

Referee Report

on “Wading the String Bog: CMB Birefringence in the Swampland”

Synopsis and overall assessment

The manuscript asks whether a future detection of cosmic microwave background (CMB) birefringence could discriminate between two broad mechanisms and, in turn, test expectations about axions and quantum gravity. The first mechanism uses a propagating ultralight axion coupled through equation (1). Equations (3)–(4) estimate how late that axion must remain mobile, while equations (5)–(8) review axion quintessence and its required field range. The authors then propose a two-stage multi-axion portal: mixing first distributes an electromagnetic anomaly coupling, and the mass matrix (15) is subsequently claimed to create a clockwork hierarchy without exponentially suppressing the light-mode coupling. Equations (17)–(18) and (27)–(29) give the resulting mass, coupling, and polarization estimates.

The second part connects this construction to the Weak Gravity Conjecture (WGC). Equations (30)–(32) apply a magnetic axion-WGC cutoff estimate to infer a lower mass bound that is compared with the upper bound (4). The manuscript then turns to a proposed alternative involving thin vacuum interfaces. After reviewing adiabatic propagation in equations (19)–(25), it solves a planar delta-function interface in equations (33)–(39). The small-jump estimate (40) is used to motivate a rotation of the observed size. Equations (41)–(42) argue that, when the last-scattering surface and the observer occupy fixed initial and final vacua, all intervening wall contributions collapse to an endpoint difference. Finally, equation (43) and Figures 2–3 discuss temperature-polarization observables, reionization, and possible tomographic tests.

The topic is timely, and the manuscript contains several appealing ideas. In particular, separating a coupling-transfer step from a hierarchy-generating step is a useful model-building question. The one-interface scattering problem is presented in a compact form, and the emphasis on combining CMB monopole, anisotropy, redshift, and radio-galaxy information is phenomenologically constructive. The paper is also candid that the WGC input is conjectural and that no exclusion theorem is being claimed.

In its present form, however, the main quantitative chain is not correct. The mass matrix (15) has an exact zero mode and no eigenvalue suppressed by q^{-N} ; its massive eigenvalues remain of order m . Thus equation (18) and the numerical portal built from it do not follow from the displayed theory. The slow-roll formula (5) has the sign of $\dot{w}$ reversed. The WGC analysis conflates the QCD decay constant, site periodicities, an effective light-mode field range, and the electromagnetic coupling scale, and it is not clear why the QCD cutoff should constrain every orthogonal light eigenmode. Finally, a homogeneous rolling axion also produces isotropic endpoint rotation, so the claimed dichotomy “anisotropic means ultralight, isotropic means walls” is not valid. The exact multi-wall problem also does not reduce solely to an endpoint difference once reflection is retained.

These issues affect the central claims in the abstract, Sections 4–7, and the summary rather than isolated details. I therefore recommend **major revision**. A substantially revised manuscript could be suitable for a high-energy theory journal, but the clockwork construction, WGC inference, and observational discriminator must first be rebuilt or appropriately weakened.

Assessment of the logical structure and principal derivations

The canonical slow-roll setup in equations (5)–(8) is an appropriate way to display the field-range problem, but differentiating $w + 1 \simeq 2\epsilon_V/3$ along the forward cosmological evolution gives

$$\dot{w} \simeq \frac{4}{3}H\epsilon_V(2\epsilon_V - \eta_V),$$

not the last expression in equation (5). For $V \propto \phi^p$ this gives $\dot{w} \simeq +(2/3)Hp^3(M_{\text{Pl}}/\phi)^4$, opposite to equation (7). The magnitude-based field-range estimate survives, but the stated dynamics does not.

The two-field positivity bound (12) is basically the determinant condition for the kinetic and mass matrices, and the hierarchical eigenvalues (13) can be obtained in the indicated corner after canonical normalization. This does not establish the many-field proposal. For the $N + 1$ fields actually appearing in equation (15), the vector

$$(\phi_0, \phi_1, \dots, \phi_N) \propto (1, q^{-1}, \dots, q^{-N})$$

is exactly massless. The other eigenvalues are

$$m_k^2 = m^2 \left(q^2 + 1 - 2q \cos \frac{k\pi}{N+1} \right), \quad k = 1, \dots, N,$$

up to the indexing convention. None scales as $m^2 q^{-2N}$. An additional explicit breaking term can lift the zero mode by a suppressed amount, but its location, normalization, anomaly vector, and radiative stability then matter for both its mass and photon coupling.

The interface calculation is useful when interpreted as a fixed-frequency spatial scattering problem. For one helicity, equations (34)–(39) give a transmission amplitude of the form $t_{\pm} = [1 \pm i\Delta\theta/2]^{-1}$, so the transmitted linear polarization is rotated by $\arctan(\Delta\theta/2)$. The small-jump limit therefore contains a factor of one half. That result supports an endpoint description only in the forward, negligible-reflection limit. A sequence of separated interfaces has transfer matrices containing propagation phases and repeated reflections; at second and higher order in the jumps it is not determined only by the initial and final values of θ .

Major comments

1. Equation (15) does not generate the hierarchy asserted in equation (18).

The displayed clockwork quadratic form contains one exact Goldstone mode and N modes with masses between approximately $m(q - 1)$ and $m(q + 1)$. It does not produce a lightest massive eigenstate with $m_{\text{ultra}}/m \sim q^{-N}$. Consequently, the inference $N \simeq 18$ in equation (18), the ultralight state used below it, and the coupling estimates (27)–(29) are not predictions of the Lagrangian shown in the paper.

The authors should provide a complete mass and kinetic matrix including the operator that lifts the exact zero mode. They should diagonalize that matrix, normalize the eigenvectors, and compute the projection of the electromagnetic anomaly vector onto the lifted mode. If the lifting is placed at a clockwork endpoint, its integer charge and the selected endpoint should be explicit; these choices determine whether the mass and coupling are suppressed or enhanced. The indexing should also be corrected: equation (15) sums over $i = 0, \dots, N - 1$ but contains fields through ϕ_N , whereas the surrounding text repeatedly calls this a system of N fields. Finally, the proposed 10^{-33} eV eigenvalue needs a technical-naturalness argument against the same gravitational and instanton effects emphasized elsewhere in the paper.

2. The assumed continuous “virialization” of axion couplings is not a complete compact-axion construction.

The discussion around equations (14)–(17) treats the mass diagonalization as a generic $SO(N)$ rotation and the photon coupling as an arbitrary vector of comparable components. For compact axions, periodicities, instanton charges, and electromagnetic anomaly coefficients form an integer charge lattice. General real rotations do not preserve that lattice, and the direction that solves strong CP cannot simply be declared orthogonal to all other nonperturbative potentials while an unrelated photon vector is distributed democratically without specifying the ultraviolet charge assignments.

At minimum, the authors should give an explicit integer charge matrix for the confining sectors and an electromagnetic anomaly vector, identify the periodic field domain, and show that the QCD direction in equation (14) is a mass eigenvector while the light lifted mode retains the coupling (17). This would also expose whether additional light modes noted below equation (18) are unavoidable and whether their energy density and isocurvature are acceptable. Without such an example, equations (16)–(17) are an ansatz, not a demonstrated portal.

3. The WGC mass window (30)–(32) is not derived for the proposed light eigenmode.

Several physically different quantities are denoted by f : the QCD axion decay constant entering equation (10), the periodicities of the dark sites, the field range of the light clockwork eigenmode, and the effective scale inferred from its photon coupling. They cannot be interchanged in the domain wall bound. This is particularly important because quintessence requires an effective range of order M_{Pl} or larger, whereas the numerical range 10^{11} – 10^{16} GeV inserted below equation (31) was inferred from the photon portal.

There is also a logical tension within the construction. Around equation (14), the QCD direction is arranged to be a mass eigenvector precisely so that it does not mix with the orthogonal dark modes. Section 6 then applies the QCD cutoff to the ultralight modes on the ground that they mix with the QCD axion and inherit its residual discrete symmetry. Electromagnetic anomaly mixing alone does not establish that implication. A magnetic WGC constraint must be applied to the particular axion/three-form pair and its canonically normalized charge, not automatically to every eigenstate carrying a small $F\tilde{F}$ coupling.

The authors should derive the bound in the full charge basis and then project it onto the light eigenmode. All decay constants and charge vectors should be distinguished. If that derivation cannot be supplied, equations (31)–(32) and the claimed collision with equation (4) should be presented only as a speculative possibility rather than evidence for a quantitative WGC tension.

4. The sign of the slow-roll evolution and the early-oscillation estimate must be corrected.

As noted above, the last relation in equation (5) and the last relation in equation (7) have the wrong sign for forward time. A monomial field rolling toward its minimum generally has $w+1$ increasing, not decreasing, within the stated slow-roll approximation. The sign error should be corrected wherever the $w-\dot{w}$ behavior is discussed.

Equation (3) also needs a derivation with the physical integration limits. If $m > H_{\text{LSS}}$, the onset time t_i precedes last scattering, so CMB rotation cannot be integrated from t_i ; it is set by $\phi(t_0) - \phi(t_{\text{LSS}})$. Moreover, the written matter-era solution has $\dot{\phi}(t_i) = -\phi(t_i)/t_i$, rather than the approximately frozen initial condition. The omitted sine solution is comparable when mt_i is order one. A reliable bound should evolve the field across radiation, matter, and late dark-energy domination, impose the initial velocity and abundance, and then evaluate the endpoint difference after last scattering. Equation (4) may remain a reasonable order-of-magnitude guide, but it is not established by the current equation (3).

5. Time-dependent axion propagation and spatial wall scattering must be derived separately.

Equation (21) is obtained by Fourier transforming in time while retaining an axion gradient, and

is then asserted to describe both a homogeneous evolving $\phi(t)$ and a static wall $\phi(z)$. A time-dependent coefficient does not admit the fixed-frequency Fourier ansatz used in that derivation. For a homogeneous axion in FRW, one instead obtains mode equations in conformal time, schematically $A''_{\pm} + [k^2 \pm gk\phi']A_{\pm} = 0$. A static spatial wall is a fixed- ω scattering problem and can produce reflection. Their WKB forward limits are related along a null ray, but the exact problems, boundary conditions, and nonadiabatic effects are different.

The manuscript should give the two derivations independently and state where the geometric-optics identification is used. There is also an internal sign problem between equations (22) and (23): with $H = \omega\mathcal{K} - (g/2)\partial_z\phi\sigma_2$, its integral contains a minus sign, whereas the text inserts a plus sign to obtain the displayed rotation matrix. The physical sign depends on conventions, but the equations must be mutually consistent.

6. The endpoint and isotropy claims are too strong, and equations (41)–(42) miss the rotation factor.

Equations (24)–(25) already show that smooth axion birefringence is an endpoint effect in geometric optics. A homogeneous rolling axion therefore produces an isotropic monopole rotation. Anisotropic rotation arises from spatial fluctuations of the axion at emission and observation, whose amplitude depends on the inflationary spectrum, mass, abundance, and isocurvature constraints. It is not necessarily order one, nor is it a mandatory signature of every ultralight model. Thus the repeated conclusion that an anisotropic signal identifies an ultralight whereas an isotropic signal identifies interfaces is not logically valid.

For walls, equations (39)–(40) imply a small-jump polarization rotation $\Delta\theta/2$ per oriented crossing, while equations (41)–(42) identify the rotation itself with $\int d\theta$ and omit this factor. Beyond leading order the angle is $\arctan(\Delta\theta/2)$ for one wall, and a multi-wall transfer matrix also contains reflected paths and separation-dependent phases. Endpoint cancellation is therefore a leading-order forward-scattering statement, not an exact theorem. The paper should correct the normalization, quantify the order- $\Delta\theta^2$ corrections, and replace “strictly” or “necessarily” isotropic language with conditions on the endpoint vacua, defects, reflection, and fluctuations. A useful comparison would state predictions for the monopole angle and for the angular spectrum $C_L^{\alpha\alpha}$ in each model rather than offering a binary diagnostic.

7. The vacuum jump and anomaly normalization need to be fixed.

Equation (40) assumes $\Delta\phi \simeq f$. For the periodic potential in equation (8), adjacent equivalent minima are separated by $2\pi f$. Depending on the branch structure of the proposed dark sector, a domain wall may connect inequivalent branches with another jump, but that jump must be specified. Together with equation (10), $\Delta\phi = 2\pi f$ gives $\Delta\theta = C_\gamma\alpha_{\text{QED}}$ before the additional factor of one half converting a helicity phase difference to a linear-polarization angle. This differs from the chain of estimates in equations (39)–(40).

The authors should define the periodicity of ϕ , the normalization and periodicity of the electromagnetic θ angle in equation (33), the integer anomaly coefficient, and which vacua the walls connect. The same conventions should be used in equations (1), (9), (10), and (33)–(42). This is necessary before claiming that an order-one anomaly coefficient naturally reproduces the quoted CMB angle.

8. Figures 2–3 do not yet provide a reproducible or correctly interpreted phenomenological calculation.

The manuscript does not state the rotation amplitude used in the curves, the baseline cosmological parameters, the Boltzmann solver and modifications, the visibility history used for “suppressed reionization,” or numerical accuracy. No data or code are supplied. Since C_ℓ^{EB} is signed, a positive logarithmic plot also requires an explanation of whether the absolute value is shown. The legend and captions should identify every curve and all units.

For a spatially uniform post-recombination rotation, the spectra should first be checked against the standard analytic E/B rotation relations, including $C_\ell^{EB} = \frac{1}{2} \sin(4\alpha)(C_\ell^{EE} - C_\ell^{BB})$ when primordial parity-odd correlations vanish. The statement below Figure 3 that high- ℓ modes correspond to short distances that are “directly affected by scattering on domain walls” is not correct: high ℓ denotes small angular scales, while a uniform interface rotation acts on every traversing ray. The transition-redshift dependence is instead controlled by whether a given polarization contribution, especially the reionization contribution, is generated before or after the rotation. The numerical section should be rewritten around this causal interpretation and supplied with enough information for independent reproduction.

9. The manuscript should distinguish new results from synthesis and sharpen its relation to the literature.

Much of the exact interface scattering, Pancharatnam-phase interpretation, and vacuum-interface proposal is summarized from references by one of the present authors. The rolling-axion rotation formula, birefringence tomography, and anisotropy literature are also established. The new contributions appear to be the proposed multi-axion portal, its WGC interpretation, the combined observational narrative, and the illustrative spectra. Because the first two currently have foundational gaps and the spectra are not reproducible, the novel result is difficult to identify.

The revision should explicitly separate review from new derivation, compare the wall scenario with the wider literature on axion defects and interface birefringence, and state which observational predictions differ from existing tomographic analyses. This would make the paper’s significance much easier to assess.

Minor comments

1. Equation (2) quotes a rotation of order 10^{-3} radians, while the 0.3 degree value is 5×10^{-3} radians in equation (26). Use one benchmark or explain the intended order-of-magnitude range.
2. The kinetic sign in equation (11) should be reconciled with equation (6) by stating the metric convention. The positivity condition (12) should read $|\mu^2| < m m_{\text{QCD}}$ if either sign of the off-diagonal entry is allowed.
3. State all approximations behind equation (13), including the role of κ , canonical normalization, and the size of terms discarded in m/m_{QCD} and ε .
4. Inserting $C_\gamma = 1$ in equation (10), $g_{\phi\gamma\gamma} = 10^{-13} \text{ GeV}^{-1}$ corresponds more nearly to $f \sim 10^{10} \text{ GeV}$ than 10^{11} GeV . The arithmetic in equations (27)–(29) should be tabulated with the chosen C_γ and Planck mass convention.
5. Equation (37) sets $A_\mp = 0$ without explicitly saying that a pure $A_\pm$ incident state is being considered. For a general incident polarization both helicity amplitudes are present and transmit independently. The symbol $\tilde{A}_\pm$ in equation (38) should also be defined.
6. Equation (43) is a real-part estimator for C_ℓ^{EB} . It would be clearer to define the ensemble spectrum and then state the finite-sky or full-sky estimator separately, including the convention for complex harmonic coefficients.
7. The summary states that an abundance bound constrains the ultralight decay constant, but no such calculation is presented. Either provide the misalignment abundance and initial-condition bound or remove this new claim from the summary.
8. Several presentation errors should be corrected: the affiliations use the label b twice; the acknowledg-

ment contains corrupted characters; the text refers to a “previous subsection” although no subsections are defined; and the Figure 3 legend appears inconsistent with that of Figure 2.

9. The paper would benefit from a compact notation table distinguishing m_{QCD} , the intermediate site mass, the lifted eigenmass, f_{QCD} , site periodicities, the light-mode effective range, $g_{\phi\gamma\gamma}$, g_{ultra} , and the dimensionless vacuum jump $\Delta\theta$.

Recommendation

The manuscript addresses an interesting intersection of axion model building, cosmic birefringence, and quantum-gravity conjectures, and parts of its phenomenological program could become useful. At present, however, the displayed clockwork theory does not yield the claimed ultralight mass, the WGC bound is not attached consistently to the relevant eigenmode, and the proposed isotropy discriminator is not valid as stated. These are central rather than cosmetic problems. I recommend **major revision** and reconsideration only after the principal derivations and claims have been substantially reworked.