

Emergent Graph Theory (EGT)

A Framework for Emergent Spacetime under Energy Conservation

Marcel Langjahr

Pantareon

Pantareon Foundations

langjahr@pantareon.io · <https://pantareon.io/>

July 2026

Abstract

We present Emergent Graph Theory (EGT), a theoretical framework that models physical reality as a discrete, dynamic graph governed by strict energy conservation. The central thesis is that stable structure—at any scale—is a configuration that has become invariant under the energy-conserving dynamics of its substrate. The grammar of this stabilisation is graph-topological and thermodynamic; the instances are the universe’s catalogue of stable forms.

The framework rests on four postulates: energy is fundamentally spectral, carrying frequency-encoded histories of its interactions; information routing through the graph is emergent from spectral overlap; time is the cumulative count of interaction events, not an external parameter; and physics involves a fifth fundamental force—the activity of the substrate itself—which standard physics has not unified because it acts outside matter. Together these postulates yield the ECQ triad, a thermodynamic account of how the dynamics of energy (E), quanta (Q), and chronos (C) drive the system toward elaborate self-organisation. The decreasing matter-energy under accumulating substrate sediment creates the thermodynamic pressure that forces structures to optimise themselves: this is the proposed mechanism of emergent complexity.

Implemented in the *Aeoncore* physics engine, the framework demonstrates that macroscopic phenomena such as gravitational metric contraction, vacuum expansion, and baryonic binding emerge from graph topology under a single calibration anchor (the proton mass), rather than from hard-coded, per-phenomenon force laws. All four standard interactions are implemented as conserving graph operators — confinement as edge tension, beta decay against the physical neutron lifetime, exact charge bookkeeping with a parameter-free radiative

back-channel, and gravitation as mass-gated densification of the metric itself — and space is not a stage but a particle species: vacuum nodes live in the same conservation ledger as matter, gravitate in the emergent Friedmann law, and carry the dark sector structurally.

On a fourteen-cell reference campaign on the audited engine (frozen declaration, mass-based counting), every cell runs the complete cosmological cycle — big bang, era sequence, cold-window nucleosynthesis, gravitational collapse, crunch, and re-seed — as one closed process, and every cell’s successor generations cycle healthily in turn: global energy drift stays at the 10^{-10} scale through entire 2×10^5 -step cycles including the bounce, the charge census closes to zero exactly — through 121 stellar collapses and 10.7 million fusion events — and the four-sector energy partition closes to 4×10^{-13} . Within that budget the engine builds a cold nucleosynthesis ladder from deuterium (hundreds of standing nuclei) through helium up to $A = 49$ across 63 distinct species, and terminates it physically: the chain’s upper edge is a measured collapse edge, progenitors of $A \approx 40$ –64 collapsing through an emergent three-pressure cascade into 120 neutron stars and, in the campaign’s densest cell, its first white dwarf. The cycle’s clock runs with the system scale — the engine’s first quantity to pass a pre-declared scale-running criterion. The inheritance of constants across cycles remains open in a sharpened form: the spectral archive is carried exactly across the bounce and is now read by a frequency-selective operator, but it does not yet steer the successor dynamics.

Source Code and Live Demonstration: The source code is not publicly available. The reference run of this paper, however, is: the complete dataset behind the numbers below can be explored as an interactive live demonstration — the universe replayed in the browser, chapter by chapter — at

<https://graphdynamics.org/>.

Scope of this Paper: Detailed mathematical derivations and quantitative results from the *Aeon-core* implementation will be released in dedicated follow-up publications. The present work establishes the conceptual framework and ontological commitments of EGT. Current simulation results (reported in the companion results paper, v0.2) characterise the emergent substrate as a non-geometric energy/interaction network rather than a metric spacetime; whether a metric, causal spacetime is recovered in a coarse-grained limit remains an open question under active investigation. The framework presented here is the theoretical hypothesis under test, not a set of claims established by the present engine.

Claim Level: EGT at its present scale makes *no claim to reproduce our universe one-to-one* — that correspondence is precisely what the scaling programme exists to pursue. What the engine claims, and what its frozen-declaration methodology tests, are structural properties: exact conservation, composition laws, the shapes and directions of emergent phenomena, and the closure of complete cosmological cycles. Quantitative comparisons against nature are reported throughout as *mapped distances* — they calibrate the scaling frontier and are expected to move as scale and operator depth grow; a magnitude gap at today’s scale is data, not failure. Two kinds of open item are kept deliberately distinct: magnitude gaps, which belong to the scaling axis, and structural absences (such as running couplings), which are defined construction stages of the architecture. Both are stated where they occur, once, with what it would take to close them.

1 Introduction

Everything we recognise as structure—a galaxy, an atomic nucleus, a stable molecular configuration—is a configuration that has become invariant under the dynamics of its substrate. The premise of Emergent Graph Theory is that the substrate, at every scale at which structure stabilises, is energy flowing through a graph of relations, and that the principles governing how structure emerges from elementary interactions are not domain-specific.

Contemporary physics faces significant challenges in reconciling the information-theoretic properties of Black Holes with the geometric nature of spacetime. Furthermore, the “Fine-Tuning Problem” suggests that the constants of nature are highly specific, yet standard models offer no mechanism for their selection other than the Anthropic Principle.

EGT proposes a computational and ontological solution: The universe is modeled as a high-dimensional graph where the metric is defined by edge tension (Energy), evolution is driven by interaction processes (Chronos), and structure is emergent (Quanta). Reality, in this view, is not a stage on which dynamics occur—it is the running execution of those dynamics. Stable structure is the answer to the question *what configuration of this graph is invariant under further evolution*, given its energetic and topological constraints.

The framework rests on four commitments that distinguish it from standard physics. First, energy carries information: it is not a scalar quantity but a spectral signature, a distribution over frequencies that encodes the causal history of interactions. The foundational postulate, that energy carries spectral information, is not a speculative leap but a formalisation of daily practice in astronomy, where spectral analysis has long served as the primary tool for decoding the history and composition of matter.

Second, time is not an external parameter against which evolution is measured. Time is the cumulative count of interaction events—of what has happened. A system at rest accumulates no time; a system in active interaction accumulates time at the rate of its event density. This relational view recovers, in graph-dynamics form, the position that time is not a container but a count.

Third, physics involves five fundamental forces, not four. Standard physics recognises electromagnetism, the weak and strong nuclear forces, and gravitation. EGT recognises a fifth: the activity of the substrate itself. This force has not been identified in standard physics as a single, unified interaction because it acts outside matter—the standard view treats space as a passive stage on which matter fields evolve. EGT brings the substrate into the dynamics as an active participant, and its activity unifies phenomena that are currently treated by separate phenomenological frameworks: vacuum energy, dark energy, the cosmological constant, vacuum fluctuations, and long-range biomolecular coupling.

In the implementation this commitment is literal, and it is the framework’s sharpest structural departure: *space is a particle species*. Vacuum nodes are citizens of the same graph and the same conservation ledger as protons and electrons — they carry energy, participate in the metric, gravitate in the emergent expansion law, and are counted, conserved, and audited by the identical bookkeeping. Expansion is vacuum-node population dynamics; the metric is edge density among them; and, as measured below, the dark sector is not an add-on but

the ledger share of this population. Nothing in the engine distinguishes “matter physics” from “space physics” except the operator set that happens to act on each type.

Fourth, structure emerges under thermodynamic pressure. The interplay of energy, quanta, and chronos—the ECQ triad—governs how a system organises itself. As substrate sediment accumulates and matter energy diminishes, the system is forced toward increasingly elaborate self-organisation. This is the proposed thermodynamic mechanism of emergent complexity, including the conditions that make biological organisation possible.

Using the *Aeoncore* engine, we simulate this topology under strict energy conservation enforced per transaction ($\Delta E < 10^{-9}$ per step). We show that when such a system is subjected to radiative and thermal stress, it spontaneously clusters into baryonic structures and undergoes a conservation-faithful collapse–reseed, providing a candidate mechanism for the inheritance of physical constants across cosmological epochs—proposed here and examined quantitatively in a dedicated companion results paper.

The present paper focuses on the cosmological and particle-physics scales as the foundational level of investigation. These scales are the substrate on which all higher-order phenomena—chemical, biological, and cognitive—presuppose the stabilisation already achieved. Establishing the framework at this level is therefore a precondition for its extension to other domains. This paper establishes the conceptual framework of EGT. Detailed derivations, simulation results, and quantitative predictions are addressed in companion documents.

2 Foundations and Postulates

The theory rests on four postulates. The first two concern the nature of energy and information transport in the graph. The third concerns the nature of time. The fourth concerns the structure of fundamental interactions.

2.1 Postulate I: The Spectral Vector

We postulate that energy is fundamentally spectral. Every quanta q possesses a spectral state vector $\mathcal{S}_q$, defined as a set of frequency-energy bands:

$$\mathcal{S}_q = \{(f_0, E_0), (f_1, E_1), \dots, (f_n, E_n)\} \quad (1)$$

The total energy is the scalar sum $E_{\text{tot}} = \sum E_i$, but the information lies in the distribution $E(f)$. This

spectrum encodes the causal history of the energy—fusion events, redshifting, and resonant interactions leave distinct frequency fingerprints that are conserved and superimposed during energy transfer.

2.2 Postulate II: Resonant Conductivity

Information routing is emergent. The conductivity σ_{uv} between two nodes u and v is a function of their spectral overlap and the edge tension:

$$\sigma_{uv}(f) \propto \frac{\sqrt{T_{uv}}}{1 + \gamma|f_{\text{source}} - f_{\text{target}}|} \quad (2)$$

This implies that energy “self-routes”: a packet of energy at frequency f_x will propagate preferentially through pathways resonant with f_x . The graph acts as a spectral filter, effectively “addressing” information to compatible structures. The underlying move—function carried by a *dynamic* graph of correlations between elements, rather than by fixed pairwise weights—has a direct precedent in theoretical neuroscience, where the correlation theory of brain function [12] holds that representation is carried by dynamic linking among units, not by static connection strength alone. Postulate II is, in this sense, the energy-conserving and scale-invariant form of correlation-based routing.

2.3 Postulate III: Time as Events

Time, in EGT, is not an external parameter against which the system evolves. It is the cumulative count of interaction processes occurring within the system:

$$C(t) = \sum_{s=0}^t N_{\text{events}}(s), \quad (3)$$

where $N_{\text{events}}(s)$ is the number of discrete interaction events at step s . A system at rest accumulates no time; a system in active interaction accumulates time at the rate of its event density.

This is a substantive claim about the nature of time, not a definition internal to the simulation. A photon that propagates without interaction accumulates no Chronos, consistent with its zero proper time in standard physics. A region of dense interaction—a stellar core, a folding protein, the early universe—accumulates Chronos rapidly. The flow of time is the rate at which something happens.

The events that contribute to C are the discrete state-changing processes: bond formation and breaking, vacuum decay, fusion, decay channels, spectral merging, and substrate reorganisation. Continuous-field dynamics—gradual changes in tension, position

updates—do not contribute, as they do not constitute distinct events.

This postulate makes Chronos in the ECQ triad (§4) an operationally definable quantity rather than an abstract "interaction cost". It also provides a natural reading of relativistic time dilation: a moving observer participates in fewer interaction events per unit of an external reference, and so accumulates Chronos more slowly relative to that reference. The relational character of time is built into the postulate.

2.4 Postulate IV: The Fifth Force

Standard physics recognises four fundamental interactions: electromagnetism, the weak nuclear force, the strong nuclear force, and gravitation. EGT recognises a fifth: the activity of the substrate itself.

The fifth force does not act between matter quanta through edges of the graph; it acts *through the substrate*, between the substrate and matter, and within the substrate's own dynamics. Its manifestations include:

- **Vacuum decay:** the spontaneous creation of matter–antimatter pairs from vacuum nodes, balanced in energy and inheriting spectral signature from the parent substrate.
- **Vacuum recombination:** the inverse process, in which substrate nodes collapse and redistribute their energy and spectral content.
- **Bridge tunnelling:** energy transport through the substrate between distant points, bypassing the matter topology entirely.
- **Substrate polarisation:** the alignment of vacuum dipoles in response to surrounding charged matter, what standard physics calls vacuum polarisation in QED.
- **Cosmological expansion:** the substrate generating new substrate, which in standard cosmology is described phenomenologically through the Friedmann equations.
- **Spatial edge discovery:** the emergent geometry feeding back on topology, allowing geometrically near but topologically distant nodes to develop new edges through the substrate.

These processes share a common structure: they operate without requiring a pre-existing matter edge between the affected nodes; they involve the substrate as an active participant; they carry spectral

information across their action; and they are universal across configurations.

Standard physics has not unified these phenomena because each is treated within a separate framework: vacuum energy in QFT, dark energy in cosmology, vacuum polarisation in QED, the cosmological constant in general relativity, vacuum fluctuations in quantum mechanics, biomolecular long-range coupling in biophysics. EGT's claim is that they are all manifestations of one fundamental interaction, distinguished from the four standard forces by the fact that it acts through the substrate rather than through matter.

This is the most consequential postulate of EGT for empirical prediction. If a single force underlies all substrate phenomena, then the relations between dark energy, vacuum fluctuations, and biomolecular coupling are not coincidences but expressions of common dynamics. The fifth force generates testable predictions across these domains, addressed in dedicated follow-up work.

In the implementation, however, the fifth force is not a residual category but a concrete operator inventory, each entry with its own ledger account and declared constant: (i) *gradient-driven substrate export* — spectral-true transport of radiation toward the south archive along the pole-potential gradient (the working heat sink; macroscopically the dark-energy-like drain); (ii) *bridge tunnelling* — mass transfer from horizon states into the archive; (iii) the *crunch drain* — the bridge-charge-driven absorption that completes a cycle once the pole potentials cross; (iv) *wedge funding of topological gravity* — conversion of a node's extractable agitation into structural tension (kinetic $\rightarrow$ structure, Jeans-like); and (v) *expansion cooling* — the redshift-like drain of propagating radiation into the archive. Every entry is conservative by construction; the ledger closes to the drift figures reported below.

3 Time, Cycles, and the Substrate Archive

The substrate is not a passive medium. Through the fifth force, it accumulates the spectral history of every interaction that has occurred within the system. This accumulation is concentrated, in the cosmological configuration, at the south pole of the substrate's polar structure.

The south pole of the cosmological substrate is the integrated record of fifth-force activity throughout the system's history. Every vacuum decay, every recombination, every bridge tunnelling event contributes its spectral signature to the south-pole

accumulation. Over the lifetime of a cosmic cycle, the south-pole spectrum becomes an increasingly dense archive of the system’s interaction history.

This archive is not metaphorical. It is the literal accumulation of spectral content in the energy concentrated at the south-pole node, transported there by the fifth-force tunnelling channel. The spectral signature carried by each tunnelled energy quantum is preserved, merged with the existing south-pole spectrum, and added to the archive.

The substrate as a whole, concentrated at the south pole, is the integrated record of the universe’s history. Black holes, in this picture, are extreme local concentrations of substrate sediment—not separate archives, but local manifestations of the same underlying mechanism. The information consumed by a black hole is preserved in its local substrate concentration and contributes to the broader south-pole accumulation through bridge tunnelling.

The mechanism by which this archive could determine physical constants in subsequent cycles, and the measured status of that claim — carried exactly, now read by a frequency-selective operator, not yet steering — is established in §11.

4 The ECQ Triad

The four postulates establish the substrate, the kinematics of information, the nature of time, and the active role of the fifth force. The ECQ triad establishes how these come together into a thermodynamic account of how systems organise themselves over time.

The triad consists of three quantities. Two of them are forms of energy; the third is time as defined by Postulate III.

4.1 The Three Quantities

E — Substrate energy. The energy concentrated at the south pole of the cosmological substrate. This is the sediment of fifth-force activity: every vacuum decay, every recombination, every bridge tunnelling event has contributed energy to this accumulation. E grows monotonically over the lifetime of a cycle, as fifth-force processes transport energy from the active matter region to the substrate archive.

Q — Matter energy. The total energy bound in matter structure: nuclear masses, kinetic agitation, radiation in matter, edge tension between matter nodes. Q is what standard physics would call the energy of “everything that is not vacuum or substrate”. As the cycle progresses, Q tends to decrease: matter loses energy to the substrate through

fifth-force processes, and the matter sector becomes increasingly impoverished in its energy content.

C — Chronos. The cumulative count of interaction events, as defined in Postulate III. C grows monotonically with the system’s activity.

A fourth quantity, the vacuum-class energy E_{vac} , stands outside the triad. It is the substrate’s remaining potential, the energy that has not yet been mobilised into either matter (Q) or sediment (E). For a closed cosmological system, the total is conserved:

$$E_{\text{tot}} = E + Q + E_{\text{vac}} = \text{const.} \quad (4)$$

4.2 The Coherence State

The interaction of the three triad quantities defines the *coherence state* of the system, a quantity that characterises how far the system has progressed through its cycle. The schematic form is:

$$\mathcal{K} = \frac{E \cdot C}{Q}. \quad (5)$$

All three contributions push $\mathcal{K}$ upward as the cycle ages: E grows, Q shrinks, C accumulates. In cosmological-configuration runs this behaviour is confirmed quantitatively: the sediment E is non-decreasing in 100% of samples, Q rises only during the brief materialisation phase (peaking within the first few percent of a generation) and declines monotonically thereafter, and $\mathcal{K}$ increases monotonically with a net growth of order 10^3 – 10^4 per generation. The acceleration is a net effect of the simultaneous growth of the numerator and decline of the denominator; it is not uniformly convex at finite sampling resolution. The four-sector partition underlying the triad (E , Q , E_{vac} , and the structural pole sector) is verified in the implementation to close against the total energy at machine precision.

4.3 The Critical Regime

The coherence state $\mathcal{K}$ is a diagnostic quantity: it tracks how far the cycle has progressed, not what happens at any particular value. As $\mathcal{K} \rightarrow 1$, the substrate sediment approaches its maximum, the matter sector approaches its minimum, and the accumulated history approaches its archive limit. The actual termination of the cycle—the Big Crunch—is governed by the physical condition described in §11: an energy-gradient inversion between the alpha-dipole’s poles. The coherence state diagnoses the approach to that condition; it does not cause it.

Persistent structures inhabit the regime $\mathcal{K} \approx 1 - \varepsilon$, with small $\varepsilon > 0$: close enough to the critical point

that the thermodynamic pressure drives elaborate self-organisation, far enough from it that the system has not yet collapsed.

The exact functional form of $\mathcal{K}$ is to be determined empirically. Equation (5) is the simplest candidate; the precise normalisation that yields the most informative diagnostic is calibrated against cosmological-configuration runs in companion work.

4.4 The Thermodynamic Engine of Self-Organisation

The qualitative behaviour of the triad encodes what we believe to be the central mechanism of EGT. As the system ages:

- Substrate sediment E accumulates—the fifth force concentrates energy at the archive.
- Matter energy Q diminishes—the matter sector loses energy to the substrate.
- Chronos C accumulates—events stack up as interaction continues.

The ratio E/Q grows accordingly: more sediment per unit of matter energy. This is the formal expression of a thermodynamic pressure: matter has less and less energy with which to maintain its structures, while the substrate has more and more accumulated history pressing on the matter.

Under this pressure, structures that survive must be increasingly efficient. Inefficient configurations do not accumulate enough Chronos through resonant pathways to persist; they are dispersed by the same pressure that selects for the resonant ones. Efficient configurations—those whose spectral signatures align with stable substrate channels—accumulate Chronos at higher rates and persist.

The result is a directed pressure toward self-organisation. We state its evidential status plainly: in the instrumented runs the coherence state rises monotonically after materialisation, with growth factors of order 10^3 – 10^4 per generation, but the rise is *not* uniformly convex, and the selection argument above — that with Q shrinking and E growing only increasingly efficient structures persist — is a mechanism hypothesis, measured as a strong tendency rather than proven as a necessity. This is the proposed mechanism of emergent complexity.

Possible biological consequences. If this thermodynamic engine drives the formation of structures generally, it should also drive the formation of structures of biological complexity, where the same principles—efficient resonance, persistent memory,

robust optimisation—are observed at a different scale. Whether the ECQ framework can predict the temporal scale of biological emergence (for instance, the age of a lifeform from its structural complexity) is an open research question. We do not pursue it here, but it is a natural extension of the present framework.

4.5 The Energetic Computational Capacity

A critical distinction must be made between the *Aeoncore* simulation and physical reality. While the simulation requires external hardware to compute topological changes, in the physical universe **energy is the hardware**. The complexity of the universe (the maximum number of nodes $|\mathcal{V}|$) is bounded by the total system energy, just as the simulation’s resolution is bounded by available computational memory. Real physical processes do not “simulate” interaction; the expenditure of energy against the resistance of Chronos accumulation *is* the computation. Physical reality is the runtime execution of the graph topology on the substrate of raw energy.

5 Scale Invariance and the Universal Grammar

The postulates introduced above make no reference to any particular physical scale. Spectral Vectors, resonant conductivity, and the ECQ-Triad are formulated as properties of nodes and edges, not as properties of cosmological or molecular regimes. This is not an incidental feature of the formalism but a substantive thesis of EGT: *the principles governing the emergence of stable structure are not domain-specific*.

What differs between systems at different scales is not the underlying grammar, but the scale of the nodes, the time horizon of the dynamics, and the set of interactions that dominate the topology. The same thermodynamic update that cools an expanding young universe is, formally, the same operation that cools any energy-conserving graph approaching its stable configuration. The same edge-tension dynamics that drive cosmological metric decay are, in their abstract form, the same dynamics that govern any system in which topology and energy density are coupled.

This claim is not reductive. It does not assert that biology reduces to physics, or that cognition reduces to neural topology. It is a claim about *structural invariance*: regularities in how complex systems stabilise hold across scales we ordinarily

treat as distinct disciplines. Acknowledging this invariance is what makes a single theoretical framework operationally meaningful at multiple scales. At the cognitive scale in particular, this connects to a long-standing position in theoretical neuroscience—that brain function rests on a dynamic graph of correlations among elements rather than on fixed connection strengths alone [12]—of which EGT’s resonant conductivity is the energy-conserving, scale-invariant counterpart.

The methodological consequence is significant. A result reproducible by the same formalism at distinct scales is stronger evidence for the underlying thesis than the same result obtained by disjoint methodologies. Unification is not asserted; it is operationalised.

The present paper develops EGT at the cosmological and particle-physics scales. These scales constitute the foundational level of physical organisation: every higher-order phenomenon presupposes the stabilisation already achieved at the level of fundamental constants, baryonic matter, and large-scale structure. Establishing the framework at this level is therefore a precondition for its extension to other domains. The applicability of EGT to systems at other scales follows from the scale-invariance of its postulates and is the subject of separate investigation; within the engine itself, however, the postulate is already measurable, and a first measured scale test is reported below.

5.1 A First Measured Scale Test: The Running Crunch Clock and the Export Efficiency

Scale invariance is not only a postulate to be exported to other domains; within the engine it is a measurable question, and it has now been measured. Four chained campaigns (internal identifiers K18–K18d in the released results kit; every trend criterion, seed set, and verdict tier declared and frozen *before* the corresponding runs) swept the initial energy over a full decade (scaling factors 1–10 on the fast cycle preset, node budgets from $\approx 5 \times 10^4$ to $\approx 5 \times 10^5$), on two dynamical arms — the present canon and a candidate rule stack under evaluation — with each quantity carrying a declared *mechanical null*: what the normalised transport rules trivially predict, so that only departures from that null count as signal.

The legacy negatives do not decide. The first sweep ($\times 1 / \times 3 / \times 10$, twelve runs) asked whether the measured negatives reported below — the helium yield, the correlation slope, the black-hole mass plateau, the isotropy index — migrate toward the

physical values at larger scale. We report the outcome as measured: *no quantity met the frozen “runs with scale” criterion, and no decidable trend toward the physical values exists over the accessible range.* The $\times 10$ cells were truncated by a compute decision after the first-generation crunch, capping the verdict machinery at two scale points (half a decade), on which every decidable quantity is consistent with stability; for several quantities (the canon correlation slope, the helium yield, the black-hole plateau) the instrument was too blunt to decide at all — and undecidable is counted as undecidable, never as stable. Finite-size rescue of the negatives is therefore neither supported nor excluded. Descriptively the decade is mixed: the candidate-arm correlation slope falls monotonically toward the physical 1.8 ($3.76 \rightarrow 3.30 \rightarrow 2.97$) but below the declared resolution floors, and a dedicated three-seed follow-up returned “not seed-robust on the fast preset” — an honest negative that closes that reading at this preset; the crunch time, by contrast, falls per-seed strictly monotonically over the full decade, *against* its scale-free mechanical null.

The crunch clock runs with the scale. A dedicated first-generation campaign made the crunch-time trend verdict-capable: six scale points $\times$ two arms $\times$ three seeds (36 points, a paired per-seed design; reused cells admitted only after exact identity-anchor replication gates — bit-identical crunch step under a pinned binary; both gates passed). The result is the engine’s first quantity to pass a frozen scale-running criterion: a *weak-tier* trend verdict (Kendall- τ criterion; not strictly monotone) on *both* arms. On the canon arm $\tau_b = -0.867$ in all three seeds (exact one-sided $P = 0.0083$ per seed), with per-seed endpoint spans of 1027/487/527 steps against a declared 100-step threshold; the seed-mean crunch time falls from 2854 to 2174 steps, -24% per energy decade (-12% on the candidate arm). Three fresh seeds later damp the rate to -8 to -16% per decade while the direction holds in all six seeds — the magnitude is more seed-dependent than the first three seeds suggested; the weak-tier verdict stands as frozen. Descriptively, larger universes crunch *hotter*: the absolute cooling anchors, where measurable, do not fall (the 50 MeV crossing drifts slightly later, with censoring caveats), while the temperature at step 100 rises from ~ 300 –410 MeV at $\times 1$ to ~ 800 –1470 MeV at $\times 10$; fine-cadence runs resolve the nominal first-generation maximum as a boot-injection transient (steps 4–12, ~ 0.7 – 5×10^5 MeV), leaving the absolute crossings as the only trustworthy cooling clock.

Mechanism localised. A descriptive diagnosis

on the same data locates the clock. In the settled phase the substrate-export channel is the practically sole feeder of the store pole (the store-energy increment equals the cumulative-export increment to $\approx 100.000\%$ in every coarse-cadence cell). The acceleration is *not* event-driven — annihilation, fusion, photodissociation, and decay rates all scale sublinearly (exponents 0.86–0.94) — but sits in the inflow itself: the store-pole inflow grows as $\text{scale}^{1.098 \pm 0.034}$ against an active-pool growth of $\text{scale}^{1.015 \pm 0.007}$, a relative excess that nominally accounts for ~ 70 – 80% of the decade effect. The excess sits in the *drainable* fraction: the rest-mass share of total energy falls with scale ($14.5\% \rightarrow 11.3\%$ at a fixed mid-generation checkpoint, with a declared phase-matching caveat), so the hotter big bang is itself the plausible accelerator — code-grounded and direction-consistent, not yet fully quantified.

A dimensionless carrier, discovered and then replicated. The diagnosis suggested a dimensionless reformulation, the *export efficiency* $\mathcal{E}_x \equiv \Delta L_{\text{ex}}/(\Delta t \cdot \bar{N})$ — cumulative substrate export per step and per unit of active-side energy. Because it was *discovered* on the campaign data, it was given no verdict there; instead a confirmatory wager was declared and frozen before measurement (a one-sided wager that the efficiency rises with scale, plus a frozen consistency band $[-0.01, +0.14]$ for the log-log slope on the canon arm) and tested on 36 virgin cells with three fresh seeds — no value known in advance, under an exact identity anchor and a pinned binary. The wager holds: *confirmed at the weak tier* — slope $+0.0612/+0.0594/+0.0535$ per seed, mean $+0.058 \pm 0.003$ inside the pre-registered band, Kendall $\tau_b \geq +0.6$ in all three seeds (the τ criterion alone carries a null probability of $\approx 3.2 \times 10^{-4}$); strict monotonicity, the strong tier, is missed. The export efficiency is thus the engine’s first dimensionless quantity to survive a discovery $\rightarrow$ pre-declared confirmatory replication chain on independent data, and it is the empirical entry point of the programme’s transmutation axis — the question whether the framework’s dimensionless quantities run with scale, which §9.5 identifies as the framework’s candidate mechanism for the missing gravitational hierarchy. The paired arm contrast replicated (canon minus candidate positive in 3/3 seeds), but the candidate arm *missed* its own declared band: its export efficiency *falls* with scale (-0.087 ± 0.042 , negative in all three seeds) — a new anomaly, reported with the same weight as the confirmation, and located below.

The running is carried by the drainable share. A descriptive decomposition on the same

cells (internal identifier K18e in the released results kit) resolves what the efficiency measures and pins down the accelerator that the mechanism diagnosis above had left code-grounded but not yet quantified. In its open-phase domain of validity the continuous export formula reproduces the measured efficiency to a few per cent (model minus measurement -0.7% to -4.5% , mean -2.4% , in all 37 cells), so the dimensionless slope is a property of the transport law rather than a boundary artefact. Resolved to first order — the covariance measured, not assumed — the running slope e_x (the log-log scaling exponent of $\mathcal{E}_x$) $= +0.058 \pm 0.003$ separates into a fill-trajectory term $e_{\text{gradient}} = +0.015 \pm 0.005$ and a drainable-fraction term $e_{\text{drain}} = +0.094 \pm 0.017$; the drainable share is the dominant carrier and in fact *over*-carries the running ($\sim 160\%$ of e_x), a negative covariance between the two trajectories (-0.028 ± 0.011 , beside a closure residual -0.023 ± 0.009) drawing the total back to the measured value. The chain “hotter big bang $\Rightarrow$ larger drainable surplus $\Rightarrow$ more efficient export” is thereby quantified rather than merely asserted. The same diagnosis *locates* the candidate-arm anomaly reported above: it lies neither in the drain populations (both rise with scale) nor in the boot phase (flat-to-positive on both arms) but in the short threshold segment just before the crunch, where the aggressive crunch-drain branch takes over; the stationary candidate export scales positive with scale, like the canon, and the negative sign is a cadence-confounded property of that pre-crunch segment alone.

From measured to derived: the engine’s first theoretical number. With the decomposition in hand, the exponent was made the target of a frozen derivation attempt (internal identifier K18f). The rule was fixed *before* the work: a strict *fit ban* under which every number entering a chain must carry a code anchor, an engine constant, or an independently measured state exponent, with no export-derived quantity admissible as an input (the ban held in all eight adversarial checks). The outcome is the engine’s first dimensionless quantity carried from measured to *derived* rather than re-measured: a predicted slope $e_x^{\text{pred}} = +0.059 \pm 0.012$ against the measured $+0.058 \pm 0.003$, at a declared depth of *consistency*. Four independent derivation paths were run; three viable ones — a self-consistent fill-rate ODE, an anchor-breaking audit integral, and a derived channel kinetics — converge centrally on $+0.059$ – $+0.066$, while the fourth, a budget partition, is reported *refuted*, its apparent band hit exposed as the mutual compensation of two errors. The mechanism is convergent across the three: the drainable

matter–radiation share per node runs with scale ($m_{\text{rad}}/N \propto \text{scale}^{+0.0998 \pm 0.0105}$) while the agitation-excess pool is scale-flat and merely dilutes, the running itself driven by count-bound budget items — rest mass tracking $n_{\text{matter}} \propto \text{scale}^{0.86}$, an edge tension of ~ 100 MeV per edge held scale-invariant, and a bind quantum of 275 — that follow the sublinear structure clock, with the absolute anchors (a CMB floor of 0.037 MeV, a radiation threshold of 0.1) numerically dead at $T/T_{\text{CMB}} \sim 10^4$. The best chain, the fill-rate ODE, predicts as a by-product the crunch-trigger time to $\leq 3\%$ without taking that time as input. The honesty is declared with the result: this is a depth-one *consistency* derivation, not a first-principles one — *why* the drainable share runs remains an input, the precision tier (a prediction inside $\pm 2\sigma$ of the measurement) is defensible by no chain, and the saturation of the local exponent across the decade ($+0.066 \pm 0.019$ at $\times 1$ – $\times 3$ falling to $+0.032 \pm 0.026$ at $\times 5$ – $\times 10$) is reproduced by none of them.

A frozen counterfactual bounds the reach.

The derivation’s mandatory out-of-sample test (internal identifier K18g) was declared as a wager before its runs: switching the agitation drain off — a radiation-only export, one override — should, on the best chain, raise the slope to $+0.080 \pm 0.011$ and lengthen the clock by $+25$ to $+28\%$. Run on 25 fresh cells under an exact identity anchor (bit-reproduced) with a charge census of zero on both arms, the point prediction is *substantially refuted*, and reported as such. The small universes ($\times 1$ – $\times 3$) no longer crunch at all — the southern fill saturates below the polarity-flip threshold of 0.5 — rather than merely running slower; the efficiency level falls by -73 to -79% instead of the predicted $\sim -20\%$; and where a crunch survives, from $\times 5$ upward, the clock lengthens by $+55$ to $+87\%$ with factor-of-two seed scatter. The frozen verdict cascade formally returned *undecidable* — crunch-less cells yield no exponent to fit — because its falsification thresholds had been written to guard only the “effect too weak” side, a named design hole, while reality broke out on the opposite, “too strong / regime change” side. The consequence is a bounded reach, honestly drawn: the agitation drain is not the flat diluter the best chain assumed but a marginally load-bearing southern feeder, so the *mechanistic* reading is downgraded, while the depth-one number — an interpolation within the measured decade ($\times 1$ to $\times 10$), and on a different arm — is left untouched. The counterfactual thus makes empirically sharp the depth-one/depth-two boundary the derivation had drawn by hand: the chain describes the band

and does not extrapolate. It also returns a new standalone finding — a sharp radiation-only crunch crossover between $\times 3$ and $\times 5$, so that cycle-ability itself becomes scale-dependent once the agitation drain is removed, the same per-node m_{rad}/N running seen from the other side. The method lesson is recorded with the rest: the retrospective structural cross-check passed while quasi-circular, and only the out-of-sample counterfactual discriminated and exposed the missing feedback physics — a derivation’s worth shows at the counterfactual, not in retrospection.

Limits. All of the above is one preset (the fast cycle), one generation (generation zero), and one energy decade. Runtime scales effectively as edges $^{-1.8}$ (the pre-declared estimate of edges $^{-0.55}$ was measurably too flat), a throughput factor of ~ 50 – 80 at $\times 10$; extending the lever beyond a decade requires event-driven (Gillespie-class) scheduling as a declared prerequisite. The candidate-arm export anomaly has since been localised to the short pre-crunch threshold segment — the stationary candidate export scaling positive with scale, like the canon — though it is cadence-confounded and not yet precisely quantified; the unlocated remainder of the candidate crunch-clock run, the depth-two question of *why* the drainable share runs, and the exponent’s unexplained saturation across the decade remain open.

A first renormalisation ladder, actually climbed. The magnitude falsification of §9.5 elevates scale-dependence from robustness check to candidate mechanism: if the missing 10^{19} hierarchy is to emerge at all, some effective coupling must *run*. This has now been measured directly (internal identifier K35, declaration — blocking rule, aggregation, the full coupling list, and kill criteria — frozen before the first analysis line). A metric Kadanoff blocking was applied to 30 ensemble snapshots (10 sink-free cells $\times 3$ epochs, $\sim 3.5 \times 10^4$ nodes and 3.3×10^5 edges late): at each level the metrically nearest pairs — the highest-density edges, density being the inverse metric distance — are contracted until the node count halves, extensive quantities summed, parallel cross-edges merged as conductances, roughly ten levels deep. Exactly one coupling runs stably: the mean connectivity, at -0.312 e-folds per halving (inter-quartile range -0.320 to -0.308 ; sign-consistent in 30 of 30 snapshots across all seeds and epochs; median $R^2 = 0.76$ for the log-linear model). The substrate’s topology is renormalisation-self-similar with a stable exponent — a constant $\sim 73\%$ of connectivity survives each halving — and the exponent is structure-specific by measurement, not by argument:

degree-preserving randomisation of the same graphs (double-edge swaps, edge attributes carried with their records) decays at only -0.17 e-folds per halving under the identical blocking, eleven inter-quartile widths from the substrate’s value, so roughly half the decay is generic to the degree sequence and rule while the substrate’s hierarchical clustering contributes the rest. The energy couplings, by contrast, do *not* run: the effective tension-per-length moves by factors above 10^2 across the ladder, but the pre-registered robustness pass resolved this against a coupling reading — the declared bundle aggregation amplifies multiplicity quadratically, and under per-member normalisation the flow loses all ladder stability (sign consistency 60%, deeply negative log-linear R^2): the apparent cold-era stiffening was substantially a bundling artefact, and no energy coupling runs on this ladder, even loosely; the binding energy internalises predominantly at the finest levels (a characteristic binding scale, not the log-uniform spectrum scale invariance would prefer); and the effective-dimension ladder is instrument-limited beyond the first level (ball growth saturates on blocked graphs; the level-zero values, 5.3–5.9, validate the instrument against the independently measured ~ 5.8). The verdict is stated plainly: the substrate carries a genuine RG-stable topological exponent — a new structural constant — but the energy-scale running that dimensional transmutation would require has not been found in the current operator set. The missing hierarchy remains open, now with the first ladder measured rather than conjectured.

6 Implementation: Aeoncore

The theory is modeled in *Aeoncore*, a scale-invariant, Rust-based simulation engine where energy flows are fully spectral and the governing laws apply identically regardless of graph size, scaling only with available computational resources. The engine does not search for low-energy configurations by gradient descent or heuristic minimisation; it lets configurations emerge from physical dynamics that respect the conservation laws of the underlying system.

Tested system scale. All results reported in this paper come from a declared and, by astrophysical standards, small regime: graphs of order 10^4 nodes (peaks near 2.5×10^4 in the reference runs), matter sectors of 10^3 – 10^4 quanta, run lengths of 2×10^4 to 1.7×10^5 steps, seed ensembles of $n = 8$, and up to three full expansion–collapse cycles within a single run; the scaling campaigns of §5.1 extend the initial energy and node budget by a further decade (node budgets to $\approx 5 \times 10^5$, seed sets of two to three per

cell, six on the crunch clock), for the first generation only; the bridge-signature campaigns of §10 add multi-generation ensemble runs of 1.2×10^5 steps (15 cells) and passive-observation cells at three substrate scales up to $\approx 1.1 \times 10^5$ nodes. Whether the quantitative gaps reported below — the helium yield, the isotropy index, the correlation slope, the black-hole mass scale — are regime-limited, i.e. whether they shift at larger initial energies and node counts, has now been swept over a decade of initial energy (§5.1): on the decidable range (half a decade) no trend toward the physical values was found, the full-decade criterion could not be triggered for the legacy negatives, and the one quantity that does run with scale — the crunch clock — is not among them; runtime grows steeply with scale (effectively edges^{-1.8}), and no result here is claimed beyond the stated range. At this size the system’s nearest laboratory reference class by size and temperature is heavy-ion physics rather than cosmology — a correspondence that has since been tested on all three of its branches and closed as a regime claim (see the note on comparison regimes below).

6.1 Thermodynamic Integrity (Zero Drift)

Unlike typical simulation engines, *Aeoncore* operates on strict conservation laws comparable to Lattice QCD simulations. The total energy of the system is tracked with double-precision floating-point accuracy:

$$E_{\text{sys}} = \sum (mc^2 + E_{\text{kin}} + E_{\text{rad}}) + \sum E_{\text{tension}} \quad (6)$$

Every fusion event, particle decay, and metric expansion is balanced as a transaction: energy is only ever moved, never created or destroyed—the same increment removed from one node is added to another. Conservation is therefore a property of the construction rather than a tolerance imposed after the fact. Each transaction balances at the machine-precision floor (per-step residual of order 10^{-10}). On the current reference baseline (§11.5) the measured *global* drift over complete 2×10^5 -step cycles — including the emergent collapse and re-seed — lies between 5×10^{-11} and 9×10^{-10} across all fourteen cells, with the charge ledger closing to zero exactly and the four-sector energy partition closing to $\leq 3 \times 10^{-13}$. (An earlier build’s 0.012% full-cycle figure was float accumulation in the hot plasma era; the audited engine retires it.) The figure is insensitive to the choice of gravitational funding rule (identical to five decimals under both variants tested), indicating that it reflects numerical accumulation rather than physics.

This places *Aeoncore* in the same conservation class as lattice-QCD simulations.

The system also tracks the ECQ triad of §4 as a global metric, partitioning the total energy into the matter component Q , the substrate sediment E , and the vacuum potential E_{vac} , with chronos C accumulated as the cumulative event count.

6.2 Calibration: The Boot Sequence

Calibration in *Aeoncore* separates two classes of quantity explicitly: *genuinely calibrated* scales, where the target observable is demonstrably sensitive to the searched parameter, and *declared anchors*, where a physical constant is taken as input and stated as such. An earlier calibration scheme was retired after an internal falsification test showed the searched parameter had no influence on the target mass; the present scheme replaces it.

1. **String-tension search (genuine):** the engine instantiates the semiclassical string hadron of §7 — three quarks bound by a Y-shaped string junction — and performs a binary search on the string tension σ until the minimised hadron energy matches the physical proton mass (938.27 MeV). The target is genuinely sensitive to the parameter ($E \sim \sqrt{\sigma}$), and the result is bracket-independent.
2. **Declared anchors (honest inputs):** the fine-structure constant $\alpha^{-1} = 137.036$ is a CO-DATA input, *not* an emergent output — an earlier circular derivation was identified and removed. The neutron–proton mass splitting (+1.293 MeV) is likewise anchored, because the current-quark string model alone predicts the wrong sign (§7); the model’s own value is printed as a diagnostic at every boot.
3. **Derived scales:** interaction thresholds (neutron lifetime, fusion barrier, metric break point, Chandrasekhar-type collapse threshold) are then derived from the calibrated and anchored values, keeping the simulation internally consistent with its own metric grammar.

6.3 Parameter Accounting

How many free parameters does the model have? The honest answer requires four distinct classes, and we state all of them. *Empirical anchors* (seven scalars, never fitted to simulation output): $\alpha^{-1} = 137.035999084$, the proton-mass calibration target for the string tension, the n – p splitting of +1.293 MeV (declared as an anchor because the

string model’s own prediction has the wrong sign), m_e , m_u , m_d , and $\tau_n = 880\text{s}$ — plus two empirical *input tables*: the semi-empirical mass formula with light-nuclide binding energies, and a chemical bond-energy table. The nuclear binding landscape is therefore imported, not emergent, and every nucleosynthesis result inherits it. *Chosen continuous constants* (approximately 26): the cycle-defining rates found by parameter scan (the export rate, the cooling rate, and the crystallisation-onset threshold Ω), the declared mode constants (τ_p , the black-hole lifetime and horizon parameters), the five hand-weighted couplings of the space equation (registered as free pending a first-principles action), the declared substrate scale of 275 MeV per space quantum, and a handful of hand-picked coefficients inside otherwise anchor-derived formulas. *Structural canon rules* (four, adopted after falsification-driven testing): exact charge conservation, apex-funded topological gravity, full-drag substrate export, and two anti-runaway topology gates that are declared scaffolding to be retired. Finally, we state plainly what a code audit will find beyond the configuration surface: the thirteen operator cadences of the discrete scheduler are themselves free parameters (the code declares them as such), and on the order of forty dimensionless rate, branching, and threshold coefficients remain hardcoded in the operator implementations — an electron-capture probability of 0.2, the fusion probability law, decay energy splits. None of these has been shown to be forced by the theory. Counted narrowly (canon-active, continuous, not empirically anchored) the model has ~ 26 free physical constants; counted as a hostile reviewer would, including operator hardcodes and cadences, the number is of order 70–80. Reducing this set — event-driven scheduling in place of fixed cadences, and a single-action derivation of the space-equation couplings — is declared future work, and the direction of travel so far has been encouraging: the recent canon steps (topology pricing, apex funding, charge conservation) each *removed* behaviour that previously had to be capped by hand, and the expansion sector has since derived one chosen constant (the matter-expansion index, an adiabatic exponent — validated) and demonstrated the feasibility of replacing a second (the expansion rate as a Friedmann-type function of the emergent energy density — feasibility confirmed, structurally separated from the sink candidate); each removes rather than adds a degree of freedom (§9.5); the derived matter-expansion index became canon with the 2026-07-11 configuration flip (below), and the Friedmann rate followed on 2026-07-18 once the

sink’s removal had dissolved the documented incompatibility — flipped, like the index, for the derived kinematics and the removed degree of freedom, with no yield claim (§9.5).

6.4 Emergent 3D Space

Aeoncore does not simulate particles moving inside a 3D coordinate system. Instead, it simulates a high-dimensional graph. The 3D positions (x, y, z) observed in visualisations are emergent projections resulting from the force-directed layout of the graph topology. Space creates itself where node density and edge tension interact, forming clusters that we perceive as spatial dimensions. Once emergent geometry has stabilised, it can act back on the topology through the fifth force—via spatial edge discovery (Postulate IV), enabling new edges between geometrically near nodes.

Two notions of dimensionality must be kept distinct here. The *embedding* used for geometry-coupled dynamics and visualisation is three-dimensional by construction. The *intrinsic* dimension of the graph topology, measured by BFS ball growth in matter-era states, is higher (of order 5–6 in cosmological runs). The claim of this section is therefore not that the substrate topology is three-dimensional, but that a stable three-dimensional geometric shadow emerges and becomes dynamically active; the relation between intrinsic graph dimension and the effective dimensionality of physics on the graph is an open, quantitatively tracked question.

7 The Hadronic Sector: Edges as Strings

The scale-invariance thesis of §5 predicts that the same edge grammar that carries cosmological structure should also carry the physics of bound matter at the smallest simulated scale. The hadronic sector is where this prediction is tested most sharply, because quantum chromodynamics provides an unusually clean effective description to compare against: at hadronic distances the colour field between quarks collapses into a flux tube of constant tension — a string. We emphasise the scope of the claim from the outset: this is *not* a fundamental string theory or a theory-of-everything claim. It is the observation that the EGT edge — an energy-carrying tension relation — and the QCD flux tube obey the same effective grammar, and that this identification is quantitatively productive.

7.1 The String Reading of the Edge

An edge with tension T and geometric length R carries energy σR with σ the tension per unit length. A baryon is modelled as three quarks joined by a Y-shaped junction of such edges; the semiclassical ground state minimises

$$E(R) = \sum_q \sqrt{m_q^2 + (\hbar c/R)^2} + 3\sigma R + E_{\text{Coul}}(R), \quad (7)$$

the relativistic kinetic term expressing localisation on scale R , the linear term the Y-junction string energy. A single binary search on σ against the proton mass fixes the confinement scale ($\sigma = 24.45 \text{ MeV}$ per length unit at confinement radius $R^* = 6.40$); everything below follows without further fitting.

7.2 Regge Trajectories

The rotating-string spectrum is the classic signature of confinement phenomenology [5]. Integrating the EGT string numerically (relativistic leapfrog on the edge grammar) yields angular momentum rising as the square of the mass,

$$J = \frac{M^2}{8\sigma}, \quad (8)$$

with the measured trajectory linear in M^2 at $R^2 = 1.000000$ and slope matching $1/(8\sigma)$ exactly; the constant offset relative to the Nambu–Goto value $1/(2\pi\sigma)$ is the known geometric factor between the Y-junction and open-string configurations. We state precisely what this does and does not establish. A linear tension relation rotating relativistically *will* produce $J \propto M^2$; the linearity is therefore a property of the model we wrote down, not a discovery, and the near-perfect R^2 measures the integrator rather than the physics. What is genuinely not injected is that the *same* edge grammar — the tension relation the cosmological runs use, with no hadron-specific term added — reproduces the confinement phenomenology at all, and that its slope is fixed by a single anchor with nothing left to tune. The falsifiable content sits in the magnitude, which is measured next and misses.

The slope in physical units, and where it lands. That the trajectory is linear is a structural result; whether its *magnitude* matches nature is a separate question, and one the calibration does not settle — σ is fixed against the proton mass, so the slope is a prediction downstream of that single anchor. Converting through the engine’s own constants ($\hbar c = 1000 \text{ MeV}\cdot\text{len}$ fixes the length unit at 0.1973 fm , with no assumption about the proton radius), the calibrated tension is

$\sigma = 0.0244 \text{ GeV}^2$ (124 MeV/fm) against the lattice value of $\approx 0.18 \text{ GeV}^2$ (900 MeV/fm), and the Regge slope is $\alpha' = 5.11 \text{ GeV}^{-2}$ against the light-meson value of $\approx 0.88 \text{ GeV}^{-2}$. The string is therefore about seven times too soft, and the trajectory correspondingly about six times too steep — the two deviations are one number seen twice, related by the 8-versus- 2π geometric factor above ($7.4/(8/2\pi) = 5.8$), which is itself a consistency check on the conversion. We record this as a characterised negative of a different class from the gravitational magnitude falsification of §9.5: there the failure was four to five decades and structural, here it is a factor of six in an otherwise exactly reproduced law. The likely locus is the same absence the neutron–proton splitting exposes below: in QCD the bulk of the nucleon mass is chiral-condensate energy rather than string tension, so a model that carries the full proton mass on a string necessarily assigns that string too little tension per unit mass. Chiral dynamics is thus not one honest gap but two, and the second one is now quantified.

7.3 Topology Selection: Y over Δ

For three quarks two string topologies compete: the Y junction and the Δ triangle, and three determinations were originally cited as agreeing on the winner. An internal code audit has since demoted one of them, and we report the demotion rather than the original triple. (i) The semiclassical energetics — the Y minimum lying below the Δ minimum — is *not* an independent determination: in the implementation the Δ energy is formed from the Y energy by construction, $E_\Delta - E_Y = 3\sigma r(\sqrt{3} - 1)$, which is positive for every tension and every radius. The comparison is an algebraic identity of the parameterisation, its verifying test unfalsifiable by construction, and it carries no evidential weight. What survives is a pair, not a triangle. (ii) Lattice QCD measures the three-quark potential to be Y-shaped [6]. (iii) The substrate itself: inserted as a live subgraph into a running cosmological simulation — full operator pipeline, no protective flags — the Y-configured hadron entered at exactly its semiclassical ledger value (938.27) and remained fully bound for over 1,800 steps in the hot plasma era, while the Δ configuration unbound completely within 300 steps, for the reason predicted *before* the run from the engine’s edge-rupture thresholds (vacuum-touching junction edges inherit the metric break point; direct quark–quark edges fall to the pair-production floor). That dynamical selection is a genuine result and agrees with the lattice; the analytic third leg was bookkeeping wearing the costume of a test.

7.4 The Uncertainty Floor and Constituent Mass

The substrate initially lacked a ground-state protection: resonant transport could drain a bound quark’s localisation energy, so the semiclassical ground state was not a fixed point of the dynamics. The remedy is an *extraction constraint*, not a new energy term: for a node whose strongest binding is confinement-class, thermal channels may not draw its agitation below the localisation floor $\hbar c/r - m$. With this constraint active, the bound state’s kinetic sector becomes a genuine substrate fixed point (bound-quark agitation settles on the floor instead of eroding below it). Two consequences deserve emphasis. First, the effective energy of a bound quark becomes $m + (\hbar c/r - m) = \hbar c/r$: constituent-quark inertia emerges from localisation alone, with no fitted constituent mass. Second, precisely because the current mass cancels, the floor cannot by itself produce the neutron–proton splitting — the model value has the wrong sign, which is the expected signature of absent chiral dynamics. In EGT terms the chiral question becomes geometric and falsifiable: a physical n – p splitting requires an *emergent radius asymmetry* between up- and down-type localisation, not a fitted constant. This reformulation, together with the declared anchor of §6.2, is stated as an honest limit of the present model.

7.5 Scope

The hadronic sector thus supplies the microphysical end of the scale-invariance argument: one grammar — edges as tension relations — yields Regge phenomenology, lattice-consistent baryon topology, and constituent inertia at the hadronic scale, while the same grammar, uncalibrated further, produces the cosmological phenomenology of the following sections. What it does not supply is flavour physics, chiral symmetry breaking, or gauge structure; these remain outside the current model and are documented as such.

8 The Four Interactions as Graph Operators

Before the cosmological mapping, we state explicitly how the four standard interactions are realised, because the realisation is itself a thesis: in EGT, forces are not exchange particles but *conserving operators* — scheduled transformations of graph state, each drawing on and paying into the same energy ledger. Every operator below is dimensioned by the boot calibration of §6.2 (one genuine fit, the string tension

against the proton mass; declared anchors stated as such) and audited by the same conservation gates.

The strong interaction is the edge itself. An edge is a tension relation; at hadronic scale it *is* the QCD flux tube, and §7 develops this bank in full: a genuinely calibrated string tension, an exactly linear Regge trajectory, lattice-consistent Y-junction topology, and constituent inertia from localisation alone. At the nuclear tier the same grammar continues as the binding account: fusion founds composite nodes against an imported binding landscape (a declared input table), and the edge-break threshold consults the same landscape. The strong sector is where the framework is quantitatively strongest and its mapped distance is precisely known (the slope magnitude, factor six, §7).

The weak interaction is the decay operator family: beta decay of the free neutron against the physical lifetime anchor $\tau_n = 880$ s, electron capture, and proton decay as the late-era clock. Each is a discrete, ledger-balanced transformation of a node's identity — charge, nucleon content, mass, and recoil all move in one transaction. This sector has recently been audited to an unusual depth: a bookkeeping defect in beta decay (the nucleon counters of an identity pair not moving together) was found by census, fixed, and the fix re-measured — with consequences reported in §11 that propagated through the collapse gate into the cycle clock itself, and that *unlocked* the deuterium channel whose absence had previously been read as physics. The weak sector is thereby the cleanest demonstration of the framework's methodology: every operator's bookkeeping is exact, checkable, and consequential.

Electromagnetism enters as three structures. First, charge is the ledger's second exact invariant: the conserved-charge census plus the declared archive account sums to zero *exactly* — not approximately — through complete cycles including collapse and re-seed, measured across the full reference campaign. Second, the radiation field is spectral: photons live as frequency-banded radiation on nodes, and the radiative back-channel — photo-disintegration — is the parameter-free Boltzmann tail $\exp(-S/T_{\text{bath}})$ of the ambient bath against each nuclide's separation energy, the operator arc of §9 having converged on it precisely because it contains nothing to tune. Third, Coulomb structure acts inside the hadron model (E_{Coul} in the string Hamiltonian) and as charge-gated behaviour of the reaction operators. What electromagnetism is *not*, at present, is a long-range gauge field on the graph — a defined construction stage, stated here rather than implied away.

Gravitation is implemented at three tiers, and none of them is a force arrow between particles. (i) *Mass curves the metric, operatively*: topological gravity is a mass-gated friend-of-friend triangulation — where two massive nodes share a neighbour, a new direct edge is funded from the parent edges' tension (with a declared fraction fundable from the apex node's extractable agitation: kinetic energy converting to structure, the Jeans mechanism at operator level). Since edge density *is* the inverse metric, attraction is literally the densification of space around mass, without any energy travelling from one body to the other. The reference campaign measured the resulting wells directly: the mean edge density on mass concentrations is $8\times$ the vacuum baseline, decaying to background within two to three hops — a gravitational potential profile, read off the substrate. A frame-dragging channel (massive apex, tenth-strength) exists in the same operator. (ii) *Collapse is emergent*: no stellar types are placed. A concentration crossing the Chandrasekhar-scale mass and critical density is resolved against three counter-pressures — electron degeneracy, neutron degeneracy, or none — yielding white-dwarf, neutron-star, or black-hole outcomes respectively, with supernova ejecta driving a kinetic burst footprint. The reference campaign's census (121 collapse events) resolves to neutron stars in all but one case — the sole white-dwarf outcome, the cascade's second pressure branch, appearing in the campaign's densest cell — from progenitors of $A \approx 40\text{--}64$, thereby terminating the nucleosynthesis ladder physically (§11). (iii) *Cosmological gravity is the Friedmann operator* of §9.5: H derived from the ledger's own energy density over the vacuum population, the volume being the vacuum-node count itself. The absolute strength of gravity — the 10^{19} hierarchy — is the framework's declared central open problem (§10); everything else in this list is running physics.

The fifth force — the substrate's own activity — has its operator inventory in Postulate IV and its cosmological reading in the next section. The table there (Table 1) assigns each graph entity its observable and its force class.

9 Cosmological Mapping

The graph topology of *Aeoncore* maps onto standard cosmological observables. The mapping is not arbitrary but follows from the postulates: vacuum substrate as the active medium, matter-class forces between connected nodes, fifth-force activity throughout the substrate.

Table 1: Graph–cosmology mapping. Items under “fifth force” are unified manifestations of substrate activity (Postulate IV).

Graph entity	Cosmological observable	Force class
Vacuum substrate	Dark matter halos	Fifth force
Substrate dynamics	Dark energy	Fifth force
Substrate expansion	Cosmological expansion	Fifth force
South-pole sediment	Vacuum energy archive	Fifth force
Bridge tunnelling	Long-range substrate transport	Fifth force
Edge density	Gravitational coupling	Matter (gravity)
Edge tension (matter)	Local field stress	Matter (electromagnetism)
Node mass clustering	Baryonic structure	Matter (gravity, EM)
Alpha-dipole bipolarity	Cosmic polarity	Substrate topology

9.1 Vacuum Substrate as Dark Matter — Now a Measured Share

The vacuum substrate carries energy in the same ledger as matter and gravitates in the Friedmann source density, but — since the lake sink left the canon — it has *no radiative export channel*: it holds energy and curves, and does not shine. That is the operational definition of dark matter, and in EGT it is not a postulated component but a structural consequence of space being a particle species. The reference campaign measured its share of the cosmological (matter-plus-vacuum) density directly from the ledger partition: 39–59% at boot, rising through the cycle to 45–100% at the approach to collapse. The gravitating, non-luminous sector *dominates* the density budget for most of cosmic history — the same qualitative statement observational cosmology makes of our own universe, obtained here with no dark component inserted.

9.2 Substrate Activity as Dark Energy — A Λ -Slot Without a Λ

The dark-energy reading is sharper still, because it is structural. The Friedmann source density is $\rho = (E_{\text{matter}} + E_{\text{vac}})/n_{\text{space}}$. The matter term dilutes with the growing volume; the vacuum term is $n_{\text{space}}\bar{e}_{\text{vac}}/n_{\text{space}} = \bar{e}_{\text{vac}}$ — the mean energy per vacuum node — which does *not* dilute volumetrically: more space brings its own energy with it. That is precisely the structure of a cosmological constant, and no Λ parameter exists anywhere in the code to produce it. The reference campaign tested the consequence: late in the cycle the measured H becomes vacuum-dominated and sits on the parameter-free floor prediction $H_{\text{floor}} = H_0\sqrt{\bar{e}_{\text{vac}}/\rho_0}$ at ratio 1.10–1.20 across the ensemble — consistent, through $\sqrt{1/\text{vacuum share}}$, with the measured 70–80% vacuum share at that epoch. What does *not* occur is a de-Sitter phase: in every cell the crunch closes the cycle before H flattens onto the floor —

in this cosmology, the cycle beats Λ . Two honest notes accompany the result: the FRW vocabulary is a reading (EGT has no Einstein tensor), and at the detail scale the measured ratio falls below unity, which is impossible if the diagnostic partition and the Friedmann source were identical — a bookkeeping-reconciliation question between two ledger views, flagged and open.

9.3 Black Holes as Substrate Concentrations

Black holes, in EGT, are extreme local concentrations of substrate sediment, formed where matter accumulates beyond a topological threshold and engages the fifth force at maximum local intensity. In the implementation this is realised without any placed black-hole object: the horizon is a *state*, re-derived each pass from local substrate concentration (mass above a nucleon-count floor and near-singular incident edge density), and carries four channels — accretion from the densest infall edge, bipolar jets perpendicular to the infall, Hawking-type evaporation with $\dot{M} \propto 1/M^2$ ending in a final burst that returns the node to ordinary space, and bridge tunnelling into the south-pole archive. The black-hole information paradox does not arise: the spectral content of consumed matter is preserved in the local substrate, contributing—through bridge tunnelling—to the broader south-pole accumulation that defines the universe’s integrated archive (§3).

Under the current canon these horizon states are measurably emergent again — they had been suppressed by the lake sink’s cycle acceleration and returned with the 2026-07-16 flip (the nucleosynthesis arc above): the sink-free ensemble campaign contains 120 horizon-node observations across 27 distinct nodes, seed-stochastic as expected. This population supports a first Bekenstein-type reading [24] (internal identifier K36, declaration frozen before the analysis; the graph-native surface of a horizon node is its set of near-singular contact

edges — discrete area quanta — its content the swallowed mass, its information the node’s spectral signature). Boundary growth correlates with swallowed content across 94 consecutive-checkpoint transitions: $\rho(\Delta A, \Delta M) = 0.22$ with bootstrap CI (+0.005, +0.428) on the saturation-weighted area measure — weakly positive, in the same confirmed-weak class as the export-efficiency exponent of §5.1, and honestly noisy at the checkpoint cadence, where accretion, evaporation and tunnelling average together. The area-law question proper — does the stored information track the boundary rather than the content? — is a characterised negative. The initial reading used the spectral signature as the information measure and found it collapsed (65% of horizon observations at ≤ 1 band); the pre-registered interpretive fork (checkpoint artefact versus genuine archive drain) was then resolved not by more simulation but by code inspection: the collapse is neither, it is the engine’s own live spectral-maintenance operator, which compacts every node’s spectrum each cadence below a fixed energy floor for performance — so live and serialised reads see the identical floored spectrum, and spectral entropy is structurally unusable as a horizon information measure here. The test was therefore redirected onto the one uncompacted proxy available, the compositional entropy of a horizon’s neighbourhood. On it the measure is well defined (27% degenerate, not 65%), and the verdict is null: horizon information tracks neither the boundary ($\rho = -0.19$ against patch count, bootstrap CI (−0.41, +0.04), crossing zero) nor the content ($\rho = +0.06$ against mass). Boundary *growth* with swallowed content stands (Q1 above); a holographic *area law* for information does not appear on the measures the engine exposes — whether a better information proxy would reveal one is left open. The Hawking law itself, being coded rather than emergent, is excluded from every claim in this paragraph.

9.4 Measured Large-Scale Structure and Honest Negatives

The cosmological configuration has been measured against the standard observables of structure formation, with mixed — and deliberately reported — results. The matter two-point correlation function follows a power law over the accessible range with $R^2 = 0.98$, i.e. the clustering is genuinely scale-free, but with a steeper index ($\gamma \approx 4$ –5) than the observed galaxy value $\gamma \approx 1.8$: the emergent structure is more compactly segregated than the real cosmic web. This scale-free regime is itself era-dependent: in the slow-cooling configuration the late cold state

consolidates into a small number of massive halos (the mass-clustering correlation rises from ≈ 0.1 to ≈ 0.7) and the power law gives way to a few-object distribution. A friends-of-friends halo catalogue exists and its cumulative mass function is itself a power law; minimal-spanning-tree statistics show the clump–filament–void signature rather than a random gas. Against these positives stand three characterised negatives. Primordial nucleosynthesis is a characterised negative in both accessible regimes, in opposite directions. The fast cycle configuration overproduces heavy elements and underproduces helium ($Y_{\text{He}} \approx 0.10$, $Z \approx 0.3$ against the observed 0.245 and ~ 0) — the fusion chains run too deep in the dense early phase. The slow-cooling configuration has now been measured and does *not* approach the physical ratio from the other side: metal overproduction essentially vanishes ($Z \approx 0.03$) but the helium yield collapses with it ($Y_{\text{He}} \approx 0.013$, $X_{\text{H}} \approx 0.95$). The helium fraction scales with the cooling rate, and the observed point — helium-rich yet metal-free — lies outside the regime accessible to either configuration. Enforcing charge conservation (now canon) roughly doubles the slow-cooling helium yield ($Y_{\text{He}} 0.013 \rightarrow 0.028$, $X_{\text{H}} 0.95 \rightarrow 0.91$) — a prediction declared *before* the measurement (surviving electrons feed electron capture and deepen the fusion throughput), moving the yield toward the observed value with no tuning; the remaining gap to 0.245 stands as reported (the formation-channel results below change the mechanistic reading of this gap without closing it quantitatively). Causal-cone isotropy in the hot regime is low (isotropy index 0.01–0.03; signal propagation strongly anisotropic), so no emergent Lorentz-type isotropy can be claimed there. The cold late era — the fairer test — has now been measured: the isotropy index rises by an order of magnitude, to 0.25–0.31, while the signal-speed dispersion falls from $\text{CV} \approx 2$ to ≈ 1 . Isotropy is thus an equilibration trend that grows as the substrate cools, but it remains far from the full isotropy an emergent Lorentz symmetry would require. Finally, the particle-count asymmetry deserves a corrected reading. The raw node-count index is negative (of order 20% hot, 40% cold — nominally antimatter-heavy), but a sector decomposition shows this to be a lepton bookkeeping effect, not antimatter dominance: the baryonic sector is essentially pure matter (baryon-number-weighted index +0.99 to +1.0; no anti-baryons survive into the cold state), while the count is dominated by positrons and antineutrinos — the Standard-Model-like by-products of the matter sector’s own processing (each exothermic pp-chain fusion emits $e^+ + \nu$, each beta decay an

$\bar{\nu}$, each proton decay an e^+). The index therefore *anti-correlates* with physical matter dominance; its cold-era deepening is the proton-decay washout, not growing antimatter, and its sign is phase-dependent (early, electron-rich phases measure positive) while lying far beyond seed noise ($\sim 40\sigma$), i.e. structural. Two honest qualifications follow. First, the baryonic matter dominance is itself declared, not emergent: the hydrogen pair-production channel creates (proton, electron) pairs without anti-partners, so the framework does not address the baryogenesis question — it assumes its outcome, and the symmetric pair channel’s output annihilates completely. Second, the positron accumulation was traced to a charge-non-conserving electron sink (electrons were consumed with their host nuclei without a compensating charge ledger). Charge conservation has since been promoted to a second ledger invariant of the canon: the graph’s conserved-charge census plus a declared archive account (black-hole accretion, Hawking bursts, the crunch archive) sums to zero exactly, measured through a full cycle including the crunch and re-seed. Under it the lepton sector normalises (electron-dominated, the stranded positrons drop by two orders of magnitude), fusion activity rises by roughly a fifth, and the emergent cycle shifts by under two percent — the anomalous count asymmetry was, in the end, the shadow of a single missing ledger entry. These negatives are reported with the same weight as the positives: they delimit what the framework does and does not currently reproduce.

A note on comparison regimes is owed here. At 10^3 – 10^4 matter quanta and a hot-era temperature near 10^{12} K (~ 90 MeV, just below the QCD hadronisation scale), the system’s nearest laboratory reference class by size and temperature scale is heavy-ion physics — the “Little Bang” of chemical freeze-out, light-cluster yields, flow anisotropy, and femtosopic correlations — rather than FLRW cosmology. Read in that class, part of the gap table above is a category mismatch of the reference rather than a dynamical failure: no fireball of this size carries cosmological helium fractions; the measured deuteron photodisintegration–refusion cycling is the light-cluster survival problem of heavy-ion physics; a system built on a polar dipole has a preferred axis by construction, making flow-style anisotropy coefficients the regime-appropriate reading of the isotropy index; and the correlation slope belongs against cluster–cluster statistics, not the galaxy two-point function. The cosmological observables quoted above remain the declared *targets of the scaling axis*; they were never the natural benchmarks of the

present regime. This reframing was itself declared testable at today’s scale, and it has now been tested on all three of its branches — equilibrium statistics, freeze-out systematics, and kinematics — with a uniformly negative or non-load-bearing outcome. We report the closure with the same weight as the reframing.

The statistical branch fails first. A nuclear-statistical-equilibrium consistency test — Saha abundances computed with the engine’s *own* binding energies along the engine’s own measured (T, ρ) trajectory, so that it probes the statistical mechanics rather than the imported landscape — returns a three-zone verdict: in the classically valid middle zone ($T \approx 12$ – 75 MeV, $n_p \lambda^3 = 0.08$ – 0.5) the deuteron population lies robustly below equilibrium in all three runs and under both volume conventions — zero deuterons where NSE expects 13–832, $P(0) \leq 2.3 \times 10^{-6}$ down to $\sim 10^{-36}$; cold ($T < 1.5$ MeV) the deuteron-to-proton ratio freezes onto run-specific plateaus at least 1.96 dex below NSE under any volume choice, without ever having been on the equilibrium curve (the multi-dex nominal magnitudes there are classical extrapolations; the direction is robust); hot ($T > 100$ MeV) the same direction holds but is count- and volume-limited, hence undecidable. There is no freeze-out in the technical sense: the pre-declared tracking criterion never emitted a freeze-out temperature in any of the six cases — what happens cold is reaction quenching without prior equilibrium. The weak sector was never in equilibrium either (n/p a factor 3–10 below $\exp(-\Delta m/T)$ already in the hot phase). The mechanism is operator-sharp, not a mystery: the forward fusion rate carries no phase-space factor and consults the binding energy only as an endothermy veto, and the back channel — photodisintegration as a hard 15 MeV threshold on node heat — is blind to binding energy, subjecting the deuteron ($B = 2.22$ MeV) and ^4He ($B = 28.3$ MeV) to the identical destruction law, so no Boltzmann ratio can form.

The freeze-out systematics confirm this at the pattern level. The heavy-ion data show a near-constant per-nucleon penalty factor (307 ± 76 in Pb–Pb at the LHC, falling to a preliminary 54.0 ± 4.7 at the lowest RHIC beam energy); the engine instead shows one large step from protons to deuterons and then essentially none — $S_{23} = 0.36$ – 0.93 across all four runs against ~ 390 in the data, a Poisson-solid rejection ($P \approx 10^{-13}$ even in the count-weakest run) — with a geometric-mean penalty of 3.8–8.7 per nucleon, a factor 6–80 below the lowest published heavy-ion value. The isospin branch is off by two

decades: the engine freezes out at $n/p = 0.0075 - 0.021$ against ≈ 1 in the baryon-free fireball, in all four runs and beyond any counting discussion. Only the cooling-rate trend is undecidable rather than negative: the sign of the comparison inverts between the two defensible mappings of the cooling axis onto beam energy, and $n = 4$ runs with the charge-conservation axis entangled in the trend do not support a directed verdict either way.

The kinematic branch — the last left open — does not carry the analogy either, though it is a partial hit rather than a clean miss. A quadrupole anisotropy analogous to v_2 is present at high significance ($> 3\sigma$) but lives in position space, not momentum space, and is axis-locked only in post-crunch generations (seed-random in generation 0); collective radial expansion appears in nine of ten generations but with no stable mass ordering and an *anti*-Hubble profile — fast inside, slow outside, the signature of overdamped, event-pumped displacement rather than pressure-driven flow; and the femtoscopic size-multiplicity scaling fails outright, the measured size range spanning a factor 16.6 where the $N^{1/3}$ lever allows only 1.13 — system size is morphology-dominated (compact-spherical versus filament-web), not multiplicity-driven.

The “Little Bang” reference class is therefore closed as a regime claim. What survives is stated plainly. First, an engine-own, reproducible kinematics — event-pumped, overdamped, polar-imprinted expansion — including one structurally heavy-ion-like candidate: a prolate, axis-near anisotropy that builds up across post-crunch re-expansions (positive in five of six windows, alignment median 0.80–0.96), which remains a candidate, not a finding, until its correlation with the crunch polarity axis is tested. Second, a single named operator as the consequence: both statistical negatives point at the same back channel, and a binding-energy-dependent photodisintegration with detailed balance is the first-principles-nearest candidate — with a prediction declared *before* any such change was made — S_{23} must move from ~ 0.5 towards > 1 , and the light-cluster hierarchy must become binding-ordered — whose outcome is reported with the back-channel arc below. The reframing thus does not stand as a rescue of the cosmological gap table; it stands as a further characterised negative that has, in exchange, localised its own cause to one operator.

The common root beneath these nucleosynthesis negatives — composite break-up behind a hard 15 MeV cut, blind to binding energy — has since been re-engineered twice, with flip criteria declared and frozen before each set of runs, and the outcome

is a chain of characterised negatives that terminates in a mechanism rather than a mystery. The first redesign (v1) set each nuclide’s break-up threshold to its minimal nucleon separation energy from the engine’s own binding landscape — 2.22 MeV for the deuteron, 5.50–6.26 MeV for $A=3$, 19.82 MeV for ^4He , zero for the di-proton — introducing no new fitted constant. Its positives are real and are reported first: the cold survival hierarchy becomes binding-ordered in every seed ($^4\text{He} \geq d \gg A=3$ in 3 of 3, against the inverted baseline ordering in 0 of 3), the anomalous $A=3$ accumulation point disappears (suppression $\times 10$ – 13 at matched temperature), the declared S_{23} prediction is met directionally on the fast regime (pooled endpoints $0.434 \rightarrow 1.875$) though the crossing is fragile and inverts on the detail scale ($0.5 \rightarrow 0.14$, stock starvation), and the di-proton artefact self-cleans (cold counts 0–1 against 13–14) with no special-case code. But the destruction does *not* freeze out with the universe’s temperature: the gate tested a node’s *own* heat, fusion births its products hot, and the measured $A=3$ destruction hazard peaks at $T = 1 - 2.2$ MeV — a factor 5–10 *below* the thresholds. What the operator delivers is exposure-time attrition, not binding-ordered freeze-out, and on the full-cycle scale it sterilises the composite sector (cold census with $n = 0$, metallicity $0.137 \rightarrow 0$; the helium gain of the charge-conservation canon step, $\times 1.8 - 2.1$, is returned in full). The frozen flip condition failed on three of five conjuncts; no bug was found — ledgers exact, energy drift at the 10^{-10} level — the design was incomplete: *no canon flip*. The second redesign (v2b) repaired exactly the diagnosed defect: the driver becomes the ambient photon bath — the radiation temperature over non-pole nodes; the sediment poles, being export-exempt archive, are not part of the interacting medium — and the break-up probability the parameter-free Boltzmann tail $\exp(-S/T_{\text{bath}})$, with no floor, no auxiliary gate, and nothing to tune, destruction now being measured channel-sharply from event logs rather than inferred from censuses (two mis-specifications — one estimator, one bath aggregate — were caught by control measurements mid-campaign and corrected before the deciding full-cycle arms, with post-hoc status explicitly labelled). The declared discriminator passed: the $A=3$ destruction hazard in the cold bath window (0.5, 2.2) MeV falls from 7–40% (v1) to 3 events in 1902 (0.16%) — the attrition defect is structurally removed and a genuine freeze-out exists, though its declared binding-ordering is only weakly evidenced (2 of 9 generations evaluable, on single-event statistics). Sterilisation nonetheless replicated on both

scales (full-cycle cold census $p = 2212$, $n = 4$, $d = 0$, ${}^4\text{He} = 0$), through two measured legs, each necessary. First, the bath itself does not freeze out: it floors at 0.48–0.66 MeV from mid-cycle to the crunch, never crossing the deuteron-safe 0.32 MeV, and the ~ 179 deuterons formed cold late in the cycle were all destroyed at the floor (deuteron lifetime there 85–100 steps). Second, the neutron race is lost: the bath crosses the deuteron threshold only ≈ 66 neutron lifetimes after the neutron peak, where real nucleosynthesis wins the same race at a ratio of 0.2–0.3 — a dimensionless timing coincidence the engine misses, and a scaling-axis question rather than an operator defect. Under the frozen honesty clause this counts as a negative against the back-channel family *as a whole*; both operators remain behind default-off flags.

The floor itself has now been diagnosed, and the diagnosis closes the arc. The ~ 0.5 MeV bath is not a mirror of the matter temperature — which falls correctly to 0.04–0.06 MeV, at its CMB-analogue anchor of 0.037 — and not decay heat, but a throughput equilibrium between a strong pump and a structurally blind sink. On the source side, annihilation carries $\sim 95\%$ of the radiative event injection (38.5 MeV per step, releasing reclaimed edge tension as heat); on the sink side, 93–96% of the ambient radiation comes to rest on vacuum-substrate nodes, which both south-pole drain paths exclude by type: the vacuum sea is export-immune, and only 3.3–5.7% of the bath is drain-visible. A steady-state balance of these measured fluxes reproduces the floor (0.48–0.78 MeV modelled against ~ 0.5 measured), and the counterfactual is equally sharp: were the sea drain-visible, the same balance yields ≈ 0.03 MeV — CMB-cold and deuteron-safe. The factor of ~ 18 between the two *is* the export exemption of vacuum radiation, verified at code level. A second, operator-own effect compounds it: the bath-coupled formation–destruction cycling pumps additional mass and binding through the annihilation channel into the sea, so the operator arm holds $9.9\text{--}11.0\times$ more vacuum agitation than its baseline and floors warm where the baseline crosses the deuteron-safe line before its own crunch — the operator heats the very bath that is supposed to freeze it out. The back-channel as a sole lever is thereby exhausted and mechanistically understood. The floor is the binding boundary condition that every subsequent destruction *or* formation candidate inherits: at the floor, full detailed balance would still destroy deuterons at $\approx 1.2\%$ per step, and the formation side does not change the bath temperature. The effective lever is a sea sink — extending export,

or Hubble-type dilution, to vacuum radiation — and since the canon ground for the exemption speaks of vacuum *agitation* remaining north (§4), the exemption of vacuum *radiation* is not semantically forced: opening it was posed as a declared canon question, not a tuning move.

That canon question has now been answered by measurement, and the constraint removed — by a pre-declared intervention, not by tuning. Opening the exemption for radiation only — one boolean flag, no new rate, no new constant; agitation and the poles remain exempt — was declared, with frozen pass/fail criteria, *before* implementation and before any run. The declared primary criterion (floor at or below 0.32 MeV in the cooling era) passed cleanly: the ambient floor falls in *every* cooling phase to 0.036–0.052 MeV, at the diagnostic’s CMB anchor (0.037) and where the steady-state estimate predicted (~ 0.03), against a baseline best of 0.282 and 0.476 for the bath-coupled photodisintegration operator alone. The ledger gates hold exactly (charge census plus archive $\equiv 0$; energy drift at machine level), the cycle survives (three crunch–re-seed generations), and the declared side-effect is reported: the sink accelerates the crunch, and at the fast scale it does so faster than any cold window can open — the result is read at the detail scale only.

With the floor thus decoupling-safe, the combined configuration — lake sink plus bath-coupled photodisintegration — was run against the sterilisation result that had previously killed the photodisintegration branch: in the cold census ($T_{\text{bath}} < 0.32$ MeV, the zone the earlier operator sterilised) the composites now appear where the reference had exactly zero, ${}^4\text{He}$ at 16–24 and $A=3$ at 36–40, present in 100% of cold snapshots and many σ above zero. Adversarial verification then *overturned the first reading* of why. The survival narrative — the floor drops below the destruction thresholds, so the nuclei live — is refuted: event logs show ${}^4\text{He}$ is never photodisintegrated in any run, and the sterilising reference would not have destroyed it at its own ≈ 0.5 MeV floor either (probability $\sim 10^{-17}$). The verified mechanism is *formation*: a sink-only control gives the ladder ${}^4\text{He} 3 \rightarrow 9 \rightarrow 24$ (canon $\rightarrow$ sink $\rightarrow$ sink+photodisintegration) — the sink lengthens the cold formation phase, and the photodisintegration’s actual role is to dissolve di-protons ($61 \rightarrow 0$), returning free protons to genuine fusion. Nucleosynthesis in this engine is formation-limited, not survival-limited. The single-run caveat was itself declared a blocker and has been closed: an independent seed reproduces the composite core in the same range (${}^4\text{He}$ 12–28, $A=3$ 44–45), while deuterium

remains seed-variable at Poisson level (0–7) — the strict three-species criterion (d , ${}^4\text{He}$, n all > 0) rides on neutron counting noise and is *not* claimed; the result carries on ${}^4\text{He}/A=3$. (All deuterium counts in this section are (Z, N) -keyed and therefore subject to the census defect retracted below; the mass-keyed re-measurement puts the true count at this scale in the same Poisson range, 12.) Adding the positional-gravity channel on top tests additivity, and the verdict is channel-specific, not general: the heavy composites do not rise further (mean ${}^4\text{He} \approx 19$ against ≈ 20 without gravity — formation-window-limited, not production-limited), while deuterium does ($d = 7$ in later generations against ≈ 0) — the two mechanisms feed different channels and do not multiply. We state the status plainly: this configuration (expansion kernel, lake sink, bath-coupled photodisintegration) became *canon* on 2026-07-11 — after the scale-test arc of §5.1 had additionally established its ledger exactness and near-scale-free fill dynamics, the defaults were flipped under the standard procedure (bit-exact gates in both directions: the new default reproduces the measured candidate trajectories exactly, and the historical physics remains reachable through explicit overrides, with pre-flip checkpoints deserializing to their historical values) — and no quantitative Y_{He} match is claimed: the yields are counts on a 10^3 -quanta system, mechanistically verified but uncalibrated, with the binding-energy landscape still imported.

The lake sink has since left the canon, removed by its own declared measurements. At the ensemble scale — the fast preset, where cycle statistics live — a five-arm component isolation (internal identifier K31, declaration and analysis frozen before the runs) attributed a previously noted cost of the 2026-07-11 configuration to the sink specifically, and not to the photodisintegration: under the full canon nothing beyond $A=4$ forms and no horizon state is ever reached, while switching off the sink alone restores heavy composites (A up to 16 in a 4×10^4 -step window, seed-stochastic at roughly one seed in three) and the first horizon states. The mechanism is cycle acceleration, not nucleosynthesis: the sink’s archive feed advances the polarity inversion so strongly that the generation length collapses (crunches every $\sim 2,300$ steps against a single long cycle without it), amputating the cold formation phase in which heavy fusion happens — (A deuterium trade-off was initially read into these arms from the standing $(Z=1, N=1)$ census; that reading has since been retracted — see the deuterium bottleneck below.) A dynamic, phase-gated sink — throttled while heavy nuclei build, open in the

cold end phase — was designed as the reconciliation, declared, built default-off, and tested on a confirmatory campaign (15 cells at 1.2×10^5 steps); its activity sensor proved structurally blind (the gate keys on a quantity read inside the export loop, which falls silent exactly in the cold phase it was meant to detect), so the treatment arm degenerated to sink-off and the gate never fired — reported as an instrument failure, not evidence. The same campaign, however, produced the deciding side finding: with the sink on, the ambient bath at this scale *never* cools — each accelerated crunch re-seeds and re-heats, holding the median matter agitation near 80 MeV — whereas the sink-free arms cool to 0.29–0.63 MeV, *below* the deuterium-safe 0.32 MeV line the sink had been introduced to reach. The operator installed as the cooling lever is, at the ensemble scale, the obstacle to cooling; and the derived matter-expansion index (§9.5) now supplies the faster matter cooling that motivated the sink before the index existed. The default was therefore flipped off on 2026-07-16 under the standard procedure (forced-on reproduces the prior canon trajectory bit-exactly, crunch at step 1,162; the new default reproduces the sink-free reference, crunch at step 2,713; energy drift at the 10^{-11} level, charge census zero), making the current canon the expansion kernel plus bath-coupled photodisintegration, without the sink. Under it, heavy composites and horizon states are again emergent (A up to ~ 39 ; 27 distinct horizon-state nodes across the ensemble campaign, §10). The detail-scale re-measurement has since been run (frozen declaration, two seeds), and the flip’s price at that scale is real — and inverts the trade. Sink-free, the cold-window heavy ladder collapses (${}^4\text{He}$ 1–9, $A=3$ at 1–8 against the sink era’s 16–24 and 36–40) at a properly cold floor (0.012–0.027 MeV) — at this scale the flip costs yield across the board rather than trading one species for another.

A retraction, and the bottleneck it exposes.

The first reading of these arms reported the opposite for deuterium — thousands of standing deuterons sink-free against a handful with the sink, read as a composition *bracket* around the real (helium-rich, deuterium-trace) point. That reading was wrong, and the error was ours: a bookkeeping defect in the beta-decay operator set the proton counter without clearing the neutron counter, so every decayed free neutron became a phantom ($Z=1, N=1$) node carrying the bare proton mass (938.27, not the 1875.6 of a real deuteron). Every deuterium census keyed on $(Z=1 \wedge N=1)$ — ours included — counted those phantoms. Re-measured on mass, which is immune to the defect, the picture is not a bracket but a

scale split: at the detail scale genuine deuterium is essentially absent (*one* deuteron in the sink-free runs, 12 in the sink era), while at the ensemble scale under the current canon it is substantial (~ 900 per snapshot, at exactly the physical deuteron mass) — with the same ensemble preset but constant expansion it collapses to 11. What the phantoms hid was therefore not a uniform wall but a genuine regime dependence: the deuterium channel is open where the cycle is fast and the expansion dynamic, and shut where it is slow. For a short interval we read the detail-scale absence as the physics — a total bottleneck, mechanically explaining a chain that ends at helium. That reading did not survive the fix either, because the defect was never merely cosmetic: the phantom (1, 1) state fed a spurious neutron-degeneracy term into the collapse gate and a non-zero nuclear binding into the edge-break threshold — the phantoms were *load-bearing*, distorting the very dynamics on which the census was taken. The clean-engine reference campaign (§11; same canon, same presets, same seeds, the book-keeping fix the only difference) settles it: standing cold-window deuterium rises from 1 to 25–533 per cell across both scales, the proton–proton ladder has its second rung, and the chain climbs — to $A = 34$ and $Z = 11$ across up to 45 distinct species at unit scale, terminated from above by the collapse cascade rather than by attrition (§8). The “total bottleneck” was the bug’s dynamical shadow. We report this sequence — a pre-declared measurement overturning our earlier reading, and the audited re-measurement then overturning the correction’s own conclusion — as the third and fourth such cases in this paper, and as exactly what the frozen-declaration procedure is for. The mechanism was then measured rather than guessed, through event-log archaeology with a pre-registered band hypothesis whose specific band placement was refuted by its own kill criterion: what the declared histograms show instead is that the sink era spread its heavy-fusion events across many temperature decades ($5\text{--}10\times$ more $A\geq 3$ events, the strongest cold bin at $0.03\text{--}0.1$ MeV), while the sink-free bath *plunges* through the fusable zone (essentially no events between 0.1 and 30 MeV). The sink was, at this scale, a descent brake, and the heavy-ladder yield is cooling-rate-limited — the bracket has a measured regulator. The sink-free detail cycle also closes later than the historical window (163–185 k steps; the preset window has been widened accordingly). Two honest qualifications remain: no configuration reaches the real composition point, and this is the second time a canon component has been partially retracted by its

own pre-declared measurements (after the v1 back-channel above), which we record as the procedure working rather than as churn.

9.5 The Expansion Rate: From Chosen Constant to Derived Law

The expansion sector has since been pushed from chosen constants towards derived law, in three steps whose evidential status we state plainly: one validated, one feasible but structurally separated, and one clearly falsified.

The adiabatic index (validated). Historically the engine cooled matter agitation and radiation at the same per-channel rate — and a structural clamp in the export channel even *forbade* matter from cooling faster than radiation. Adiabatic expansion fixes the ratio: with $T \propto a^{-3(\gamma-1)}$, radiation ($\gamma = 4/3$) redshifts at $1 \times H$ while non-relativistic matter ($\gamma = 5/3$) cools at $2 \times H$. The matter-expansion index of 2.0 — applied in the redshift channel, where expansion cooling lives — is therefore an adiabatic exponent, derived rather than chosen; it replaces a chosen unit coupling and removes one degree of freedom. The prediction, declared *before* the A/B, was faster matter cooling and a deeper, longer cold formation floor, and both have now been measured on the two reference regimes (seed-888 detail runs, 1.65×10^5 steps): matter cools measurably faster in the hot phase (matched-step temperature ratio 0.897), and the cold floor deepens from 0.075 to 0.048 MeV in the canon configuration and from 0.055/0.040 to 0.035/0.028 MeV across the stationary generations of the sink configuration, with the ledger exact in both arms (charge zero exactly, energy drift at the 10^{-10} baseline). The nucleosynthesis yield is honestly modest — ${}^4\text{He}$ 40 $\rightarrow$ 45 and deuterium 2 $\rightarrow$ 7 in the stationary census of the sink configuration, no dramatic lift — and part of the $2\times$ asymmetry was already present empirically at index 1.0 through fusion binding losses and the export drain, which is why the observable is modest. The index is kept at its derived value for the correct kinematics and the removed degree of freedom, *not* for a match improvement, and is not adjusted.

Friedmann feasibility (confirmed, then structurally separated). A second step makes the rate itself emergent: $H = H_0 \sqrt{\rho/\rho_0}$, with ρ the ledger energy density per vacuum node, so that the vacuum node count acts as an emergent scale factor and the previously chosen cooling constant is demoted to an initial-value anchor — again with no new fitted constant. The operator behaves as FRW kinematics requires: H genuinely dilutes (from $3.5\times$ to $0.70\times$ the anchor across the first generation,

monotonically with vacuum growth) and gives the dimensionless BBN timing invariant $H\tau_n$ an epoch dependence for the first time (0.49 at the hot maximum, falling to 0.10 at the deuterium epoch, against the physical $\sim 2.2\text{--}3.4$); under the constant rate it had none. It is nevertheless *not* part of the current nucleosynthesis candidate: it is structurally incompatible with the substrate-sea sink, because both operators act on the same vacuum-radiation account — the sink exports it to the archive while the Friedmann density counts it — producing a seed-robust fragmentation runaway (edge-to-node ratio $\sim 98\times$ by the second generation), and every redefinition of ρ tested either restored the runaway or neutralised H ($\sim 100\times$ too weak). This is a documented operator incompatibility, not a numerical defect, and it was resolved by composition at the time: Friedmann expansion stood as a separate, default-off prediction module. **The incompatibility died with the sink, and the module is now canon.** After the sink’s removal (the 2026-07-16 flip recorded with the nucleosynthesis arc), the re-test was run under a frozen declaration rather than presumed (two arms, two seeds, detail scale, plus two further closure seeds at 3×10^5 steps): the fragmentation runaway is gone — the edge-to-node trajectory of the Friedmann arm sits at $0.99\times$ the sink-free canon against the historical $98\times$ — H dilutes by more than $8\times$ across the generation, every cycle closes, and the ledgers are exact throughout. The race also moves the right way: the bath’s deuteron-safe crossing advances from 82/118 to 45/109 neutron lifetimes after the neutron peak in the two paired seeds. A pre-registered out-of-sample *yield* prediction, frozen before its runs completed, then failed — the cold-bin fusion enhancement inverted in one of two fresh seeds, and the helium counts sit at Poisson level — so no nucleosynthesis benefit is claimed, and none enters the flip rationale. On exactly the basis the adiabatic index established — the derived kinematics and the removed degree of freedom, not a match improvement — the module became canon on 2026-07-18 under the standard procedure (suite green; forced-off reproduces the constant- H canon bit-exactly, crunch at step 2,713; the new default reproduces the pre-flip override anchor bit-exactly, crunch at step 1,493; drift 4×10^{-11} , charge zero). The previously chosen cooling constant is thereby demoted, in the canon itself, to an initial-value anchor; the *absolute* magnitude question is untouched by this and remains open below. At the fast ensemble scale the dynamic H shortens the cycle to $\sim 55\%$ (the boot-spike regime the feasibility study had already noted there) — a characterised throughput

change, not a physics claim.

The absolute magnitude (falsified). The remaining anchor is the overall scale of H . The Planck-analogue construction $G = \hbar c/M_{\text{fund}}^2$ was implemented with the pre-declared candidate $M_{\text{fund}} = \sigma r_p = 156.35 \text{ MeV}$ (string tension times confinement radius, the engine’s one genuinely calibrated energy), and with the asymmetric reading fixed *before* the run: only a result outside the band $H\tau_n \in [1, 10]$ is informative. The result is a clear falsification: $H\tau_n = 2.4 \times 10^5$ for σr_p , and 3.9×10^4 and 3.7×10^4 for the two alternative candidates (m_p and the metric-break scale) — four to five decades above the physical value, for the entire candidate family. The reading is structural rather than parametric: every scale the engine possesses is hadronic (156 MeV–1 GeV), whereas real gravity is weak because $(M_{\text{Pl}}/m_p)^2 \sim 10^{38}$; a gravitational coupling built from hadronic constants is therefore *necessarily* decades too strong at this invariant. The engine has no Planck hierarchy, and the absolute expansion magnitude is not derivable from EGT as it stands. We declare this open rather than patch it: the only route consistent with the framework’s own postulates is an emergent scale far above m_p , generated dynamically by scale-invariant running over many decades (dimensional transmutation) rather than inserted — which is what elevates the scale-dependence measurements reported under the scale-invariance heading (§5.1) from a robustness check to the framework’s candidate mechanism for the missing hierarchy.

10 Bridge Signatures: Measured Steps Between the String Sector and Geometry

The framework’s self-conception is that of a connecting structure: a hadronic string sector on one bank (§7 — edges as strings, an exact Regge trajectory, a genuinely calibrated tension) and FRW-type kinematics with emergent horizons on the other (§9.5, §9.3). The magnitude falsification quantified what is missing between the banks: the 10^{19} hierarchy. Rather than assert the bridge, we measured which of its load-bearing signatures the substrate actually carries. Three were declared — induced-gravity dilution, renormalisation running, and horizon thermodynamics — each with frozen declarations, pre-registered kill criteria, and the falsification thresholds written on both sides (the two-sided lesson of the K18g counterfactual above). This section reports the first, and collects the verdicts of all three.

10.1 Mass-Coupled Metric Attraction, and the Dilution That Is Not a Power Law

The Sakharov reading of induced gravity [2] — G small not through a small coupling but through dilution over the substrate’s degrees of freedom — makes two separable claims: that the substrate exhibits a mass-coupled attraction at all, and that its effective strength dilutes with substrate size. Both were measured (internal identifier K34), and they came apart.

The measurement had to become passive first, and that is itself a finding. Four successive intervention designs — inert test masses, pressure-supported test masses, native-particle clusters, clusters placed in empty space outside the substrate sea — failed for the same structural reason, each diagnosed to its operator chain: the engine *metabolises* foreign structure. Inert mass of nucleon scale is recycled, ledger-exactly and era-independently, into native protons and neutrons through the orphan-absorption and pair-materialisation channels; the hot era destroys any inserted structure within a hundred steps (the Jeans statement at operator level); the cold era condenses proton clusters into nuclei within the same window; and outposts wired to the sea drain into it by discovery-edge fusion. A substrate whose operators integrate every foreign object into their own dynamics cannot be probed by placement — measurement must be observational, on structures the physics itself built. The parallel to the quantum-mechanical measurement problem is not decorative: the simulation enforces the same epistemology.

Mass coupling: confirmed. The passive design tracks natural pairs in the cold window — nuclei against electrons, a mass span of ~ 1 to $\sim 8 \times 10^3$ MeV at identical charge magnitude — matched on initial metric distance ($\pm 25\%$) within the same universe, the observable being the pairwise approach rate in the emergent metric, $R = -\Delta d_{\text{metric}}/\Delta t$. Heavy pairs approach faster than their matched light twins in 8 of 9 cells (bootstrap CI above zero in each), 167 of 217 matched pairs overall (77%); the single null cell, at the largest scale, is reported rather than absorbed. The substrate exhibits a genuine, robustly mass-coupled metric attraction between structures it built itself — consistent with the continuous mass term in the metric-contraction operator, whose per-edge form is additive in the masses, and with heat acting as anti-gravity in the same law (the operator-level Jeans criterion). Proton pairs are censored at 70–90% by cold-era fusion — the condensation channel at work — so the comparison rests on the

two stable classes, as declared.

Sakharov dilution: excluded as a power law by its own third scale. Across substrate sizes of 7×10^3 and 2.2×10^4 vacuum nodes the mass-normalised differential G_{eff} fell with size, and the two-scale exponent, $\alpha = 0.62 \pm 0.56$ for $G_{\text{eff}} \propto N^{-\alpha}$, was still compatible with the Sakharov value $\alpha = 1$. The pre-registered third scale (6.8×10^4 nodes, one further tripling) disciplined the claim: the pooled estimate collapses to $\alpha = 0.14 \pm 0.29$, the flow is non-monotone — the dilution seen from $\times 1$ to $\times 3$ does not continue to $\times 9$ — and $\alpha = 1$ is excluded over the measured decade. The verdict repeats the method lesson of the scale-test arc: two points describe, the third discriminates. A simple dilution over degrees of freedom is not the mechanism of the missing hierarchy; whatever generates it must be sought in running, not in counting.

10.2 The Ladder and the Horizon: Verdicts in Place

The renormalisation measurement (K35) is reported where the framework’s own text pointed for it, under the scale-invariance heading (§5.1): one topological coupling runs stably (-0.312 e-folds of connectivity per halving, invariant across seeds and epochs — a new structural constant of the substrate), while the energy couplings move era-coupled rather than ladder-stable, so the transmutation channel the hierarchy requires is, on the current operator set, *not found*. The horizon measurement (K36) is reported with the black-hole state (§9.3): boundary growth tracks swallowed content weakly but positively (a first Bekenstein signature), while the area law proper is a characterised negative: the spectral poverty that first blocked it was traced by code inspection to the engine’s own live spectral-maintenance operator (not a serialisation artefact), and on the one uncompact information proxy available — neighbourhood composition — horizon information tracks neither boundary nor content.

10.3 What the Bridge Carries, and What It Does Not

One qualification governs how the negatives below may be read, and a code audit forced it on us. Every coupling in the engine is resolved once at boot and then frozen for the run: there is no renormalisation scale, no cutoff, and no operator anywhere in the core that recomputes a coupling as a function of scale or energy. A quantity that cannot vary cannot be observed to run. The Sakharov exponent’s collapse to $\alpha = 0.14 \pm 0.29$ and the absence

of a running energy coupling on the renormalisation ladder are therefore *structurally forced by the present implementation*, not empirical verdicts on the underlying idea — we measured our own architecture, and the honest content of both results is that the architecture lacks the degree of freedom the hypothesis requires. This is a construction stage, and a named one: giving the operator set a scale-dependent coupling is a defined next build, after which the same two measurements become genuine tests. The topological exponent (which *does* run, being a property of the graph rather than of a frozen constant) and the horizon results are unaffected by this qualification.

Stated as a ledger, with the negatives at equal weight. The substrate *carries*: a mass-coupled metric attraction among its own structures (robust, 8/9); boundary growth correlated with swallowed content at its horizons (weak, CI above zero); a renormalisation-self-similar topology with a stable exponent (sharp, 30/30); and, from the earlier arcs, an exactly linear Regge trajectory on one bank (its slope a factor six from the physical value, §7) and derived FRW cooling kinematics on the other. The substrate *lacks*: power-law dilution of the attraction ($\alpha = 1$ excluded); a cleanly running energy coupling (not found at the first ladder); and an area law for horizon information (a characterised negative on the measures the engine exposes, the boundary *growth* notwithstanding). The bridge therefore stands as measured scaffolding on both banks with the span — the hierarchy — still missing, and we close it as we closed the magnitude question: open by measurement, not defect by construction. The framework’s falsifiable commitment is unchanged: if the hierarchy is to emerge, some coupling must run, and the ladder that would show it is now built and climbable.

11 The Alpha-Cycle

The cosmological cycle of EGT is the long-term consequence of the postulates: substrate accumulation at the south pole, matter energy depletion, chronos accumulation as events occur. Eventually, the energy distribution between the poles inverts, and the cycle terminates.

11.1 The Alpha-Dipole

The substrate has a polar topology. Its two distinguished nodes are the alpha-dipole:

- α_{active} (north pole): the active region of the universe, where matter dynamics dominate.

- α_{store} (south pole): the substrate archive, where fifth-force activity accumulates as spectral sediment.

The two are connected by a topological bridge—an Einstein–Rosen-like channel through the substrate—through which energy and spectral information are transported by the fifth force (Postulate IV).

11.2 The Big Crunch as Polarity Reversal

The Big Crunch is a physical consequence of an energy-gradient inversion between the alpha-dipole’s poles. Throughout the cycle, fifth-force activity transports energy from the active region (north pole) to the substrate archive (south pole). As long as $E_{\text{active}} > E_{\text{store}}$, the bridge supports net flow toward the south, and the matter sector continues to function. When E_{store} accumulates sufficiently to exceed E_{active} by a substantial margin, the gradient reverses: the south pole becomes the dominant energy source, the bridge supports flow in the opposite direction, and the topological structure of the north pole collapses through the bridge.

The triggering condition is a property of the substrate’s polar energy distribution, not of any abstract coherence quantity. The Big Crunch is what the system *does* when its polar gradient inverts.

This triggering condition has now been tested directly against its natural stellar alternative — that the crunch is a virial event, a collapse fired when inward gravitational binding overtakes outward thermal support (a virial ratio R crossing unity from below, as in stellar core collapse). The hypothesis is falsified, and proxy-free: counting only crunch events against $R=1$ crossings — a statistic invariant under any monotone redefinition of the binding term — unity is neither *necessary* (one configuration crunches 18 times while R never crosses unity, remaining at 0.001–0.42 throughout) nor *sufficient* (another sits at $R \approx 1.8$, nominally over-bound, for 84% of its runtime and crunches once). Where the tension-proxied ratio crosses unity at all it does so from *above* — the opposite direction to the stellar reading — because radiation grows toward the crunch as the stagnant vacuum bath fills while tension shrinks: cooling here fills a reservoir rather than removing a support pressure. The south-pole fill fraction remains the tighter predictor (CV 0.18 against 0.49 for R over 24 crunches), and an operator cross-check argues the same way: two independently added operators that shift the crunch time in opposite directions crunch at markedly dif-

ferent ratios ($R \approx 0.75$ versus $0.15\text{--}0.44$) — no fixed virial threshold exists, while both shift the crunch exactly as a fill switch predicts. The proxy caveat has since been retired by re-measuring with the true position-computed gravitational binding (eight runs, 60 crunches): that ratio does rise from below in 60/60 crunches, flutter-free — the hypothesised direction — but its value at crunch is a cooling clock, not an invariant (CV 4.58, a spread of four orders of magnitude; correlation of $\log R$ with cooling time 0.84), and the rise is denominator-driven: the kinetic term falls with cooling ($1.9 \times 10^5 \rightarrow 6 \times 10^3$) while the gravitational binding at crunch is stable (CV 0.18–0.24) and the matter’s spatial extent does not contract at all before the crunch. The trigger is the south-fill partition switch of the polar gradient stated above; the crunch is *not* hydrostatic — converging with the independent finding that the expansion is substrate-driven rather than pressure-supported.

What the falsification removes is exactly one claim — *trigger equals pressure-balance crossing*. The broader stellar life-cycle phenomenology survives it and continues to carry: the two-stage termination (crystallisation onset, then polarity inversion) is a collapse–bounce–ejection sequence of supernova type; an extended cooling era precedes the collapse; local, density-primary metric singularities — the emergent black holes — provide object-scale collapse in parallel to the global crunch; and the cooling–crunch coupling is thermodynamically real (cooling routes radiation through tension into south-pole sediment and so fills the pole), but the switch it throws is the polarity, not $2T + U$.

The hydrostatic collapse channel itself, however, has been measured to be *feasible*, and remains design-open. The engine carries a full position-based gravitational term (mass-weighted attraction with a Pauli hard core, energy-neutral by construction) that is gated off in the canon. Switched on as a pure configuration override — no code change — across a strength sweep, the ledger holds at every strength (global drift $\sim 10^{-10}$; the energy-neutrality claim survives the addition of a full force term), and structure clusters gravitationally and strength-scaled: at matched era, +36% in the clumping ratio, +27% in close pairs, and –12% in bounding volume against the gravity-off baseline, with the fusion count rising +11% at the mild setting and +29% at the strong one. In the detail regime total fusions rise by +76% (deuterons $\times 4$, ${}^3\text{He} \times 2.7$): position gravity acts as a density-driven *production* amplifier — clumping raises local density, density raises fusion — which is stellar physics. Three honest limits are stated with

it: the baseline already clusters strongly from the bond dynamics alone, so the term amplifies gravitational clustering rather than creating it; the yield gain is flat (no ${}^4\text{He}$ gain, no metals, a di-proton artefact strongly amplified, and the free neutrons consumed); and gravity alone does not deliver a stellar collapse — it accelerates the existing south-fill crunch. A local Jeans-instability trigger on this now gravitationally dynamic base — the global crunch emerging from cascading local collapse — is the declared next design step, not yet built.

11.3 The Role of the ECQ Triad

The ECQ triad of §4 is the diagnostic counterpart of this dynamics. As the cycle progresses, E (substrate sediment) grows, Q (matter energy) shrinks, C (chronos) accumulates, and the coherence state $\mathcal{K}$ rises. The critical regime $\mathcal{K} \rightarrow 1$ is the diagnostic signature of the system approaching the polarity-reversal threshold: it tracks how far the cycle has progressed, but does not itself trigger the reversal. The reversal is the physical event; ECQ is the measurement that anticipates it. This distinction has been tested directly in seed ensembles: the collapse onset brackets the pole-potential equality tightly across seeds (last stable samples 0.48–0.57 of the pole ratio, followed by self-amplifying inversion within $\sim 10^2$ steps), while $\mathcal{K}$ at onset varies between seeds by large factors and its terminal spike occurs *during* the collapse, as a consequence of $Q \rightarrow 0$. No critical value of $\mathcal{K}$ exists; the polar condition does.

11.4 The Big Bang as Spectral Re-seeding

The Big Bang is the inverse of the Big Crunch. Once the polar collapse has completed, the accumulated spectral content of $\mathcal{S}_{\text{store}}$ is released through the bridge, structuring the next cycle’s emergent dynamics.

In the *Aeoncore* implementation the full cycle — accumulation, polarity reversal, collapse, and re-seeding — runs end to end as one closed, energy-conserving process, and it *recurs*: up to three complete generations have been observed in a single run. Two earlier qualifications of this framework have since been resolved by adding missing physics rather than by tuning. First, the collapse is no longer a placed threshold: the cycle terminates in a two-stage process in which a configurable crystallisation onset Ω triggers the late-era phase transition, while the Big Crunch itself is the *emergent* polarity inversion of the alpha dipole at equal pole potentials — a condition of the substrate’s polar energy distribu-

tion, not a control knob. Second, recurrence failed originally because heat had no outlet; once the fifth-force export channel was priced into the ledger, the cycle closed of its own accord. The spectral archive $\mathcal{S}_{\text{store}}$ is carried through the bridge into the seed of the next generation with its frequency content preserved (full archive returned; global drift 0.012% over a 170,000-step cycle including the bounce). As a by-product, the late era reproduces the ordering of standard late-universe chronology — a radiation-, then matter-, then decay-dominated era whose clock is the proton lifetime, with the measured scaling law (decay-era depletion rate $\propto 1/\tau_p$).

The further claim of §3 — that the eigenfrequencies dominating the previous archive re-emerge as the fundamental constants of the next cycle — has now been tested interventionally, in three stages, and we state all outcomes plainly. *First*, the spectrum-collapse test declared above has been run: collapsing the inherited multi-band form to the single boot frequency at re-seed (energy-exact, no ledger change) leaves the successor generation identical in every counter — nodes, edges, matter count, fusion totals — with a first divergence of exactly 1 ULP (relative 1.2×10^{-16}) in the aggregate coherence diagnostic, traceable to three pre-mapped floating-point partition channels; a re-partition control (splitting bands into exact halves) remains bit-identical throughout, showing that the engine clears band partitions exactly and only the frequency *form* survives to processing. Passive spectral inheritance is therefore semantically *inert*: no operator in the engine reads the archived form, and the earlier map-level finding (“implemented but unread”) is upgraded to an interventional proof. *Second*, the frequency-selective operator defined by that follow-up has been built — Postulate-II native: the existing resonant transfer rate is modulated by the energy-weighted band overlap of the coupled nodes, reusing the existing merge tolerance, with zero new fitted constants and the empty-spectrum limit reducing exactly to the historical scalar coupling. With the reader active, the same collapse intervention now produces *semantic* divergence at re-seed (node counts 1131 vs. 1180, $\sim 4\%$; edge-age statistics at relative deviation 0.137) against the 1-ULP signature without it: the operator demonstrably reads the inherited band content, making the archive dynamical for the first time. As a by-product, the reader *delays* the first crunch by roughly a factor of eight (step 30,000 vs. 3,600 in the reader-off baseline) while the generation remains productive — a longer stable era, not immortality; together with the independently measured drain operator that *accelerates* the crunch,

the two move the crunch in opposite directions — support up, longer stable era; support down, earlier collapse — exactly as the south-fill partition switch of the polarity-reversal subsection above predicts; the virial reading of the same pair was tested and rejected there (the two operators crunch at markedly different virial ratios). *Third*, the directedness prediction — declared *before* measurement, that a reading operator would make second-generation structure shifts sign-consistent across seeds — has been *refuted* in the fast-cycle regime (6 seeds $\times$ reader-on/off, ~ 19 –21 generations per arm, forced-crunch threshold identical in both arms): the only borderline effect ($p = 0.047$) reverses sign under an alternative endpoint definition ($p = 0.83$); the signal is weakest, and wrong-signed, precisely in the spectral channel the operator reads (permutation $p = 0.063$); and the thesis-faithful monotone test (Kendall’s τ over generations) is null in *both* arms (τ 0.07–0.15, p 0.25–0.41) — the dynamics is a cyclic sawtooth, not a progression, and the apparent directedness of the reader-off baseline is a sediment-drift artefact of the coarse metric. The memory claim thus splits cleanly: the archive is carried exactly, and is now *read* — but it does not yet *steer*. This is an honest negative against spontaneous cross-cycle learning in the tested regime (an effect over long, unforced generations is not excluded, but is not shown); the reader remains a default-off flag pending the full canon procedure, and any surviving inheritance of constants would have to be carried by an imposed self-consistency map rather than by emergent directedness.

The cycle machinery on which these tests run has been audited for internal coherence across multi-operator, multi-generation runs. Total energy is conserved to $\sim 10^{-10}$ across every individual bounce (not merely end to end; the earlier-quoted 0.012% full-cycle figure is float accumulation over the long detail cycle’s hot era, not per-bounce loss), the charge ledger closes to 0.0 bit-exactly at every crunch, the ECQ invariant holds to $\leq 2.2 \times 10^{-14}$ over all generations, and the re-seed is clean — sediment reset exact, no cumulative drift, seed-reproducible. The single named exception is the boot generation: operator-laden configurations crunch their first generation while matter is still hot (20–50 MeV), because the accelerating operators drive the polar fill past the crunch threshold before the boot heat has dissipated — a property of the arbitrary genesis initial condition, not of the cycle or of any conservation law. Generation 0 is accordingly *declared* the ignition generation, and cycle coherence is asserted for the stationary cycles

(generation 1 onward), where all regimes are fully coherent.

11.5 The Cycle as Measured: The Clean-Engine Baseline

The entire cycle claim has been re-established from scratch on the audited engine — twice, because the audits themselves demanded it. After the weak-sector audit (§8) a first fourteen-cell campaign froze the baseline; its successor-generation physics then exposed a second bookkeeping defect (a fusion eligibility gate that could consume an ejecta-fattened lepton, leaking one elementary charge — found by a per-operator charge audit on a deterministic replay, repaired by typing fusion as a strictly baryonic operator), and the campaign was re-run in full on the repaired canon: fourteen cells spanning three regimes (ten ensemble cells, three full-detail cells of 2×10^5 steps, one triple-scale cell), every census mass-keyed, under a frozen declaration. This is the reference baseline of the framework, and it delivers the cycle narrative end to end — now including the generations *after* the first bounce:

Every cell runs the complete cycle — and every successor generation cycles healthily. Fourteen of fourteen cells traverse big bang, era sequence, collapse, crunch, and re-seed, and all fourteen successor generations run structurally healthy (edge-to-node medians 5.5–8.5); two seeds close second full cycles in-window, one runs a healthy third generation. The corrected bookkeeping shifts cycle clocks only where its bugs had fired — cells without stellar collapse reproduce their previous crunch steps to the exact step.

The era sequence is instrumented. With the epoch markers serialised, one detail cell reads: recombination by step 12,000; first $Z \geq 2$ nuclei at 44,800; first $Z \geq 3$ at 46,300; the deuteron-safe cold window opening at 52,000; collapse-driven crunch at 97,200; re-seed. The ambient bath floors at 0.04–0.23 MeV across all cells — deuteron-safe everywhere, the decoupling the operator arc of §9 was built to reach.

Cold nucleosynthesis builds the catalogue, and collapse terminates it. In the cold window the mass-keyed ladder stands at, e.g., 255 deuterons, 141 ^3He , 153 ^4He , and 147 nuclei of $A \geq 5$ in a single unit-scale cell — and at 844/721/475/503 in the triple-scale cell, which carries 63 distinct species up to $A = 49$, $Z = 13$. Above that, the ladder does not fade — it is *harvested*: 121 collapse events across the campaign resolve through the three-pressure cascade to 120 neutron stars and one white dwarf (the cascade’s second pressure branch, observed in the

densest cell), from progenitors of $A \approx 40$ –64 — the cascade removing exactly the mass range beyond the surviving census. The upper edge of the periodic table at this scale is a collapse edge, not a failure to fuse. Most of that stellar activity lives in the *successor* generations (24 collapses before the repair opened them, 121 after): in this cosmology, generation two is the stellar era. No $Z \geq 26$ forms in any cell: the iron frontier is the ladder’s next mapped distance on the scaling axis.

The ledger holds through everything above. Global drift stays at the 10^{-10} scale over complete multi-generation runs *including* collapse and re-seed (worst cell 1.3×10^{-9} — the most event-rich detail cell, its accumulation tracking activity, with charge and partition exactness confirming the numerical reading); the charge census closes to zero exactly in all fourteen cells — *through* all 121 stellar collapses, the very regime whose latent leak the audit had caught; the four-sector ECQ partition closes to $\leq 4 \times 10^{-13}$; and 10.7 million fusion events balance transaction-exact. The abstract’s conservation claim is this table.

The generation after the bounce: measured, repaired, closed. The first campaign delivered the first systematic look past the re-seed and found it bimodal: six cells rebuilt healthy second-generation cycles while eight entered a chemically poor densification regime. The mechanism was localised to a race — the per-generation re-anchoring of the Friedmann density ρ_0 , implemented as a lazy insert, could fire before the re-seeded medium was re-populated, locking the anchor at its numerical floor and the expansion rate at its ceiling; the floored anchor predicted the regime in 14 of 14 cells. The repair — no anchor on an empty medium, a zero-knob guard reusing the floor constant that already existed — was tested A/B over all nine affected-and-control seeds (six of six floored seeds healed, three of three healthy seeds unchanged) and promoted to canon under the standard bit-exact flip procedure. Extended control runs confirm the stakes: under the old behaviour a floored universe never cycles again (10^5 steps without a second crunch); under the repaired canon, all fourteen do. The successor-generation claim at the top of this section is that repair, measured.

11.6 The Initial State

If physical constants are inherited from previous cycles, what governed the first cycle? We posit that at $N = 0$, the spectral distribution $\mathcal{S}_{\text{initial}}$ was characterised by white noise—a uniform distribution of energy across all frequencies with maximum entropy.

Without pre-existing resonances, structures formed through stochastic optimisation: random fluctuations established the first stable resonant pathways, and these survived to seed subsequent cycles. The progression from $N = 0$ onward is the gradual sharpening of the substrate’s spectral content: each cycle adds discrimination, each cycle’s archive is more structured than the previous one’s, and over many cycles the universe converges on an increasingly specific set of physical constants.

12 Implications for Quantum Computation

The EGT framework offers a natural reading of quantum computation that does not require ontological multiplication of realities. In standard interpretations, the apparent need to account for 2^n simultaneous computational states is often invoked as evidence for branching universes, on the grounds that the state space of a sufficiently large quantum register exceeds the atom count of the observable cosmos.

In EGT, this argument rests on a category confusion. The dimensionality of a computational state space is not a count of physical realities, but a measure of the spectral degrees of freedom available within a coherent subgraph. A quantum register of n qubits corresponds to a subgraph of n spectrally coupled nodes, each carrying its own Spectral Vector $\mathcal{S}_q$. The state space lives in the frequency domain of this subgraph, not in a multiplicity of separate universes.

Quantum interference, the mechanism by which incorrect computational paths cancel and correct paths reinforce, follows directly from Postulate II. Resonant pathways minimize the Chronos cost C and survive; non-resonant pathways encounter high impedance and dissipate. Decoherence, in this view, is the loss of spectral isolation through unwanted edge couplings to the environment—precisely the failure mode that cryogenic isolation in physical quantum computers is engineered to prevent.

This interpretation is ontologically parsimonious: a single graph, locally coherent subgraphs, no parallel realities. A detailed treatment of superposition, entanglement, and measurement as subgraph dynamics is left to future work.

Declared Predictions and Falsification Tests

The framework’s internal discipline is falsification-driven, and we list its record explicitly. *Passed*: the edge-runaway falsification — removing topology pricing must reproduce the runaway (measured: energy per node 141 without the ledger entry, 1.0 with it); one prediction declared *before* its measurement — that enforcing charge conservation would raise the slow-cooling helium yield (confirmed, $Y_{\text{He}} 0.013 \rightarrow 0.028$); the spectrum-collapse test at re-seed, now run: replacing the inherited multi-band form by the boot frequency leaves every counter of the successor generation identical, with a single ULP of floating-point divergence — passive spectral inheritance proved semantically inert by intervention; and the export-efficiency wager, declared and frozen *before* measurement and confirmed at the weak tier on 36 virgin cells (slope $+0.058 \pm 0.003$ inside the pre-registered band; §5.1) — the record’s second before-declaration, and the first to survive a discovery $\rightarrow$ pre-declared replication chain on independent data; and, built on that same measurement, the export-efficiency exponent has since been *derived* rather than re-measured — a closed chain from the engine’s code mechanics and independently measured state exponents to $e_x^{\text{pred}} = +0.059 \pm 0.012$ (measured $+0.058 \pm 0.003$), under a strict no-fit rule confirmed in all eight adversarial checks — reported as derived only at the consistency tier (depth one), the drainable share’s own scale-running taken as input and the $\pm 2\sigma$ precision defensible by no chain; the engine’s first theoretically-carried dimensionless quantity (internal identifier K18f in the released results kit). *Falsified and replaced*: the wealth-gate protection hypothesis (degree-priced edges protecting rich structures) failed measurement and led to the apex-funding rule; and the virial reading of the Big Crunch trigger — collapse at a pressure-balance crossing — failed a proxy-free event-count test ($R=1$ neither necessary nor sufficient; §11) and led to the position-gravity feasibility measurement of the hydrostatic channel. *Falsified*: the reader-directedness prediction, declared *before* measurement — that a frequency-selective operator reading the archive would make cross-generation structure shifts sign-consistent across seeds; the operator was built and demonstrably *reads* (semantic divergence under spectrum collapse; first crunch delayed $\sim 8\times$), but directedness failed on 6 seeds ($\sim 19\text{--}21$ generations per arm, fast-cycle regime) — an honest negative against spontaneous cross-cycle learning, with the reader retained as a default-off flag; the Planck-analogue derivation of the gravitational coupling, and with it the absolute expansion magnitude, across the entire hadronic candidate family ($H \tau_n$ four to five decades above the physical value; §9.5); and the heavy-ion reference-class reframing, closed as a regime claim on all three of its declared branches (equilibrium statistics, freeze-out systematics, kinematics); and the frozen counterfactual of the export-efficiency derivation (§5.1), declared *before* its runs — that disabling the agitation drain would raise the exponent to $+0.080$ and lengthen the cycle clock by $+25$ to $+28\%$ — substantially refuted on 25 fresh cells, where the small universes ($\times 1\text{--}\times 3$) cease to crunch altogether and, where collapse survives (from $\times 5$), the clock runs $+55$ to $+87\%$ longer with factor-two seed scatter: the agitation drain proves a marginally load-bearing southern feeder, not the flat diluter the mechanistic reading supposed, so that reading is downgraded while the depth-one interpolated value stands untouched (the frozen thresholds guarded only the effect-too-weak side, a named one-sided design hole, and formally returned *undecidable*, not an acquittal; internal identifier K18g). *Passed, campaign-scale*: the clean-engine reference baseline (§11.5) — fourteen cells under one frozen declaration, every structural gate (drift, charge — through 121 stellar collapses — partition closure, determinism protocol, and successor-generation health) passed in every cell, the deuterium supersession above emerging from it as measurement rather than hope. *Open and defined*: the neutron–proton mass splitting, which in this framework requires an emergent radius asymmetry between constituent configurations; the functional form of the coherence state, where $E \cdot C/Q$ is the declared simplest candidate and alternatives are discriminable against the recorded run histories; the post-bounce generation is now closed — the measured bimodality traced to a floored density re-anchor (regime-predictive in 14 of 14 cells), repaired by a zero-knob guard, and re-measured to 14-of-14 healthy successor generations (§efsec:baseline). All of these are internal tests of the model against its own dynamics; an external, observational discriminator remains the programme’s principal open obligation.

13 Scaling Outlook

Every result in this paper was produced on a single desktop CPU, the engine single-threaded per run for bit-exact determinism. That is the honest frame for what remains open: under the claim level stated at the outset, the open items are not defects but a defined, instrumented frontier — each with a named

experiment and a known cost.

What the measured curves already say. The cycle clock runs with scale (the first frozen scale-running verdict); the export efficiency survives a discovery-to-replication chain and has been carried to a depth-one derivation; the nucleosynthesis ladder, once the bookkeeping was audited, climbed from helium to $A = 49$ without a single new operator. The pattern across the programme’s history is consistent: measured obstacles have repeatedly resolved into bookkeeping, operator composition, or scale — not into failures of the framework’s grammar.

The scaling axis. The iron frontier ($Z \geq 26$), the helium mass fraction, and the composition point are magnitude questions on declared observables; the instruments exist, and each next decade of scale is a compute purchase, not a research risk. Runtime currently grows steeply with scale, and event-driven (Gillespie-class) scheduling is the declared engineering step that converts the next decade from weeks to days.

The architecture axis. Two construction stages are defined: running couplings (the engine’s constants are frozen at boot; the renormalisation ladder that would carry a running one is built and climbable, and it is the framework’s candidate route to the 10^{19} hierarchy), and the post-bounce anchor experiment of §11.5 (one flag, one campaign). Both are bounded programmes with pre-registered verdicts waiting to be run.

Cost asymmetry. The engine has no gradient pass, no optimiser state, and no accelerator requirement; its cost is linear in events with a bounded working set. Scaling this system buys *experiment throughput* — more cells, more seeds, more decades of the levers above — which is what a falsification-driven programme converts directly into verdicts.

14 Application Surface

The properties measured above — determinism, exact conservation, full-cycle cosmology, mass-keyed auditability — map onto application surfaces that conventional simulation stacks do not serve natively:

A falsification instrument for emergent-physics hypotheses. The engine’s working discipline — frozen declarations, pre-registered kill criteria, bit-exact reproduction, ledger-exact accounting — is built into the tooling, not layered on top. Any hypothesis expressible in the operator grammar can be tested with the same rigor documented in this paper’s record.

Deterministic generative cosmology. A seed is a universe: every run is a complete, replayable cosmological history — eras, nucleosynthesis, supernovae, collapse, and re-birth — reproducible bit-for-bit from one integer. This already powers an interactive exploration surface in production, where users run, inspect, and share engine universes; the same property serves procedural content, education, and outreach at marginal cost.

An auditable physics core. Because every transaction balances and every trajectory replays, the engine is inspectable end to end — a property of interest wherever simulation results must be attributable, from research infrastructure to regulated settings.

One grammar, two substrates. The same conserving operator grammar that produces baryonic structure here ingests natural language in the programme’s language mode, where it yields measured sequence prediction beyond its count controls on identical hardware. The bet the two instruments share — that one energy-conserving graph grammar underlies structure formation across domains — is the programme’s core intellectual property.

Pantareon offers the technology for licensing, acquisition, or funded co-development.

15 Conclusion

EGT offers a framework where physics and information theory are indistinguishable. By modelling reality as a spectral energy graph, we have argued that the universe does not merely increase in entropy; it actively compiles itself into an increasingly structured spectral archive. Through fifth-force activity—the substrate’s own dynamics—the system preserves its own history: each cosmological cycle is spectrally primed by the last — a priming that, as measured above, the ledger carries faithfully and a Postulate-II reading operator now demonstrably consumes, but which does not yet *steer*: directed cross-cycle evolution remains refuted in the tested regime.

Thermodynamic consistency. While the system creates local order (increasing spectral complexity), it strictly adheres to global energy conservation ($\Delta E_{\text{tot}} = 0$ in the closed cosmological configuration). The emergence of structure in EGT is dissipative self-organisation in a closed system with internal gradients. The resonant conductivity of Postulate II acts as a selection filter: chaotic spectral configurations face high impedance and dissipate, while resonant structures persist. Under the thermodynamic pressure of the ECQ triad—the simultaneous accumulation of substrate sediment E , depletion of matter energy Q , and growth of chronos C —the system is driven toward elaborate self-organisation: measured in the instrumented runs as a strong, monotone tendency, not asserted as a necessity.

Position within the Pantareon Foundations programme. The present work establishes EGT at the cosmological and particle-physics scales. These scales constitute the foundational substrate on which all higher-order organisation depends. The scale-invariance of the postulates indicates that the same formalism may be brought to bear on systems at other scales of physical organisation. Such extensions—including the question of whether the ECQ engine drives biological complexity in a way that allows the temporal scale of biological emergence to be predicted from structural complexity—are deferred to separate investigations within the Pantareon Foundations programme.

Whether this model corresponds to physical reality or is a useful computational metaphor remains to be determined by experiment. The programme’s stated trajectory is to graduate from the internal falsification record documented above to *external* prediction: quantitative observables — a helium yield, an isotropy index, a correlation slope — computed from the substrate dynamics at scales where they can confront data. The present paper documents the instrument standing at its reference baseline: a machine that runs complete, energy-exact cosmological cycles — big bang to re-birth, with the four interactions as audited operators, space as a conserved particle population, a nucleosynthesis ladder terminated by emergent stellar collapse, and the dark sector carried structurally by the vacuum ledger — together with the measured distances between that machine and our universe, and the named experiments that close them. The distances are the scaling programme; the experiments are declared; what they require is iteration speed. The frontier is limited by compute and defined construction stages, not by ideas. Empirical results of our simulation are released alongside as the versioned results kit.

References

- [1] Verlinde, E., "On the Origin of Gravity and the Laws of Newton", *JHEP* **04**, 029 (2011).
- [2] Sakharov, A. D., "Vacuum quantum fluctuations in curved space and the theory of gravitation", *Dokl. Akad. Nauk SSSR* **177**, 70 (1967).
- [3] Wolfram, S., "A Class of Models with the Potential to Represent Fundamental Physics", *arXiv:2004.08210* [hep-th] (2020).
- [4] Dowker, F., "Causal sets and the deep structure of spacetime", *arXiv:gr-qc/0508109* (2005).
- [5] Chew, G. F., Frautschi, S. C., "Principle of Equivalence for All Strongly Interacting Particles Within the S-Matrix Framework", *Phys. Rev. Lett.* **7**, 394 (1961).
- [6] Takahashi, T. T., Matsufuru, H., Nemoto, Y., Suganuma, H., "Three-Quark Potential in SU(3) Lattice QCD", *Phys. Rev. Lett.* **86**, 18 (2001).
- [7] Rovelli, C., *Quantum Gravity*, Cambridge Monographs on Mathematical Physics, Cambridge University Press (2004).
- [8] Jacobson, T., "Thermodynamics of Spacetime: The Einstein Equation of State", *Phys. Rev. Lett.* **75**, 1260 (1995).
- [9] Padmanabhan, T., "Thermodynamical Aspects of Gravity: New insights", *Rep. Prog. Phys.* **73**, 046901 (2010).
- [10] Ambjørn, J., Jurkiewicz, J., Loll, R., "Spectral Dimension of the Universe", *Phys. Rev. Lett.* **95**, 171301 (2005).
- [11] Friston, K., "The Free-Energy Principle: A Unified Brain Theory?", *Nat. Rev. Neurosci.* **11**, 127–138 (2010).
- [12] von der Malsburg, C., "The Correlation Theory of Brain Function", in *Models of Neural Networks II*, E. Domany, J. L. van Hemmen, K. Schulten (eds.), Physics of Neural Networks, Springer, New York, ch. 2, pp. 95–119 (1994); orig. MPI für Biophysikalische Chemie, Internal Report 81-2 (1981).
- [13] Wheeler, J. A., "Information, Physics, Quantum: The Search for Links", in *Complexity, Entropy and the Physics of Information*, W. H. Zurek (ed.), Addison-Wesley, Redwood City, CA, pp. 3–28 (1990).
- [14] Hagberg, A. A., Schult, D. A., Swart, P. J., "Exploring Network Structure, Dynamics, and Function using NetworkX", *Proceedings of the 7th Python in Science Conference (SciPy 2008)*, pp. 11–15 (2008).
- [15] Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., Teller, E., "Equation of State Calculations by Fast Computing Machines", *J. Chem. Phys.* **21**(6), 1087–1092 (1953).
- [16] Hastings, W. K., "Monte Carlo Sampling Methods Using Markov Chains and Their Applications", *Biometrika* **57**(1), 97–109 (1970).
- [17] Ambjørn, J., Görlich, A., Jurkiewicz, J., Loll, R., "Nonperturbative Quantum Gravity", *Phys. Rep.* **519**(4–5), 127–210 (2012).
- [18] Diestel, R., *Graph Theory*, 5th Edition, Springer (2017).
- [19] Watts, D. J., Strogatz, S. H., "Collective Dynamics of 'Small-World' Networks", *Nature* **393**, 440–442 (1998).
- [20] Shannon, C. E., "A Mathematical Theory of Communication", *Bell Syst. Tech. J.* **27**(3), 379–423 (1948).
- [21] Landauer, R., "Irreversibility and Heat Generation in the Computing Process", *IBM J. Res. Dev.* **5**(3), 183–191 (1961).
- [22] Penrose, R., *The Emperor's New Mind*, Oxford University Press (1989).
- [23] Hawking, S. W., "Particle Creation by Black Holes", *Commun. Math. Phys.* **43**(3), 199–220 (1975).
- [24] Bekenstein, J. D., "Black Holes and Entropy", *Phys. Rev. D* **7**(8), 2333–2346 (1973).