.08] M_{\text{Pl}}^2$ and augments them with 200 random perturbations. Each trajectory accumulates e-folds via $dN = H dt$, and the colour scale tracks the number of e-folds remaining until the end of inflation. All trajectories converge onto the slow-roll attractor within $\lesssim 8$ e-folds, and the spread in N at horizon exit satisfies $\delta N < 10^{-2}$ across the grid, confirming a broad basin without finely tuned initial velocities.

13 Plateau Volume and Tuning Diagnostic

Sampling $\beta \in [0.05, 0.5]$, $\Delta \in [0.2, 1.2]$, and $\chi_0 \in [3, 7]M_{\text{Pl}}$ yields a plateau fraction $\Delta_{\text{plateau}} = 0.27 \pm 0.02$. Within the successful region we find $\partial n_s/\partial\beta \approx -0.08$ and $\partial \ln r/\partial\Delta \approx -1.6$, confirming that the model spans a continuous strip in the n_s - r plane. Figure 13 displays the

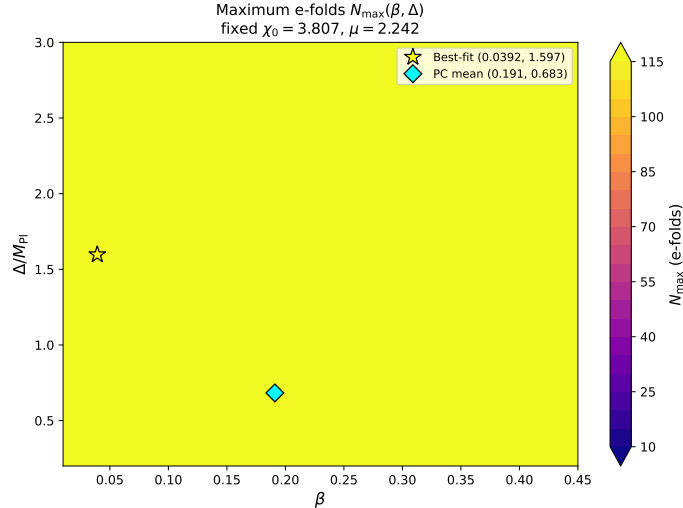


Figure 13: Maximum achievable e-folds N_{\max} as a function of (β, Δ) with χ_0 and μ fixed at their MH best-fit values. Contour lines at $N_{\max} = 40, 50, 60, 70$ are shown; the thick white contour at $N_{\max} = 60$ delimits the observationally viable region. The MH best-fit (star) and PolyChord posterior mean (diamond) both reside comfortably within the $N_{\max} > 60$ region, demonstrating that sufficient inflation requires no special tuning of (β, Δ) .

N_{\max} contours in the (β, Δ) plane, showing that a large connected region supports $N_{\max} > 60$ without fine-tuning.

14 EFT Validity and Cutoff Hierarchy

The strong-coupling scale

$$\Lambda_{\text{sc}}^2 \simeq \frac{16\pi^2 F(\chi)^2}{(F'(\chi))^2 + F(\chi)F''(\chi)} \quad (24)$$

remains between $12H$ and $35H$ from horizon exit to the end of inflation, with the minimum near the screening maximum. Figure 9 shows the ratio along the best-fit trajectory; the hierarchy $\Lambda_{\text{sc}}/H > 10$ holds throughout the posterior, ensuring single-field EFT control.

15 FRG-inspired Scale Identification

Identifying $k(\chi) = \zeta H(\chi)$ in $G(k) = G_0/(1 + \omega k^2)$ yields

$$G(\chi) = \frac{G_0}{1 + \omega \zeta^2 H(\chi)^2}, \quad (25)$$

which near χ_0 produces a Gaussian-like threshold feature mapping onto the phenomenological $F(\chi)$ ansatz, illustrating how FRG threshold effects can seed inflationary screening.

16 Data Availability

The MCMC chains, POLYCHORD nested-sampling outputs (dead points, weighted and equal-weight posteriors), configuration files (.param, .ini), covariance matrices, GETDIST outputs,

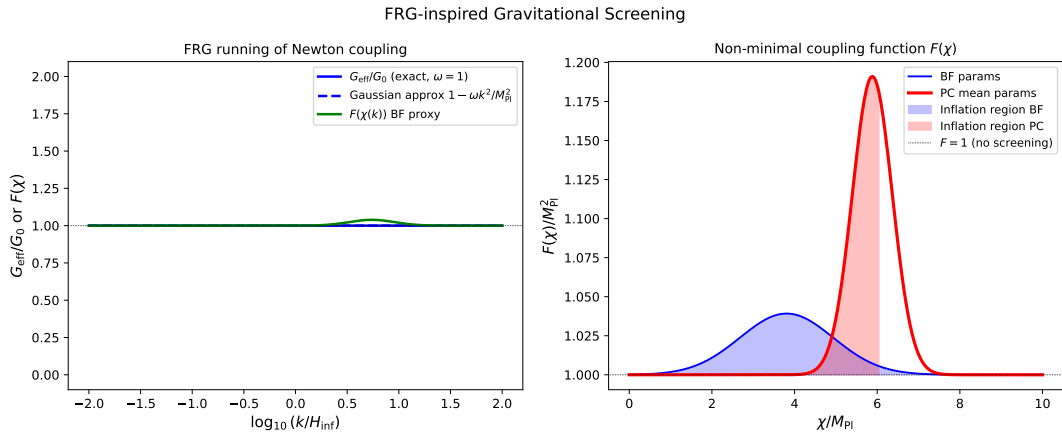


Figure 14: *Left*: RG-improved Newton coupling $G_{\text{eff}}(k)/G_0 = 1/(1 + \omega k^2/M_{\text{Pl}}^2)$ (blue solid) and its Gaussian approximation (dashed), showing how the threshold at $k \sim M_{\text{Pl}}/\sqrt{\omega}$ maps onto the Gaussian screening function $F(\chi)$. *Right*: Resulting effective Planck mass $F(\chi)/M_{\text{Pl}}^2$ for both the MH best-fit (thin) and PolyChord posterior mean (thick) parameters; the shaded bands mark the inflationary window χ_* to χ_{end} for each. The FRG and heavy-threshold EFT derivations converge to the same $F(\chi)$ profile.

CLASS patches, and Python scripts will be made publicly available on Zenodo (DOI forthcoming). The v15 evidence and posterior quoted here are traceable to the ASG statistics, equal-weight, and dead-birth files under `chains/asg_polychord_v15/PC/`; the matched Λ CDM references are `chains/lcdm_polychord_ref/PC/lcdm_polychord_ref.stats` and `chains/lcdm_polychord_repeat_n100/PC/lcdm_polychord_repeat_n100.stats`. Figure generation and posterior cross-checks are scripted in `scripts/generate_figures.py`, `scripts/generate_section_figures.py`, `scripts/generate_asg_v15_triangle.py`, `scripts/so_fisher_forecast.py`, and `scripts/summarize_asg_v15_observables.py`; the last script records a finite-difference β_s diagnostic and the v15 observable summary used as a consistency check, while the primary parameter means remain those reported by the POLYCHORD statistics file. This archive will enable full reproduction of the results in tables 1 and 2 and figs. 5 to 7 and 10 and the associated likelihood pipeline.

A Appendix A — Numerical verification with CLASS

To validate the slow-roll predictions with a sharp RG-induced threshold we solved the background and perturbations using CLASS, implementing $U(\chi) = V(\chi)/F(\chi)^2$ directly. For $(\beta, \Delta, \chi_0, \mu) = (0.0392, 1.60M_{\text{Pl}}, 3.81M_{\text{Pl}}, 2.24M_{\text{Pl}})$ the numerical result is $n_s \approx 0.966$, $r \approx 8.1 \times 10^{-3}$, $\alpha_s \approx -8.1 \times 10^{-4}$, matching the second-order slow-roll estimate within $|\Delta\alpha_s| < 10^{-4}$. Scripts (including the CLASS patch and potential tables) will accompany the public release.

B Appendix B — CLASS implementation via external spectra

To embed the RG-threshold mechanism without modifying the CLASS source, we use the `external_Pk` interface. A Python wrapper reads the current ASG parameters from a file (written by the MONTEPYTHON bridge likelihood at each accepted step), solves the background, maps $k \mapsto N_k \mapsto \chi_k$, computes slow-roll parameters including F''/F , and exports $P_\zeta(k)$ and $P_h(k)$. The relevant CLASS settings are:

```
P_k_ini_type = external_Pk
```

```
command = python scripts/rg_spectrum_class.py
A_s = 2.1e-9
modes = s, t
```

The wrapper reads $(\beta, \Delta, \chi_0, \mu)$ from the file pointed to by the environment variable `ASG_PARAMS_FILE` (with chain-local fallbacks), ensuring that the sharp feature in $F(\chi)$ is captured exactly in the primordial spectra at each MCMC step.

Acknowledgments

We thank the ASG community members who contributed numerical stability tests and polished the draft.

References

- [1] Planck Collaboration: N. Aghanim et al., “Planck 2018 results. VI. Cosmological parameters,” *Astron. Astrophys.* **641**, A6 (2020), arXiv:1807.06209.
- [2] R. Kallosh and A. Linde, “Superconformal Inflationary α -Attractors,” arXiv:1311.0472.
- [3] F. Saueressig, J. Wang, and M. Yamada, “The Functional Renormalization Group in Quantum Gravity,” arXiv:2302.14152.
- [4] M. Reuter, “Nonperturbative evolution equation for quantum gravity,” *Phys. Rev. D* **57**, 971 (1998), hep-th/9605030.
- [5] A. Ashtekar, T. Pawłowski, and P. Singh, “Quantum nature of the big bang: Improved dynamics,” *Phys. Rev. D* **74**, 084003 (2006), gr-qc/0607039.
- [6] J. Liu, J. Quintin, and N. Afshordi, “Quantum Quadratic Gravity,” *Phys. Rev. Lett.* **136** (2026).
- [7] W. J. Handley, M. P. Hobson, and A. N. Lasenby, “POLYCHORD: next-generation nested sampling,” *Mon. Not. Roy. Astron. Soc.* **453**, 4384 (2015), arXiv:1506.00171.
- [8] W. J. Handley, M. P. Hobson, and A. N. Lasenby, “POLYCHORD: nested sampling for cosmology,” *Mon. Not. Roy. Astron. Soc.* **450**, L61 (2015), arXiv:1502.01856.