Cosmological Horizons as Epistemic Bounds of Conservation-Congruent Encodings

Peter David Fagan, School of Informatics, University of Edinburgh

May 13, 2026

Abstract

The foundations of standard cosmology rely on modelling reality with continuous macroscopic field equations. This note identifies an observer-resource assumption hidden by such equations and analyzes it using the Conservation-Congruent Encoding (CCE) framework. Within CCE, a projection Π is a coarse-graining of physical reality whose erasure, refinement, and maintenance carry energetic or informational costs. The note argues that several horizon-like limits also admit an observer-indexed operational reading within this ledger. Forward in time, heat death is read as *Predictive Dissipation*: the irreversible loss of macroscopic signal when a projection truncates the metric exhaust required to hold that signal distinct from the bath. Backward in time, the Big Bang singularity is read as *Retrodictive Divergence*: the divergence of the Landauer-scale cost required to re-instantiate erased branch distinctions under an assumption of costless resolution. Cosmological expansion and redshift are treated as standard geometric identities with an additional CCE bookkeeping interpretation, while gravitational singularities mark lower area-capacity limits for physically instantiated projections. The result is not a new dynamics, but a stricter operational license for using continuous models: observers do not access an infinite continuous universe at arbitrary precision, but work within a finite epistemic bubble bounded by their physical embedding.

1 Introduction

The foundational equations of relativistic cosmology are normally written as if their variables were continuously available to an ideal observer. In the Friedmann–Lemaître–Robertson–Walker description, in the Raychaudhuri equation, and in stress–energy closure assumptions, quantities such as the scale factor $a(t)$, the expansion scalar θ, and the stress–energy tensor $T_{\mu\nu}$ are treated as real-valued objects whose precision can be increased without cost. This is a useful convention, but it hides a blind spot: these equations contain no internal variable that tracks their own divergence from an underlying physical reality and do not account for observer effects. They track curvature, density, expansion, and focusing, but not the lossiness of the macroscopic projection that made those quantities available in the first place.

That omission matters because a macroscopic equation is also a compression map. It discards unresolved degrees of freedom in order to become usable. Since the equation has no term for its own lossiness, it will continue to output formal values after the projection that supports those values has ceased to be physically maintainable. Pushed past its operational domain, the same formalism that gives precise predictions over its operational domain will happily return infinite densities, zero-area congruences, or unbounded future dilution, because the equation has no register for the precision of its own predictive or retrodictive power; it fails to identify the point at which its outputs become operationally unsupported for that projection.

This hidden assumption of macroscopic equations is referred to here as the *Platonic Observer Fallacy*: the implicit assumption that mathematical limits such as $t \to +\infty$, $t \to -\infty$, $r \to 0$, $A \to 0$, or $A \to \infty$ can be taken without accounting for the physical reality of observation itself which is required to track those limits. It assumes a projection operator Π that sits outside the universe, with infinite thermal capacity, unlimited entropy export, and zero metric-sector exhaust. It therefore grants a physical observer the privileges of a costless mathematical viewpoint.

This fallacy has two arrow-dependent forms. Forward in time it ignores dissipation: the *truncation of metric exhaust* treats a resolved macrostate as though it can be stabilized without paying the continuous thermal cost of holding it distinct from the bath. Backward in time it ignores multiplicity: the *assumption of costless resolution* treats erased or marginalized alternatives as though they can be re-instantiated without Landauer cost. The resulting failures are, respectively, *Predictive Dissipation* and *Retrodictive Divergence*.

The Platonic Observer Fallacy leads directly to the *Principle of Epistemic Bounding*: macroscopic equations of reality are operationally licensed only within the domain of the observer. A formal limit becomes operationally undefined for that observer when maintaining the required projection exceeds the capacity of the observer. The issue is not whether a symbol can be pushed to a limiting value on paper; it is whether the distinctions represented by that symbol are supported by a physically realizable observer.

The Conservation-Congruent Encoding (CCE) framework resolves this by treating the projection Π as a physically instantiated, resource-bound act. A projection is an embodied process of selection and compression: it must be carried by physical memory, stabilized against ambient fluctuations, and continually updated through irreversible operations. While Landauer's principle establishes the strict thermal cost of erasing information, and reversible computation defines the theoretical limit of preserving it [16, 4, 5, 15], CCE abstracts this necessity into a generalized physical ledger. A macroscopic field becomes physically meaningful only if the observer possesses the capacity to support the distinctions it represents [11].

Applying this ledger suggests an operational reading of several standard cosmological endpoints. Heat death becomes the limiting case of Predictive Dissipation, where macroscopic distinctions dissipate entirely into the unresolved bath, leaving the observer with no distinguishable signal. Backward extrapolation toward the Big Bang becomes Retrodictive Divergence, where the thermodynamic cost of importing erased branch distinctions exceeds any finite observer capacity and, in the formal limit, diverges without bound. Cosmological expansion and redshift retain their standard geometric meaning, while CCE assigns an additional metric-sector bookkeeping cost to the stabilized source world-tube. Finally, gravitational singularities are treated not as observer-accessible infinitely dense points, but as limits where tracking a geometric congruence can exceed a projection's area-capacity floor.

The resulting framework does not replace the dynamics of general relativity; it issues a stricter, physical license for their use. Continuous geometry remains formally pristine, but its encoding is bound by conservation laws. The remainder of this note structures that bounding. Section 2 formally defines the conservation-congruent ledger, linking phase-space measures, projection channels, relative entropy, and metriplectic evolution. Section 3 applies this ledger to time, recasting heat death and the Big Bang as Predictive Dissipation and Retrodictive Divergence. Section 4 treats space, retaining the standard redshift, time-dilation, CMB blackbody, and surface-brightness identities while assigning CCE bookkeeping costs to them. Section 5 applies the same ledger to gravity, placing observer-indexed capacity and sensitivity contours around the Raychaudhuri limiting regimes. Together, these bounds describe a finite, physically constrained epistemic bubble inside the broader continuous description.

2 Quantifying the Epistemic Bound

The introduction stated the physical rule; this section gives the bookkeeping that the Platonic Observer Fallacy suppresses. CCE treats each macrostate as the output of a projection channel, evolves that output by a metriplectic ledger, and bounds the resulting distinctions by a finite thermodynamic substrate.

2.1 The Projection Channel and the Measure-Theoretic Cutoff

In the classical formulation, let (Ω_t, ω_t) be the instantaneous state space with symplectic form ω_t and dimensionless Liouville measure

$$d\mu_t(x) = \frac{\omega_t^n / n!}{h^n}, \tag{1}$$

where $2n = \dim \Omega_t$. A statistical state is a density $p_t(x)$ with respect to $d\mu_t$. The observer's projection is a coarse-graining channel, written as a Markov kernel

$$K_t(dm|x) = P(M_t \in dm | X_t = x), \qquad p_t(m) = \int_{\Omega_t} K_t(dm|x) p_t(x) d\mu_t(x). \tag{2}$$

The deterministic projection $m_t = \Pi_t(x_t)$ is the special case $K_t(dm|x) = \delta_{\Pi_t(x)}(dm)$. Throughout this note, Π denotes the physically chosen projection class, while K_t denotes the probabilistic channel that implements it at time t. In a quantum version, the analogous object is a completely positive trace-preserving coarse-graining map from microscopic density operators to observer records [23]. The macrostate is therefore a physical channel output, represented by a probability measure over the model space, not a free annotation attached to the microstate.

Standard continuum physics often permits formal pointwise limits or arbitrarily fine refinements, even when the corresponding operational measurement would require a finite-resolution detector and a physical stabilization protocol. Under CCE, this limit is not physically available for free. Evaluating a field at a mathematical point would require $p_t(m)$ to approach a singular Dirac delta measure against the unresolved background. The physical strain of maintaining this instantaneous distinction against continuous thermal fluctuations of the bath is captured by

$$I_\Pi(X_t; M_t) = D_{\mathrm{KL}}(p_t(x, m) \,\|\, p_t(x) p_t(m)). \tag{3}$$

This quantity is the instantaneous information load of the projection: the nats of micro–macro correlation that must be stored, stabilized, refreshed, or discarded for the macrostate to remain meaningful. It also avoids counting points in a continuum: an effective branch number is the exponential of an entropy or relative phase volume [7].

Because the KL divergence between a singular delta measure and an absolutely continuous background measure diverges, isolating a perfect mathematical point requires unbounded informational strain and, under the CCE ledger, unbounded free-energy support. Therefore CCE imposes a strict measure-theoretic cutoff. The continuous limit $V \to 0$ is physically truncated at the scale where I_Π saturates the observer's instantaneous thermal capacity. The projection operator Π cannot output exact points; it can only output finite, minimum-variance macroscopic cells bounded by the observer's physical ledger.

2.2 Kinematics vs. Dynamics: The Metriplectic Drift

The projection channel becomes a cosmological equation only when one asks how its outputs evolve. Standard cosmological models assume that if a macroscopic field is accurately measured at an instant

t_0, its governing differential equation can be rolled forward indefinitely without loss of physical fidelity. This conflates two distinct operations: the instantaneous state mapping and the dynamical evolution of the projection. Under CCE, these operations carry two separate physical ledgers.

2.2.1 The First Ledger: Static Instantiation Cost

The first ledger is the static projection cost: the instantaneous physical strain of resolving the field into a finite macroscopic cell. Evaluating this resolved variable at a specific time requires the projection $K_t(dm|x)$ to establish a stable statistical distinction against the background measure of the bath. The corresponding kinematic load is

$$\mathcal{C}_{\Pi}^{\text{stat}}(t) \equiv I_{\Pi}(X_t; M_t) = D_{\text{KL}}\big(p_t(x,m) \,\|\, p_t(x)p_t(m)\big). \tag{4}$$

This static instantiation cost dictates what must be physically stored or stabilized for the finite macroscopic cell to remain meaningful at t, but it does not guarantee that the formal differential equation will remain bound to the physical world-tube as time evolves.

2.2.2 The Second Ledger: Dynamical Divergence

The second ledger tracks dynamical drift. Let $p_t^{\text{phys}}(m)$ be the channel output obtained by evolving the underlying physical state and then applying K_t. Let $q_t(m)$ be the macro-distribution obtained by rolling forward the macroscopic equation of state from the same initial record, $q_{t_0} = p_{t_0}^{\text{phys}}$. The drift between the formal equation and the physical world-tube is

$$\mathcal{D}_{\Pi}(t) = D_{\text{KL}}\big(p_t^{\text{phys}}(m) \,\|\, q_t(m)\big). \tag{5}$$

Standard continuous dynamics implicitly set $\mathcal{D}_{\Pi}(t) = 0$ for all t at no cost. CCE instead treats growth in $\mathcal{D}_{\Pi}$ as the strict signature that the idealized equation of state is drifting away from the underlying physical reality.

2.2.3 The Metriplectic Closure

To see how this drift enters physical observables, let $\mathcal{O}_t(m)$ be a resolved macroscopic observable, such as a density, scale factor, or expansion scalar, with expectation

$$\langle \mathcal{O} \rangle_t = \int \mathcal{O}_t(m) p_t^{\text{phys}}(m) dm. \tag{6}$$

In a lossless continuum idealization, the observable is transported only by the conservative Hamiltonian sector,

$$\frac{d\langle \mathcal{O} \rangle_t}{dt} = \langle \partial_t \mathcal{O} + \{\mathcal{O}, H\} \rangle_t \,, \tag{7}$$

or, when $\mathcal{O}$ has no explicit time dependence, simply $d\langle \mathcal{O} \rangle_t/dt = \langle \{\mathcal{O}, H\} \rangle_t$. This is the formal image of the Platonic Observer Fallacy: the macroscopic projection is assumed to remain resolved while its dynamical drift and unresolved exhaust are ignored.

CCE replaces that truncation with a metriplectic closure. To keep the formal equation bound to the physical world-tube, the conservative Poisson sector must be supplemented by a symmetric, positive metric bracket that encodes the required dissipative correction [21, 22]:

$$\frac{d\langle \mathcal{O} \rangle_t}{dt} = \langle \partial_t \mathcal{O} + \{\mathcal{O}, H\} \rangle_t + \langle [\mathcal{O}, S_{\Pi}]_{\Pi} \rangle_t \,. \tag{8}$$

Here H denotes the conservative generator appropriate to the chosen continuum closure, S_Π is the generalized entropy associated with the projection, and $[\cdot,\cdot]_\Pi$ is determined by the unresolved bath and the physical substrate used to maintain K_t. For this closure to be conservation-congruent in a closed resolved-plus-bath ledger, the dissipative sector must be degenerate on the conserved generator and any other protected invariants; in particular $[H, S_\Pi]_\Pi = 0$. If the resolved sector exchanges energy or entropy with an external bath, the same equation is instead an open-system balance and the missing exchange must be carried by an explicit bath port so that the combined ledger conserves the relevant total quantity. In local resolved coordinates, a representative closed-ledger metric bracket has the form

$$[F, G]_\Pi(m, t) = \nabla_m F^\mathsf{T} G_\Pi(m, t) \nabla_m G, \qquad G_\Pi(m, t) \succeq 0, \qquad G_\Pi(m, t)\nabla_m H = 0, \tag{9}$$

though CCE requires only the existence of a positive dissipative sector with the appropriate conservation degeneracies or, in an open setting, an explicitly balanced port representation.

2.2.4 The Master Bookkeeping Inequality

The metric bracket extracts a continuous metric-sector exhaust: the irreversible entropy-production rate paid by the projection,

$$\dot\Sigma_\Pi^{\text{irr}}(t) = \langle [S_\Pi, S_\Pi]_\Pi \rangle_t \geq 0, \tag{10}$$

measured in nats per unit time when S_Π is dimensionless. While the static load $\mathcal{C}_\Pi^{\text{stat}}$ is the stock of resolved correlation, $\dot\Sigma_\Pi^{\text{irr}}$ is the thermodynamic flow required to refresh that stock and correct the drift. The following bound is therefore used as a CCE admissibility condition for a specified maintenance protocol, rather than as a theorem of the metric bracket alone:

$$\dot\Sigma_\Pi^{\text{irr}}(t) \geq \left[\frac{d}{dt}\mathcal{C}_\Pi^{\text{stat}}(t)\right]_+ + \Gamma_\Pi(t)\mathcal{C}_\Pi^{\text{stat}}(t) + \left[\frac{d}{dt}\mathcal{D}_\Pi(t)\right]_+, \qquad [u]_+ = \max\{u, 0\}, \tag{11}$$

where $\Gamma_\Pi(t)$ is the minimum refresh or error-correction rate imposed by the ambient noise model and the tolerated error probability of the projection. The three terms are additive when the implementation realizes them as distinct integration, maintenance, and correction operations; if a particular protocol couples them, the right-hand side should be read as the corresponding protocol-level lower ledger rather than a double count. Under that convention, the continuous metric exhaust pays for three operations:

- *Integration cost*: $[d\mathcal{C}_\Pi^{\text{stat}}(t)/dt]_+$, paying for newly built micro–macro correlation.
- *Maintenance cost*: $\Gamma_\Pi(t)\mathcal{C}_\Pi^{\text{stat}}(t)$, paying to preserve existing correlation against diffusion into the bath.
- *Correction cost*: $[d\mathcal{D}_\Pi(t)/dt]_+$, paying for the accumulating divergence between the rolled-forward equation and the physical state it claims to track.

Additional erasure terms may enter when a particular computation discards distinctions, but they contribute to $\dot\Sigma_\Pi^{\text{irr}}$ only after the irreversible operation and the bath that receives the exported entropy have been specified.

Standard cosmological equations, such as ideal FLRW evolution or the Raychaudhuri equation, are recovered only by assuming that the static cell remains valid and that the dynamical exhaust vanishes. CCE refuses that shortcut. An equation of state is a lossy projection whose physical validity requires the observer to afford the continuous metric exhaust accumulated while computing and stabilizing its macroscopic variables.

2.3 The Observer's Finite Capacity

The physical validity of a macroscopic projection terminates at the exact contours where the demands of the metric bracket exceed the observer's finite thermodynamic substrate. Along any chosen parameterization, such as coordinate time t, redshift z, or affine parameter λ, define the accumulated metric exhaust

$$\Sigma_{\Pi}(s_1, s_2) = \int_{s_1}^{s_2} \dot{\Sigma}_{\Pi}^{\text{irr}}(s) ds. \tag{12}$$

A CCE boundary is reached where the static stock $\mathcal{C}_{\Pi}^{\text{stat}}$, the dynamical drift $\mathcal{D}_{\Pi}$, or the accumulated exhaust Σ_{Π} exceeds the memory, free-energy, or entropy-export capacity of the observer implementing K_t. These boundaries define the observer's epistemic bubble and appear in two arrow-dependent modes of projection failure.

2.3.1 Predictive Dissipation: Truncating Metric Exhaust

As a physical system evolves forward in time or extends across a spatial parameter, its underlying micro-histories naturally mix and dissipate into the unresolved bath. This objective physical mixing causes the idealized macroscopic equation of state to continuously drift from the underlying physical reality. To keep the formal equation precise, the CCE ledger requires the observer to continuously extract metric exhaust to correct this drift.

The standard forward oversight is the truncation of this exhaust: treating the mathematical macrostate as though it remains perfectly sharp without acknowledging the continuous objective dissipation of the physical signal. The resulting operational failure is *Predictive Dissipation*. The universe objectively diffuses the signal, draining the usable static correlation $\mathcal{C}_{\Pi}^{\text{stat}}$. The projection mathematically fails not because the universe ceases to exist, but because this objective dissipation eventually leaves the observer unable to distinguish the finite macroscopic cell from the background noise.

This establishes the observer's forward and outward epistemic bounds. These limits are projection- and apparatus-indexed contours where the conserved signal density of the tracked field drops below the observer's specified sensitivity threshold. Beyond this boundary, the idealized equation may continue to output formal values that are divorced from the physical reality.

2.3.2 Retrodictive Divergence: Costless Resolution

Conversely, as a macroscopic equation of state is retrodicted backward in time or traced inward toward a focal point, it must map a single present macrostate to an exponentially proliferating tree of compatible micro-histories. The standard backward oversight is the assumption of costless resolution: treating these erased or marginalized distinctions as though they can be perfectly re-instantiated without supplying the necessary Landauer work. The resulting operational failure is *Retrodictive Divergence*.

Reversing a dissipative channel requires an ongoing integration cost to physically label and track the diverging branches. Because the observer's encoding capacity is finite, this branch load and its associated thermodynamic cost grow until they strictly exceed the hardware's absolute limits. The projection mathematically fails not because the physical universe collapses, but because the observer runs out of the physical memory required to compute the necessary distinctions.

This divergence mechanism establishes the observer's backward and inward epistemic bounds. These limits are projection- and apparatus-indexed contours where the informational demands of recovering erased dynamics saturate the observer's total informational capacity. Beyond these bounds, the formal equations may continue to output mathematical values—often diverging to infinite densities or singular points—but these infinities are merely the mathematical artifacts of an unbacked ledger. The equation has not respected the physical constraints of observation itself.

3 Temporal Horizons

The temporal application has a narrow purpose: to show why moving a projection through time carries an information cost measured in nats. Prediction asks how many nats of a present record remain capable of resolving future variables once metric exhaust has been paid. Retrodiction asks how many compatible past alternatives must be distinguished to recover a specific earlier state once erased distinctions are re-instantiated. Both are channel questions whose limits are governed by the metriplectic ledger.

3.1 Prediction: Predictive Dissipation

Suppose the observer resolves variables R_t and ignores bath variables B_t, so that $x_t = (R_t, B_t)$. A dissipative update is represented by a stochastic or coarse-grained evolution kernel

$$p_{t+\Delta t}(x') = \int P_{t+\Delta t,t}(dx'|x) p_t(x) d\mu_t(x). \tag{13}$$

After this update, the projection channel $K_{t+\Delta t}$ produces the future macro-record. Distinct micro-histories that differ only in B_t can therefore induce overlapping future records. The formal model may treat this as ordinary statistical merging, but the CCE ledger treats it as Predictive Dissipation: resolvable correlation is irreversibly dissipated into the bath unless the projection pays the metric exhaust required to maintain it.

For a given projection, define the predictive information

$$\begin{aligned} I_+(t) &= I(M_{t_0}; R_t) \\ &= D_{\mathrm{KL}}(p(m_{t_0}, r_t) \| p(m_{t_0}) p(r_t)) \,. \end{aligned} \tag{14}$$

This is the number of nats by which the present record constrains the future resolved state. If the projection requires a predictive noise floor $\mathcal{I}_\Pi^{\mathrm{noise}}(t)$ nats to distinguish that state from bath fluctuations and model drift, then the predictive deficit is

$$\Delta I_\Pi^+(t) = [\mathcal{I}_\Pi^{\mathrm{noise}}(t) - I_+(t)]_+, \qquad [u]_+ = \max\{u, 0\}. \tag{15}$$

Restoring this missing margin is a physical operation, not a symbolic one. Let $T_\Pi(t)$ be the local reservoir temperature of the apparatus or bath used by the projection. When restoration requires logically irreversible erasure or reinitialization of missing distinctions, Landauer's bound gives the lower reset scale

$$W_+^{\min}(t) \geq k_{\mathrm{B}} T_\Pi(t) \, \Delta I_\Pi^+(t). \tag{16}$$

The point is not that every predictive update literally performs an irreversible erasure. Measurement, reversible storage, calibration, and error correction have their own ledgers. The displayed bound applies to the reset part of the protocol: any lost distinction that must be physically reinitialized costs at least one nat of memory or negentropy per nat restored, with $k_{\mathrm{B}} T_\Pi(t)$ setting the ideal thermal scale for that irreversible step.

The forward epistemic bound is therefore the first time at which the available predictive information falls to the required margin,

$$t_{+} = \inf\{t > t_0 : I_{+}(t) \leq \mathcal{I}_{\Pi}^{\text{noise}}(t)\}. \tag{17}$$

Beyond this contour, the equation may still output formal values, but the observer no longer has the nats required to keep those values resolved as macroscopic distinctions. In cosmological language, heat death is the limiting case of Predictive Dissipation: the projection's future signal has dissipated below its predictive noise floor.

3.2 Retrodiction: Retrodictive Divergence

Retrodiction runs the same bookkeeping in the opposite direction. A dissipative channel that is simple to use forward is many-to-one in the forward direction and one-to-many when inverted. The Platonic error is to assume costless resolution: erased or marginalized alternatives are treated as though they can be recovered without physically reinstating the distinctions that the forward channel discarded. Given a present model state m_{t_0}, compatible past states are therefore described by a posterior distribution,

$$p(x_t|m_{t_0}) = \frac{P(m_{t_0}|x_t)p_t(x_t)}{P(m_{t_0})}, \qquad P(m_{t_0}|x_t) = \int K_{t_0}(dm_{t_0}|x_{t_0})\, P(dx_{t_0}|x_t). \tag{18}$$

Before this load can be compared with a finite apparatus capacity, the projection must fix a retrodictive partition $\mathcal{P}_{\Pi}(t) = \{C_j\}$, a Liouville-cell normalization, a support convention, and an error tolerance. With

$$p_j(t|m_{t_0}) = \int_{C_j} p(x_t|m_{t_0}) d\mu_t(x_t), \tag{19}$$

the retrodictive branch load is the coarse-grained conditional entropy

$$\mathcal{B}_{\Pi}(t; m_{t_0}) = H_{\mathcal{P}_{\Pi}}(X_t|M_{t_0} = m_{t_0}) = -\sum_{C_j \in \mathcal{P}_{\Pi}(t)} p_j(t|m_{t_0}) \ln p_j(t|m_{t_0}). \tag{20}$$

It is measured in nats, and $\exp \mathcal{B}_{\Pi}$ is the effective number of compatible past branches only relative to this specified partition, support, and tolerance. Changing those choices changes the contour. Exact retrodiction requires either selecting one branch, storing the unresolved alternatives, or importing an external record that performs the same discrimination.

The formal reason this cost can grow in irreversible regimes is that the forward process exports records into inaccessible degrees of freedom. Stochastic thermodynamics expresses this time-asymmetry as a path-space relative entropy,

$$D_{\text{KL}}(P_{\text{F}}[\gamma] \| P_{\text{R}}[\tilde{\gamma}]) = \left\langle \ln \frac{P_{\text{F}}[\gamma]}{P_{\text{R}}[\tilde{\gamma}]} \right\rangle_{\text{F}} = \frac{\langle \Delta s_{\text{tot}} \rangle_{\text{F}}}{k_{\text{B}}}. \tag{21}$$

This identity motivates, but does not by itself derive, a branch-entropy law for a particular coarse graining [9]. The CCE retrodictive postulate used below is narrower: when the bath entropy between t and t_0 is the inaccessible record needed to disambiguate the present macrostate, and when no external record supplies those labels, the missing branch information is bounded at the scale

$$\Delta \mathcal{B}_{\Pi}(t) \gtrsim \frac{\Delta S_{\text{bath}}(t)}{k_{\text{B}}}. \tag{22}$$

This is a CCE admissibility assumption for such projections, not a theorem of stochastic thermodynamics alone. Under that assumption, reversing compression requires paying for distinctions that the forward projection discarded. When those distinctions proliferate under backward extrapolation, the cost becomes a Retrodictive Divergence.

Let $\mathcal{I}_\Pi^{\text{cap}}$ be the observer's total macroscopic-distinction capacity in nats. The capacity deficit for exact retrodiction is

$$\Delta I_\Pi^-(t) = [\mathcal{B}_\Pi(t; m_{t_0}) - \mathcal{I}_\Pi^{\text{cap}}]_+. \tag{23}$$

If the observer must instantiate those excess labels through irreversible memory reset operations coupled to the local reservoir $T_\Pi(t)$, then

$$W_-^{\text{min}}(t) \geq k_\text{B} T_\Pi(t)\, \Delta I_\Pi^-(t). \tag{24}$$

If the labels are instead supplied by measurement, reversible storage, or an external record, those resources must be charged separately against the apparatus capacity and calibration ledger. The backward epistemic bound is the first lookback time at which the required branch load exceeds capacity,

$$t_- = t_0 - \inf\{\tau > 0 : \mathcal{B}_\Pi(t_0 - \tau; m_{t_0}) > \mathcal{I}_\Pi^{\text{cap}}\}. \tag{25}$$

For $t < t_-$, exact branch recovery is unsupported by this observer's finite thermodynamic substrate. Coarse retrodictions, parameter fits, and early-universe model comparisons may still be possible; what fails is unlimited backward refinement at zero thermodynamic cost.

Thus the temporal epistemic window of a projection is bounded by two nat budgets. The forward bound is Predictive Dissipation: predictive mutual information falls below the noise threshold after metric exhaust is paid. The backward bound is Retrodictive Divergence: conditional branch entropy exceeds macroscopic-distinction capacity because costless resolution is physically unavailable. A different observer, projection, or thermodynamic resource can move these contours, but no finite observer obtains them for free.

4 Spatial Horizons: CCE Bookkeeping for Cosmological Observables

Space has analogous horizons. A distant source is not just a point on a passive manifold; it is a world-tube whose emitted distinctions must be maintained, transported, decoded, and retrodictively assigned to a source model. The same metric-sector exhaust that limits temporal prediction also limits spatial observability. Across large separations, maintaining a coherent projection carries an accumulated energetic toll.

Set $c = 1$ and write the relevant cosmological patch in Friedmann–Lemaître–Robertson–Walker form,

$$ds^2 = -dt^2 + a^2(t)\left[d\chi^2 + S_k^2(\chi) d\Omega^2\right]. \tag{26}$$

For a photon with momentum p^μ emitted by a comoving source at t_em and received by a comoving observer at t_0, the redshift is

$$1 + z = \frac{a(t_0)}{a(t_\text{em})} = \frac{(u_\mu p^\mu)_\text{em}}{(u_\mu p^\mu)_\text{obs}}. \tag{27}$$

The standard kinematic redshift gives the logarithmic stretch along the null channel. Choose the path parameter λ to increase from emission to observation in the expanding FLRW patch, so $d\ln a(t(\lambda))/d\lambda \geq 0$; with the opposite orientation the following oriented derivative must be sign-adjusted. Then

$$\ln(1+z) = \int_{t_{\rm em}}^{t_0} H(t)dt = \int_{\lambda_{\rm em}}^{\lambda_{\rm obs}} \frac{d}{d\lambda}\ln a(t(\lambda))d\lambda. \tag{28}$$

CCE assigns a bookkeeping exhaust to this standard stretch through a positive calibration density $\eta_\Pi(\lambda)$, interpreted as projection-dependent nats per logarithmic metric stretch:

$$\dot{\Sigma}^{\rm irr}_{\Pi,z}(\lambda) := \eta_\Pi(\lambda)\frac{d}{d\lambda}\ln a(t(\lambda)), \qquad \Sigma_{\Pi,z}(\lambda_{\rm em},\lambda_{\rm obs}) := \int_{\lambda_{\rm em}}^{\lambda_{\rm obs}} \dot{\Sigma}^{\rm irr}_{\Pi,z}(\lambda)d\lambda. \tag{29}$$

Equivalently,

$$\widehat{\Sigma}_{\Pi,z} \equiv \int_{\lambda_{\rm em}}^{\lambda_{\rm obs}} \eta_\Pi^{-1}(\lambda)\dot{\Sigma}^{\rm irr}_{\Pi,z}(\lambda)d\lambda = \ln(1+z), \tag{30}$$

and, when η_Π is constant over the channel,

$$\Sigma_{\Pi,z} = \eta_\Pi \ln(1+z), \qquad 1+z = \exp(\Sigma_{\Pi,z}/\eta_\Pi). \tag{31}$$

These equations are definitions of the CCE redshift ledger unless an independent detector-and-bath model fixes η_Π. The total accumulated exhaust Σ_Π may also include detector response, foreground confusion, lensing scatter, and retrodictive source-assignment work, so $\Sigma_\Pi \geq \Sigma_{\Pi,z}$. Within this calibration, redshift acts as the invariant geometric input to the minimum adiabatic bookkeeping cost of transporting the source world-tube through the FLRW metric sector. CCE keeps this kinematic statement but changes its operational role. A redshifted observation is a channel from source phase-space distributions to detector records. If $Y_{\rm em}$ labels a resolved source macrocell and Ξ_z is the detector record, write

$$\begin{aligned} q_z(\xi_{\rm det}|y)d\xi_{\rm det} &= P(\Xi_z \in d\xi_{\rm det}|Y_{\rm em}=y), \\ q_z(\xi_{\rm det}|y)d\xi_{\rm det} &= \int K_{\rm det}(d\xi_{\rm det}|\xi_0)\,U_z(d\xi_0|\xi_{\rm em})\,E(d\xi_{\rm em}|y). \end{aligned} \tag{32}$$

Here E prepares the emitted photon distribution, U_z is the Liouville/geodesic transport kernel along the null congruence, and $K_{\rm det}$ includes detector response, foregrounds, unresolved bath variables, and lensing scatter. Noise is not an additive scalar on a world-tube; it is the stochastic part of the covariant channel from emission to record.

Two compatibility conditions make this precise. First, neighboring source macrostates must remain statistically distinguishable after transport. Because the observer does not have direct access to the source manifold, assigning a received detector record Ξ_z to a specific source coordinate y^a is fundamentally a parameter-estimation problem. The fidelity of this assignment is governed by the Fisher information carried by the received-record family $q_z(\xi_{\rm det}|y)$,

$$\mathcal{J}_{ab}(z;y) = \int q_z(\xi_{\rm det}|y)\,\partial_a \ln q_z(\xi_{\rm det}|y)\,\partial_b \ln q_z(\xi_{\rm det}|y)\,d\xi_{\rm det}. \tag{33}$$

The Fisher line element $ds_F^2 = dy^a\mathcal{J}_{ab}dy^b$ is the local distinguishability metric on source assignments. Under the usual regularity and local-unbiasedness assumptions, the Cramér–Rao bound gives $\mathrm{Cov}(\hat{y}) \succeq \mathcal{J}^{-1}$. With a prior density $\pi(y)$, the Bayesian Cramér–Rao, or van Trees, inequality gives schematically

$$\mathrm{BayesMSE}(\hat{y}) \succeq \left(\mathbb{E}_\pi\mathcal{J}(z;Y_{\rm em}) + \mathcal{J}_\pi\right)^{-1}, \qquad (\mathcal{J}_\pi)_{ab} = \int \pi(y)\,\partial_a\ln\pi(y)\,\partial_b\ln\pi(y)\,dy. \tag{34}$$

Thus redshift-driven Fisher loss inflates posterior uncertainty for the same statistical agent used in temporal retrodiction [28]. To make the local cutoff operational, fix a source chart, a prior $\pi(y)$, a source point y or prior-averaging rule, and a tolerated error probability or Bayes-risk threshold. Let $\mathcal{S}_\Pi \subset T_yY$ be the nonzero source-coordinate separations that the projection elects to treat as distinct macrostates, measured in that chart and prior model. Let $\alpha_\Pi > 0$ be the resulting minimum squared Fisher distance required by that threshold. In a CCE model that converts posterior uncertainty into source-assignment branch load, the local resolution condition is

$$\delta y^a \mathcal{J}_{ab}(z;y)\delta y^b > \alpha_\Pi \quad \text{for every} \quad \delta y \in \mathcal{S}_\Pi, \qquad \text{(Condition I)} \tag{35}$$

The mutual-information condition $I_\Pi(Y_{\rm em};\Xi_z) > \mathcal{I}_\Pi^{\rm noise}(z)$ is the integrated version, with $\mathcal{I}_\Pi^{\rm noise}$ fixing a tolerated error probability rather than a hard cliff. Different choices of $(y, \pi, \mathcal{S}_\Pi, \alpha_\Pi, \mathcal{I}_\Pi^{\rm noise})$ define different, explicitly parameterized spatial contours.

Second, the transport statement needs a frame convention. For photons, let $f(x,p)$ be the distribution on the future null mass shell $p^\mu p_\mu = 0$, $p^0 > 0$, evaluated in a local orthonormal frame on a chosen spacelike slice. The invariant content needed below is collisionless transport of this on-shell distribution, equivalently conservation of phase-space occupation along the geodesic flow:

$$\frac{df}{d\lambda} = 0, \qquad \frac{I_\nu}{\nu^3} = \text{constant along a ray}, \qquad \text{(Condition II)} \tag{36}$$

The more detailed mass-shell measure may be written in a chosen frame, but the subsequent optical identities use only this radiative-transfer invariant. Thus CCE adds no new optical mechanism. It reads standard redshift identities as bookkeeping rules for a finite projection: energy is assigned to the metric port, source-assignment support is specified by the chosen statistical model, and phase-space occupation numbers remain invariant [20, 25, 8].

Let $\mathcal{B}_\Pi(z)$ be the CCE branch-load functional for assigning the received record to a source cell at redshift z, with the source cell, prior, support measure, and tolerance fixed by the projection. The spatial epistemic horizon for those fixed choices is then an operational contour,

$$\begin{aligned} z_\Pi(y;\pi,\mathcal{S}_\Pi,\alpha_\Pi,\mathcal{I}_\Pi^{\rm noise},\mathcal{I}_\Pi^{\rm cap}) = \inf\{z > 0 : \min_{\delta y \in \mathcal{S}_\Pi} \delta y^a \mathcal{J}_{ab}(z;y)\delta y^b \le \alpha_\Pi \text{ or} \\ \mathcal{B}_\Pi(z) > \mathcal{I}_\Pi^{\rm cap}\}. \end{aligned} \tag{37}$$

Crossing this contour means that maintaining the source assignment requires more integration time, calibration information, memory, or free energy. It is not a discontinuity in the null geodesic.

Three standard observables are then faces of the same transfer map.

4.1 Adiabatic World-Tube Deformation: Supernova Time Dilation

Type Ia supernova light curves broaden by the same factor that redshifts their spectra. In CCE terms, this is an adiabatic deformation of the source world-tube. For a resolved photon packet,

$$\nu_{\rm obs} = \frac{\nu_{\rm em}}{1+z}, \qquad \Delta E_{\rm obs} = h\nu_{\rm obs} = \frac{\Delta E_{\rm em}}{1+z}. \tag{38}$$

The time stretching is the operational response forced by that same null geometry once the projection ledger is included. Continuous metric exhaust does not scatter the packet, as in tired-light accounts, and it does not act as a passive stretching of an empty vacuum. It steadily lowers the Fisher information carried by the received phase marker relative to the source phase. To keep adjacent phase markers statistically distinguishable, the observer must widen the temporal acceptance window until the phase-estimation covariance again satisfies the Cramér–Rao bound. For two neighboring source phase markers emitted from the same comoving world-line and received by the same observer, the null channel therefore imposes

$$\int_{t_{\rm em}}^{t_0} \frac{dt}{a(t)} = \int_{t_{\rm em}+\Delta\tau_{\rm em}}^{t_0+\Delta\tau_{\rm obs}} \frac{dt}{a(t)}. \tag{39}$$

For small intervals this gives

$$\frac{\Delta\tau_{\rm obs}}{a(t_0)} = \frac{\Delta\tau_{\rm em}}{a(t_{\rm em})}, \qquad \Delta\tau_{\rm obs} = (1+z)\Delta\tau_{\rm em}. \tag{40}$$

For a photon packet the product

$$\Delta E_{\rm obs}\Delta\tau_{\rm obs} = \Delta E_{\rm em}\Delta\tau_{\rm em} \tag{41}$$

is action-cell bookkeeping: the null channel conservatively delivers the packet's action cell, so the CCE-mandated widening of $\Delta\tau_{\rm obs}$ forces the conjugate energy scale to fall by the same factor. It is not a separate quantum-uncertainty derivation of light-curve dilation. If $L_{\rm em}(\tau;\Theta)$ is a rest-frame light-curve template with source parameters Θ, the received bolometric flux has the schematic form

$$F_{\rm obs}(t;z,\Theta) = \frac{1}{4\pi D_{\rm L}^2(z)} L_{\rm em}\left(\frac{t-t_*}{1+z};\Theta\right), \tag{42}$$

where $D_{\rm L}$ is the luminosity distance and t_* the observed epoch of a chosen phase marker. Any intrinsic template width $w_{\rm em}$ is observed as

$$w_{\rm obs} = (1+z)w_{\rm em}. \tag{43}$$

Supernova time dilation is therefore not an added optical effect. It is the standard coordinate stretch of the source world-tube under cosmological redshift, read by CCE as conservation-congruent bookkeeping for the transported record [12].

4.2 Adiabatic Metric Exhaust: The CMB Blackbody

The cosmic microwave background gives the same point distributionally. A blackbody spectrum remains blackbody under cosmological redshift because metric exhaust is continuous and phase-space preserving, not scattering off a material medium. For a Planck distribution at emission,

$$f_{\rm em}(\nu_{\rm em}) = \frac{1}{\exp(h\nu_{\rm em}/k_{\rm B}T_{\rm em})-1}. \tag{44}$$

Liouville transport gives $f_{\rm obs}(\nu_{\rm obs}) = f_{\rm em}((1+z)\nu_{\rm obs})$, hence

$$f_{\rm obs}(\nu_{\rm obs}) = \frac{1}{\exp\left[h\nu_{\rm obs}/k_{\rm B}(T_{\rm em}/(1+z))\right]-1}. \tag{45}$$

The observed spectrum is again Planckian, with

$$T_{\rm obs} = \frac{T_{\rm em}}{1+z}. \tag{46}$$

The same result appears in intensity form. Since I_ν/ν^3 is invariant,

$$I_\nu^{\rm obs}(\nu_{\rm obs}) = (1+z)^{-3} I_\nu^{\rm em}((1+z)\nu_{\rm obs}). \tag{47}$$

Substituting the Planck intensity gives $I_\nu^{\rm obs} = B_\nu(T_{\rm em}/(1+z))$. The CMB temperature law is the thermal signature of adiabatic metric-sector exhaust: phase-space cell occupancy is preserved while the metric-sector energy scale relaxes [19].

4.3 Retrodictive Source-Assignment Load: Tolman Surface Brightness

The Tolman surface-brightness relation shows the cost of assigning a faint extended source to a resolved spatial history. Let a small proper source element dA_{em} have emitted bolometric surface brightness

$$S_{\mathrm{em}} = \frac{dL_{\mathrm{em}}}{4\pi\, dA_{\mathrm{em}}}. \tag{48}$$

The observed flux from that element and its observed angular area are

$$dF_{\mathrm{obs}} = \frac{dL_{\mathrm{em}}}{4\pi D_{\mathrm{L}}^2}, \qquad d\Omega_{\mathrm{obs}} = \frac{dA_{\mathrm{em}}}{D_{\mathrm{A}}^2}. \tag{49}$$

Using Etherington reciprocity, $D_{\mathrm{L}} = (1+z)^2 D_{\mathrm{A}}$, the observed surface brightness is

$$S_{\mathrm{obs}} = \frac{dF_{\mathrm{obs}}}{d\Omega_{\mathrm{obs}}} = S_{\mathrm{em}}\frac{D_{\mathrm{A}}^2}{D_{\mathrm{L}}^2} = S_{\mathrm{em}}(1+z)^{-4}. \tag{50}$$

Equivalently, the logarithmic attenuation is

$$-\ln\frac{S_{\mathrm{obs}}}{S_{\mathrm{em}}} = 4\ln(1+z). \tag{51}$$

For the CCE ledger, the split

$$4\ln(1+z) = \underbrace{\ln(1+z)}_{\text{energy exhaust}} + \underbrace{3\ln(1+z)}_{\text{source-assignment load}} \tag{52}$$

is a bookkeeping convention tied to the standard radiative-transfer identity, not an independent derivation of Tolman dimming. Since I_ν/ν^3 is invariant,

$$I_\nu^{\mathrm{obs}}(\nu_{\mathrm{obs}}) = (1+z)^{-3} I_\nu^{\mathrm{em}}((1+z)\nu_{\mathrm{obs}}). \tag{53}$$

CCE assigns the redshift-only minimum of this spectral-intensity attenuation to a source-assignment branch load by defining

$$e^{-\mathcal{B}_\Pi(z)} := (1+z)^{-3}, \qquad \mathcal{B}_\Pi(z) := 3\ln(1+z), \qquad N_{\mathrm{br}}(z) = e^{\mathcal{B}_\Pi(z)} = (1+z)^3. \tag{54}$$

Equivalently, for a uniform emitted cell nested in a uniform retrodictive support, if a particular source-assignment model chooses X_{ret} with volume $(1+z)^3$ times the emitted source cell X_{em}, then

$$\begin{aligned} D_{\mathrm{KL}}(P_{\mathrm{cell}} \,\|\, P_{\mathrm{ret}}) &= \ln\frac{d^3x_{\mathrm{ret}}}{d^3x_{\mathrm{em}}} \\ &= 3\ln(1+z) = \mathcal{B}_\Pi(z). \end{aligned} \tag{55}$$

This KL expression should be read as the explicit source-support convention for that CCE model. It is not a claim that a two-dimensional surface-brightness observable alone forces a unique three-dimensional spatial support volume. Integrating over observed frequency adds the remaining factor $d\nu_{\mathrm{obs}} = (1+z)^{-1} d\nu_{\mathrm{em}}$, so

$$S_{\mathrm{obs}} = e^{-\mathcal{B}_\Pi(z)}(1+z)^{-1} S_{\mathrm{em}} = S_{\mathrm{em}}(1+z)^{-4}. \tag{56}$$

Thus

$$S_{\mathrm{obs}} = S_{\mathrm{em}} \exp\left[-\ln(1+z) - \mathcal{B}_\Pi(z)\right]. \tag{57}$$

In this CCE reading, the surface is not merely dimmer; it is harder to hold as a resolved object. If $S_{\mathrm{noise},\Pi}(z)$ is the surface-brightness floor of the observer's spatial projection, the Tolman-limited horizon is

$$z_{\mathrm{Tol},\Pi} = \inf\{z > 0 : S_{\mathrm{em}} e^{-\ln(1+z) - \mathcal{B}_\Pi(z)} \le S_{\mathrm{noise},\Pi}(z) \text{ or } \mathcal{B}_\Pi(z) > \mathcal{I}_\Pi^{\mathrm{cap}}\}. \tag{58}$$

At sufficient redshift, the formal source remains in the equations, but the model-assigned branch load can exceed the observer's spatial macroscopic-distinction capacity [27, 10].

5 Gravitational Horizons as Epistemic Bubble Boundaries

The spatial examples replace an unbounded geometric reading with a finite port ledger. Gravitational horizons apply the same move to congruence geometry. Standard cosmology allows the cross-sectional area $A(\lambda)$ to approach two formal extremes: $A \to 0$, associated with singular focusing, and $A \to \infty$, associated with de Sitter-like dilution and heat death. Both smuggle in a Platonic observer. The first assumes costless resolution as label density grows without bound; the second truncates metric exhaust by assuming infinite sensitivity as signal density is diluted.

CCE replaces these limits with a finite geometric window. A projection can follow a congruence only while its cross-section lies between a lower capacity bound and an upper sensitivity bound,

$$A_\Pi^{\min} \le A(\lambda) \le A_\Pi^{\max}. \tag{59}$$

In this reading, singularity and heat death are not opposite exceptions. They are the inward and outward failures of the same encoded congruence: Retrodictive Divergence at the capacity floor and Predictive Dissipation at the sensitivity floor.

5.1 The Convergence Bound: Retrodictive Capacity

The singularity theorems sharpen the issue by using smooth congruences to turn local focusing into global geodesic incompleteness. For a null congruence with affine parameter λ, tangent k^μ, cross-sectional area A, expansion

$$\theta = \frac{1}{A}\frac{dA}{d\lambda}, \tag{60}$$

shear $\sigma_{\mu\nu}$, and twist $\omega_{\mu\nu}$, the Raychaudhuri equation is

$$\frac{d\theta}{d\lambda} = -\frac{1}{2}\theta^2 - \sigma_{\mu\nu}\sigma^{\mu\nu} + \omega_{\mu\nu}\omega^{\mu\nu} - R_{\mu\nu}k^\mu k^\nu. \tag{61}$$

For hypersurface-orthogonal congruences, $\omega_{\mu\nu} = 0$. If the null convergence condition $R_{\mu\nu}k^\mu k^\nu \ge 0$ holds, and since $\sigma_{\mu\nu}\sigma^{\mu\nu} \ge 0$, one obtains

$$\frac{d\theta}{d\lambda} \le -\frac{1}{2}\theta^2. \tag{62}$$

Thus an initially converging bundle, $\theta(\lambda_0) < 0$, focuses in finite affine parameter, no later than $\Delta\lambda = 2/|\theta(\lambda_0)|$ in the comparison solution. Equivalently, the formal area variable is driven toward $A \to 0$. The timelike version replaces $1/2$ by $1/3$ and affine parameter by proper time; the operational issue is the same. Penrose, Hawking, and successors use this focusing structure to show that, under stated global assumptions, spacetime geodesics cannot be extended indefinitely [26, 24, 14, 20].

CCE does not refute that result. It changes its operational interpretation. The Raychaudhuri equation evolves a mathematical congruence whose neighboring geodesics remain distinguishable down to arbitrarily small area. A physical observer must encode the labels that distinguish them, while semiclassical gravity bounds the entropy that can pass through a light-sheet. Here the h^3 phase-space ledger changes regimes: when bulk phase-space cells approach gravitational collapse, thermodynamic volume capacity gives way to holographic area capacity $A/(4\ell_{\rm P}^2)$. For a light-sheet L generated from area A, Bousso's covariant entropy bound gives

$$\frac{S[L]}{k_{\rm B}} \leq \frac{A}{4\ell_{\rm P}^2}. \tag{63}$$

CCE uses this as an admissibility postulate for projections whose branch labels must be physically encoded in degrees of freedom crossing the light-sheet: those labels cannot require more entropy capacity than the sheet can support. If an external apparatus supplies independent records, its memory and calibration costs must be charged separately against $\mathcal{I}_{\Pi}^{\rm cap}$ and cannot be hidden inside the light-sheet area. For labels carried on the sheet, the available capacity in nats is then

$$C_A(\lambda) = \frac{A(\lambda)}{4\ell_{\rm P}^2}. \tag{64}$$

For a projection whose branch-label load is $\mathcal{I}_{\Pi}^{\rm req}$, admissibility demands

$$C_A(\lambda) \geq \mathcal{I}_{\Pi}^{\rm req}, \qquad A(\lambda) \geq A_{\Pi}^{\rm req}, \qquad A_{\Pi}^{\rm req} := 4\ell_{\rm P}^2 \mathcal{I}_{\Pi}^{\rm req}. \tag{65}$$

The active lower CCE area for this projection is therefore

$$A_{\Pi}^{\rm min} \equiv A_{\Pi}^{\rm req}. \tag{66}$$

This distinguishes the load demanded by the chosen projection, $\mathcal{I}_{\Pi}^{\rm req}$, from the maximum physical capacity of the apparatus that implements it. Admissibility requires $\mathcal{I}_{\Pi}^{\rm req} \leq \mathcal{I}_{\Pi}^{\rm cap}$, while the saturated apparatus envelope is

$$A_{\Pi}^{\rm cap} := 4\ell_{\rm P}^2 \mathcal{I}_{\Pi}^{\rm cap}. \tag{67}$$

A larger apparatus raises the possible envelope $A_{\Pi}^{\rm cap}$ for more demanding projections; it does not force every smaller-load projection to stop at that larger area. For an observer apparatus $\mathcal{O}_{\Pi}$ with gravitating energy E_{Π}, circumscribing radius R_{Π}, and enclosing area $A_{\Pi}^{\rm app}$, the stored distinction budget is objectively bounded by

$$\mathcal{I}_{\Pi}^{\rm req} \leq \mathcal{I}_{\Pi}^{\rm cap} \leq \mathcal{I}_{\Pi}^{\rm phys} := \min\left\{ \frac{2\pi E_{\Pi} R_{\Pi}}{\hbar c}, \frac{A_{\Pi}^{\rm app}}{4\ell_{\rm P}^2} \right\}. \tag{68}$$

The first term is Bekenstein's entropy bound for a finite system, expressed in nats, while the second is the corresponding holographic area bound for a strongly gravitating enclosure. If only a finite proper time $\Delta\tau$ is available for the retrodictive computation, the usable ledger is further constrained by the Margolus–Levitin/Lloyd operation budget,

$$N_{\rm ops}(\Delta\tau) \leq \frac{2E_{\Pi}\Delta\tau}{\pi\hbar}. \tag{69}$$

Thus different instruments can have different operational cutoffs, but those cutoffs are apparatus-indexed physical facts fixed by energy, size, area, runtime, and error tolerance, not arbitrary observer preferences [3, 18, 17].

An absolute envelope is obtained by taking the apparatus to be the accessible universe. Lloyd estimates that the universe can have performed of order 10^{120} elementary operations on roughly 10^{90} ordinary bits, with a gravitational/holographic ceiling of order 10^{120} bits. In nat units this gives the generous universal bound

$$\mathcal{I}_{\Pi}^{\text{cap}} \leq \mathcal{I}_{U}^{\max} \sim 10^{120} \ln 2 \simeq 7 \times 10^{119} \text{ nats}. \tag{70}$$

Consequently no physically embedded observer can push the CCE capacity parameter beyond a universal finite ceiling; equivalently, any projection-required area floor obeys

$$A_{\Pi}^{\text{req}} \leq A_{\text{cap}}^{U} = 4\ell_{\text{P}}^{2}\mathcal{I}_{U}^{\max} \simeq 7 \times 10^{50}\,\text{m}^{2}, \qquad \sqrt{A_{\text{cap}}^{U}} \simeq 3 \times 10^{25}\,\text{m}. \tag{71}$$

This enormous number should not be read as the floor for every observer. It is the objective upper envelope on all possible apparatus-level capacity cuts. For ordinary projections one must use the required load of the chosen projection and the much smaller Bekenstein or holographic capacity of the actual apparatus.

The scale matters. Taking $\ell_{\text{P}} = 1.616 \times 10^{-35}\,\text{m}$ and an illustrative projection-level macroscopic-distinction load $\mathcal{I}_{\Pi}^{\text{req}} \sim 10^{30}$ nats gives

$$A_{\Pi}^{\min} = A_{\Pi}^{\text{req}} \simeq 4(1.616 \times 10^{-35}\,\text{m})^{2}10^{30} \simeq 1.0 \times 10^{-39}\,\text{m}^{2}, \qquad \sqrt{A_{\Pi}^{\min}} \simeq 3.2 \times 10^{-20}\,\text{m}. \tag{72}$$

This is far below ordinary macroscopic resolution and even below a nuclear length scale, though still about 2×10^{15} Planck lengths across. The CCE floor should therefore not be oversold as a laboratory-scale cutoff or replacement for quantum gravity. For this projection it occurs before $A \to 0$, but perhaps after the classical fluid, optical, or matter model defining the congruence has already failed. Conversely, raising the floor to $1\,\text{m}^{2}$ would require $\mathcal{I}_{\Pi}^{\text{req}} \sim (4\ell_{\text{P}}^{2})^{-1} \simeq 10^{69}$ nats and an apparatus with $\mathcal{I}_{\Pi}^{\text{cap}}$ at least that large. CCE becomes predictive only after Π, its physical implementation, and its error scale are fixed; otherwise $A_{\Pi}^{\min}$ is a projection-required contour inside an objective apparatus-capacity envelope, not a freely adjustable constant.

The area floor is therefore an apparatus-indexed but projection-specific CCE cutoff: the coordinate where maximum covariant entropy allowed by semiclassical general relativity falls below the minimum entropy the observer's chosen encoding is required to instantiate, subject to the apparatus capacity bounds above [6, 2, 1, 13].

Choose λ to increase toward the focal endpoint of the converging congruence and assume the first crossing is taken along this oriented path. The CCE stopping parameter is

$$\lambda_{\Pi}^{-} = \inf\{\lambda > \lambda_{0} : A(\lambda) \leq A_{\Pi}^{\min}\}. \tag{73}$$

When $\lambda \geq \lambda_{\Pi}^{-}$ along that oriented path, the observer cannot maintain the congruence's macroscopic-distinction load. Neighboring geodesics may remain formal curves, but they are no longer separately encodable states for that observer. The CCE cutoff occurs before $A \to 0$: convergence becomes operationally under-resolved once the area ledger cannot carry the labels required by Π.

In this sense, CCE places an operational cutoff on the singular limit. It does not alter the Raychaudhuri equation or truncate the singularity theorems. It changes only the license to treat the limiting projection as physically available information. The singularity is the lower boundary at which costless resolution fails and Retrodictive Divergence saturates the available area ledger.

5.2 The Expansion Bound: Predictive Dissipation

The same congruence variable also captures the opposite endpoint. Choose λ to increase along the expanding bundle. If $\theta > 0$, then

$$\frac{dA}{d\lambda} = \theta A > 0, \tag{74}$$

and the formal cross-section grows along that oriented path. In an asymptotic de Sitter regime, the classical description lets the world-tube expand toward $A \to \infty$. This is the geometric image of heat death: a fixed macroscopic distinction spreads over a growing cross-section until local contrast disappears.

CCE treats this as sensitivity failure, not as an infinite empty future. To make the bound dimensional, let E_{signal} be the free-energy-equivalent signal load carried by the tracked distinction, measured in joules. Equivalently, one may divide by $k_{\text{B}}T$ and work entirely in nats at a specified bath temperature. Its operational surface density is

$$\rho_{\text{signal}}(\lambda) = \frac{E_{\text{signal}}}{A(\lambda)}, \tag{75}$$

with units of joules per unit area in the free-energy convention. Let ρ_{noise} be the corresponding metric-sector noise floor, also in joules per unit area, including detector integration, foreground confusion, and the cost of holding the world-tube as a protected distinction. The signal-density requirement is

$$\rho_{\text{signal}}(\lambda) > \rho_{\text{noise}} \tag{76}$$

for the source to remain a macroscopic distinction of Π. The maximum admissible cross-section in this calibrated ledger is

$$A_{\Pi}^{\max} = \frac{E_{\text{signal}}}{\rho_{\text{noise}}}. \tag{77}$$

The heat-death stopping parameter is the first hitting time along the expanding path,

$$\lambda_{\Pi}^{+} = \inf\{\lambda > \lambda_0 : A(\lambda) \geq A_{\Pi}^{\max}\}. \tag{78}$$

For $\lambda \geq \lambda_{\Pi}^{+}$ along that path, the formal geometry may keep expanding, but the protected world-tube has dissolved into unresolved degrees of freedom. The endpoint is not an observable infinite void; it is where the congruence falls below the observer's signal-to-noise threshold.

Together, the convergence and divergence bounds form a calibrated geometric window,

$$A_{\Pi}^{\min} \leq A(\lambda) \leq A_{\Pi}^{\max}. \tag{79}$$

Relativistic geometry remains formally useful inside this interval. Outside it, the limiting congruence requires impossible label capacity at small area or impossible sensitivity at large area, relative to the chosen projection and apparatus.

5.3 Trapped Surfaces as Capacity Surfaces

This area-capacity reframing applies directly to Penrose trapped surfaces. A closed spacelike two-surface $\mathcal{S}$ is trapped when both future-directed null congruences orthogonal to it have negative expansion,

$$\theta_{(k)} < 0, \qquad \theta_{(\ell)} < 0. \tag{80}$$

Classically, both outgoing and ingoing light fronts decrease in area. Combined with the convergence condition and global assumptions, this yields geodesic incompleteness [24, 14].

In CCE language, a trapped surface is also a two-sided compression device. The two future light-sheets generated by the null normals k^μ and ℓ^μ both have non-positive expansion, so the covariant entropy bound applies to each:

$$\frac{S[L_k]}{k_{\rm B}} \le \frac{A(\mathcal{S})}{4\ell_{\rm P}^2}, \qquad \frac{S[L_\ell]}{k_{\rm B}} \le \frac{A(\mathcal{S})}{4\ell_{\rm P}^2}. \tag{81}$$

For branch labels carried on the two light-sheets, CCE assigns separate capacities,

$$C_k(\mathcal{S}) = \frac{A(\mathcal{S})}{4\ell_{\rm P}^2}, \qquad C_\ell(\mathcal{S}) = \frac{A(\mathcal{S})}{4\ell_{\rm P}^2}. \tag{82}$$

Along the generators these obey

$$\frac{dC_k}{d\lambda_k} = \theta_{(k)} C_k, \qquad \frac{dC_\ell}{d\lambda_\ell} = \theta_{(\ell)} C_\ell. \tag{83}$$

Because both expansions are negative, either light-sheet can be driven toward the projection-required load,

$$C_k(\mathcal{S}_{\lambda_k}) < \mathcal{I}_\Pi^{\rm req} \quad \text{or} \quad C_\ell(\mathcal{S}_{\lambda_\ell}) < \mathcal{I}_\Pi^{\rm req}. \tag{84}$$

The corresponding CCE capacity cut for the chosen projection is

$$\mathcal{T}_\Pi = \{\mathcal{S} : A(\mathcal{S}) = 4\ell_{\rm P}^2 \mathcal{I}_\Pi^{\rm req}\} = \{\mathcal{S} : A(\mathcal{S}) = A_\Pi^{\rm min}\}. \tag{85}$$

At this cut, a formal null generator need not cease to exist. Rather, any continuation requiring separately tracked branches demands more covariant entropy capacity than the projection can instantiate.

For black holes, this does not replace the classical event horizon as a global causal boundary. It adds an operational layer. To an exterior finite observer, the horizon is the interface where attempted refinements of interior histories are not returned as recoverable macroscopic distinctions but represented on a finite area ledger. To an infalling observer, the relevant CCE boundary is the later capacity cut where that observer's projection can no longer carry its distinctions. The event horizon is therefore supplemented by observer-indexed capacity horizons.

6 Conclusion

The foundations of relativistic cosmology typically treat the extremes of the universe as absolute geometric or temporal facts. This note has argued that several such limits also have observer-indexed operational boundaries dictated by the physical thermodynamics of projection. Heat death, the Big Bang singularity, the fading of highly redshifted observables, and black-hole trapped surfaces remain distinct standard phenomena, but within the CCE ledger they can be read as directional stress tests of the same finite projection operator Π.

When a macroscopic equation of state is stripped of the Platonic Observer Fallacy and forced to pay its thermodynamic ledger, the infinite limits of continuous geometry are no longer automatically licensed as physically available records. The failures manifest in two distinct modes. In the first, *Predictive Dissipation*, forward in time and outward in space, the fallacy is the truncation of metric exhaust: the projection ignores the continuous thermal cost of holding a macrostate distinct from the bath. Predictive contrast irreversibly dissipates into unresolved degrees of freedom ($A \geq A_{\Pi}^{\max}$), and transported spatial world-tubes lose Fisher distinguishability. In the second, *Retrodictive Divergence*, backward in time and inward under gravitational focusing, the fallacy is the assumption of costless resolution: exact retrodiction demands the Landauer-priced re-instantiation of erased branches. The required thermodynamic import diverges in the formal limit, and the label density of converging congruences hits the covariant area-capacity floor ($A \leq A_{\Pi}^{\min}$).

This framework does not alter the differential equations of general relativity; it intercepts their operational interpretation. The central principle is that no observer can use a continuous mathematical projection whose informational obligations exceed its physically instantiated thermodynamic ledger.

Operationally, the observer should not treat an infinite spacetime manifold as available at arbitrary precision. They work inside a unified, thermodynamically bounded epistemic bubble. The underlying physical universe may vastly exceed any single encoding, but any observer-indexed physics occurs within the finite thermodynamic limits of what that observer can physically afford to resolve.

References

[1] James M. Bardeen, Brandon Carter, and Stephen W. Hawking. The four laws of black hole mechanics. *Communications in Mathematical Physics*, 31(2):161–170, 1973.

[2] Jacob D. Bekenstein. Black holes and entropy. *Physical Review D*, 7(8):2333–2346, 1973.

[3] Jacob D. Bekenstein. Universal upper bound on the entropy-to-energy ratio for bounded systems. *Physical Review D*, 23(2):287–298, 1981.

[4] Charles H. Bennett. Logical reversibility of computation. *IBM Journal of Research and Development*, 17(6):525–532, 1973.

[5] Antoine Bérut, Artak Arakelyan, Artyom Petrosyan, Sergio Ciliberto, Raoul Dillenschneider, and Eric Lutz. Experimental verification of Landauer's principle linking information and thermodynamics. *Nature*, 483(7388):187–189, 2012.

[6] Raphael Bousso. A covariant entropy conjecture. *Journal of High Energy Physics*, 1999(07):004, 1999.

[7] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley-Interscience, Hoboken, 2 edition, 2006.

[8] Harald Cramér. *Mathematical Methods of Statistics*. Princeton University Press, Princeton, 1946.

[9] Gavin E. Crooks. Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences. *Physical Review E*, 60(3):2721–2726, 1999.

[10] I. M. H. Etherington. On the definition of distance in general relativity. *Philosophical Magazine*, 15(100):761–773, 1933.

[11] Peter David Fagan. Towards a physical theory of intelligence. Companion manuscript distributed with the draft materials, 2026. Background source for the broader CCE framework referenced in this note.

[12] Gerson Goldhaber et al. Timescale stretch parameterization of type Ia supernova B-band light curves. *The Astrophysical Journal*, 558:359–368, 2001.

[13] Stephen W. Hawking. Particle creation by black holes. *Communications in Mathematical Physics*, 43(3):199–220, 1975.

[14] Stephen W. Hawking and Roger Penrose. The singularities of gravitational collapse and cosmology. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 314(1519):529–548, 1970.

[15] Yonggun Jun, Momčilo Gavrilov, and John Bechhoefer. High-precision test of Landauer's principle in a feedback trap. *Physical Review Letters*, 113(19):190601, 2014.

[16] Rolf Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3):183–191, 1961.

[17] Seth Lloyd. Computational capacity of the universe. *Physical Review Letters*, 88(23):237901, 2002.

[18] Norman Margolus and Lev B. Levitin. The maximum speed of dynamical evolution. *Physica D: Nonlinear Phenomena*, 120(1–2):188–195, 1998.

[19] John C. Mather et al. Measurement of the cosmic microwave background spectrum by the COBE FIRAS instrument. *The Astrophysical Journal*, 420:439–444, 1994.

[20] Charles W. Misner, Kip S. Thorne, and John Archibald Wheeler. *Gravitation*. W. H. Freeman, San Francisco, 1973.

[21] Philip J. Morrison. Bracket formulation for irreversible classical fields. *Physics Letters A*, 100(8):423–427, 1984.

[22] Philip J. Morrison. A paradigm for joined hamiltonian and dissipative systems. *Physica D: Nonlinear Phenomena*, 18(1-3):410–419, 1986.

[23] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge, 2000.

[24] Roger Penrose. Gravitational collapse and space-time singularities. *Physical Review Letters*, 14(3):57–59, 1965.

[25] C. R. Rao. Information and the accuracy attainable in the estimation of statistical parameters. *Bulletin of the Calcutta Mathematical Society*, 37:81–91, 1945.

[26] Amal Kumar Raychaudhuri. Relativistic cosmology. i. *Physical Review*, 98(4):1123–1126, 1955.

[27] Richard C. Tolman. On the estimation of distances in a curved universe with a non-static line element. *Proceedings of the National Academy of Sciences*, 16(7):511–520, 1930.

[28] Harry L. Van Trees. *Detection, Estimation, and Modulation Theory, Part I.* John Wiley & Sons, New York, 1968.