Expansion as Boundary Printing: One Printer for Inflation, Dark Energy and Dark Matter, and a Primordial-Tensor Null

David Elliman
Neuro-symbolic Ltd
dave@neusym.ai

Revised canon snapshot, 2026-06-25

Abstract

In the finite quantum-error-correction (QEC) substrate programme [5, 8], space is the crystallised, self-correcting phase of a discrete register network. We develop the consequence that cosmic expansion is not the stretching of a metric but *Holographic Boundary Crystallization*: the printing of fresh, zero-entropy cells at the causal horizon to keep the patch under its Bekenstein bound [1, 10]. Under this one mechanism inflation and dark energy are the same printer at different clock rates. We report three results. (i) The *scalar* sector — the amplitude A_s, the tilt n_s, the e-fold count, and the dark-energy equation of state w_0 — is now reduced to a local single-clock queue-balance theorem: the scalar readout is the colour-restoring post-decoder current, the nonzero horizon modes are white only if the serial-clock/fixed-total and de Sitter-homogeneous local covariance structure excludes a separate horizon-mode service covariance, and saturated constant-H printing must lock the colour-restoring channel load so that $N_{\text{shell}}\alpha_0^4 = C_F = 4/3$. This gives the sharp candidate $A_s = (3/4)\alpha_0^4 = 2.13 \times 10^{-9}$, within 1σ of Planck [16], conditional on the channel-lock/critical-latch and spatial-whitening identities. (ii) The *tensor* sector is sharply different. The naive de Sitter estimate $r \simeq 23$ rests on a metric-stretching premise the model explicitly denies; the printer is a *scalar* process whose graviton — the substrate's transverse (shear) phonon — is absent at the pre-rigid printing front, giving a primordial-tensor **null**, $r_{\text{linear}} = 0$, with only a scalar-induced floor $r_{\text{induced}} \sim 2 \times 10^{-9}$ (no observable B-modes). (iii) The physical boot cooling law is the printer's own dilution of frustration, $\dot{F} = -nHF$, not a freely chosen annealing schedule. The K04 Metropolis parameter $\gamma = 0.995$ is retained only as a simulation proxy; the physical residuals are the crystallisation epoch, the measured Kibble–Zurek exponent μ, and the lattice-to-physical energy bridge. The picture is falsifiable: a B-mode detection at $r \sim 10^{-3}$–$10^{-1.4}$ would refute it. We tier honestly throughout — the scalar *form*, the local-amplitude candidate, and the tensor-null *mechanism* are robust within their stated premises; CMB completion and absolute dark-sector normalisation retain named gates — and every number is reproduced by a short self-checking program.

Contents

1 Stretching versus printing

Standard inflationary cosmology stretches a metric: a scale factor $a(t)$ grows and everything rides along. A discrete substrate cannot do this. Stretching the lattice spacing would change the speed of light and dissolve the framework's ultraviolet cutoff; inserting new cells into the bulk would violate the local Euler characteristic and cascade defects. The only consistent option is to add cells *at the boundary.* As the substrate's error-correction radiates Landauer waste heat [13] outward to the edge of the causal patch, the patch entropy approaches the Bekenstein–Hawking bound $S_{\max} = A/4\ell_P^2$ [1, 10]; to stay under it the horizon must grow its area by precipitating new $\mathbb{Z}^3 \otimes Q_3$ cells *at the boundary. The universe does not stretch; it prints.* This is Holographic Boundary Crystallization (HBC).

This note takes HBC seriously as a dynamical mechanism and asks what it predicts for the things we actually observe: the primordial scalar perturbations, the primordial tensors, the dark-energy equation of state, and the dark-sector bookkeeping. The striking outcome is that one printer, with a finite service stop rule and a physical cooling law, organises all of them — and that the scalar and tensor sectors come apart in a sharply testable way.

2 One printer, two clock rates

The event unit. Each print event lays down one boundary cell and writes one irreversible record. When the shell count is translated into a de Sitter entropy this is the ordinary one-bit record unit, $u_{\text{event}} = \ln 2$. The scalar amplitude itself, however, is not a Landauer-heat coefficient: it is a normalized service-current covariance, and the heat unit cancels from the amplitude ledger.

The stop rule. The high-bandwidth (inflationary) phase prints at the capacity of the gauge-projected scalar queue. The scalar readout is the post-decoder, colour-restoring topology current, not the all-channel entropy current. The finite channel geometry has 14 weight-4 supports, 28 transverse service labels, and six colour-restoring scalar labels; each scalar label trips 8/12 strain checks and therefore carries load $2(8/12) = C_F = 4/3$. A coherent constant-H printer has both sustainability and no-idle saturation,

$$\lambda_{\text{shell}} \le C_F, \qquad \lambda_{\text{shell}} \ge C_F.$$

Thus the stop rule is the queue-balance condition

$$\lambda_{\text{shell}} = N_{\text{shell}}\, \alpha_0^4 = C_F = \tfrac{4}{3}, \tag{1}$$

with $\alpha_0 = 1/137$ the substrate coupling and C_F the $SU(3)$ fundamental Casimir. This is not a derivation from total Bekenstein–Hawking area saturation alone; it is a local single-clock scalar-printer theorem. Adding a second scalar source or a nonlocal horizon-mode service operator would be new physics and would reopen both the tilt and amplitude audits. Inflation is this printer running at saturation during boot; dark energy is the same printer at steady state once the boot has been pruned to baseline. They are not two phenomena but two clock rates of one machine.

3 The scalar sector

Treating the perturbations as the shot noise of discrete print events, the saturation identity (1) gives the scalar-amplitude candidate, and the 28-channel boundary clock fixes the tilt and the equation of state. All four scalar observables then follow once the channel-lock/critical-latch and spatial-whitening identities are accepted:

Observable	Printer reading	Value	vs. Planck
Scalar amplitude A_s	$\frac{3}{4}\alpha_0^4 = \alpha_0^4/C_F$	2.13×10^{-9}	$+1.0\sigma$
Spectral tilt n_s	$1 - \frac{1}{28} = 27/28$	0.9643	-0.15σ
e-folds N_e	2×28	56	in $[50, 60]$
Dark-energy w_0	$-1 + \frac{1}{28} = -27/28$	-0.964	DESI-ward

The candidate amplitude is the inverse shell count $A_s = 1/N_{\text{shell}} = \alpha_0^4/C_F$; the tilt is the first eigenvalue 1/28 of the serial 28-channel boundary clock acting on the power ledger (so $n_s = 27/28$, and by the same gap $w_0 = -n_s = -27/28$ for the steady-state phase, with $w_a = -1/28$ in the Chevallier–Polarski–Linder parametrisation [3, 14]); $N_e = 2 \times 28$ is the Markov relaxation length. The present DESI preference for $w_0 > -1$ [4] is the direction this predicts.

Honesty about tiers. The *forms* (amplitude = inverse shell count; tilt = clock gap) are robust, and the four numbers land at $\leq 1\sigma$ without retuning. The amplitude gate has been narrowed sharply, but not removed: all-channel entropy loading can keep total printed entropy fixed while changing the colour-restoring channel projection, and an allowed horizon-mode service covariance can change $S_j(k = aH)$ while leaving the homogeneous clock intact. The resulting value $A_s = (3/4)\alpha_0^4$ is therefore a conditional saturated-printer candidate, not a fully discharged theorem. The remaining locks are channel-lock/critical-latch and spatial-whitening. These caveats limit how firmly the value is derived as a statement about nature; they do not affect the scalar/tensor contrast of §4, which is the load-bearing result.

4 The tensor null

The same event unit fixes the inflationary scale. Writing the shell count as the de Sitter entropy per event, $N_{\text{shell}} = S_{\text{dS}}/u_{\text{event}}$ with $S_{\text{dS}} = 8\pi^2 \bar{M}_{\text{P}}^2/H_*^2$, gives

$$H_* = \sqrt{\frac{6\pi^2}{\ln 2}}\, \alpha_0^2\, \bar{M}_{\text{P}} \simeq 1.2 \times 10^{15} \text{ GeV}. \tag{2}$$

For *any* ordinary inflationary background this scale is fatal: the de Sitter graviton vacuum would give a tensor-to-scalar ratio

$$r_{\text{naive}} = \frac{(2/\pi^2)(H_*/\bar{M}_{\text{P}})^2}{A_s} \simeq 23, \tag{3}$$

about 650× the bound $r < 0.036$ [2]. The resolution is that the estimate is a *category error*.

The premise mismatch. $r_{\rm naive}$ is computed by quantising the transverse-traceless metric perturbation — the graviton — in a *smoothly stretching* de Sitter background and squeezing its adiabatic vacuum. HBC denies exactly that premise: the universe prints, it does not stretch. The standard tensor formula is being applied to a process the model says is not happening.

The printer is scalar. The printer's perturbations are the shot noise of discrete cell printing — a density (scalar) source. At linear order the scalar, vector and tensor sectors decouple, so a scalar source produces *no* transverse-traceless mode. Equivalently, the transverse-traceless projector annihilates every scalar source of the form $a\delta_{ij} + b\hat{k}_i\hat{k}_j$. The linear primordial tensor ratio is therefore

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

The only tensor channel is second-order scalar-induced gravitational waves, at the floor

$$r_{\rm induced} = C_{\rm SIGW} A_s,$$

where $C_{\rm SIGW}$ is the standard second-order transfer coefficient. With the natural unit coefficient this is $r_{\rm induced} \simeq 2.1 \times 10^{-9}$; changing $C_{\rm SIGW}$ by an order of magnitude still leaves the signal far below near-term B-mode reach.

No graviton to squeeze (the no-squeezing theorem). One might still worry that the emergent coarse-grained metric squeezes a graviton vacuum. It does not, because in this substrate the graviton *is* the crystal's transverse (shear) phonon, and a freshly-printed front has no shear rigidity. Concretely: the printed degree-3 cell is, by Maxwell constraint counting [15], shear-*floppy* (coordination $z = 3$, below the rigidity threshold $2d = 6$; six internal zero modes) while remaining compression-stiff — so it carries the scalar (compression) mode but no transverse mode. Shear rigidity, and with it a propagating graviton, appears only once crystallisation supplies the bond-angle (next-nearest) constraints that lock the cell into its Q_3 geometry. This is not a finite-size artefact: by rigidity percolation [11] the macroscopic shear modulus is an order parameter, exactly zero and scale-invariant below the crystallisation threshold, while compression percolates earlier — so a partially-crystallised bulk carries scalars but no long-wavelength shear. A mode crossing the horizon while its region is still being printed therefore finds no graviton to squeeze, and its tensor amplitude is set only when the region crystallises, at the lower rate $H_*^{\rm xtal}$:

$$r = r_{\rm naive} \left(\frac{H_*^{\rm xtal}}{H_*} \right)^2. \tag{4}$$

A modest crystallisation lag — a $\sim 25\times$ drop in H, i.e. $\sim$ 2–3 e-folds of post-printing evolution — already brings r below the bound; the printer's own (scalar, emergent-gravity) structure pushes it to the second-order floor $r \sim 2 \times 10^{-9}$.

The prediction. Boundary printing predicts an exact linear tensor null and therefore *essentially zero observable primordial tensors*: $r_{\rm linear} = 0$, with only a second-order scalar-induced floor of order 2×10^{-9}. This is the note's sharpest, most falsifiable statement. A primordial B-mode detection at $r \sim 10^{-3}$ to the current bound would mean the graviton vacuum *was* squeezed at H_*, refuting either the no-squeezing mechanism or the one-Landauer-bit event unit.

Tier. The premise mismatch and the scalar/vector/tensor decoupling are rigorous; the graviton-as-shear-phonon mechanism (floppy front, crystallisation bracing, rigidity percolation) is rigorous as linear elasticity on the cell and as percolation theory; the one physical input is that the emergent metric does not squeeze the graviton on CMB scales during printing, which the discrete-printing premise supplies but which is not yet a closed continuum theorem.

5 The cooling schedule fixes itself

The remaining quantity is the Kibble–Zurek freeze-out correlation length $\xi(R)$ [12, 17], which controls the frozen defect density and the crystallisation lag of §4. It has been easy to confuse the simulation protocol with the physical law. The K04 Metropolis dynamics is scale-invariant: $(w_4, w_6, \lambda, T) \mapsto s(w_4, w_6, \lambda, T)$ leaves all acceptances unchanged, so neither w_6/Λ nor the physical ramp can be derived from K04 alone. The printer supplies the physical cooling law instead: *the printer is the cooler.* The fresh $S = 0$ cells minted to drive expansion dilute the frustration density, $\dot{F} = -nHF$, so the cooling rate is nH and the quench time $\tau_Q = 1/(nH_c) \sim H_c^{-1}$ is printer-set, not a chosen anneal ramp. Hence

$$\xi = \xi_0\,(\tau_Q/\tau_0)^{\mu}, \qquad \mu = \frac{\nu}{1+\nu z}, \tag{5}$$

a function of the crystallisation-epoch Hubble rate and one universality exponent μ. A direct consequence is that the frozen-defect density and the tensor ratio are *co-monotonic* in the one cooling rate (faster cooling $\to$ smaller ξ $\to$ more frozen defects *and* a shorter lag $\to$ larger r), so the dark-sector bookkeeping and the tensor bound are a single joint constraint on the boot. The identification of $H_*^{\rm xtal}$, and whether the metric (inflation) and matter (QCD) crystallisations are one event or two, are flagged open.

6 The universality exponent μ

Because ξ is hyper-sensitive to μ (it spans ~ 9 orders of magnitude across $\mu \in [\frac{1}{4}, \frac{3}{4}]$), μ is the single decisive number left. We measure it the framework's own way: a Kibble–Zurek quench of the $K04$ embedded ensemble at several cooling rates R, fitting the trapped-defect density $d \sim R^{-\mu}$. A finite-size-scaling run on $L = 8, 12, 16, 24$ (24-core cluster; $4 \times 4 \times 3 = 48$ quenches) shows a clean Kibble–Zurek signal at every L, and the $1/L$ extrapolation gives

$$\mu(\infty) \simeq 0.22 \in [0.21, 0.24], \tag{6}$$

the bracket coming from the small-R saturation systematic. This sits close to the mean-field Kibble–Zurek value $\mu = \nu/(1+\nu z) = \frac{1}{4}$ (with $\nu = \frac{1}{2}, z = 2$) and *below* the 2D random-bond-Ising / Nishimori class value 0.375: the 3D $K04$ ordering transition is therefore most likely *not* in the RBIM class — its exponent is its own, near mean-field. With μ pinned, ξ collapses from the ~ 9-OOM ambiguity to ~ 1 OOM: the cooling schedule is now a *measured* quantity, not a free dial. The exponent carries genuine error bars (four R-points, three reps, d near its ceiling at small R), so it is $\mu \approx 0.22 \pm 0.03$, not a four-digit constant; a wider R-range and more reps would sharpen it.

7 What is settled, what is conditional, what is open

Settled (robust). Expansion as boundary printing; inflation and dark energy as one printer at two clock rates; the scalar/tensor contrast — a scalar printer with no linear tensor source —

and the consequent primordial-tensor null as a *falsifiable* prediction; the graviton-as-shear-phonon mechanism (Maxwell floppiness, crystallisation bracing, rigidity percolation); the printer-is-cooler reduction $\xi = \xi(H_c, \mu)$ and the frozen-defect $\leftrightarrow$ tensor co-monotonicity.

Conditional (form firm, value premise named). The absolute scalar amplitude has a sharp local saturated-printer candidate, $A_s = (3/4)\alpha_0^4$. Its residual is not a hidden numerical fit but two operator identities: channel-lock/critical-latch and spatial-whitening. A new nonlocal scalar source or second clock would reopen the amplitude, while the 28-clock tilt can survive a scale-independent amplitude shift. The inflation scale H_* uses the one-bit entropy translation; the exact crystallisation lag that sets r within the null remains a dynamical printer/crystallisation question.

Open. Sharper precision on μ ($\mu(\infty) \approx 0.22 \pm 0.03$ is now measured; a wider R-range and more reps would tighten it, and test the near-mean-field reading); whether the metric and matter crystallisations are one transition or two; the lattice-to-physical energy bridge ($w_6 \leftrightarrow \Lambda$), where the selected ramp-start bridge gives $w_6/\Lambda = 1/6$ but K04-internal Metropolis dynamics cannot derive an absolute scale; and a closed continuum proof that the emergent metric does not squeeze the graviton during printing. None of these affect the central falsifiable claim: $r \sim 2 \times 10^{-9}$.

CMB completion is conditional, not solved by K04 debris. The old claim that the printed defect abundance by itself completes the CMB is retired. K04 debris is a pinned fossil/static branch, not a mobile recombination-era cold component. The current CMB route uses an effective R4 zero-mode reservoir: a pressureless rest-count component with $w = c_s^2 = 0$, $\Omega_{\text{zero}} h^2 = 0.096708$, total dark $\Omega_{\text{dark}} h^2 = 0.120885$, and $z_{\text{eq}} = 3430.3$. A CAMB check gives the right peak/equality shape at roughly one-percent level once this dust slot is inserted. The live caveat is severe: documented active R4 dynamics do not derive a conserved massive dust charge, so the reservoir remains a conditional/AeST-class premise rather than a substrate theorem. The remaining CMB gates are explicit: the selector-locked value $H_0 = 67.3$ gives $100\theta_*$ high by 0.290% (a fit wants $H_0 \simeq 66.3$, $\Omega_k \simeq +0.002$, or a phantom-flipped $w(a)$ coefficient), and the halo ledger must choose between a CDM-like zero-mode halo branch and an active MOND/R4 branch with a new $> 95\%$ zero-mode galaxy depletion or screening theorem. This paper therefore treats the CMB as conditionally supplied at the effective Boltzmann/equality level, not fully Locked.

A Reproducibility

Every numerical and logical statement is produced by a self-asserting program in the repository [9] (each exits 0 only if its internal checks pass):

Program	Role
python_code/boundary_printing_ledger.py	The scalar sector (§3) and the tensor exposure: A_s, n_s, N_e, w_0 vs Planck.
python_code/boundary_printing_tensor_theorem.py	The premise mismatch and tensor-null mechanism (§4).
python_code/boundary_printing_tensor_prediction_audit.py	TT-projector audit: $r_{\text{linear}} = 0$, induced floor $C_{\text{SIGW}} A_s$.
python_code/boundary_printing_no_squeeze.py	Graviton = shear phonon; the printed cell is shear-floppy, crystallisation braces it.
python_code/boundary_printing_rigidity_percolation.py	Shear-floppiness is a scale-invariant bulk phase (rigidity percolation).
python_code/boundary_printing_q3_bracing.py	The bracing is the Q_3 cell's bond-angle locking, not a stand-in.
python_code/boundary_printing_suppression_lag.py	The suppression factor $r = r_{\text{naive}}(H_*^{\text{xtal}}/H_*)^2$ and the required lag.
python_code/boundary_printing_cooling_schedule.py	The printer-is-cooler reduction $\xi = \xi(H_c, \mu)$; frozen-defect $\leftrightarrow$ tensor.
python_code/boundary_printing_mu_measurement.py	The Kibble–Zurek exponent μ at $L = 8$ (§6).
python_code/boundary_printing_mu_fss.py	FSS extrapolation $\mu(L) \to \mu(\infty) \approx 0.22$ from the $L = 8$–24 grid (§6); data `mu_fss_results.jsonl`.
python_code/item131_hbc_stop_rule_proof.py	Finite service-ledger candidate for $\lambda_{\text{shell}} = C_F$ under the local single-clock scalar-printer premise.
python_code/item131_hbc_channel_whitening_closure.py	Tests the local whitening route; later audits keep channel-lock and spatial-whitening as live amplitude locks.
python_code/k04_absolute_scale_cooling_law_audit.py	Splits w_4/w_6, w_6/Λ, protocol γ, and physical boundary-printer cooling.
python_code/item123_cmb_zero_mode_theorem.py	Conditional pressureless zero-mode reservoir/effective-dust target for the CMB dark slot.
python_code/item123_cmb_theta_halo_completion_gate.py	CAMB theta-star and halo non-double-counting gate for the CMB completion.

The cosmological-engine, gravity and dark-sector mechanisms, with full claim tiers, are stated in the sector papers [5–8].

References

[1] Jacob D. Bekenstein. Black holes and entropy. *Phys. Rev. D*, 7:2333–2346, 1973.

[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.

[3] Michel Chevallier and David Polarski. Accelerating universes with scaling dark matter. *Int. J. Mod. Phys. D*, 10:213–223, 2001.

[4] DESI Collaboration. DESI 2024 VI: cosmological constraints from the measurements of baryon acoustic oscillations. arXiv:2404.03002, 2024.

[5] D. G. Elliman. Cosmology, dark energy, and inflation in the finite-QEC substrate. Zenodo preprint, `https://doi.org/10.5281/zenodo.20671860`, 2026.

[6] D. G. Elliman. Foundations and methodology for a finite-QEC substrate: code, crystallisation, ledgers, and audit protocol. Zenodo preprint v3.0, https://doi.org/10.5281/zenodo.20682178, 2026.

[7] D. G. Elliman. Gravity, horizons, and black holes in the finite-QEC substrate. Zenodo preprint, https://doi.org/10.5281/zenodo.20671864, 2026.

[8] D. G. Elliman. A finite-QEC substrate program for particle physics and cosmology: current canon and audit methodology. Zenodo preprint v5.0, https://doi.org/10.5281/zenodo.20682177, 2026.

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

[10] Gary W. Gibbons and Stephen W. Hawking. Cosmological event horizons, thermodynamics, and particle creation. *Phys. Rev. D*, 15:2738–2751, 1977.

[11] D. J. Jacobs and M. F. Thorpe. Generic rigidity percolation: The pebble game. *Physical Review Letters*, 75(22):4051–4054, 1995.

[12] Tom W. B. Kibble. Topology of cosmic domains and strings. *J. Phys. A*, 9:1387–1398, 1976.

[13] R. Landauer. Irreversibility and heat generation in the computing process. *IBM J. Res. Dev.*, 5:183–191, 1961.

[14] Eric V. Linder. Exploring the expansion history of the universe. *Phys. Rev. Lett.*, 90:091301, 2003.

[15] J. Clerk Maxwell. On the calculation of the equilibrium and stiffness of frames. *Philosophical Magazine*, 27(182):294–299, 1864.

[16] Planck Collaboration. Planck 2018 results. VI. cosmological parameters. *Astron. Astrophys.*, 641:A6, 2020.

[17] Wojciech H. Zurek. Cosmological experiments in superfluid helium? *Nature*, 317:505–508, 1985.