A Primordial-Tensor Null from Boundary Printing: Pre-registering $r_{\text{linear}} = 0$ and a Scalar-Induced Floor

David Elliman
Neuro-symbolic Ltd
dave@neusym.ai
https://neusym.ai

Prediction note: 30 June 2026

Abstract

This short note pre-registers a near-term primordial-tensor prediction from the finite-QEC boundary-printing branch. In this branch inflation is not a smooth de Sitter stretching of an already-existing metric. It is the printing of new boundary cells by a scalar service process. The consequence is sharp: at linear order the printer has no transverse-traceless source, hence

$$r_{\text{linear}} = 0.$$

The often-quoted scale $r \sim 2 \times 10^{-9}$ is not a primordial squeezed tensor amplitude. It is only the expected second-order scalar-induced floor,

$$r_{\text{induced}} = C_{\text{SIGW}} A_s, \qquad A_s = (3/4)\alpha_0^4 \simeq 2.1290 \times 10^{-9},$$

where C_{SIGW} is an external second-order radiation-transfer coefficient, not a fitted substrate parameter. The prediction-grade near-term statement is therefore a null: no primordial B-mode detection at $r \geq 10^{-3}$ after dust, lensing, and systematics are removed. A robust primordial signal at that level would falsify the scalar boundary-printer / no-squeezing branch. The note also records the post-hoc rules: the denominator or branch cannot be adjusted after LiteBIRD or CMB-S4; a hidden squeezed graviton vacuum is a new branch, not this prediction.

Plain summary

Standard inflation predicts that quantum fluctuations of the gravitational field can be stretched into a background of primordial gravitational waves. Those waves would leave a special curl-like pattern, called B-modes, in the polarisation of the cosmic microwave background. The usual summary number is r, the tensor-to-scalar ratio.

The boundary-printing picture says something different. During the boot phase the universe is not stretching an already-existing continuum. It is printing fresh boundary cells. That printer is a scalar counting process. A scalar source does not produce a transverse-traceless gravitational wave at linear order. In the framework's own language, the graviton is the shear phonon of the crystallised substrate; the pre-rigid printing front has compression but no shear rigidity, so there is no squeezed primordial tensor vacuum to amplify.

The prediction is therefore a clean null:

$$r_{\text{linear}} = 0.$$

There can still be a much smaller second-order background induced by scalar perturbations. Its natural scale is the scalar amplitude itself,

$$r_{\text{induced}} \sim A_s \simeq 2.13 \times 10^{-9},$$

far below the reach of near-term B-mode experiments. Thus a convincing primordial detection at $r \gtrsim 10^{-3}$ would not be a small correction. It would kill this branch.

1 The Prediction

The pre-registered prediction is

$$\boxed{r_{\text{linear}} = 0}. \tag{1}$$

The secondary floor is

$$\boxed{r_{\text{induced}} = C_{\text{SIGW}} A_s}, \qquad A_s = (3/4)\alpha_0^4 \simeq 2.1290 \times 10^{-9}. \tag{2}$$

The coefficient C_{SIGW} is not claimed here as a new substrate derivation. It is the ordinary second-order scalar-induced gravitational-wave transfer coefficient, depending on the cosmological transfer calculation and the observational convention used to quote an effective r. The framework fixes the scalar amplitude A_s; it does not tune C_{SIGW}. Note that $A_s = (3/4)\alpha_0^4$ reproduces the observed scalar amplitude [1] and was fixed before this note, so it enters here as a postdiction-grade input rather than the live claim; the live claim is the tensor null $r_{\text{linear}} = 0$.

For $C_{\text{SIGW}} = 1$, Eq. (2) gives

$$r_{\text{induced}} \simeq 2.1290 \times 10^{-9}.$$

Even the deliberately broad bracket $C_{\text{SIGW}} \in [0.1, 10]$ gives

$$2.1 \times 10^{-10} \leq r_{\text{induced}} \leq 2.1 \times 10^{-8},$$

which is many orders of magnitude below a near-term detection threshold $r \simeq 10^{-3}$.

2 Why the Standard Tensor Estimate Does Not Apply

The standard slow-roll estimate for primordial tensors quantises the transverse-traceless metric perturbation in a smoothly stretching de Sitter background. In reduced Planck units the tensor power is schematically

$$P_t = \frac{2}{\pi^2}\left(\frac{H_*}{\bar{M}_{\text{P}}}\right)^2, \qquad r = P_t/A_s.$$

If the framework's high boot scale $H_* \simeq 1.2 \times 10^{15}$ GeV is inserted into this formula, the result is catastrophic:

$$r_{\text{naive}} \simeq 23,$$

far above the current bound $r_{0.05} < 0.036$ from BICEP/Keck [2]. This is a real exposure only if the branch contains a squeezed de Sitter graviton vacuum.

The branch denies that premise. Expansion is boundary printing, not metric stretching. The print event is a scalar cell-counting event, so the source has the form

$$S_{ij} = a\,\delta_{ij} + b\,\hat{k}_i\hat{k}_j.$$

The transverse-traceless (TT) projector — the operator that extracts the genuine spin-2 gravitational-wave part of a source and discards its scalar (trace and longitudinal) pieces — annihilates this source:

$$\Lambda^{\rm TT}_{ij,kl} S_{kl} = 0.$$

This is the core of the prediction. It is not a small prefactor. It is the absence of a linear tensor source in the scalar printer branch.

3 The No-Squeezing Mechanism

The framework's no-squeezing script gives the same result in elastic language. The bare printed Q_3 cell has central-force bonds. It is compression-stiff but shear-floppy: longitudinal density modes have a restoring force, while transverse shear modes do not. Crystallisation adds the bracing constraints that turn shear rigidity on. Thus the graviton, interpreted as the transverse-traceless shear phonon of the crystallised substrate, is not present at the pre-rigid printing front.

The two branches are therefore mutually exclusive:

Stretching branch. A squeezed graviton vacuum exists during the high scale phase. Then $r_{\rm naive} \simeq 23$, unless the one-bit event unit is abandoned for a much larger effective event unit.

Printing branch. The one-bit boundary printer is retained. The source is scalar; the TT projector kills it; $r_{\rm linear} = 0$.

The current finite-QEC canon adopts the second branch.

4 Falsification Rules

The prediction is falsified by a robust primordial tensor signal at

$$r \geq 10^{-3}$$

after foregrounds, lensing, and systematics are removed. This is the relevant near-term scale because LiteBIRD targets $\delta r \sim 10^{-3}$ [3], and CMB-S4 has a comparable no-detection target in the $r < 10^{-3}$ range [4]. A detection at that level would not be explained by the scalar-induced floor in Eq. (2). It would force a squeezed tensor branch and would refute the boundary-printer null.

The following post-hoc moves are explicitly disallowed:

- changing the prediction from $r_{\rm linear} = 0$ to a neighbouring small nonzero primordial value after the data arrive;
- replacing the scalar printer by a hidden squeezed-graviton vacuum while calling it the same branch;
- fitting $C_{\rm SIGW}$ to a near-term B-mode detection;
- changing $A_s = (3/4)\alpha_0^4$ or the one-bit printer premise to save the tensor prediction.

The following moves are allowed:

- record branch failure;

- open a new labelled tensor branch with a separate date and status;
- compute a sharper C_{SIGW} using an external second-order Boltzmann calculation, provided it is not fitted to a claimed primordial signal.

5 What Would Count as Support

A future non-detection down to $r \simeq 10^{-3}$ would not prove the framework. Many models predict unobservably small tensors. It would, however, keep the boundary-printer branch alive and would distinguish it from high-scale slow-roll branches whose natural tensor amplitude lies within reach.

A positive tensor detection at $r \gtrsim 10^{-3}$ would be much more decisive than a null. It would say that the early universe contained a propagating, squeezed, transverse-traceless gravitational vacuum at a level incompatible with the scalar printer. That is why this is a useful near-term prediction: it is easy for the sky to kill.

6 Reproducibility

The numerical audit is intentionally small; all scripts are in the project reproducibility repository [5]:

Script	Purpose
`python_code/boundary_printing_tensor_prediction_audit.py`	Applies the TT projector to scalar printer sources, verifies $r_{\mathrm{linear}} = 0$, evaluates $A_s = (3/4)\alpha_0^4$, brackets $r_{\mathrm{induced}} = C_{\mathrm{SIGW}} A_s$, and computes the incompatible squeezed-branch event-unit requirement.
`python_code/boundary_printing_tensor_theorem.py`	Records the premise mismatch between a smoothly stretching de Sitter tensor formula and the boundary-printing branch.
`python_code/boundary_printing_no_squeeze.py`	Checks the cell-level rigidity mechanism: the printed cell is compression-stiff but shear-floppy until crystallisation adds bracing.

The machine-readable pre-registration accompanying this note is in the same folder, under the file name `primordial_tensor_prediction_preregistration.json`.

7 Conclusion

The finite-QEC boundary-printer branch predicts a primordial tensor null:

$$r_{\mathrm{linear}} = 0.$$

The only remaining tensor background is a second-order scalar-induced floor of order A_s, naturally around 2.1×10^{-9}, far below the reach of LiteBIRD and CMB-S4. The prediction is risky in the useful sense: a robust primordial B-mode detection at $r \geq 10^{-3}$ would falsify this branch. The value is recorded here before the relevant experiments converge.

References

[1] Planck Collaboration. Planck 2018 results. VI. cosmological parameters. *Astron. Astrophys.*, 641:A6, 2020. doi: 10.1051/0004-6361/201833910.

[2] BICEP/Keck Collaboration. Improved constraints on primordial gravitational waves using Planck, WMAP, and BICEP/Keck observations through the 2018 observing season. *Phys. Rev. Lett.*, 127:151301, 2021. doi: 10.1103/PhysRevLett.127.151301.

[3] LiteBIRD Collaboration. Probing cosmic inflation with the LiteBIRD cosmic microwave background polarization survey. *Prog. Theor. Exp. Phys.*, 2023(4):042F01, 2023. doi: 10.1093/ptep/ptac150.

[4] CMB-S4 Collaboration. CMB-S4 science case, reference design, and project plan. arXiv:1907.04473, `https://doi.org/10.48550/arXiv.1907.04473`, 2019. Primordial tensor and delensing sensitivity reference.

[5] David Elliman. It-from-bit: Reproducibility repository. `https://github.com/dgedge/itfrombit`, 2026. Code and reproducibility repository for the canon snapshot series.