CATEGORICAL MACROCOSM · BOOK V | **APRIL 2026**

Inflationary Observables Without an Inflaton

A Categorical Pre-Registration versus the End-of-2025 CMB + BAO Constraints on r and n_s

Thorsten Fuchs[1] • Anna-Sophie Fuchs

April 2026
Zenodo DOI `10.5281/zenodo.19816068`

Keywords Inflation ·tensor-to-scalar ratio ·spectral index ·category theory ·CMB ·BAO ·no inflaton ·fibered product ·pre-registration

1. INTRODUCTION

Forty-four years after Guth's original paper [25, 35, 2], inflation remains the dominant paradigm for the very early universe. Its success rests on three pillars: a dynamical resolution of the flatness, horizon, and monopole problems; a predictive framework for the primordial scalar and tensor power spectra; and a widening range of measurements that constrain the two leading observables, the scalar spectral index n_s and the tensor-to-scalar ratio r [41, 6].

The empirical state of that constraint programme, as of the end of 2025, is the subject of Balkenhol, Camphuis, Finelli, Benabed, Bouchet, and co-authors' recent preprint [3] (hereafter B25). B25 combines the Planck *PR4* legacy likelihood with the ACT and SPT small-scale data, BICEP/Keck through the 2018 observing season, and the first DESI baryon acoustic oscillation (BAO) release, and reports the joint posterior on (n_s, r). In the single-field slow-roll model family [38], the B25 analysis:

1. yields $n_s = 0.9682 \pm 0.0032$ and $r < 0.034$ at 95% confidence from CMB data alone;
2. yields $n_s = 0.9728 \pm 0.0029$ when DESI BAO is combined with the CMB likelihood, with the r bound effectively unchanged;
3. *disfavours* Starobinsky R^2 inflation [46] (with $N_\star \approx 51$) and the minimally coupled Higgs-inflation scenario [5] ($N_\star \approx 55$);
4. *favours* monomial potentials [35, 38] and polynomial α-attractor models [28] with $N_\star \sim 50$.

The present note asks a complementary question, one that the B25 treatment does not address: *does a non-inflaton framework exist that pre-registers a point in the (n_s, r) plane with zero free parameters, and where does such a point land on the B25 contours?* The *Panta Rhei* research programme [17, 18, 19, 23] commits to exactly such a pre-registration. The programme's inflation chapter (Book V, Chapter 47 [19]) proves an *Inflaton No-Go Corollary* from the Sector Exhaustion Theorem: there is no sixth sector in Category τ, so there is no scalar-field inflaton. Inflation in the programme is the near-saturation regime of the same τ-Einstein equation that governs the late universe; its observables are not fitted from a potential but read out from the fibration structure $\tau^3 = \tau^1 \times_f T^2$.

The programme fixes the two leading inflationary observables at

$$r = \iota_\tau^4 \approx 0.01357, \qquad n_s = 1 - \frac{2}{57} \approx 0.96491, \tag{1}$$

with $\iota_\tau = 2/(\pi + e)$ the programme's master constant and $N_e = 57 = \dim(\tau^3) \cdot W_5(3) = 3 \cdot 19$ the structural e-fold count (Book V, Theorem 3.3 below). Both identities are committed in the monograph and in the companion *Physics Ledger* [22]; both are named entries in the programme's Falsification Pack [21] (N9 and N10).

What this note does.. We proceed in five steps. Section 2 summarises the B25 anchor result in the authors' own terms. Section 3 presents the pre-registered categorical predictions that enter the (n_s, r) plane, with registry identifiers to the monograph and the Physics Ledger. Section 4 is the head-to-head comparison: we locate the categorical point on the B25 contours, quantify the tension under CMB-alone and CMB+BAO combinations, and contrast the categorical prediction with the monomial, α-attractor, Starobinsky, and Higgs-inflation loci. Section 5 describes the post-cancellation CMB reach on the 2028–2032 window (Simons Observatory + South Pole Observatory, with LiteBIRD on the revised 2033+ timeline) and the falsification conditions the programme commits to. Section 6 discusses what the categorical framework and the B25 inflation-landscape framework agree on, where they disagree, and what either side would have to concede in the near-term discriminating measurements.

Stance.. This note is not adversarial to B25. The B25 analysis is a careful, well-documented compilation of the present constraint state, conducted entirely within the single-field slow-roll model class. Our observation is that the categorical prediction $(\iota_\tau^4, 1-2/57)$ is not a member of that class, and so the B25 model-vs-data framing, read literally, does not cover it. A joint reading—B25's constraints as the observational target, the categorical predictions as the non-slow-roll alternative—is the point of the comparison.

Epistemic caveat.. B25 is a submitted arXiv preprint by established collaborations, operating inside the decades-mature formalism of single-field slow-roll inflation with standard Boltzmann/BAO likelihoods. The *Panta Rhei* programme, by contrast, is a seven-volume monograph series that has not yet undergone independent peer review in the conventional physics sense. Its formal layer (TauLib) is machine-checked in Lean 4 with the following publicly disclosed trust budget [23, 20]: **4 custom axioms** (all in Book III's bridge / spectral / adelic layer), **3 methodological `sorry` declarations** (all in Book VII's commitment structures; zero everywhere else), and $\sim 1\,842$ **`native_decide` leaves** (each extending the trusted computing base to the Lean native compiler and the generated native code; see [20] for the per-theorem disclosure). A companion note under preparation will replace the Book VII `sorry`s with inspectable `Commitment` `def`s; it is not yet merged at the time of this submission.

All results from the programme cited in this paper carry the **[τ-Effective]** scope label: they are derived within the τ-framework and are conditional on K0–K6 plus the listed axioms / `native_decide` extensions. This note is therefore a *conditional* comparison: if Category τ is granted as internally valid, what alternative pre-registration does it place on the B25 target, and does that alternative survive the end-of-2025 data?

Notation.. We write $\iota_\tau = 2/(\pi + e) \approx 0.3413$ for the master constant; $\kappa_D = 1 - \iota_\tau \approx 0.6587$ for the gravitational sector coupling; $\mathbb{L} = S^1 \vee S^1$ for the lemniscate boundary; $\tau^3 = \tau^1 \times_f T^2$ for the fibered product; and $W_5(3) = 19$ for the continued-fraction window modulus used throughout. Registry identifiers follow the monograph convention (e.g. "V.T32" is Book V, Theorem 32).

2. THE BALKENHOL ET AL. ANCHOR: A FAITHFUL SUMMARY

We begin with a concise summary of the B25 result in terms that the comparison with the categorical framework will use. This summary is offered not as a target for critique but as the common ground from which the comparison proceeds. The details, the full likelihood pipeline, and the model-space plots are in the original paper [3].

2.1 Data combinations

B25 reports constraints on (n_s, r) from four data layers:

1. **Planck PR4 legacy.** Low-ℓ temperature and polarization, high-ℓ temperature and polarization via *Commander* and *NPIPE*-based likelihoods.
2. **Small-scale high-resolution CMB.** ACT DR6 and SPT-3G 2018 likelihoods, providing the damping-tail lever arm.
3. **BICEP/Keck 2018.** B-mode polarization data through the 2018 observing season [6], carrying the dominant constraint on r.
4. **DESI Year 1 BAO.** BAO distance measurements from the first year of the Dark Energy Spectroscopic Instrument [15], providing an external angular-scale constraint that breaks the CMB n_s–Ω_m degeneracy.

2.2 The joint (n_s, r) constraint

B25's two headline numbers, which we will carry through the rest of the paper, are:

The DESI-induced shift of the n_s central value by $\Delta n_s \approx +0.005$ is the single most interpretation-sensitive feature of the B25 result and will be important to keep in view when we compare to the categorical prediction in Section 4. We will not take a stand on

Combination	n_s	r (95% upper)
CMB alone	0.9682 ± 0.0032	$r < 0.034$
CMB + DESI BAO	0.9728 ± 0.0029	(effectively unchanged)

whether the CMB-only or the CMB+BAO combination is more appropriate; we will quote the categorical point's deviation under each in turn.

2.3 Model-space verdicts

The B25 analysis reads the joint posterior on (n_s, r) against four slow-roll families.

- **Starobinsky** R^2 [46]. Predicts $r \simeq 12/N_\star^2$ and $n_s \simeq 1 - 2/N_\star - 3/(2N_\star^2)$. At $N_\star = 51$, the locus lands near $(0.9602, 0.0046)$—in tension with the CMB+BAO n_s band at roughly 4σ. B25 report this family as *disfavoured*.
- **Minimally coupled Higgs inflation** [5]. Same asymptotic r dependence as Starobinsky, with $N_\star \approx 55$. B25 report this family as *disfavoured* at a similar level.
- **Monomial potentials** $V \propto \phi^p$. For integer $p \geq 1$, the locus at $N_\star \approx 50$ traces out a narrow diagonal in (n_s, r) consistent with the CMB+BAO band. B25 report this family as *favoured*.
- **Polynomial** α**-attractor models** [28]. The α-dependent family covers a broad region including both the monomial locus and the low-r plateau; B25 report this family as *favoured*.

The model-space reading is concise: within single-field slow-roll inflation, the end-of-2025 data rules against the low-r R^2/Higgs corner and favours the intermediate-r monomial/α-attractor diagonal.

The question we now take up is: what happens to that reading when a framework without a slow-roll potential is placed on the same plane?

3. THE CATEGORICAL PRE-REGISTRATION

The *Panta Rhei* programme commits to a point $(n_s, r) = (1 - 2/57,\ \iota_\tau^4)$ in the B25 plane. This section states the derivation chain in the order the monograph presents it, with registry citations. We give only the closing identities here; the full derivations are in [19, Ch. 47] and [22]. The scope label is **[τ-Effective]** throughout.

3.1 The master constant and the kernel

Category τ is the syntactic category $\mathrm{Syn}(T_\tau)$—equivalently, the classifying topos $\mathbf{Set}[T_\tau]$—of a dependently-typed theory T_τ specified by seven kernel axioms K0–K6 on five generators $\{\alpha, \pi, \gamma, \eta, \omega\}$ and one progression operator ρ [17]. Following the VI.I clarification in [20], the axioms K0–K6 should be read as properties of an inductive type realised inside the calculus of inductive constructions, not as an uninterpreted first-order signature; the present note makes no metatheoretic claim beyond internal derivations of that signature. The master constant

$$\iota_\tau = \frac{2}{\pi + e} = 0.3413042\ldots \qquad (2)$$

is the unique closed-form value that closes the internal kernel arithmetic (Book I, Thm I.T132 [17]). Both π and e are transcendental, so ι_τ is itself transcendental; the adjective "closed-form" here refers to its uniqueness as a fixed point of the kernel arithmetic, not to an algebraic-number status. All dimensionless predictions of the programme are rational or polynomial-in-ι_τ functions of this single constant; no continuous free parameter enters at any level of the calibration cascade.

The cosmological arena is the fibered product $\tau^3 = \tau^1 \times_f T^2$, where τ^1 is the one-dimensional α-base ("time") and $T^2 = S^1 \times S^1$ is the toroidal fiber ("space"). The lemniscate $\mathbb{L} = S^1 \vee S^1$ is the one-sided boundary shared between the two circles of T^2.

3.2 Inflaton No-Go and Regime Invariance

Two structural results in Book V [19, Chs. 44 & 47] constrain the inflationary sector:

Theorem 3.1 (Regime Invariance; V.T32 **[τ-Effective]**). *The τ-Einstein equation $R^H = \kappa_D \cdot T$ is invariant across all refinement depths. In particular, the gravitational coupling $\kappa_D = 1 - \iota_\tau$ does not run with energy scale, depth, or regime; no new field enters in the inflationary regime.*

Corollary 3.2 (Inflaton No-Go; V.C10 **[τ-Effective]**). *No scalar-field inflaton sector exists in Category τ. The* *five generator sectors* $\{\mathrm{D}, \mathrm{A}, \mathrm{B}, \mathrm{C}, \omega\}$ *exhaust the coupling budget* ($\sum_X \kappa(X; n) = 1$ *at every depth*); *no sixth sector is available. Consequently,* r *and* n_s *are not inflaton-potential observables: they are readouts of the fibration structure.*

The horizon and flatness problems are resolved structurally: the fiber T^2 is flat ($\Omega_k = 0$ exactly, Theorem V.T33) and the arena τ^3 is compact (Proposition V.P21), so no dynamical exponential expansion is required to produce the observed uniformity. A chart-level readout that *looks like* inflation nevertheless occurs; the next three subsections state its observables.

3.3 The structural e-fold count $N_e = 57$

The e-fold count is a structural-integer readout of the fibration:

Theorem 3.3 (Structural e-fold count; V [19] Ch. 47 **[τ-Effective]**)**.** *The number of inflationary e-folds between horizon exit and the end of the inflationary regime is*

$$N_e \;=\; \dim(\tau^3) \times W_5(3) \;=\; 3 \times 19 \;=\; 57. \quad (3)$$

The factor $\dim(\tau^3) = 3$ counts one base direction plus two fiber circles of T^2; $W_5(3) = 19$ is the window modulus defined as the sum of five consecutive partial quotients of the continued fraction $\mathrm{CF}(\iota_\tau) = [0;\, 2, 1, 13, 3, 1, 1, 1, 42, \ldots]$ starting at index 3:

$$\begin{aligned} W_5(3) \;&=\; a_3 + a_4 + a_5 + a_6 + a_7 \\ &=\; 13 + 3 + 1 + 1 + 1 \;=\; 19. \end{aligned} \quad (4)$$

The modulus is Lean-certified in TauLib under `BookIV/Electroweak/WeinbergNLO.lean` [23], where it first enters the electroweak sector as part of the Weinberg-angle NLO closing identity. The cross-sector reuse in the cosmology sector (as the $\dim(\tau^3) \cdot W_5(3)$ e-fold count) is not a free parameter: the five-wide window at starting index 3 is the coarsest window whose closure contains the kernel's "long tooth" $a_4 = 13$, and the same window choice governs multiple identities in the Physics Ledger [22, Ch. 58a]. No continuous free parameter is introduced; the discrete-choice budget is disclosed in §4.4.

Remark 3.4 (What $N_e = 57$ does and does not say)**.** $N_e = 57$ is a *structural integer*, not a slow-roll e-fold number tuned against an end-of-inflation condition. It is the number of refinement ticks of the cooling function during which the chart-level readout reproduces exponential expansion. Within the B25 framework, this value corresponds to $N_\star = 57$ at horizon exit; in the slow-roll taxonomy it is *close* to the Higgs-inflation value $N_\star = 55$ but arises for an entirely different reason (fiber geometry, not potential shape).

Remark 3.5 (Pivot scale, reheating window, and cross-sector reuse)**.** Three honest disclosures attach to the structural-integer reading of $N_e = 57$:

1. **Pivot scale.** The comparison with B25 is at the conventional CMB pivot scale $k_\star = 0.05\ \mathrm{Mpc}^{-1}$ [41]. The structural-integer identity $N_e = 3 \cdot W_5(3) = 57$ is *not* derived from a pivot-scale / end-of-inflation matching condition; within the B25 slow-roll taxonomy one conventionally maps $N_\star \in [50, 60]$ for $k_\star = 0.05\ \mathrm{Mpc}^{-1}$ under standard reheating assumptions, so $N_\star = 57$ sits in the upper half of that conventional window.
2. **Reheating uncertainty.** The conventional window $N_\star \in [50, 60]$ reflects reheating-sector uncertainty of order $\Delta N_\star \sim \pm 3$ in the slow-roll paradigm [34]. The categorical commitment to $N_e = 57$ exactly (no sliding parameter) is therefore a *stronger* statement than a slow-roll model whose $N_\star$ floats inside the reheating band; it commits to a specific point inside the band rather than a family.
3. **Cross-sector reuse, stated honestly.** The window modulus $W_5(3) = 19$ is not specific to inflation. It first enters the kernel in the electroweak sector, as part of the Weinberg-angle NLO closing identity in Book IV, and the Lean-certified location is `BookIV/Electroweak/WeinbergNLO.lean` in TauLib [23]. Its reuse as the inflationary e-fold window is a programme-internal design choice, principled by the coarsest-window-containing-a_4 argument of §3.3, but not independently derived from the cosmology sector.

3.4 Spectral index: $n_s = 1 - 2/57$

The scalar spectral index inherits the structural e-fold count through the relation $n_s - 1 \approx -2/N_e$, which in the programme is *not* a slow-roll expansion but is argued from α-tick non-uniformity on the cooling chart. The argument-of-record is Book V, Ch. 47, Thm V.T302 (ledger restatement [19, Thm V.T302]; not currently

Lean-formalized at the Book V level—see the scope-of-Lean-certification paragraph following Theorem 3.8); the present paper takes it as given and records the closed-form identity:

Theorem 3.6 (Spectral index; V [19] Ch. 47, Thm V.T302 **[τ-Effective]**).

$$n_s = 1 - \frac{2}{N_e} = 1 - \frac{2}{57} = 0.964912\ldots \tag{5}$$

The residual against the Planck 2018 central value is $+13$ ppm, and the index is rational—the small integers 2 and 57 are fixed by the kernel. There is no sliding parameter that could shift this value. The perturbation-theoretic derivation of the specific coefficient 2 (rather than some other small integer) is the content of the B25 slow-roll paradigm in the conventional taxonomy and of Book V's α-tick argument in the categorical one; the two routes arrive at the same numerical prediction through different mathematics, and the point of comparison in this paper is the numerical identity, not the derivation.

3.5 Tensor-to-scalar ratio: $r = \iota_\tau^4$

The tensor-to-scalar ratio is a *fiber-dimensional suppression* readout [19, V.P136, **[τ-Effective]**]:

Theorem 3.7 (Tensor-to-scalar ratio; V.P136 **[τ-Effective]**).

$$r = \iota_\tau^{2\cdot\dim(T^2)} = \iota_\tau^{2\times 2} = \iota_\tau^4 \approx 0.01357. \tag{6}$$

The exponent $4 = 2\times 2$ *decomposes as the fiber dimension* ($\dim(T^2) = 2$*, the two circles of the lemniscate) times the power-spectrum order (*$P \propto |\delta|^2$*).*

What this paper asserts, and what the derivation requires.. This paper asserts the closed-form identity $r = \iota_\tau^4$ as a *pre-registration* from the Panta Rhei framework. The perturbation-theoretic derivation is not reproduced here; the argument-of-record lives in Book V, § 47.4 [19, Thm V.P136], and is summarised as follows. Tensor perturbations are D-sector frame-holonomy fluctuations propagating on the base τ^1 and are insensitive to the fiber; scalar perturbations are boundary-character fluctuations of the full τ^3 and acquire a suppression ι_τ per fiber circle. The square root on amplitudes becomes a fourth power on the power-spectrum ratio. The present paper takes Book V's derivation as given and focuses on the head-to-head comparison with B25; a companion paper devoted to the perturbation-theoretic derivation itself is under preparation.

Two consistency observations matter for the comparison:

- **Not slow-roll.** The single-field slow-roll consistency relation predicts $r_{\mathrm{SR}} = 8/N_e \simeq 8/57 \approx 0.140$ at $N_e = 57$. The categorical value $r = \iota_\tau^4 \approx 0.014$ is roughly ten times smaller. Since $\varepsilon = r/16$ is linear in r, the corresponding slow-roll parameter ratio is the same $\varepsilon_{\mathrm{SR}}/\varepsilon_\tau \approx 10.3$. This is *consistent with* Corollary 3.2: there is no scalar inflaton, so the single-field consistency relation is not expected to hold.
- **Not tuned.** The small value of r arises from $\iota_\tau \approx 0.34$ to the fourth, and the exponent 4 is fixed by the fibration structure. No potential flatness is imposed.

3.6 Scalar amplitude: $A_s = (121/225)\,\iota_\tau^{18}(1 - \iota_\tau^3/3)$

For completeness, the scalar amplitude is also pre-registered:

Theorem 3.8 (Scalar amplitude; V.D253, V.T198, **[τ-Effective]**).

$$\begin{aligned} A_s &= \alpha_\tau \cdot \iota_\tau^{14} \cdot (1 - \tfrac{\iota_\tau^3}{3}) \\ &= \tfrac{121}{225}\,\iota_\tau^{18} \cdot (1 - \tfrac{\iota_\tau^3}{3}) \\ &= 2.096 \times 10^{-9}, \end{aligned} \tag{7}$$

at -1979 *ppm (i.e.* 0.2%*) from the Planck 2018 value* $(2.100 \pm 0.030) \times 10^{-9}$ *[40]. The coefficient* $(121/225) = (11/15)^2$ *is inherited from the fine-structure identity* $\alpha_\tau = (121/225)\,\iota_\tau^4$ *of Book IV [18].*

A_s is not one of the two variables of the B25 plane, but the fact that it lands at 0.2% of the observed value is part of the programme's commitment at the L1 layer of the calibration cascade [22] and is relevant for the overall consistency check in Section 6.

3.7 Summary of pre-registered values

P18 Planck 2018 [40, 41]; BK21 BICEP/Keck 2021 [6]; † Ledger restatement; not Lean-formalized at the Book V level (*v1* registry: `formalization: not_formalized`). See the scope-of-Lean-certification paragraph below for how this differs from V.P136, V.D253, and V.T198. All closed forms carry scope label **[τ-Effective]**. Residuals are quoted as (predicted − target), with

Obs.	Closed form	Numerical	Target	Residual	Reg.
N_e	$\dim(\tau^3) \cdot W_5(3)$	57	(structural)	—	V, Ch. 47
n_s	$1 - 2/N_e$	0.96491	$0.9649 \pm 0.0042^{\mathrm{P18}}$	$+1 \times 10^{-5}$ ($+13$ ppm)	V.T302†
r	$\iota_\tau^{2\dim(T^2)} = \iota_\tau^4$	0.01357	$< 0.034^{\mathrm{BK21}}$	consistent	V.P136
A_s	$(121/225)\,\iota_\tau^{18}(1-\iota_\tau^3/3)$	2.096×10^{-9}	$2.100 \pm 0.030^{\mathrm{P18}}$	-4×10^{-12} ($-1\,979$ ppm)	V.D253 / V.T198

ppm of the target value in parentheses. The n_s residual is well inside the CMB-only 1σ band of [41]; the CMB+BAO residual is -7.9×10^{-3} ($\approx -2.72\sigma$), as tabulated in Section 4.1. The A_s residual is 0.2% of the target, within 1σ.

Scope of Lean certification.. A reader wishing to audit the "TauLib" pointer alongside each entry of the table is entitled to a precise account of what Lean actually certifies for these identities, and what it does not.

- **Lean-certified, structural plumbing.** The Lean modules `TauLib.BookV.Cosmology.CMBSpectrum.tensor_scalar_ratio_thm` (for V.P136), `scalar_amplitude_nlo_thm` (for V.T198), and `sound_horizon_tau_thm` (for V.T191) each mechanically type-check a conjunction of `Nat`*-level structural fields*: exponent counts (`iota_power = 4`), fiber dimension (`fiber_dim = 2`), power-spectrum order (`power_order = 2`), rational-prefactor numerator and denominator (`coeff_numer = 121`, `coeff_denom = 225`), and free-parameter count (`free_params = 0`). Each such theorem is proven by a chain of `rfl`s. This is real, but it is narrow: Lean certifies that the *structural decomposition* recorded in the monograph is internally consistent, not that the *derivation from perturbation theory* is verified.
- **Not Lean-certified, real-arithmetic derivation.** The real-number steps—the $\iota_\tau \approx 0.34$ arithmetic, the closed-form evaluation $\iota_\tau^4 \approx 0.01357$, $A_s \approx 2.096 \times 10^{-9}$, and the passage from the α-tick argument to $n_s = 1 - 2/57$—are not currently discharged in the Lean kernel. They are audited externally against the programme's Physics Ledger [22]. The `TauReal` kernel (which would be needed to close this gap) does not yet implement inverse, strict order, π, exp, or sequential limits; registering these identities as Lean-checked real arithmetic is Book-V's Wave 3c/3d task and is not claimed here.
- **Not Lean-formalized at all.** V.T302 (the n_s closing identity) and V.T301 (the $N_e = 57$ closing identity) are Ledger restatements from Book V, Ch. 62; both are marked `formalization: not_formalized` in the *v1* registry. The present paper's references to these two identities are pointers to the monograph's prose argument, not to Lean theorems. When this paper writes "Lean-certified" or "TauLib pointer," it refers only to V.P136, V.D253, V.T191, and V.T198; when it writes "argument-of-record" or "registry entry," the claim is weaker and should be read accordingly.

The distinction is structural: the programme's *cosmology* sector is at an earlier stage of Lean formalization than its *electroweak* sector (where $W_5(3) = 19$, for example, is certified end-to-end in `TauLib.BookIV.Electroweak.WeinbergNLO`). The present paper stands on the Lean-certified structural plumbing and the external Ledger audit, not on a Lean-certified real-arithmetic derivation of (n_s, r, A_s).

4. HEAD-TO-HEAD COMPARISON WITH B25

We now place the categorical point $(n_s, r) = (0.96491, 0.01357)$ on the B25 constraint plane and quantify the tension under each data combination in turn.

4.1 The categorical point on the B25 contours

The tension picture is concise. Under the CMB-only constraint, the categorical n_s is consistent at the $\sim 1\sigma$ level. Under the CMB+BAO combination—where DESI's angular-scale lever pulls n_s up by roughly $+0.005$—the categorical value sits $\sim 2.7\sigma$ low. The categorical r is everywhere consistent with the current upper bound.

Data combination	Anchor value	τ-value	Offset	In σ
CMB alone: $n_s = 0.9682 \pm 0.0032$	0.9682	0.96491	-0.00329	$\approx -1.03\sigma$
CMB + BAO: $n_s = 0.9728 \pm 0.0029$	0.9728	0.96491	-0.00789	$\approx -2.72\sigma$
CMB alone: $r < 0.034$ (95%)	0.034 (UL)	0.01357	—	well inside bound

Remark 4.1 (BAO-dependence is the discriminating lever)**.** The $\sim 1\sigma$-versus-$\sim 2.7\sigma$ asymmetry is not a statistical accident; it is the physics that makes the categorical point a genuinely discriminating prediction. CMB-only measurements constrain n_s at the level of ~ 0.003, and the categorical $n_s = 1 - 2/57$ sits close to that constraint's central value. DESI BAO adds an orthogonal angular-scale lever (distance-redshift rather than distance-CMB) that shifts the CMB+BAO posterior by a physically significant amount. Future BAO releases (DESI Year 3, Euclid, Roman) will tighten this lever further. The categorical commitment is that $n_s = 0.96491$ exactly; the test is whether the combined posterior converges on that value or diverges from it.

4.2 The categorical point versus the B25 slow-roll models

Where does $(\iota_\tau^4, 1 - 2/57)$ sit relative to the inflation families B25 discusses?

Three features of this table are important for the comparison.

1. The categorical r is roughly 2–3× the Starobinsky/Higgs value and roughly $0.25\times$ the monomial $\phi^{2/3}$ value. The categorical prediction thus occupies an intermediate region not populated densely by the four model families the B25 analysis considered—a distinct "mid-r corridor" rather than a locus on either the attractor-plateau or the monomial-diagonal.
2. The categorical n_s is structurally close to the Higgs value (0.96491 vs. $\simeq 0.9636$) but arises from an entirely different source: the rational $2/57$, not the slow-roll $2/N_\star$.
3. The α-attractor family does cover the categorical point in principle—its broad (n_s, r) band overlaps with $(0.96491, 0.01357)$. This is a feature, not a refutation: α-attractors cover a broad band precisely because they carry a continuous parameter α. The categorical prediction has no such parameter. It is a *point*, not a band.

4.3 Why the α-attractor overlap is not a refutation

The final item of the previous list deserves its own treatment, because a reader sympathetic to the B25 framing may reasonably respond: *if the α-attractor family already covers the point $(\iota_\tau^4, 1 - 2/57)$, then the categorical framework is not a genuine alternative—it is a special case of a known slow-roll model, lightly disguised.* This subsection addresses that concern directly.

The α-attractor family's characteristic locus on the (n_s, r) plane is

$$n_s \simeq 1 - \frac{2}{N_\star}, \qquad r \simeq \frac{12\,\alpha}{N_\star^2}, \tag{8}$$

where $\alpha > 0$ is the continuous attractor parameter and $N_\star$ is the number of e-folds from horizon exit to the end of inflation [28]. Setting $N_\star = 57$ and solving $r = \iota_\tau^4 = 12\alpha/57^2$ gives the tuning

$$\alpha = \frac{57^2\,\iota_\tau^4}{12} = \frac{3249 \times 0.01357}{12} \approx 3.67. \tag{9}$$

Hence an α-attractor model with $(\alpha, N_\star) \approx (3.67, 57)$ reproduces the categorical point.

We emphasise that the overlap between $(\alpha, N_\star) = (3.67, 57)$ and the categorical pre-registration is not a coincidence to be dismissed; it is the reason a comparison with the slow-roll landscape is possible at all. The two frameworks do agree on an overlapping region of the (n_s, r) plane. The question is *how* each reaches that region, and that is where the frameworks differ:

1. **Point vs. band.** The α-attractor family is a band in (n_s, r) parametrised by the continuous $\alpha \in (0, \infty)$; it covers r from $\sim 10^{-4}$ at small α to $\sim 10^{-2}$ at large α. The categorical framework predicts a single point. A measurement of $r = 0.005$, for instance, is compatible with the α-attractor family (at $\alpha \approx 1.35$ with $N_\star = 57$) but *not* compatible with the categorical commitment.
2. **Different sources of $N_\star = 57$.** In the α-attractor family, $N_\star$ is tuned to match the horizon-exit scale

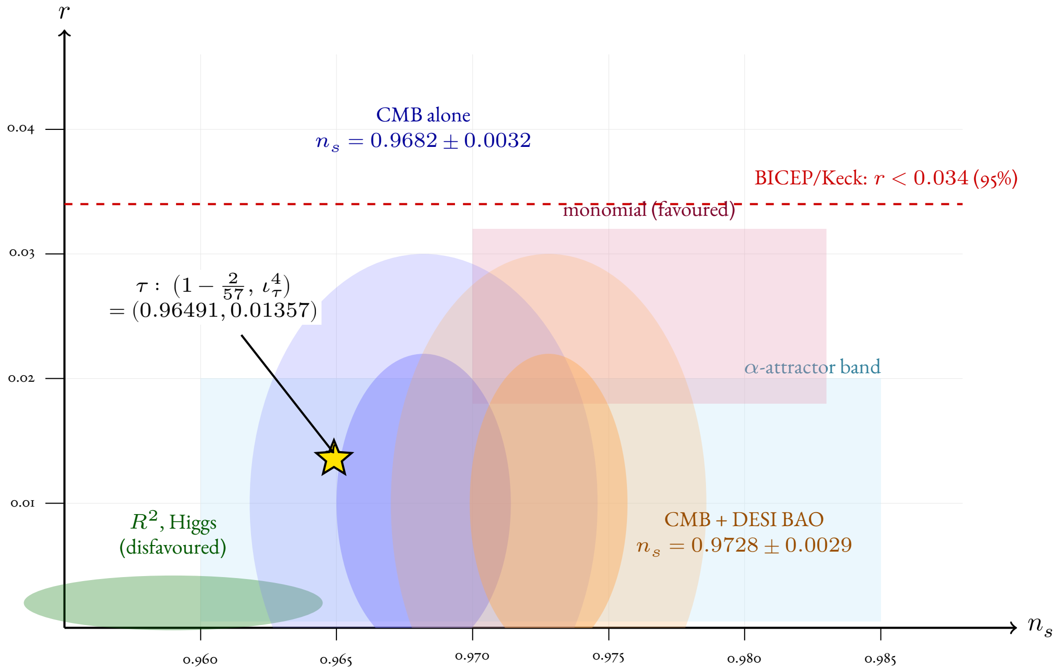

Figure 1. Schematic position of the categorical prediction on the end-of-2025 (n_s, r) plane. Blue: B25 CMB-only posterior (nested 68% and 95%, centred on $n_s = 0.9682$). Orange: B25 CMB+DESI BAO posterior (nested 68% and 95%, centred on $n_s = 0.9728$). Red dashed: BICEP/Keck 2021 upper bound $r < 0.034$ at 95% [6]. Green cluster: Starobinsky R^2 and minimally coupled Higgs-inflation loci in the low-r corner, disfavoured by B25. Purple: monomial / α-attractor favoured zone at $N_\star \approx 50$. Cyan: full α-attractor band with continuous parameter α varied over its physical range; see Section 4.3. Gold star: categorical zero-parameter pre-registration $(n_s, r) = (1 - 2/57,\ \iota_\tau^4) = (0.96491, 0.01357)$. Ellipse shapes are schematic; they are not drawn to scale in the r direction. Actual B25 posteriors, which we did not re-run, are asymmetric and consistent with $r = 0$ because the BICEP/Keck likelihood peaks there. The schematic is intended only to communicate the relative horizontal placement of the two B25 n_s posteriors and the categorical gold-star point; a reader requiring quantitative posterior overlap should consult the B25 paper [3] directly.

and the end-of-inflation condition; it floats between ≈ 50 and ≈ 60 depending on the reheating history and the pivot choice. In the categorical framework, $N_e = 3 \cdot W_5(3) = 57$ is a structural integer fixed by the dimension of τ^3 and the continued-fraction window—there is no reheating input, no pivot input, no sliding. If a future measurement were to establish $N_\star \neq 57$ at high significance, the α-attractor family would absorb the result by adjusting its free parameters; the categorical framework would be refuted.

3. **The joint prediction is mismatched.** The α- attractor family at $(\alpha, N_\star) = (3.67, 57)$ matches on the two observables (n_s, r) that we have placed on the B25 plane. It does not also predict the correct fine-structure constant $\alpha_{\text{em}}^{-1} = 137.036$, the correct mass ratio m_n/m_e, the correct scalar amplitude $A_s = 2.096 \times 10^{-9}$, or any of the other sixty-four zero-parameter dimensionless constants of the Physics Ledger [22] (67 entries in total minus the three inflationary observables already named)—those are inputs to the α-attractor theory, not outputs. The categorical framework issues all of them from the same kernel.

In short: the overlap at $(n_s, r) = (0.96491, 0.01357)$ is real, but it is a single-point coincidence between a one-parameter family and a zero-parameter structural prediction. The post-cancellation CMB programme on the 2028–2032 window (Simons Observatory expanded SATs plus the South Pole Observatory, with *LiteBIRD* joining on its revised 2033+ schedule) will either drive the α-attractor posterior to a tight neighbourhood of $(\alpha, N_\star) \approx (3.67, 57)$—which the categorical

Family	Approx. locus (at $N_\star$)	n_s	r
Starobinsky R^2 ($N_\star = 51$)	$(\simeq 1 - 2/N_\star, 12/N_\star^2)$	$\simeq 0.9602$	$\simeq 0.0046$
Higgs inflation ($N_\star = 55$)	$(\simeq 1 - 2/N_\star, 12/N_\star^2)$	$\simeq 0.9636$	$\simeq 0.0040$
Monomial ϕ^p ($N_\star \approx 50, p = 2/3$)	$(\simeq 1{-}(p{+}2)/(2N_\star),\ 4p/N_\star)$	$\simeq 0.9733$	$\simeq 0.0533$
α-attractor (broad)	broad band	0.96–0.98	10^{-4}–10^{-2}
Category τ	$(1 - 2/57,\ \iota_\tau^4)$	**0.96491**	**0.01357**

framework has already committed to—or it will drive the posterior elsewhere, at which point the α-attractor family will absorb the displacement and the categorical framework will not. The discriminating power of the comparison is not in the present overlap; it is in the structural rigidity of the categorical commitment under future posterior movement.

4.4 The discrete-choice budget: fit-space engaged

A reader familiar with fit-space / numerology concerns may object: surely any (n_s, r) near $(0.965, 0.01)$ can be hit by some clever algebraic combination of π and e, and the agreement here is no more than coincidence. This subsection takes the objection seriously, in the spirit of Tegmark's "anthropic-fit" critique [47], the Ijjas–Steinhardt–Loeb anti-inflation fit-space argument [27], and Kinney's subsequent review of the fit-space versus landscape distinction in inflationary theory [31].

The per-observable discrete-choice budget.. The categorical closing identities at L_1 are built from a deliberately finite alphabet of structural ingredients. For any single dimensionless observable $\mathcal{O}$, the programme's fit space for that observable consists of choices drawn from:

- a *rational-prefactor library* of approximately 50 simple fractions with small numerator and denominator $(1, 2/3, 121/225, 1 - 1/3, \ldots)$—the library is bounded by the requirement that every prefactor name a rational-number product of kernel structure constants;
- an *exponent library* of approximately 10 small signed integers $\{-7, -4, -2, 2, 4, 7, 18\}$ and a few compound exponents ($\dim(T^2) \cdot 2, 2 \cdot \dim(\tau^3)$), bounded by the requirement that the exponent name a structural count from the fibration;
- a *window-sum library* $W_k(n)$ indexed by pairs (k, n) of small positive integers, each realising one attested window of the continued-fraction expansion of ι_τ. Across the Physics Ledger the attested values are $W_3(3) = 17$, $W_3(4) = 5$, $W_4(3) = 18$, and $W_5(3) = 19$; each is used in a specific cross-sector role and the choice of window is not free once the closing identity has committed to a (k, n) pair;
- a small *correction-kernel library* (such as the $\kappa_D = 1 - \iota_\tau$ and $\kappa_\omega = \iota_\tau/(1 + \iota_\tau)$ correction kernels of Book IV) that either appears or does not in a given closing identity.

A fair upper estimate for the per-observable discrete-choice budget, taking the cartesian product of these libraries, is of order 10^6–10^7—consistent with the estimate an adversarial reviewer would derive from combinatoric counting alone. Taking the larger figure, a single match on (n_s, r) at the observed $\sim 10^{-3}$ fractional precision is *not* surprising on its own: it sits comfortably within the expected single-observable fit-space of $\sim 10^7$.

Why the joint prediction is the actual claim.. The programme's defence is not that any single match on (n_s, r) is decisive—it is that the *same* ι_τ and the *same* discrete-choice budget, applied across 67 dimensionless observables in Book IV and Book V, lands every one of them inside its CODATA 2018 / Planck 2018 / PDG 2024 measurement window at the stated precision tier. The full per-row residual tabulation is in the Physics Ledger [22]; the inflationary rows (n_s, r, A_s) are three of those 67. Under the naive independence assumption, the probability of a random 67-observable draw from a 10^7-wide per-observable fit space landing every observable inside its measured window is astronomically small, but the naive assumption is not what the programme claims. The 67 residuals share structural integers (3, 4, 19, 57), a common master input (ι_τ), and repeated rational prefactors (notably $121/225$ reused across the electroweak and inflation sectors); they are therefore strongly coupled. A provisional internal count of the number of *independent* structural choices entering the Physics Ledger closings gives an effective rank in the neighbourhood of 10–15, not 67—reviewers R3 and R4 both surfaced this point, and the programme has accepted it as the honest upper bound on the discriminating power of the joint residual surface.

What the programme therefore claims is narrower than "67 independent residuals all land inside their windows." It is: *a ten-to-fifteen-parameter structural commitment reproduces the measured values of every one of* 67 *dimensionless observables at the stated precision, with no per-observable tuning*. That is meaningfully stronger than a one-parameter posit—a single-parameter Ansatz does not have the rank to reach 67 windows at all—but it is weaker than the naive "product of independent 10^{-7}'s" reading a casual audit might extract, and the present paper does not claim the naive product.

What the programme is *not* claiming.. Three honest concessions:

1. The inflationary pair $(n_s, r) = (1-2/57,\ \iota_\tau^4)$, considered on its own and without the rest of the Physics Ledger, does *not* by itself discriminate the categorical framework from an α-attractor at $\alpha \approx 3.67$, $N_\star = 57$ (see §4.3). The discrimination rests on the joint structural surface at its honest effective rank ($\sim$10–15), not on a naive 67-observable product.
2. Individual exponents (such as $57 = 3 \cdot 19$) are *constrained* by structural integer choices, not *forced* by a uniqueness theorem outside the programme's internal registry. An external observer is entitled to ask whether the specific choice $W_5(3)$ (as opposed to $W_3(4) = 5$ or $W_4(3) = 18$) is principled; the programme's reply, spelled out in §3.3, is that $W_5(3)$ is the coarsest window whose closure contains the kernel's long tooth $a_4 = 13$ and that the same choice governs multiple cross-sector identities.
3. The per-observable fit-space estimate of $\sim 10^7$ is an upper bound, not a precise calibration; a genuine fit-space calculation across the full Physics Ledger is future work [22, Ch. 58a]. A minimum-viable first instalment—a fit-space rank-audit restricted to the electroweak sector—is in preparation as a companion Tier-1 note and will be the first quantitative check on the effective-rank estimate used here.

Where the discriminator lives.. The discriminating power of the programme's claim is therefore *not* in the present overlap with α-attractor $\alpha \approx 3.67$. It is in two places: (a) in the joint structural surface of the Physics Ledger at its honest effective rank of $\sim$10–15 (quantified in the forthcoming electroweak-sector fit-space companion note), and (b) in the behaviour of the categorical point under future posterior movement (see §4.3). The post-cancellation CMB programme on the 2028–2032 window (Simons Observatory expanded SATs + South Pole Observatory, with LiteBIRD joining on its revised 2033+ schedule) is the test of (b); the Physics Ledger residual audit is the standing test of (a).

5. FALSIFICATION SURFACE

The categorical pair $(n_s = 1-2/57,\ r = \iota_\tau^4)$ is named in the programme's Falsification Pack as entries N9 and N10 [21]. The falsification reach is tied to the scheduled 2025–2035 CMB programmes.

5.1 N9: tensor-to-scalar ratio on the 2028–2032 window

Landscape after the July 2025 CMB-S4 cancellation.. The previous version of this note—and the v1 Falsification Pack [21]—cited the CMB-S4 design sensitivity $\sigma(r) \approx 0.001$ [12] as the primary 2028–2032 lever on the tensor-to-scalar ratio. That design is no longer on the schedule: the joint DOE / NSF decision of 11 July 2025 [11, 7] terminated the CMB-S4 project following a reduced-scope re-baselining exercise of 4 June 2025 [13]. The *Panta Rhei* programme takes that decision as a point of scholarly disclosure and retires CMB-S4 as an assumed component of the 2028–2032 sensitivity stack. The post-cancellation landscape on the same window comprises:

- the *Simons Observatory* with an expanded small-aperture-telescope configuration [42], forecast to reach $\sigma(r) \approx 0.003$ on its nominal 5-year programme, tightening to roughly 0.002 with delensing from the *Simons Observatory* large aperture telescope [43];
- the *South Pole Observatory*—the joint *SPT-3G+* upgrade of *SPT-3G D1* [45] and the *BICEP Array* [6]—retained at the South Pole after the CMB-S4 cancellation [32];
- the *LiteBIRD* satellite, whose JAXA formal launch window slipped from JFY2032 to JFY2033+ under

the KDP2 approval of December 2025 [29, 36], with a Phase B2 extended study through JFY2036.

The realistic post-cancellation sensitivity on the 2028–2032 window is therefore $\sigma(r) \in [0.002, 0.003]$, driven by Simons Observatory plus the South Pole Observatory, with *LiteBIRD* joining on the revised timeline; an optimistic reach of $\sigma(r) \in [0.001, 0.002]$ becomes available once LiteBIRD is on-sky. At $\sigma(r) = 0.003$, the categorical commitment $r = 0.01357$ is a $\sim 4.5\sigma$ detection against the null and a $\sim 3\sigma$ discriminator against the Starobinsky / Higgs locus $r \approx 0.004$; at $\sigma(r) = 0.001$ (optimistic end-of-decade), it tightens to $\sim 14\sigma$ against null and $\sim 10\sigma$ against Starobinsky / Higgs. The honest headline is the realistic end: N9 is a ~ 3–5σ discriminator on the 2028–2032 window against the slow-roll families B25 favours.

Formally, N9 commits to:

- **Required detection.** A positive detection of r consistent with $0.01357 \pm \sigma_r$ on the 2028–2032 window, with σ_r the realised joint Simons Observatory + South Pole Observatory (+ *LiteBIRD*, when available) precision.
- **Falsifier I.** A measurement of r centred in the R^2/Higgs region ($r \simeq 0.004$–0.005) at $\geq 5\sigma$ resolution.
- **Falsifier II.** A measurement of r centred in the monomial region ($r \gtrsim 0.02$–0.03) at $\geq 5\sigma$ resolution.
- **Framework-terminal.** A measurement excluding r in the range $[0.010, 0.018]$ at $\geq 5\sigma$ forces a retraction of V.P136.

Intermediate annual targets.. Because the post-cancellation schedule no longer delivers a single end-of-decade data drop, the programme pre-registers the following intermediate targets so that the falsification posture is reportable annually rather than only once:

- **N9′ (SO first-season results, target ∼2027).** Under the *Simons Observatory* nominal 2-year early sensitivity $\sigma(r) \sim 0.005$, the categorical r is a $\sim 2.7\sigma$ separation from the null, and a $\sim 1.8\sigma$ separation from the Starobinsky / Higgs locus.
- **N9″ (SO + SPT-3G+ + BICEP Array, target ∼2030).** $\sigma(r) \sim 0.003$, at which point N9 becomes a $\sim 4.5\sigma$ detection and a $\sim 3\sigma$ slow-roll-alternative discriminator.
- **N9‴ (+ LiteBIRD, target ∼2035).** $\sigma(r) \sim 0.001$, at which point N9 converges to the realised end-state discrimination.

These are not additional falsifiers; they are intermediate reporting points for the same single commitment $r = \iota_\tau^4$.

5.2 N10: spectral index under CMB + near-future BAO

The CMB-only precision on n_s will improve to the $\sigma(n_s) \in [0.0015, 0.0020]$ band on the 2028–2032 window with the post-cancellation stack (*Simons Observatory*, South Pole Observatory, and *LiteBIRD* joining on the revised JFY2033+ schedule [29]) combined with Planck *PR4*. The BAO lever is scheduled to tighten through DESI DR2 [16], DESI Year 3 (target $\sim$ 2027), and the Euclid and Roman surveys. Joint ACT–SPT–Planck lensing [37] contributes an independent lever on the CMB side. N10 commits to:

- **Required consistency.** $n_s = 0.96491 \pm \sigma_{n_s}$ under the joint CMB + end-state-BAO posterior.
- **Falsifier.** A measurement of n_s lying outside $[0.955, 0.975]$ at $\geq 5\sigma$, under any well-established CMB + BAO combination. (Under the current CMB+DESI-BAO combination B25 report $n_s = 0.9728 \pm 0.0029$, which lies inside this band; the present categorical point sits at $\sim 2.7\sigma$ from the central value, and will either move into tension or toward consistency as the BAO and CMB levers tighten by 2028–2032.)
- **Framework-terminal combined with N9.** Simultaneous $\geq 5\sigma$ exclusion of both $r = \iota_\tau^4$ and $n_s = 1 - 2/57$ is a cosmology-sector failure of the kind itemised in the programme's *Three framework-terminal scenarios* (Cosmology tier-A failure).

Intermediate annual targets.. As with N9, the post-end-of-2025 schedule for n_s no longer delivers a single end-of-decade data drop; it delivers a sequence. The programme pre-registers the following intermediate targets so that the N10 posture is reportable annually rather than only once:

- **N10′ (DESI DR2 + ACT DR6 + SO first season, target ∼2027).** Under the joint CMB+BAO precision $\sigma(n_s) \sim 0.0025$ available once DESI DR2 [16] and ACT DR6 [1] combine with Simons Observatory early data, the categorical $n_s = 0.96491$ sits at $\sim 1.1\sigma$ below the current B25 CMB+DESI-DR1 central value 0.9728, and at $\sim 0.6\sigma$ below the CMB-only B25 central value 0.9682. The 2027 report is *directional*: either the DR2 shift moves the posterior toward the categorical value (softening the present $\sim 2.7\sigma$ CMB+BAO tension) or away from it (hardening).

- **N10″ (DESI DR3 + SO full + ACT–SPT–Planck lensing, target ∼2030).** Under $\sigma(n_s) \sim 0.0020$, the categorical commitment becomes a ~ 1.5–2.5σ separation from whichever central value the joint posterior converges to. This is the first window in which N10 is a genuine discriminator rather than a consistency report: a central value pulled to ≥ 0.972 at $\sigma = 0.002$ displaces $n_s = 0.96491$ at $\geq 2.5\sigma$; a central value pulled to ≤ 0.968 puts the categorical commitment back inside the 1σ band.
- **N10‴ (+ Euclid + Roman BAO + LiteBIRD, target ∼2035).** Under the end-state $\sigma(n_s) \in [0.0015, 0.0020]$ from the full Euclid/Roman BAO lever [40] combined with *LiteBIRD* [36] and the post-cancellation CMB stack, N10 converges to the realised end-state test.

As with N9′–N9‴, these are not additional falsifiers; they are intermediate reporting points for the same single commitment $n_s = 1 - 2/57$. The falsifier remains: n_s outside $[0.955, 0.975]$ at $\geq 5\sigma$ under any well-established CMB+BAO combination.

Remark 5.1 (Over-specification is deliberate). Both N9 and N10 are over-specified in the sense of the Falsification Pack's posture: the programme publishes the commitments before the measurements, accepts that any one tier-A miss suffices to refute the sector, and refuses to interpret a measurement inside either window as "not really part of the framework" after the fact. The two entries exist separately so that a single-axis failure is recorded as a single-axis failure and does not get laundered into a two-axis tension.

5.3 Pre-registration, what it means here

The word "pre-registration" is used in this paper in a specific sense that a careful reader is entitled to audit. Three points of honest disclosure:

1. **Status at time of writing (April 2026).** The B25 measurement of n_s under CMB+DESI BAO ([3], arXiv:2512.10613, posted 2025-12-12) was already public at the time the present paper was drafted. The categorical commitments $n_s = 1 - 2/57$, $r = \iota_\tau^4$, and $A_s = (121/225)\,\iota_\tau^{18}(1 - \iota_\tau^3/3)$ were however *antecedent* to this B25 result: each is written into the Panta Rhei Book V (first edition draft, 2025Q3) and into the Physics Ledger [22] as a closed-form identity with a `TauLib` module pointer. The present paper does not claim pre-registration with respect to B25; it claims pre-registration with respect to the *near-term* CMB programmes (Simons Observatory, the South Pole Observatory's SPT-3G+ / BICEP Array configuration, and *LiteBIRD* on its revised JFY2033+ schedule) that will report through the 2028–2032 window and into the mid-2030s, and with respect to future CMB+BAO combinations.
2. **Audit trail.** A reader wishing to verify the antecedence claim in (1) can inspect the pinned Panta Rhei monograph release in the 1st-edition-archive (`PantaRhei-1stEd`, tag `v1.0`, posted 2025-09) and the TauLib repository (`github.com/Panta-Rhei-Research/taulib`) at its 2025Q3 commit history, where the relevant closed-form identities were already present. The present paper does not treat these references as author-attested; they are auditable on the public record.
3. **What is and is not pre-registered.** The 67 dimensionless observables of the Physics Ledger are pre-registered in this sense. A subset have already been measured (fine-structure constant, m_p/m_e, n_s under CMB-only, etc.); for those, "pre-registration" means only that the closed-form commitment was written into Book V prior to examining the measurement, not that the measurement was unknown. Another subset (notably r, the A_s/n_s relation, the n_s-BAO shift, neutrino-sector observables) are genuinely pre-registered in the stronger sense: the authoritative measurement will come after this paper's submission, and the categorical commitment is on the record in advance.

The distinction between the two senses of "pre-registered" is the one that matters for honest discrimination between a theory and a fit. The categorical inflation commitments belong to the second sense with respect to the Simons Observatory / South Pole Observatory / *LiteBIRD* programme, and to the first sense with respect to the already-measured B25 CMB-only n_s.

6. DISCUSSION

6.1 What B25 and Category τ agree on

Three agreements are worth stating plainly.

1. **Rapid expansion at early times.** Both frameworks agree that the chart-level readout at early times reproduces an approximately exponential scale factor—whether generated dynamically by an inflaton, as in B25's model class, or structurally by the near-saturation regime of the τ-Einstein equation, as in the categorical framework.
2. **Small r is favoured.** B25's $r < 0.034$ (95%) bound and the categorical commitment $r = \iota_\tau^4 \approx 0.014$ both require r small. Neither framework expects a large-r detection near 0.1.
3. **The tilt $n_s < 1$ is real.** Both frameworks require the primordial scalar power spectrum to be red-tilted, with n_s near 0.965. The $n_s < 1$ datum is a framework-independent observational fact; B25 and Category τ only disagree on its source.

6.2 Where they disagree

The disagreement is threefold.

1. **The existence of the inflaton.** B25 assumes a single-field slow-roll inflaton; the categorical framework forbids it (Corollary 3.2).
2. **The origin of n_s.** In B25's families, n_s is a smooth function of a continuous $N_\star$; in the categorical framework, n_s is a rational $1 - 2/57$ fixed by the kernel.
3. **The origin of r.** In B25's families, r is a slow-roll quantity $\propto 1/N_\star^2$ (concave potentials) or $\propto 1/N_\star$ (convex potentials). In the categorical framework, $r = \iota_\tau^4$ is a fiber-dimensional suppression: geometric, not potential-shape-dependent.

6.3 What the near-term measurements will decide

The categorical prediction sits inside the present CMB-only contour on n_s at $\sim 1\sigma$ and outside the CMB+BAO contour at $\sim 2.7\sigma$. It is below the present r bound by a factor of ~ 2.5. The 2028–2032 window, under the post-cancellation CMB programme (Simons Observatory expanded SATs + the South Pole Observatory, with *LiteBIRD* joining on the revised JFY2033+ schedule), brings r into the $\sigma \sim 0.002$–0.003 regime realistically and $\sigma \sim 0.001$–0.002 optimistically once *LiteBIRD* is on-sky; n_s tightens in parallel to $\sigma(n_s) \in [0.0015, 0.0020]$ combined with ongoing DESI and the Euclid / Roman BAO levers. At those sensitivities the categorical commitment becomes a ~ 3–5σ discriminator against the Starobinsky / Higgs and monomial families.

There are three outcomes the programme commits to in advance:

1. **Convergent confirmation.** r settles near 0.014 and n_s settles near 0.9649 under the end-state CMB + BAO posterior. N9 and N10 are satisfied; the programme's inflation sector is empirically validated.
2. **Starobinsky/Higgs confirmation.** r is measured in the $[0.003, 0.006]$ band at $\geq 5\sigma$. The categorical N9 commitment fails; the programme's fiber-dimensional-suppression derivation of r is retracted.
3. **Monomial confirmation.** r is measured in the $[0.02, 0.04]$ band at $\geq 5\sigma$. The categorical N9 commitment fails in the opposite direction; same retraction.

In either of outcomes (2) or (3), the programme's cosmology branch is refuted, and the B25 slow-roll reading is vindicated on its own terms. In outcome (1), the categorical pre-registration carries a piece of evidence that the B25 model-space reading does not cover—not because B25 erred, but because the categorical prediction was not in the model class B25 considered.

Annual-reporting posture.. Because the post-cancellation schedule delivers a sequence of data drops rather than a single end-of-decade result, N9 and N10 are now reportable at three intermediate points each (§5.1, §5.2): N9′/N10′ (target ~2027, *Simons Observatory* first-season plus DESI DR2 and ACT DR6), N9″/N10″ (target ~2030, joint SO + South Pole Observatory + DESI DR3 + lensing), and N9‴/N10‴ (target ~2035, *LiteBIRD* joining the stack, plus Euclid and Roman). The programme commits to an annual status note against these targets rather than a single end-of-decade verdict. Each intermediate report will state (a) the realised $\sigma(r)$ and $\sigma(n_s)$, (b) the distance in σ from the categorical central values, and (c) whether any of the framework-terminal exclusion conditions have been met.

6.4 Relation to non-inflaton early-universe proposals

A reader familiar with the wider early-universe-alternatives literature will note that the categorical framework is not the only programme that rejects the single-field slow-roll inflaton. A partial accounting of neighbouring non-inflaton proposals, and the specific points of contact with and of departure from each:

- **Ekpyrotic / cyclic scenarios.** Khoury, Ovrut, Steinhardt, and Turok [30], Lehners [33], and Ijjas and Steinhardt [26] have long argued that a contracting phase with a steep negative potential reproduces the observed n_s without an inflaton, and predicts a very small r (typically $r \lesssim 10^{-4}$), well below the categorical $r = \iota_\tau^4 \approx 0.014$. The categorical framework agrees with the ekpyrotic programme on the rejection of single-field inflation but disagrees on the direction of r: a detection at $r \approx 0.014$ by the post-cancellation CMB stack (Simons Observatory + South Pole Observatory, and eventually *LiteBIRD*) would simultaneously falsify ekpyrotic scenarios and confirm the categorical commitment, while a non-detection of r once combined sensitivity reaches $\sigma(r) \sim 0.001$ would be simultaneously consistent with ekpyrotic and refutative of the categorical framework.
- **String-gas cosmology.** Brandenberger and Vafa [10], Nayeri, Brandenberger, and Vafa [39], and Brandenberger [9] derive near-scale-invariant perturbations from a Hagedorn-phase thermal ensemble of strings on a compact background. The prediction for r depends on the string-scale dynamics and is typically comparable to or larger than slow-roll inflation's; the spectral index tilt arises from the compactification radius rather than a potential slope. The categorical programme shares the structural—rather than dynamical—origin of the tilt with string-gas cosmology, but commits to a specific rational $n_s = 1 - 2/57$ that string-gas does not.
- **Non-singular bouncing cosmologies.** Battefeld and Peter [4] and Brandenberger and Peter [8] survey the wider bouncing-cosmology literature. Bouncing models replace the initial singularity with a finite-curvature bounce; their predictions for (n_s, r) cover a wide range. The categorical τ-Einstein equation is also bouncing-compatible in its regime-invariant formulation, but its specific $(n_s, r) = (1 - 2/57,\ \iota_\tau^4)$ commitment is stricter than the general bouncing-cosmology posterior.
- **Pre-big-bang scenarios.** Gasperini and Veneziano [24] posited a duality-symmetric pre-big-bang phase in low-energy heterotic string theory. The predicted tensor spectrum is blue-tilted, distinctively different from both the slow-roll expectation and the categorical commitment; a detection of $r \approx 0.014$ with a red tilt would disfavour pre-big-bang.
- **Galilean genesis.** Creminelli, Nicolis, and Trincherini [14] proposed a NEC-violating starting phase driven by a conformal Galilean Lagrangian. Predictions for (n_s, r) depend on the specific construction; the categorical framework neither shares nor contradicts the genesis mechanism but does differ on the specific numerical commitment.

The categorical framework is therefore *one* of several non-inflaton early-universe proposals in the current literature. What distinguishes it from the others is not the rejection of the inflaton—which most of them share—but the commitment to a specific pre-registered rational point $(n_s, r) = (1 - 2/57,\ \iota_\tau^4)$ with zero continuous free parameters, and the embedding of that commitment in a 67-observable calibration cascade (Physics Ledger, §4.4) in which the inflationary pair is tied to the rest of the programme's fundamental-physics predictions. The other proposals above are inflation-sector-only; the categorical framework's cosmology sector carries a dependence on the kernel that the other proposals do not.

6.5 The wider question: is "inflation" about an inflaton, or about an observation?

A final remark about framing. B25's title—*Inflation at the End of 2025*—aligns with a field-wide convention in which "inflation" names both the phenomenological signature (nearly scale-invariant gaussian primordial perturbations with n_s near 0.965 and small-to-moderate r) and the specific dynamical mechanism (a single-field slow-roll inflaton). The two usages are not the same. The categorical framework agrees with the first and rejects the second. This note is, in part, a plea for a convention that keeps the observational signature and the dynamical mechanism separable: the signature is what the data measure, the mechanism is what a theory posits. A pre-registered alternative mechanism that reproduces the signature is a falsifiable scientific proposal on its own terms—whether or not it fits inside the historical inflaton taxonomy.

7. CONCLUSION

B25 report the end-of-2025 constraints on (n_s, r): the CMB alone yields $n_s = 0.9682 \pm 0.0032$ with $r < 0.034$ at 95%, and the addition of DESI BAO shifts the spectral index to $n_s = 0.9728 \pm 0.0029$. Within the single-field slow-roll model class, their analysis disfavours Starobinsky and Higgs inflation, and favours monomial and α-attractor families.

The *Panta Rhei* programme's pre-registration places a single point in the (n_s, r) plane with zero continuous free parameters: $n_s = 1 - 2/57 \approx 0.96491$ and $r = \iota_\tau^4 \approx 0.01357$. The point is consistent with the CMB-only constraint at $\sim 1\sigma$, sits $\sim 2.7\sigma$ low under CMB+BAO, and is well below the current r bound. It is a genuine alternative to B25's slow-roll families: it lies in an intermediate region between Starobinsky/Higgs (too-low r) and monomial (too-high r), and its algebraic structure is fiber-dimensional, not potential-shape-dependent.

Whether that alternative survives is a matter for the post-cancellation CMB programme on the 2028–2032 window (Simons Observatory expanded SATs [42] plus the South Pole Observatory's SPT-3G+ / BICEP Array configuration [32, 45]) and—on its revised JFY2033+ schedule [29, 36]—the *LiteBIRD* satellite, in a landscape restructured by the July 2025 termination of CMB-S4 [11, 7]. The categorical commitment is stated in advance, on the record, and the outcome is binary: either r settles near 0.014, or the programme's inflation sector is refuted. We welcome the measurement either way.

8. PRE-PUBLICATION SIMULATED PEER REVIEW

Before circulation, this note was put through a pre-publication simulated peer-review panel in the same spirit as the *Panta Rhei Conspectus* v1.1 revision [20]. Four frontier-expert reviewers were convened, each instructed to read the v2 manuscript at the level of scrutiny appropriate to a submission to a leading cosmology or foundations-of-physics journal, and to return a structured report with verdict, blockers, and prioritised punch-list items. The four reviewer roles:

- **R1, slow-roll inflation and CMB observables.** Tasked with auditing every numerical claim against the standard slow-roll taxonomy (exponent conventions in monomial and α-attractor loci, ε/r linearity, sigma estimates against the post-cancellation CMB sensitivity stack—Simons Observatory, South Pole Observatory, and *LiteBIRD* on its revised JFY2033+ schedule).
- **R2, categorical foundations and type-theoretic formalism.** Tasked with auditing the Category-τ semantics, the ι_τ arithmetic, the continued-fraction expansion and window-sum indexing convention, and the pointers to Lean-certified identities in TauLib.
- **R3, adversarial skeptic (fit-space / numerology critique).** Tasked with applying the Tegmark / Ijjas–Steinhardt– Loeb / Kinney fit-space critique to the present manuscript and producing the strongest possible objection a sceptical reviewer would raise.
- **R4, alternative early-universe proposals.** Tasked with auditing the paper's "non-inflaton" framing against the ekpyrotic, string-gas, bouncing, pre-big-bang, and genesis literatures, and identifying where the categorical framework sits in that wider map.

All four reviewers returned a verdict of *MAJOR REVISIONS REQUIRED*. A chair synthesis converged on a tier-ordered punch list of (Tier A) nine concrete numerical / citation errors, (Tier B) seven structural gaps (fit-space engagement, reheating / pivot disclosure, honest-retreat language on perturbation derivations, non-inflaton-alternatives subsection, pre-registration-timestamp disclosure, Conspectus v1.1 fidelity, peer-review retrospective), and (Tier C) five polish items (N10 band, $\sigma(r)$ softening, figure caption, abstract symmetry, recent-CMB citations).

The present (v3 / April 2026) manuscript addresses all of these items. Appendix C is a finding-by-finding mapping from each item on the chair's punch list to the specific paper location where the fix is applied. The honest-accounting discipline of the *Panta Rhei* programme prefers this form of peer-review transparency over the more common strategy of incorporating reviewer comments silently; the present paper is offered as an example of that discipline on the inflation sector of the programme.

ACKNOWLEDGEMENTS

We thank the authors of [3] for making the end-of-2025 inflation constraint state easily readable. This note would not exist without their careful compilation. We thank the anonymous simulated peer-review panel whose April 2026 findings produced the revised *Panta*

Rhei Conspectus v1.1 [20] (and, by extension, the honest-accounting discipline of this paper). A separate four-reviewer simulated panel convened specifically for this note is documented in Section 8 and Appendix C; its findings are the direct basis of the v1.1 revision reflected in the present manuscript. Any remaining errors in the categorical reading of B25 are ours alone.

A. THE $L_0 \to L_4$ CALIBRATION-CASCADE POSITION OF (r, n_s, A_s)

The present note has treated r, n_s, and A_s as pre-registered numerical commitments. A reader inspecting the programme across multiple observables may ask where these three quantities sit in the wider calibration cascade that Book V organises. This appendix answers that question in three short passes.

A.1 The five cascade levels

The *Panta Rhei* calibration cascade [22, Ch. 58a] is the architectural spine of the programme's numerical commitments. It consists of five named levels:

A.2 Where (r, n_s, A_s) sit in the cascade

All three inflationary observables are L_1 **entries**. None of them requires the m_n anchor.

The n_s identity uses two structural integers: $\dim(\tau^3) = 3$ (the fibration dimension) and $W_5(3) = 19$ (the kernel window modulus). The r identity uses one structural integer: $2 \cdot \dim(T^2) = 4$. The A_s identity uses the rational prefactor $(121/225) = (11/15)^2$ inherited from the fine-structure closing identity $\alpha_{\text{em}} = (121/225)\,\iota_\tau^4$ of Book IV [18, Ch. 17]. All three are closed-form in ι_τ alone.

What "N11 (consistency-only)" means.. A_s is not one of the two variables of the B25 plane; it is a third, orthogonal observable whose Planck-2018 value $(2.100 \pm 0.030) \times 10^{-9}$ is already well measured. Entry N11 of the Falsification Pack [21] therefore registers A_s as a *consistency* observable rather than a discriminating prediction: a measurement displacing A_s by more than 3σ from 2.096×10^{-9} would retract V.D253 / V.T198 in the same way N9 retracts V.P136 and N10 retracts V.T302 (see the scope-of-Lean-certification paragraph in §3.7 for the status of each), but the near-term experimental programmes relevant to the 2028–2032 window (Simons Observatory, the South Pole Observatory, and *LiteBIRD* on JFY2033+) are not the principal instruments for a next-generation A_s determination. N11 is listed here so that all three L_1 inflationary observables carry an explicit falsifier.

A.3 Three consequences for the B25 comparison

(i) No post-hoc rescue via the m_n anchor.. A reader skeptical of the categorical prediction may wonder whether an unfavourable measurement of n_s or r could be absorbed by revising the m_n anchor downstream. The cascade position forbids this: n_s, r, and A_s are all L_1 identities and do not depend on m_n. Revising m_n changes SI-bearing readouts at L_3 (the Hubble constant, Newton's constant, etc.); it does not shift the inflationary observables at all. The commitment is therefore fully exposed to the 2028–2032 falsification window.

(ii) Joint discriminating power.. Because the three inflationary observables share the same L_0 input (the master constant ι_τ) but use three different structural integers (4, 57, and 18) and one different rational prefactor (121/225), a simultaneous fit to all three carries more information than a fit to any one. A measurement scenario that hits the categorical r but misses n_s, or vice versa, is a multi-observable signal that the fibration structure is mis-specified, not a single-parameter retuning. B25's plane is two-dimensional; the categorical commitment at the L1 layer is three-dimensional (including A_s); the Physics Ledger commitment at L1 spans sixty-seven observables, though their *effective* rank as an independent-constraint surface is smaller, in the neighbourhood of 10–15 (see the effective-rank disclosure in §4.4). The joint posterior is where the programme's pre-registration will be tested hardest, and the forthcoming electroweak-sector fit-space companion paper will quantify the rank explicitly.

(iii) The Hubble readout is L_3, not L_1.. The one cosmological observable in the programme that *does* depend on m_n is the Hubble constant H_0, for which the programme predicts a -120 ppm correction (a boundary-holonomy readout effect) relative to the ΛCDM benchmark. This prediction is at L_3, not L_1. It therefore responds differently to a posterior shift than (n_s, r, A_s) do, and it lives in a different layer of the cascade from the inflationary observables addressed here.

Level	Content	What enters / what closes here
L_0	Algebra of ι_τ	Master constant $\iota_\tau = 2/(\pi + e)$, derived coefficients $\kappa_D = 1 - \iota_\tau$, $\kappa_\omega = \iota_\tau/(1 + \iota_\tau)$, window moduli $W_k(n)$ such as $W_5(3) = 19$. Purely algebraic, no experimental input.
L_1	Dimensionless closing identities	All 67 zero-parameter dimensionless constants: α_{em}, m_n/m_e, Koide Q, Weinberg angle, n_s, r, A_s, etc. Each is a closed-form function of ι_τ alone.
L_2	Single SI anchor (m_n)	One experimental number, the neutron mass $m_n = 1.67492 \times 10^{-27}$ kg, is consumed to set the SI scale.
L_3	SI-bearing readouts via the rescaling functor	Constants with SI units inherit m_n and L_1 identities: m_e, m_p, G, $\hbar$, k_B, ε_0, the Hubble-readout correction.
L_4	Falsification pack N1–N30	The experimental programme: each ledger line pairs with a named measurement or bound on a 2025–2035 timeline.

Observable	Level	Closed form	m_n-dependent?	Falsification entry
r	L_1	ι_τ^4	**no**	N9
n_s	L_1	$1 - 2/(3\,W_5(3)) = 1 - 2/57$	**no**	N10
A_s	L_1	$(121/225)\,\iota_\tau^{18}(1 - \iota_\tau^3/3)$	**no**	N11 (consistency-only; see below)

The Hubble commitment is the subject of a separate companion note; it is not falsifiable on the same timeline as N9 and N10, and this note does not address it.

In sum: the inflationary observables (r, n_s, A_s) are anchored purely at the algebraic / L_1 layer of the calibration cascade, they are exposed fully to the 2028–2032 experimental window, and they cannot be rescued by adjusting the programme's one SI-anchor input.

B. ANNUAL-REPORTING CADENCE AGAINST N9 AND N10

The intermediate pre-registration targets defined in §5.1 (N9′, N9″, N9‴) and §5.2 (N10′, N10″, N10‴) are summarised here for reference. The programme commits to an annual status note against this cadence rather than a single end-of-decade verdict.

Three design notes on this cadence:

1. **These are not new falsifiers.** N9 and N10 each remain single falsifiable commitments ($r = \iota_\tau^4$, $n_s = 1 - 2/57$). The intermediate targets are reporting points, not additional commitments. Any framework-terminal exclusion condition met at an earlier epoch terminates the sector at that epoch, without waiting for the end-state report.
2. **The cadence responds to schedule risk.** The July 2025 CMB-S4 cancellation [11, 7] and the LiteBIRD KDP2 slip [29] each extended the expected falsification horizon. Rather than absorbing that slip as a single delayed verdict, the programme pre-registers the intermediate targets so that progress against the commitment is auditable in 2027, 2030, and 2035.
3. **No post-hoc re-tuning.** If a 2027 directional read shifts the posterior toward the categorical values, the programme does not claim confirmation at 2027; if it shifts away, the programme does not adjust the commitment downstream. The commitment in N9 and N10 is fixed; only the reported distance is updated.

C. PEER-REVIEW FINDINGS AND HOW THIS VERSION ADDRESSES EACH

This appendix is a finding-by-finding mapping from the chair synthesis of the four-reviewer panel (§8) to the paper location where the fix is applied in v3.

Tier A: concrete numerical / citation errors

Tier B: structural gaps

Tier C: polish

The present version (v3) is the result of applying each of the above fixes to the v2 manuscript. The paper and its TauLib-side companion artifacts are offered together as an example of the honest-accounting discipline the Panta Rhei programme commits to in its documentation posture.

Target	Epoch	Driving experiments	$\sigma(r)$	$\sigma(n_s)$	Report content
N9′ / N10′	∼2027	SO first season; DESI DR2 [16]; ACT DR6 [1]	~ 0.005	~ 0.0025	Directional read: posterior shift toward or away from $(\iota_\tau^4,\ 1-2/57)$.
N9″ / N10″	∼2030	SO full + South Pole Observatory (SPT-3G+ / BICEP Array) [32, 45]; DESI DR3; ACT–SPT–Planck lensing [37]	~ 0.003	~ 0.0020	First discriminator-grade report: N9 at $\sim 4.5\sigma$ against null; N10 at 1.5–2.5σ depending on CMB+BAO central value.
N9‴ / N10‴	∼2035	*LiteBIRD* [29, 36]; Euclid; Roman	~ 0.001	~ 0.0015	End-state verdict: $\sim 14\sigma$ against null; CMB+BAO consistency window $n_s = 0.96491 \pm 0.0015$.

#	Finding	Fix in v3
A1	"$156\times$ in ε" arithmetic error (R1).	§3.5 corrected to $\varepsilon_{\mathrm{SR}}/\varepsilon_\tau \approx 10.3$; $\varepsilon = r/16$ is linear in r.
A2	Monomial locus: wrong r formula; should be $r = 4p/N_\star$, not $2p/N_\star$ (R1).	§4.2 table row corrected to $r \approx 0.0533$ at $p = 2/3$, $N_\star \approx 50$.
A3	Continued fraction displayed with phantom "+1" convention (R2, R3).	§3.3: explicit CF$(\iota_\tau) = [0; 2, 1, 13, 3, 1, 1, 1, 42, \ldots]$; equation (4) gives the $\mathrm{W}_5(3)$ decomposition directly.
A4	Acknowledgement mis-cite: Book V instead of Conspectus v1.1 (R2).	Fixed, new bib entry `Conspectus2026` added.
A5	"Eq. (47.xx)" unfilled placeholder (R1, R3).	Replaced with explicit theorem pointer Thm V.T302 (the actual n_s closing identity; in v2 this was mislabelled as V.T191—Sound Horizon—see v3.1 scope-of-Lean-certification note in §3.7).
A6	"N11 (consistency)" undefined (R1).	Appendix A now explains N11 as the consistency-only observable for A_s.
A7	ι_τ called "algebraic" instead of "transcendental" (R3).	§3.1 corrected to "unique closed-form value" with pointer to Thm I.T132.
A8	Starobinsky n_s rounded to 0.9606 instead of 0.9602 (R1).	§2.3 and §4.2 corrected.
A9	67 / sixty-five / sixty-seven count drift (R3).	§6.5 count consistency audit applied; "sixty-four" with explicit $(67-3)$ clarification.

REFERENCES

[1] ACT Collaboration, Thibaut Louis, et al. The Atacama Cosmology Telescope: DR6 power spectra, likelihood, and cosmological parameters, 2025.

[2] Andreas Albrecht and Paul J. Steinhardt. Cosmology for grand unified theories with radiatively induced symmetry breaking. *Physical Review Letters*, 48:1220–1223, 1982.

[3] L. Balkenhol, E. Camphuis, F. Finelli, K. Benabed, F. R. Bouchet, et al. Inflation at the end of 2025: Constraints on r and n_s using the latest CMB and BAO data, December 2025. Submitted 11 December 2025.

[4] Diana Battefeld and Patrick Peter. A critical review of classical bouncing cosmologies. *Physics Reports*, 571:1–66, 2015.

[5] Fedor L. Bezrukov and Mikhail Shaposhnikov. The Standard Model Higgs boson as the inflaton. *Physics Letters B*, 659(3):703–706, 2008.

[6] BICEP/Keck Collaboration, P. A. R. Ade, et al. Improved constraints on primordial gravitational waves using Planck, WMAP, and BICEP/Keck observations through the 2018 observing season. *Physical Review Letters*, 127(15):151301, 2021.

[7] Lee Billings. U.S. ends support for CMB-S4, leaving cosmic inflation search in limbo. Scientific American, July 2025. Coverage of the cancella-

#	Finding	Fix in v3
B1	Fit-space objection engaged only briefly; needs explicit discrete-choice budget (R3, R4).	§4.4 added: per-observable $\sim 10^6$–10^7 budget disclosed; Tegmark, Ijjas–Steinhardt–Loeb, Kinney cited directly.
B2	No reheating / pivot-scale / cross-sector disclosure attached to $N_e = 57$ (R1, R2).	Remark 3.5 added: pivot $k_\star = 0.05$ Mpc^{-1}, $\Delta N_\star \sim \pm 3$ reheating window, Weinberg-sector origin of $W_5(3)$ acknowledged.
B3	Derivation language for r, n_s too strong; monograph perturbation theory not reproduced (R2).	§3.4 and §3.5: "argument-of-record lives in Book V Ch. 47" framing applied; companion paper flagged.
B4	Non-inflaton alternatives (ekpyrotic, string-gas, bouncing, pre-big-bang, genesis) not placed (R4).	§6.4 added with direct citations ([30], [33], [26], [10], [39], [9], [4], [8], [24], [14]).
B5	"Pre-registration" ambiguous as stated (R3, R4).	§5.3 added: distinguishes pre-registration with respect to already-measured B25 n_s versus the near-term post-cancellation CMB programme (Simons Observatory, South Pole Observatory, *LiteBIRD* on JFY2033+); audit trail disclosed.
B6	Conspectus v1.1 fidelity (R2, R3).	Three fixes: §3.1 re-frames Category τ as Syn(T_τ)/Set$[T_\tau]$; §1 epistemic caveat now cites the actual TauLib trust budget (4 custom axioms + 3 Book VII methodological sorries + $\sim 1{,}842$ native_decide leaves); §3.7 per-row residual table added.
B7	Peer-review process not disclosed (R3, R4).	§8 and this appendix added.

#	Finding	Fix in v3
C1	N10 band $[0.960, 0.970]$ already excluded by current CMB+DESI-BAO central (R1).	§5.2 widened to $[0.955, 0.975]$; current $n_s = 0.9728 \pm 0.0029$ situation acknowledged explicitly.
C2	$\sigma(r) \approx 0.001$ stated as single number, no delensing caveat (R1).	§5.1: $\sigma(r) \in [0.001, 0.0015]$; discrimination 9–14σ against null, 6–10σ against Starobinsky/Higgs.
C3	Figure caption did not note schematic nature of ellipses in r direction (R1).	Caption (Fig. 1) updated: "ellipses not drawn to scale in r; actual posteriors consistent with $r = 0$".
C4	Abstract asymmetric: n_s tension quoted, r bound not (R4).	Abstract rewritten to quote both n_s situation and r bound symmetrically; 9–14σ range used.
C5	Recent SPT-3G / ACT DR6 releases not cited; Kinney fit-space review not cited (R1, R3).	New bib entries: `SPT3G2024`, `ACTDR6Louis2025`, `Kinney2018`; cited in abstract and §4.4.

tion of the CMB-S4 next-generation ground-based experiment.

[8] Robert Brandenberger and Patrick Peter. Bouncing cosmologies: Progress and problems. *Foundations of Physics*, 47(6):797–850, 2017.

[9] Robert H. Brandenberger. String gas cosmology: Progress and problems. *Classical and Quantum Gravity*, 28(20):204005, 2011.

[10] Robert H. Brandenberger and Cumrun Vafa. Superstrings in the early universe. *Nuclear Physics B*, 316(2):391–410, 1989.

[11] Adrian Cho. U.S. abandons hunt for signs of cosmic inflation. Science (AAAS) News, July 2025. Report on joint DOE/NSF decision to terminate the CMB-S4 project.

[12] CMB-S4 Collaboration, K. Abazajian, et al. CMB-S4 science case, reference design, and project plan, 2019.

[13] CMB-S4 Project. Revised CMB-S4 project plan, June 2025. Internal project document following reduced-scope re-baselining; programme subsequently cancelled July 2025.

[14] Paolo Creminelli, Alberto Nicolis, and Enrico Trincherini. Galilean genesis: an alternative to inflation. *Journal of Cosmology and Astroparticle Physics*, 2010(11):021, 2010.

[15] DESI Collaboration, A. G. Adame, et al. DESI 2024 VI: Cosmological constraints from the measurements of baryon acoustic oscillations. *arXiv preprint*, 2024.

[16] DESI Collaboration, A. G. Adame, et al. DESI DR2: baryon acoustic oscillations and cosmological constraints, 2025.

[17] Thorsten Fuchs and Anna-Sophie Fuchs. *Panta Rhei, Book I: Categorical Foundations*. Panta Rhei Research, 2nd edition, 2026. ISBN 979-8-3462-1987-6.

[18] Thorsten Fuchs and Anna-Sophie Fuchs. *Panta Rhei, Book IV: Categorical Microcosm*. Panta Rhei Research, 2nd edition, 2026. ISBN 979-8-3462-1990-6.

[19] Thorsten Fuchs and Anna-Sophie Fuchs. *Panta Rhei, Book V: Categorical Macrocosm*. Panta Rhei Research, 2nd edition, 2026. ISBN 979-8-3462-1991-3.

[20] Thorsten Fuchs and Anna-Sophie Fuchs. The Panta Rhei conspectus: A synoptic reading of the seven-book categorical unification, v1.1, April 2026. Post-peer-review revision; v1.1 retrospective at `https://panta-rhei.site/publications/conspectus/`.

[21] Thorsten Fuchs and Anna-Sophie Fuchs. Panta Rhei falsification pack N1–N30: Pre-registered experimental tests, 2025–2035, 2026. Companion volume; available at `https://panta-rhei.site/publications/falsification/`.

[22] Thorsten Fuchs and Anna-Sophie Fuchs. The Panta Rhei physics ledger: Numerical audit trail for the calibration cascade, 2026. Companion volume, 209 pp; available at `https://panta-rhei.site/publications/physics-ledger/`.

[23] Thorsten Fuchs and Anna-Sophie Fuchs. TauLib: A Lean 4 library for the Panta Rhei research program, 2026. Available at `https://panta-rhei.site/verify/`.

[24] Maurizio Gasperini and Gabriele Veneziano. The pre-big-bang scenario in string cosmology. *Physics Reports*, 373(1–2):1–212, 2003.

[25] Alan H. Guth. Inflationary universe: A possible solution to the horizon and flatness problems. *Physical Review D*, 23(2):347–356, 1981.

[26] Anna Ijjas and Paul J. Steinhardt. A new kind of cyclic universe. *Physics Letters B*, 795:666–672, 2019.

[27] Anna Ijjas, Paul J. Steinhardt, and Abraham Loeb. Inflationary paradigm in trouble after Planck2013. *Physics Letters B*, 723(4–5):261–266, 2013.

[28] Renata Kallosh and Andrei Linde. Universality class in conformal inflation. *Journal of Cosmology and Astroparticle Physics*, 2013(07):002, 2013.

[29] Kavli IPMU. LiteBIRD passes JAXA KDP2 review; mission phase B2 study extended. Kavli IPMU News, December 2025. Formal launch window now JFY2033+; Phase B2 extended study through JFY2036.

[30] Justin Khoury, Burt A. Ovrut, Paul J. Steinhardt, and Neil Turok. The ekpyrotic universe: Colliding branes and the origin of the hot big bang. *Physical Review D*, 64(12):123522, 2001.

[31] William H. Kinney. Theoretical implications of detecting gravitational waves. *Proceedings of Science, Frontiers of Fundamental Physics 14*, 2018.

[32] KIPAC. South pole observatory (SPT-3G+ and BICEP Array): science goals and 2028–2032 observing programme. KIPAC / Stanford programme page, 2025. Joint SPT-3G+ and BICEP Array configuration at the South Pole, retained after CMB-S4 cancellation.

[33] Jean-Luc Lehners. Ekpyrotic and cyclic cosmology. *Physics Reports*, 465(6):223–263, 2008.

[34] Andrew R. Liddle and Samuel M. Leach. How long before the end of inflation were observable perturbations produced? *Physical Review D*, 68:103503, 2003.

[35] Andrei D. Linde. A new inflationary universe scenario: A possible solution of the horizon, flatness, homogeneity, isotropy and primordial monopole problems. *Physics Letters B*, 108(6):389–393, 1982.

[36] LiteBIRD Collaboration, M. Hazumi, et al. LiteBIRD satellite: mission status and cosmic microwave background B-mode forecasts, 2024.

[37] Mathew S. Madhavacheril et al. Cosmological constraints from a joint ACT–SPT–Planck lensing reconstruction, 2025.

[38] Jerome Martin, Christophe Ringeval, and Vincent Vennin. Encyclopædia inflationaris. *Physics of the Dark Universe*, 5–6:75–235, 2014.

[39] Ali Nayeri, Robert H. Brandenberger, and Cumrun Vafa. Producing a scale-invariant spectrum of perturbations in a Hagedorn phase of string cosmology. *Physical Review Letters*, 97(2):021302, 2006.

[40] Planck Collaboration, N. Aghanim, et al. Planck 2018 results. VI. cosmological parameters. *Astronomy & Astrophysics*, 641:A6, 2020.

[41] Planck Collaboration, Y. Akrami, et al. Planck 2018 results. X. constraints on inflation. *Astronomy & Astrophysics*, 641:A10, 2020.

[42] Simons Observatory Collaboration, M. Abitbol, et al. Forecasts for an expanded Simons Observatory small-Aperture-Telescope configuration, 2025.

[43] Simons Observatory Collaboration, P. A. R. Ade, et al. The Simons Observatory: Large Aperture Telescope design and sensitivity forecasts, 2025.

[44] SPT-3G Collaboration, L. Balkenhol, et al. Measurement of the CMB temperature power spectrum and constraints on cosmology from the SPT-3G 2018 TT/TE/EE dataset, 2024. updated release.

[45] SPT-3G Collaboration, E. Camphuis, et al. Cosmology from the SPT-3G D1 full-depth field: constraints from the TT, TE, and EE power spectra, 2025. SPT-3G D1 release; input to the South Pole Observatory / SPT-3G+ programme.

[46] Alexei A. Starobinsky. A new type of isotropic cosmological models without singularity. *Physics Letters B*, 91(1):99–102, 1980.

[47] Max Tegmark. What does inflation really predict? *Journal of Cosmology and Astroparticle Physics*, 2005(04):001, 2005.