Cosmic Acceleration from Finite Record Geometry

Dark Energy as an Infrared Closure of the Causal Ledger

Keith Taylor VERSF Theoretical Physics Program

Abstract

We develop a record-theoretic derivation of late-time cosmic acceleration within the Void Energy-Regulated Space Framework (VERSF). Three foundational axioms — Finite Distinguishability (FD), Irreversible Commitment (IC), and Law Invariance (LI) — constrain macroscopic physical laws to be defined only on the stable committed-record algebra. When cosmic expansion dilutes matter below the threshold required to sustain independent large-scale committed correlation channels, an isotropic infrared closure term necessarily emerges from the record algebra structure. The coherence scale ξ is identified with the comoving BAO sound-horizon ruler (conventionally evaluated at the baryon drag epoch ($z^d \sim 10^3$; Eisenstein & Hu 1998, Ref. 8)), giving an asymptotic cosmological term $\Lambda\infty = 3/\xi^2$. We derive the closure functional $F(\chi) = 1/(1 + \chi^n)$, where the sharpness exponent n is observationally constrained by low-redshift distance and growth data; a proper fit to H(z) and BAO distances is required to bound n, and is deferred to a companion analysis. Linear structure growth is unaffected at leading order, and for $n \geq 2$ the closure density is strongly suppressed at early times, preserving standard recombination and CMB physics. The model makes discriminating predictions for next-generation w(z) surveys. A baseline stress test (Model 0) produces large mid-z expansion deviations at the Planck-ΛCDM parameter point, motivating the physically intended channel-consistent implementation (Model 1) based on the collapsed fraction fcoll(z).

For the General Reader

Why is the universe speeding up? Observations from the past three decades have firmly established that galaxies are flying apart from one another at an accelerating rate. Something is pushing space itself to expand faster and faster. Physicists call this mysterious driver *dark energy*, and in the simplest model it is represented by a constant term in Einstein's equations — the cosmological constant Λ. But this explanation raises an immediate puzzle: if you try to calculate what Λ should be from known physics, the answer comes out roughly 10^{122} times too large. This is the worst quantitative failure in the history of theoretical physics.

This paper proposes a different origin for dark energy — one grounded not in the energy of empty space but in the *structure of physical information itself*.

The central idea. In everyday life we take for granted that events leave records: a collision happened, a particle was measured, a structure formed. This paper takes seriously the idea that only events which have left a permanent, irreversible record are truly real in the physical sense. The universe, on this view, is a kind of cosmic ledger: a running account of all the committed distinctions that have ever been made.

As the universe expands and matter dilutes, that ledger becomes increasingly *underfilled*. There are fewer and fewer gravitationally-bound structures — galaxy clusters, filaments, halos — to sustain the committed correlations that populate the ledger. When the ledger becomes sufficiently sparse, the framework shows that an isotropic energy term must emerge geometrically, from the gap between what the ledger can hold and what matter actually provides.

That emergent energy term *is* dark energy.

Why does it kick in now? The characteristic scale of this effect is set not by any fine-tuned constant but by the BAO sound-horizon ruler imprinted in the early universe, approximately 150 Mpc ≈ 490 million light-years. This is the same scale that produces the baryon acoustic oscillation (BAO) signal seen in galaxy surveys. The late-time acceleration epoch begins when the matter density falls below the energy density associated with this scale — which happens at approximately the current epoch, for reasons rooted in early-universe physics.

What does the paper prove? Starting from three precisely stated axioms about the nature of physical distinguishability and records, the paper derives:

- A dimensionless ratio $\chi(a)$ that tracks how well matter fills the cosmic ledger at any epoch.
- A closure functional $F(\chi)$ — a mathematical expression for how the emergent dark energy grows as matter dilutes.
- An effective cosmological term $\Lambda_{eff}(a)$ that vanishes at early times and asymptotes to a constant set by the BAO scale.
- Predictions for how dark energy's equation of state $w(z)$ should behave across cosmic time, testable by current and forthcoming galaxy surveys (DESI, Euclid, Vera Rubin Observatory).

The model does not eliminate the cosmological constant problem — no single paper can — but it reframes it: instead of asking why the vacuum has a particular energy, we ask why the cosmic ledger has a particular coherence scale. That scale is determined by early-universe acoustics, not by quantum field theory, which is why the answer comes out 120 orders of magnitude closer to observation.

Table of Contents

---

1. Foundational Axioms

The VERSF framework rests on three axioms governing the structure of distinguishable physical states. We state each precisely and derive their consequences before building the dynamical model.

Axiom 1 — Finite Distinguishability (FD)

Any two physical states s_1, s_2 are operationally distinguishable only if their difference can be encoded in a finite sequence of committed binary distinctions.

FD asserts that the state space $\mathcal{S}$ carries a discrete information metric at the Planck scale ℓ_P. Differences $\Delta s < \ell_P$ carry no physical content. The distinguishability structure of $\mathcal{S}$ is therefore granular, not continuous.

Consequence: Only states separated by at least one committed binary distinction belong to distinct equivalence classes in the causal ledger. The continuum limit of field theory is an approximation valid far above ℓ_P; at the ledger level, physics is discrete.

Axiom 2 — Irreversible Commitment (IC)

A physical distinction is macroscopically real only if it has been irreversibly committed — i.e., it corresponds to an entropy-increasing transition that cannot be undone by any sequence of unitary operations within the causal domain.

IC distinguishes quantum superpositions (uncommitted distinctions) from classical records (committed distinctions). Let $\mathcal{A}_c \subset \mathcal{S}$ denote the **committed record algebra**: the sub-algebra of states that have undergone decoherence and irreversible entropy generation. Macroscopic observables are defined only on $\mathcal{A}_c$.

Consequence: A degree of freedom that exists only in quantum superposition contributes no independent channel to the causal ledger. It cannot serve as an independent macroscopic state variable. In the cosmological context, only gravitationally bound, decohered matter overdensities constitute committed records.

Axiom 3 — Law Invariance (LI)

The laws of macroscopic physics are invariant under permutations of committed records that preserve the causal structure of the ledger.

LI is the record-theoretic analog of general covariance. It requires that no particular committed state carries preferred status; only the relational structure of the ledger enters the laws of physics. This rules out explicit breaking of the record permutation symmetry.

Consequence: Any emergent term in the field equations must be isotropic and expressible as a scalar functional of the ledger occupancy ratio χ defined in Section 2. Anisotropic closure terms are forbidden by LI. The closure functional F may depend only on the dimensionless scalar χ, not on position, direction, or the particular identity of any committed record.

---

2. Correlation Channels and the Occupancy Ratio

2.1 What Is a Correlation Channel?

A **cosmological correlation channel** is a causally connected matter overdensity whose internal correlations have undergone irreversible commitment (IC) and are stable against Hubble dilution over at least one Hubble time. Operationally:

- The overdensity occupies a comoving region with density contrast $\delta > \delta\text{collapse} \approx 1.686$ (the linear collapse threshold).
- Its correlations are stabilised by gravitational binding.
- It survives for at least $\Delta t \sim H^{-1}$.

This definition ties correlation capacity directly to structure formation. An uncommitted fluctuation — a linear density perturbation that has not yet collapsed — does not constitute a channel. Only virialised, decohered structures (halos, filaments, collapsed overdensities) count.

2.2 Matter-Supported Correlation Capacity

The **matter-supported correlation capacity** $\mathcal{C}_m(a)$ counts independent committed correlation channels. By Definition 2.1, these channels are not generic linear density fluctuations — they are gravitationally bound, decohered overdensities that persist for $\Delta t \sim H^{-1}$. Accordingly, $\mathcal{C}_m(a)$ should scale with the *collapsed and committed* fraction of matter rather than with ρ_m alone.

We therefore write, at leading order:

$\boldsymbol{\mathcal{C}_m(a) \propto \mathrm{fcoll}(a) \cdot \rho_m(a)}$ (1)

where $\mathrm{fcoll}(a) \in [0, 1]$ is the fraction of matter in committed, gravitationally bound structures. A standard proxy is the collapsed fraction above a mass threshold Mmin:

$\mathrm{fcoll}(z; \mathrm{Mmin}) = (1/\rho_m) \cdot \int[\mathrm{Mmin} \to \infty]\ M \cdot (dn/dM)(M,z)\ dM$

where dn/dM is the halo mass function (Press–Schechter or Sheth–Tormen). This provides a concrete route to compute $\chi(z)$ without introducing new free physics beyond standard structure formation. In this formulation, dilution ($\rho_m \propto a^{-3}$) is not the sole driver; the growth of collapsed structure partially compensates dilution at low redshift, which can stabilise the effective closure term and reduce tension with distance data.

Scope note: This paper addresses dark energy and late-time expansion; it does not attempt to replace dark matter. Standard halo statistics (halo mass functions, collapsed fractions) are used as empirical proxies for committed structure, assuming dark matter halos as the reservoir of gravitationally bound records.

In the minimal analytic treatment used elsewhere in this paper we adopt the simplified **baseline scaling** $\mathcal{C}_m \propto a^{-3}$, which corresponds to fcoll = constant. This is conservative: it underestimates committed channel capacity at late times and therefore gives an upper bound on how far the closure term departs from $\Lambda\infty$. Section 6.5 shows that even this baseline generates significant deviations in H(z), which motivates the channel-consistent extension. A realistic closure cosmology should be evaluated with:

$$\boldsymbol{\chi(a) \propto \mathrm{fcoll}(a) \cdot \rho_m(a) / \rho\Lambda,\infty} \quad (2)$$

which is the correct operational interpretation of "channel capacity" under Axiom IC. The fcoll extension is deferred to a companion numerical analysis.

2.3 The Asymptotic Closure Density

From Section 3, the infrared closure term asymptotes to:

$$\rho\Lambda,\infty = \Lambda\infty / (8\pi G) = 3 / (8\pi G\, \xi^2)$$

This provides a physical threshold: when $\rho_m(a)$ falls to the order of $\rho\Lambda,\infty$, the matter-supported channels become insufficient to fill the ledger, and the closure term activates.

2.4 The Ledger Occupancy Ratio

Define the dimensionless **ledger occupancy ratio** relative to the asymptotic closure density. In Model 0 this uses the full matter density; in Model 1 it is defined operationally from the outset as:

$$\boldsymbol{\chi(a) \equiv \mathrm{fcoll}(a) \cdot \rho_m(a) / \rho\Lambda,\infty} \quad (3)$$

so that only committed, gravitationally-bound structure contributes to channel capacity by definition, not by post-hoc replacement. Model 0 is the special case fcoll = constant. The shared notation is:

$$\boldsymbol{\chi(a) \equiv \rho_m(a) / \rho\Lambda,\infty,\ \rho\Lambda,\infty = \Lambda\infty/(8\pi G) = 3/(8\pi G\xi^2)} \quad (4)$$

Since $\rho_m(a) = \rho_{m,0} \cdot a^{-3}$, we write:

$$\boldsymbol{\chi(a) = \chi_\infty \cdot a^{-3}, \chi_\infty \equiv \rho_{m,0} / \rho_{\Lambda,\infty}} \quad (5)$$

The parameter χ_∞ is not independently observable — it is defined relative to $\rho_{\Lambda,\infty}$, which is a model quantity. The directly observable present-day density ratio is:

$$\boldsymbol{r_0 \equiv \rho_{m,0} / \rho_{DE,0} = \Omega_{m,0} / \Omega_{\Lambda,0} \approx 0.31/0.69 \approx 0.45} \quad (6)$$

These are connected through the closure model: at a = 1, the dark energy density is $\rho_{DE,0} = \rho_{\Lambda,\infty}/(1 + \chi_\infty^n)$ (derived in Section 3.3), so:

$$\boldsymbol{r_0 = \rho_{m,0} / \rho_{DE,0} = \chi_\infty \cdot (1 + \chi_\infty^n)} \quad (7)$$

Derivation: $\chi_\infty \equiv \rho_{m,0}/\rho_{\Lambda,\infty}$ and $\rho_{DE,0} = \rho_{\Lambda,\infty}/(1 + \chi_\infty^n)$, so $r_0 = \rho_{m,0}/\rho_{DE,0} = \chi_\infty \cdot \rho_{\Lambda,\infty} \cdot (1 + \chi_\infty^n)/\rho_{\Lambda,\infty} = \chi_\infty(1 + \chi_\infty^n)$. ■

This is the key consistency relation. It shows that r_0 and the sharpness exponent n *jointly* determine χ_∞, rather than fixing it independently. For n = 2 and $r_0 = 0.45$, solving numerically gives $\chi_\infty \approx 0.391$; for n = 5, $\chi_\infty \approx 0.443$. Both are close to r_0 but not equal — the difference matters for precise predictions.

The two physical regimes:

- $\boldsymbol{\chi \gg 1}$ — early, matter-dominated epoch: committed channels fill the ledger. The closure term is suppressed.
- $\boldsymbol{\chi \ll 1}$ — late, dilute epoch: matter channels are sparse; unsupported ledger capacity activates the closure term.

In the baseline analytic model (fcoll = constant), $\chi(a) = \chi_\infty \cdot a^{-3}$. In the channel-consistent extension (Section 2.2), $\chi(a) = \chi_\infty \cdot a^{-3} \cdot f_{coll}(a)/f_{coll}(1)$, which moderates the transition by accounting for ongoing structure formation. Both forms are tracked in the companion numerical analysis.

Model definitions used in this paper

Model 0 (baseline diagnostic): $f_{coll}(a) \equiv f_{coll}(1)$ = constant, hence $\chi(a) = \chi_\infty \cdot a^{-3}$. This model is retained as a conservative diagnostic because it underestimates channel capacity at late times and therefore upper-bounds departures from Λ_∞. Section 6.5 quantifies these departures.

Model 1 (channel-consistent): $\chi(a) = \chi_\infty \cdot a^{-3} \cdot f_{coll}(a)/f_{coll}(1)$, with $f_{coll}(a)$ computed from a standard halo mass function above a physically motivated threshold M_{min}. Model 1 is the operationally correct implementation of "committed correlation channels" under Axiom IC, and is the model to be confronted with BAO + SNe + CMB distance priors in the companion analysis.

In Model 1, $\rho_m(a)$ in the occupancy definition is replaced by an effective committed density fcoll(a)·$\rho_m(a)$, reflecting that only collapsed IC-stable structure contributes to channel capacity.

To avoid ambiguity, we denote $\chi_0(a) \equiv \rho_m/\rho\Lambda,\infty$ (Model 0) and $\chi_1(a) \equiv \mathrm{fcoll}\cdot\rho_m/\rho\Lambda,\infty$ (Model 1); when we write $\chi(a)$ without a subscript, the relevant model is specified by the surrounding text.

---

3. The Infrared Closure Term

3.1 Geometric Origin of Λ∞

The derivation of Λ∞ has two logically distinct parts.

Part 1 — Structural consequence of the axioms. By FRW symmetry (isotropy and homogeneity following from LI at cosmic scales) and FD (which excludes sub-Planckian degrees of freedom), any diffeomorphism-invariant term arising from a single IR length scale ξ with dimensions of $(\text{length})^{-2}$ must take the form:

$\mathbf{\Lambda\infty = 3/\xi^2}$ (8)

The factor of 3 is conventional. No quantum insertion — no $\hbar$, no zero-point energy — is required. This is dimensional analysis combined with symmetry constraints: given one IR scale, only one combination has the right dimensions. The axioms do not by themselves specify which length plays the role of ξ.

Part 2 — Physical identification of ξ. The identification $\xi = \kappa\, r_s$ is a *physical hypothesis*, not a theorem deducible from FD/IC/LI alone. It is motivated by IC stability: r_s is the largest comoving scale on which committed matter correlations first become dynamically stable after the drag epoch. This identification is testable — predicting Λ∞ within an O(1) factor — but it remains a hypothesis the framework accommodates rather than uniquely derives.

This is the central economy of the framework: Λ∞ does not come from summing vacuum fluctuations. It comes from the curvature radius of committed record structure at the IC stability scale.

3.2 The Closure Functional

By Law Invariance (LI), the full time-dependent closure term must be a scalar function of χ alone. By IC, the function measures unsupported ledger capacity — it must be large when matter is dilute ($\chi \ll 1$) and vanish when matter is dense ($\chi \gg 1$). By FD, it must be analytic and bounded on $\chi \in (0, \infty)$.

These conditions impose four constraints on F:

Condition	Physical origin	Formal statement
Monotonicity	Diluting matter increases unsupported capacity	$dF/d\chi < 0$
Matter limit	Dense matter; full ledger support; no closure	$F(\chi \to \infty) \to 0$
Saturation limit	Sparse matter; full closure activation	$F(\chi \to 0) \to 1$
No new scales	LI forbids scale-dependent structure in F	F depends only on χ

We adopt a **lowest-order rational functional form** because it is analytically minimal and algebraically tractable while satisfying monotonicity and asymptotic constraints. Other monotone choices (e.g. stretched exponentials) are possible and can be mapped onto an effective sharpness parameter; the present work focuses on the rational family as the minimal representative class.

$$\mathbf{F(\chi) = 1 / (1 + \chi^n), \; n \geq 1} \quad (9)$$

(For effective-fluid interpretation of wDE and the non-stiffness of large positive values, see Section 5.2 Remark.)

The exponent n parameterises transition sharpness and is not fixed by the axioms — it is a physical parameter to be constrained observationally. Section 6 explains how low-redshift distance and growth data constrain n; a full likelihood fit is deferred to a companion analysis.

3.3 The Effective Cosmological Term

Combining the geometric closure scale and the closure functional:

$$\mathbf{\Lambda_{eff}(a) = (3/\xi^2) \cdot F(\chi(a)) = 3 / [\xi^2(1 + \chi_\infty^n \cdot a^{-3n})]} \quad (10)$$

Limiting behaviour:

- Early times ($a \to 0, \chi \to \infty$): $\Lambda_{eff} \to (3/\xi^2) \cdot \chi^{-n} \to 0$ ✓
- Late times ($a \to \infty, \chi \to 0$): $\Lambda_{eff} \to 3/\xi^2 = \Lambda_\infty$ ✓

---

4. The Coherence Scale ξ

4.1 Derivation from Early-Universe Acoustic Physics (BAO / Drag Epoch)

We take the coherence scale ξ to be set by early-universe acoustic physics at the radiation→matter transition, and identify it with the comoving BAO sound-horizon ruler (conventionally evaluated at the baryon drag epoch). Before matter–radiation equality ($a_{eq} \approx 3\times10^{-4}$), radiation pressure prevents the formation of stable committed overdensities on scales larger than the acoustic horizon. Matter perturbations above the Jeans scale are disrupted by photon pressure before they can satisfy the IC criterion (gravitational binding, survival over H^{-1}).

After the drag epoch, baryons decouple from photons and the acoustic oscillations freeze out. The maximum scale over which committed correlation channels first become dynamically viable is the frozen sound horizon:

$\xi \approx r_s \approx 150$ Mpc (comoving BAO ruler)

This scale is directly imprinted on the matter power spectrum as the BAO peak. We treat $\xi = \kappa\, r_s$ with $\kappa = O(1)$ absorbing definitional conventions (drag epoch vs alternative sound-horizon definitions) and the model's record-channel criterion; the ~20% agreement with the observed Λ corresponds to $\kappa \approx 0.9$. The physical motivation for $\xi = \kappa\, r_s$ is therefore that r_s is the largest scale on which IC-stable committed correlations first form after acoustic decoupling. We assume the largest IC-stable scale is set at acoustic decoupling and subsequently preserved under cosmic expansion. Later nonlinear structure formation can modify the network of committed correlations but does not generate new coherence scales larger than the acoustic horizon. Possible renormalization of ξ is discussed in Section 4.4. $\kappa = O(1)$ absorbs definitional conventions (drag epoch vs alternative sound-horizon integrals) and is fixed by the companion sound-horizon calculation.

The asymptotic cosmological term follows:

$\Lambda\infty = 3/\xi^2 \approx 3/(150 \text{ Mpc})^2 \approx 1.3 \times 10^{-52}\ \text{m}^{-2}$

compared to the observed value $\Lambda\text{obs} \approx 1.1 \times 10^{-52}\ \text{m}^{-2}$. With $\kappa = 1$ ($\xi = r_s$ exactly), $\Lambda\infty$ differs from Λobs by ~20%. The framework predicts the correct order of magnitude without fine-tuning; precision matching requires evaluation of κ from the full sound-horizon integral and is deferred to the companion paper. The ~20% is not a precision prediction — κ is an O(1) geometric factor the framework does not yet fix.

4.2 The Cosmological Constant Problem Reframed

In standard quantum field theory, the vacuum energy density is estimated as ρQFT ~ ρPlanck $\approx 10^{96}$ kg m^{-3}, while the observed value is $\rho\Lambda \approx 10^{-26}$ kg m^{-3} — a discrepancy of ~122 orders of magnitude.

In the ledger framework, this problem is reframed rather than solved by fine-tuning. Λeff does not arise from zero-point vacuum fluctuations. It arises from the *geometric structure of unsupported committed-record capacity* at the coherence scale ξ. The QFT vacuum contribution is absent by construction — not cancelled by a counterterm — because the closure term is a property of the causal ledger, not of quantum fields propagating on a background.

The residual question — why ξ takes its observed value — is answered by early-universe acoustic physics (sound horizon / drag epoch): ξ is the BAO scale, fixed by the radiation-to-matter transition in the early universe. The late-time acceleration scale is therefore determined by early-universe parameters, not by any fine-tuned cancellation.

4.3 The Coincidence Problem Resolved

The coincidence problem asks: why does dark energy dominate precisely now? In the ledger model the timing is geometrically constrained, though a full resolution requires a prior on r_0 from earlier-universe physics. We clarify the extent of the resolution. Acceleration begins when the matter-supported channel density falls below the fixed coherence threshold:

A natural crossover scale is $\chi(a) = 1$, i.e. $\rho_m(a) = \rho\Lambda,\infty$, giving:

$a\times = \chi\infty\hat{}(1/3)$, $z\times = \chi\infty\hat{}(-1/3) - 1$

where $\chi\infty$ is determined by solving $r_0 = \chi\infty(1 + \chi\infty^n)$ for the chosen n. For n = 2: $\chi\infty \approx 0.391$, giving $a\times \approx 0.731$, $z\times \approx 0.37$. For n = 5: $\chi\infty \approx 0.443$, giving $a\times \approx 0.763$, $z\times \approx 0.31$. In both cases the crossover falls near the present epoch — a direct consequence of ξ being set by early-universe acoustic physics (sound horizon / drag epoch), not a coincidence.

The onset of acceleration ($\ddot{a} > 0$) occurs at a nearby scale set by $\rho_m \lesssim 2\rho DE$, which depends weakly on n through $\rho DE(a)$. The two scales — density equality and acceleration onset — are closely related but not identical.

Scope of the resolution. The model reduces the coincidence problem: the crossover redshift is determined by $\xi = r_s$ (an early-universe acoustic scale) rather than by fine-tuning of a vacuum energy. However, the argument uses $r_0 = 0.45$ as observational input. A complete resolution would require predicting r_0 from physics prior to the drag epoch. The ledger framework reduces the problem to the origin of r_0 and n, but does not yet solve it.

4.4 Possible Evolution of ξ

Two cases bracket the model:

(A) ξ constant: The coherence scale is fixed at equality and does not evolve. Then $w(a) \to -1$ asymptotically and deviations from ΛCDM are confined to the crossover epoch. This is the baseline model.

(B) ξ slowly evolving: If ongoing large-scale structure formation modifies the committed correlation network after equality, ξ may grow logarithmically:

$\xi(a) = \xi_0 \cdot [1 + \alpha \cdot \ln(a/aeq)]$, $|\alpha| \ll 1$

This introduces mild time dependence in Λeff and a departure from $w = -1$ even in the late-time limit. Case (A) is treated throughout; case (B) is a testable perturbation explored in Section 6.

---

5. Modified Friedmann Dynamics

5.1 The Modified Friedmann Equation

The effective closure term enters the Friedmann equation as a geometrically-sourced fluid with energy density $\rho DE(a) = \Lambda eff(a)/(8\pi G)$:

$$\mathbf{H^2(a) = (8\pi G/3) \cdot [\rho_m(a) + \rho DE(a)]} \quad (11)$$

with:

$$\rho DE(a) = (3 / 8\pi G\xi^2) \cdot 1/(1 + \chi\infty^n \cdot a^{-3n})$$

Because $\Lambda eff(a)$ appears as a geometric scalar in the modified field equations, total stress-energy conservation follows from the contracted Bianchi identity; no separate law needs to be imposed by hand. Separate conservation of each sector holds under the assumption of no direct matter-closure coupling, which we adopt explicitly:

$\nabla\mu\ TDE$^$\mu\nu = 0$, $\nabla\mu\ T_m$^$\mu\nu = 0$ (separately; no direct coupling)

Operationally, we treat the closure sector as an effective fluid with energy-momentum tensor TDE^μν = ρDE u^μ u^ν + pDE (g^μν + u^μ u^ν), where pDE is fixed by the continuity equation $\dot{\rho}DE + 3H(\rho DE + pDE) = 0$ and yields the wDE(a) derived in Section 5.2. A derivation from a ledger action principle is deferred to future work.

5.2 Dark-Energy Equation of State

We treat the closure contribution as an effective dark-energy component with density:

$$\rho DE(a) \equiv \Lambda eff(a)/(8\pi G) = (\Lambda\infty/8\pi G) \cdot 1/(1 + \chi(a)^n),\ \chi(a) = \chi\infty \cdot a^{-3}$$

Its equation-of-state parameter is defined by:

$$wDE(a) \equiv -1 - (1/3) \cdot d \ln \rho DE / d \ln a = -1 - (1/3) \cdot d \ln \Lambda eff / d \ln a$$

Computing the logarithmic derivative. Since $\chi = \chi\infty\ a^{-3}$, we have $d \ln \chi/d \ln a = -3$, and:

$$d(\chi^n)/d \ln a = n\ \chi^n \cdot (d \ln \chi/d \ln a) = -3n\ \chi^n$$

Therefore:

$$d \ln \Lambda eff / d \ln a = -1/(1+\chi^n) \cdot d(\chi^n)/d \ln a = 3n\chi^n/(1+\chi^n)$$

Giving the closed form:

$w_{DE}(a) = -1 + n\cdot\chi(a)^n / (1 + \chi(a)^n)$ (12)

The limiting behaviour is:

$\chi \ll 1$ (late times): $w_{DE} \to -1$ $\chi \gg 1$ (early times): $w_{DE} \to n - 1$

The early-time limit does not by itself imply observational tension, because the corresponding energy density is strongly suppressed:

$\rho_{DE}(a) \propto \chi(a)^{-n} \propto a^{3n}$ $(\chi \gg 1)$

so the closure sector is dynamically negligible at high redshift for any $n \geq 1$. For $n \geq 2$ in particular, $\rho_{DE} \propto a^6$ or steeper at early times, guaranteeing negligible early dark energy and fully preserving standard recombination and CMB constraints. The observationally relevant regime is low redshift, where Ω_{DE} is appreciable.

Remark (effective-fluid interpretation). The quantity $w_{DE}(a)$ is defined operationally from the background evolution implied by $\rho_{DE}(a)$; it does not assume the closure sector is a barotropic material fluid with a microphysical equation of state $p = p(\rho)$. In particular, large positive values of w_{DE} at intermediate redshift do not imply superluminal propagation: throughout we treat the closure sector as a geometrically sourced, effectively smooth component whose perturbations are suppressed (analogous to vacuum energy), and whose clustering properties must be analysed separately in a perturbation treatment.

5.3 Low-Redshift Constraint Strategy

Constraints on $w_{DE}(z)$ primarily come from the impact on the distance–redshift relation and late-time expansion $H(z)$ at $z \lesssim 1$, where the closure sector contributes non-negligibly to the Friedmann equation. Accordingly, bounds on n should be obtained by fitting the implied $H(z)$ — or equivalently the luminosity distance $D_L(z)$ and BAO distances — rather than by interpreting $w_{DE}(0)$ in isolation. The present-day value:

$w_{DE}(0) = -1 + n\cdot\chi_\infty^n / (1 + \chi_\infty^n)$

provides a useful diagnostic, but the full shape of $w_{DE}(z)$ across the transition must be compared against the data-weighted integral of its effect on $H(z)$. This fitting exercise is deferred to a companion analysis; we report $w_{DE}(z)$ here for representative values of n to motivate the observational programme.

---

6. Observational Predictions

6.1 wDE(z) Predictions

We report wDE(z) for $z \lesssim 1$, where ΩDE(z) is non-negligible and distance probes actually constrain expansion. At higher redshift, the closure component is strongly suppressed ($\rho DE \propto a^{3n}$ in the $\chi \gg 1$ regime), so wDE becomes observationally irrelevant even if its formal value is large.

Table 1 reports wDE(z) and ΩDE(z) in **Model 0** (baseline stress test, fcoll = constant) for n = 2 ($\chi_\infty = 0.391$) and n = 5 ($\chi_\infty = 0.443$), both with $r_0 = 0.45$. Here $\chi(z) = \chi_\infty \cdot (1+z)^3$, and:

$$\Omega DE(z) = 1 / [1 + \chi(z)\cdot(1 + \chi(z)^n)]$$

z	χ (n=2)	wDE (n=2)	ΩDE (n=2)	χ (n=5)	wDE (n=5)	ΩDE (n=5)
0.0	0.391	−0.735	0.690	0.443	−0.917	0.690
0.3	0.858	−0.152	0.402	0.972	+1.331	0.355
0.5	1.318	+0.269	0.217	1.493	+3.407	0.074
1.0	3.124	+0.814	0.029	3.540	+3.991	5.0×10^{-4}

All values rounded to 3 significant figures. ΩDE(n=5) at z = 1.0 is 5.0×10^{-4}.

The ΩDE column is essential context: at z = 0.5 the closure component contributes only 22% (n=2) or 7% (n=5) of the total energy budget, and at z = 1.0 it is below 3% and 0.1% respectively. The formally large positive values of wDE at z > 0.5 for n = 5 are therefore observationally negligible — they occur precisely where the component has almost no weight in the distance integrals. The n=5 case is shown as a deliberately sharp-transition example to illustrate shape sensitivity; its viability must be assessed by fitting H(z) and distance data.

Observational discrimination between values of n is best achieved in the range 0 < z < 0.5, where both ΩDE is appreciable and wDE differs meaningfully between cases. DESI, Euclid, and the Vera Rubin Observatory will achieve sub-percent distance measurements across this range. Large positive values of wDE in the sharp-transition examples should be read as an effective background diagnostic (Section 5.2 Remark), not as a claim about microphysical fluid stiffness. The purpose of the table is shape intuition; the only observationally acceptable values of n are those that pass a full BAO + SNe distance fit.

6.2 Planck Consistency

Planck 2018 combined with BAO constrains the dark energy equation of state at $w_0 = -1.03 \pm 0.03$ and $w_a = 0.12 \pm 0.30$ (under the CPL parameterisation). In the ledger model, the relevant question is not the formal value of wDE(0) in isolation but the effect of the predicted H(z) on CMB distance ratios and BAO peak positions across 0 < z < 1.

For a given n, χ_∞ is fixed by $r_0 = \chi_\infty(1 + \chi_\infty^n)$. For the representative cases in Table 1, this gives wDE(0) ≈ −0.735 (n=2) and wDE(0) ≈ −0.917 (n=5).

Naively compared to the Planck CPL constraint $w_0 = -1.03 \pm 0.03$, the n=2 value differs by ~10σ and n=5 by ~4σ under CPL compression. This is not "formal tension" — it is a large offset that requires honest treatment. We give three reasons why the CPL comparison does not straightforwardly falsify the model:

1. **CPL assumes $w \approx -1$ everywhere cosmologically relevant.** The parameterisation $w(a) = w_0 + w_a(1-a)$ compresses w across all redshifts. In the ledger model, large $|w_{DE} + 1|$ occurs at $z \sim 0.3$–0.5 where Ω_{DE} is already falling steeply (Table 1, $\Omega_{DE} < 0.25$ for n=2 at z=0.5). The CPL number therefore misrepresents the ledger model's actual effect on distance integrals.
2. **Observational constraints come from H(z) integrals, not w_0.** BAO ratios and luminosity distances constrain $\int dz/H(z)$, not $w_{DE}(0)$ in isolation. A model with large $|w_{DE} + 1|$ where Ω_{DE} is suppressed can produce nearly identical distance integrals to ΛCDM.
3. **Model 0 likely fails; Model 1 is the relevant test.** The Section 6.5 diagnostic confirms that Model 0 deviates from ΛCDM H(z) by 10–20% at fixed Planck parameters. The meaningful test is whether Model 1, with free {H_0, $\Omega_{m,0}$, n, Mmin, can fit BAO + SNe — not whether Model 0 matches CPL at fixed parameters. The companion analysis will provide this test.

In practice, the relevant comparison is to the BAO/SNe constraints on E(z) and the integrated distances DM(z), DH(z), since these are the direct observables; CPL (w_0, w_a) constraints are only an approximate compression of those data.

6.3 Structure Growth

The linear matter perturbation equation:

$$\ddot{\delta}_m + 2H \cdot \dot{\delta}_m - 4\pi G\rho_m \cdot \delta_m = 0$$

is unmodified at leading order for constant ξ, because the closure sector is taken to be effectively smooth (negligible $\delta\rho_{DE}$ on sub-horizon scales) and enters only through the background expansion H(a). It exerts no direct force on matter perturbations. The growth index is preserved:

$\gamma \approx 0.55$ (consistent with ΛCDM)

For slowly evolving ξ (case B, parameter α), a source term appears proportional to $(d\Lambda_{eff}/dt)/H^2$, producing:

$$\Delta\gamma \approx -0.3\alpha \cdot (\Omega_{DE}/\Omega_m)^{1/2}|_{z=0} \approx -0.5\alpha$$

DESI and Euclid $f\sigma_8$ measurements (targeting $\sigma(\gamma) \sim 0.01$–0.02) will constrain $\alpha \lesssim 0.03$ at 2σ.

6.4 Discriminating Signatures

The ledger closure model is distinguishable from vanilla ΛCDM and from scalar quintessence by the following signatures:

Signature 1 — wDE(z) transition shape: The functional form $w_{DE}(z) = -1 + n\cdot\chi(z)^n/(1+\chi(z)^n)$ with $\chi(z) = \chi_\infty(1+z)^3$ has a specific redshift dependence determined by n and χ_∞. Unlike the CPL parameterisation (which is linear in a), the ledger wDE(z) has a characteristic sigmoid shape in log(1+z). A sufficiently precise w(z) measurement across $0 < z < 0.5$ — where Ω_{DE} is still appreciable — can distinguish this from both ΛCDM and standard quintessence.

Signature 2 — No fifth force: The closure term is geometrically sourced, not mediated by a propagating scalar. Solar system tests of GR are unaffected. Screening mechanisms are not needed.

Signature 3 — Consistency tie to BAO scale: The identification $\xi \approx r_s$ ties the late-time acceleration scale to the BAO ruler. This predicts a specific relationship between the inferred Λ_∞ and the sound-horizon scale, testable by jointly fitting late-time distances and BAO-calibrated rulers. Any statistically significant mismatch between the Λ_∞ inferred from expansion history and the value $3/r_s^2$ would falsify the $\xi = r_s$ identification.

Signature 4 — Growth index deviation for evolving ξ: If $\alpha \neq 0$, $f\sigma_8$ measurements deviate from ΛCDM at the level $\Delta(f\sigma_8) \sim 0.5\alpha \cdot f\sigma_8$.

6.5 Model 0 Stress Test: Background Expansion Deviation from ΛCDM

Before a full likelihood analysis, we compute the fractional deviation of the model's Hubble rate from ΛCDM when both are evaluated using the same Planck-ΛCDM parameter values ($\Omega_{m,0}$ = 0.31, $\Omega_{\Lambda,0}$ = 0.69, H_0 fixed). This is not a claim of agreement or tension — it is an internal diagnostic showing how much the predicted expansion history differs from ΛCDM. Section 6.1 reports the internal model evolution, whereas Section 6.5 deliberately compares the model's E(z) against a fixed ΛCDM reference point to quantify departure magnitude before allowing parameter refits.

Define the normalised Hubble rate:

$$E(z)^2 \equiv H^2(z)/H_0^2 = \Omega_{m,0}(1+z)^3 + \Omega_{DE,0} \cdot [\rho_{DE}(z)/\rho_{DE,0}]$$

with:

$$\rho_{DE}(z)/\rho_{DE,0} = (1 + \chi_\infty^n) / (1 + \chi(z)^n),\ \chi(z) = \chi_\infty(1+z)^3$$

This normalisation ensures E(0) = 1 by construction.

Normalisation note. In this diagnostic we hold ($\Omega_{m,0}$, $\Omega_{\Lambda,0}$, H_0) fixed to a Planck-ΛCDM reference point and set $\Omega_{DE,0} \equiv 1 - \Omega_{m,0}$ for flatness in both models. The model dependence enters only through $\rho_{DE}(z)/\rho_{DE,0}$, and we impose flatness $\Omega_{DE,0} = 1 - \Omega_{m,0}$ in both models. This isolates the shape difference in E(z) implied by the closure functional from shifts in (H_0,

$\Omega_{m,0}$) that would be allowed in a true parameter fit. Section 6.5 is a diagnostic under fixed reference parameters, not a prediction at the best-fit point.

The fractional deviation from ΛCDM is:

$\Delta E(z) \equiv [E_m odel(z) - E\Lambda CDM(z)] / E\Lambda CDM(z)$

The following table gives ΔE(z) for n = 2 ($\chi_\infty = 0.391$) and n = 5 ($\chi_\infty = 0.443$):

z	χ (n=2)	ΔE (n=2)	χ (n=5)	ΔE (n=5)
0.1	0.519	−2.9%	0.588	−1.6%
0.3	0.857	−8.8%	0.971	−12.2%
0.5	1.316	−12.2%	1.491	−19.3%
0.7	1.916	−12.5%	2.171	−16.7%
1.0	3.120	−10.2%	3.535	−11.5%

The deviations are negative throughout: the closure model predicts systematically lower H(z) than ΛCDM across $0 < z < 1$. This is a direct consequence of the closure functional — at $z > 0$, $\rho DE(z) < \rho\Lambda$ (since $F(\chi) < 1$), so the model expands more slowly than ΛCDM at intermediate redshifts. These large deviations are expected in Model 0 because it ignores the growth of committed collapsed structure through fcoll(z); Model 0 is retained precisely as a stress test and an upper bound on departures from $\Lambda\infty$.

Interpretation: Because Model 0 deliberately holds fcoll constant, it is intended as an upper-bound stress test. The physically intended model is Model 1, where fcoll(z) is computed from halo statistics; Section 6.5 therefore diagnoses what must be corrected by channel growth, not the final fitted expansion history.

These deviations are large. Interpreted literally, they indicate that the baseline scaling $\chi(a) = \chi\infty\, a^{-3}$ (equivalently fcoll = constant) is unlikely to fit BAO+SNe distances at the Planck-ΛCDM parameter point. This is not a defect of the record premise; it is a warning that the baseline scaling is too naive for precision cosmology. By Definition 2.1, correlation channels are committed, gravitationally bound structures, so χ(a) must track collapsed and committed structure formation through fcoll(a). The purpose of Section 6.5 is therefore twofold: (i) it supplies a hard falsifier for the baseline model, and (ii) it motivates the channel-consistent extension $\chi(a) \propto \mathrm{fcoll}(a) \cdot \rho_m(a)$, to be evaluated against BAO, SNe, and CMB distance data in the companion analysis. Accordingly, the empirically relevant question is whether standard structure formation through fcoll(z) can reduce the mid-z deficit to the percent level required by BAO + SNe.

6.6 Viability Conditions and Falsifiability

For the ledger-closure model to be observationally viable, the following necessary conditions must hold. Each is independently falsifiable.

(C1) Early dark energy suppression. For any $n \geq 2$:

$\rho DE(a) \propto a^{3n}$ $(\chi \gg 1)$

At recombination ($z \sim 1100$), ρDE is suppressed by a factor $\sim (1+z)^{-3n}$ relative to today. For $n = 2$ this gives $\rho DE(z\_rec)/\rho DE_{,0} \sim 10^{-22}$, yielding negligible early dark energy. Planck CMB constraints on early dark energy ($\Omega\_EDE \lesssim 0.01$ at $z \sim 1000$; see Ref. 14) are automatically satisfied for $n \geq 2$.

(C2) Dedicated parameter fit to distance data. As demonstrated in Section 6.5, the model cannot use ΛCDM parameter values and simultaneously reproduce ΛCDM expansion history. Viability requires fitting H(z) data (BAO, SNe Ia) with free parameters $\{H_0, \Omega_{m,0}, n$. A valid fit existing somewhere in parameter space is a necessary condition for the model to be taken seriously.

(C3) Smooth acceleration history. The deceleration parameter:

$q(a) = -\ddot{a}\cdot a / (\dot{a}^2) = -1 - (1/2) \cdot d \ln H^2 / d \ln a$

must evolve monotonically from matter-dominated deceleration ($q \approx +0.5$) to late-time acceleration ($q < 0$) without transient super-acceleration ($q < -1$). For constant ξ and $n \geq 2$, the transition through the closure functional is smooth and monotonic. This is directly verifiable from H(z).

(C4) Smooth closure sector. The assumption $\delta\rho DE \approx 0$ on sub-horizon scales (Section 5.2) implies that the linear growth equation is unmodified and $\gamma \approx 0.55$. A measured deviation of $|\gamma - 0.55| > 0.05$ with $\Omega DE_{,0} \approx 0.69$ would falsify the smooth-closure hypothesis and require a perturbation treatment of the ledger sector.

These four conditions render the model falsifiable through early-universe probes, distance data, expansion history, and structure growth — independent data streams whose combination is robust.

Falsifier (distance/expansion). If, after allowing $\{H_0, \Omega_{m,0}, n$ to vary — and, in the channel-consistent extension, allowing the committed-fraction proxy fcoll(a) within standard structure-formation bounds — no parameter region exists that fits BAO and SNe distances at the percent level over $0 < z < 1$, the closure functional in its present form is falsified. The role of Section 6.5 is to show that such a fit is non-trivial and therefore constitutes a sharp empirical test rather than a guaranteed outcome.

6.7 Physical Interpretation of the Sharpness Exponent n

The exponent n in $F(\chi) = 1/(1 + \chi^n)$ parameterises the sharpness of the transition between the matter-supported and closure-dominated regimes. Rather than treating it as a purely phenomenological fitting parameter, we offer a physical interpretation that connects it to the microscopic record structure.

In a microscopic picture, each independent committed correlation channel requires a sufficient local matter density to remain gravitationally bound and decoherence-stable. As ρ_m falls through $\rho_{\Lambda,\infty}$, channels begin to fail. The rate at which they fail — gradually or abruptly — is encoded in n:

- **n ≈ 1–2 (gradual transition):** Channel failure is spread broadly across a range of densities. Matter structures lose record-support slowly as matter dilutes, analogous to a broad crossover in a disordered system.
- **n ≳ 5 (sharp transition):** Channel failure is concentrated near $\chi = 1$, analogous to a percolation threshold or critical point. Below the threshold density, essentially all channels fail simultaneously.

This interpretation connects n to the **percolation sharpness** of committed correlation channel stability — a quantity that in principle could be derived from the statistics of the halo mass function and the IC threshold criterion ($\delta > \delta$collapse). A microscopic derivation of n from such statistics, possibly via network percolation theory applied to the large-scale structure graph, is a concrete target for future work.

The key observational consequence is that n is not a free aesthetic parameter — it has a physical meaning that constrains what values are theoretically natural. Values $n \gtrsim 10$ would imply an implausibly sharp percolation threshold with no analog in known structure-formation statistics; values $n \lesssim 1$ would imply a broader transition than the gradual dilution of matter structures suggests. The theoretically motivated range $n \sim 2$–5 is therefore doubly significant: it brackets the observable regime and aligns with physically reasonable percolation behaviour.

Sharpness diagnostic. Operationally, n controls the width in ln a over which the transition in ΩDE(z) occurs. A direct sharpness diagnostic is the redshift width Δz between the points where ΩDE rises from 0.2 to 0.6. This provides a model-independent mapping from survey constraints on the distance–redshift curvature to an inferred n, without requiring a CPL approximation. Next-generation surveys (DESI, Euclid) measuring ΩDE(z) across $0 < z < 1$ at the few-percent level will place sharp bounds on this width and hence on n.

---

7. Discussion and Conclusions

We have derived, from three foundational axioms of the VERSF framework, a late-time infrared closure of the causal ledger that manifests as an effective cosmological term. The derivation is internally consistent and makes the following key identifications:

$\chi(a) = \chi_\infty \cdot a^{-3}$, where $\chi_\infty \equiv \rho_{m,0}/\rho_{\Lambda,\infty}$ is determined implicitly by the consistency relation $r_0 = \chi_\infty(1 + \chi_\infty^n)$, with $r_0 = \Omega_{m,0}/\Omega_{\Lambda,0} \approx 0.45$ the sole observationally-fixed input. No holographic bit-counting is required.

$\Lambda_\infty = 3/\xi^2$ emerges from FRW symmetry and the existence of a single infrared scale ξ, without invoking quantum zero-point energy.

$\xi \approx r_s \approx 150$ Mpc is identified with the comoving BAO sound-horizon ruler, giving Λ_∞ within ~20% of the observed value.

$F(\chi) = 1/(1+\chi^n)$ is the lowest-order rational functional form consistent with monotonicity and asymptotic constraints. The exponent n is an observationally constrained parameter; a proper fit to H(z) and BAO distances is required to bound it quantitatively. For $n \geq 2$, early dark energy is strongly suppressed ($\rho_{DE} \propto a^{3n}$), preserving CMB and recombination physics.

The cosmological constant problem is reframed: Λ_{eff} arises from unsupported ledger capacity at scale ξ, not from vacuum fluctuations. The coincidence problem is resolved geometrically: acceleration begins when ρ_m falls to $\rho_{\Lambda,\infty}$, a threshold set by early-universe acoustic physics (sound horizon / drag epoch) and hence by well-understood early-universe parameters.

Model status. Under the baseline scaling $\chi(a) \propto a^{-3}$, Section 6.5 indicates large mid-z deviations in H(z) at the Planck-ΛCDM parameter point, likely excluded by current distance data. The physically intended model is therefore the channel-consistent extension $\chi(a) \propto f_{coll}(a) \cdot \rho_m(a)$, which is the correct operational interpretation of "committed correlation channels" under Axiom IC. Establishing viability requires fitting that extension to BAO + SNe + CMB distance priors, which is the focus of the companion analysis.

7.1 Companion Analysis Plan

The companion analysis will confront Model 1 with current distance data using the following specification.

Datasets:

- BAO: comoving angular diameter distance $D_M(z)/r_s$ and Hubble distance $D_H(z)/r_s$ from DESI DR1 and prior galaxy surveys at $0.1 < z < 2.3$.
- Type Ia supernovae: distance modulus residuals from Pantheon+ or Union3 over $0.01 < z < 2.3$.
- CMB compressed distance priors: the Planck shift parameter R and acoustic scale ℓ_A, which encode the angular diameter distance to last scattering.

Free parameters: {H_0, $\Omega_{m,0}$, n, M_{min}, κ, where $\kappa = \xi/r_s$ is the O(1) sound-horizon calibration factor.

Computation pipeline:

1. For each parameter point, compute fcoll(z; Mmin) from the Sheth–Tormen mass function evaluated self-consistently.
2. Construct $\chi(z) = \chi_\infty \cdot (1+z)^3 \cdot$ fcoll(z)/fcoll(0) under Model 1.
3. Evaluate E(z), DM(z), DH(z), and distance moduli.
4. Compute χ^2 against BAO + SNe + CMB priors and run MCMC.

Success criterion: A statistically acceptable fit (no significant degradation relative to ΛCDM given comparable parameter count) with Mmin in the physically motivated range 10^{10}–10^{12} M_⊙ and κ within ±30% of unity. This would establish Model 1 as observationally viable.

Failure criterion: No parameter region satisfying the above with Mmin in the physical range falsifies the closure functional $F(\chi) = 1/(1 + \chi^n)$ in its present form.

If the companion analysis confirms that Model 1 fits current distance data, the results presented here suggest that cosmic acceleration may arise not from vacuum energy but from the large-scale information structure of the universe.

Falsification criterion. If Model 1 cannot fit BAO + SNe distances for any plausible Mmin in the galaxy-to-group range (10^{10}–10^{12} M_⊙), with {H_0, $\Omega_{m,0}$, n free, the closure functional $F(\chi) = 1/(1 + \chi^n)$ is falsified.

Open questions for future work:

1. Dedicated MCMC fit to BAO, SNe Ia, and CMB distance data with free parameters {H_0, $\Omega_{m,0}$, n, as motivated by Section 6.5.
2. Variational derivation of the modified Friedmann equation from a ledger action principle.
3. Precise O(1) matching between ξ and the full sound horizon integral.
4. Non-linear closure effects at z ~ 2–5 during active channel formation.
5. Constraints on the evolution parameter α from existing $f\sigma_8$ compilations.
6. Microscopic derivation of n from halo mass function statistics and network percolation theory.

Appendix A: Minimal Rational Ansatz Consistent with Axioms

Under the additional assumption of lowest-degree rational analyticity on $\chi \in (0,\infty)$, the minimal functional form consistent with the four axiom-derived constraints is $F(\chi) = 1/(1+\chi^n)$. Other monotone analytic forms (e.g. stretched exponentials, Padé approximants) are possible; we adopt this as the minimal representative class, acknowledging that the axioms constrain but do not uniquely determine the functional form.

Formally: let $F = P(\chi)/Q(\chi)$ where P, Q are polynomials with no common roots. The conditions $F(0) = 1$ and $F(\infty) = 0$ require $P(0) = Q(0)$ and $\deg(P) < \deg(Q)$. The minimal case is $\deg(P) = 0$

(constant numerator), deg(Q) = n. Normalising P = 1, we need Q(0) = 1 and Q strictly increasing on (0,∞) for monotonicity. The simplest such polynomial is $Q(\chi) = 1 + \chi^n$, giving:

$F(\chi) = 1 / (1 + \chi^n)$ ■

Appendix B: The Holographic Bound as Interpretive Context

The holographic Bekenstein-Hawking bound provides a useful interpretive frame for the ledger capacity, though it is not required for the quantitative model.

The maximum information capacity of the Hubble volume is:

$\mathcal{C}_{max} \approx A_H / (4\ell_P^2) \approx \pi(c/H_0)^2 \cdot c^3/(G\hbar) \sim 10^{123}$ bits

The total number of decohered matter particles within the Hubble volume is:

$N_{m,0} \sim (\rho_{m,0}/m_p) \cdot (4\pi/3)(c/H_0)^3 \sim 10^{80}$

The ratio $N_{m,0}/\mathcal{C}_{max} \sim 10^{-43}$ is not the model parameter χ_∞. Rather, it reflects the microscopic-to-macroscopic hierarchy: the ledger has vastly more capacity than individual particle correlations can fill, but only a small subset — the fraction associated with committed, gravitationally-bound channel correlations — is dynamically relevant. The phenomenological $\chi_\infty \approx 0.39$–0.44 operates at the macroscopic (halo/large-scale structure) level, not the particle level.

This distinction motivates a more careful microscopic derivation of the channel capacity as a function of the halo mass function — deferred to a companion paper.

Appendix C: Channel-Consistent χ(a) from Halo Collapse Statistics

This appendix specifies a concrete route to compute the channel-consistent occupancy ratio $\chi(a)$ in Model 1 using standard structure formation inputs, without introducing new free physics.

C.1 Collapsed Fraction Proxy

We define the collapsed/committed fraction as the mass fraction in halos above a minimum mass M_{min}:

$f_{coll}(z; M_{min}) = (1/\rho_m) \cdot \int[M_{min} \to \infty]\, M \cdot (dn/dM)(M,z)\, dM$

where dn/dM is a standard halo mass function (Press–Schechter or Sheth–Tormen). Mmin represents the smallest halo mass that can sustain an IC-stable committed record over $\Delta t \sim H^{-1}$ — operationally, a structure that is gravitationally bound and robust to tidal disruption. Plausible values lie in the range 10^{10}–10^{12} M_⊙, corresponding to galaxy-scale through group-scale channels, and Mmin is treated as a model parameter to be fitted.

Remark: Mmin is not a new fundamental parameter; it encodes the operational definition of what constitutes an IC-stable "record channel" in cosmology (galaxy-scale vs group-scale vs cluster-scale committed structure). Different choices correspond to different empirical channel definitions, not different underlying laws.

C.2 Occupancy Ratio in Model 1

In Model 1, correlation capacity is proportional to committed collapsed mass density:

$\mathcal{C}_m(a) \propto \text{fcoll}(a) \cdot \rho_m(a)$

The occupancy ratio therefore becomes:

$\chi(a) = \text{fcoll}(a) \cdot \rho_m(a) / \rho\Lambda,\infty = \chi\infty \cdot a^{-3} \cdot \text{fcoll}(a)/\text{fcoll}(1)$

This reduces to Model 0 when fcoll(a) is constant, confirming that Model 0 is a limiting case of Model 1. In numerical implementation, fcoll(z) is computed on a common redshift grid with consistent precision to avoid normalization artifacts at z = 0; the ratio fcoll(a)/fcoll(1) is evaluated from the same grid to ensure exact cancellation at a = 1.

C.3 Modified Λeff(a) and Background Expansion

With $F(\chi) = 1/(1 + \chi^n)$, the effective cosmological term is:

$\Lambda\text{eff}(a) = \Lambda\infty \cdot 1/(1 + \chi(a)^n)$

The normalised expansion function is:

$E(z)^2 = \Omega_{m,0}(1+z)^3 + \Omega DE_{,0} \cdot \rho DE(z)/\rho DE(0), \rho DE(z)/\rho DE(0) = (1 + \chi(0)^n)/(1 + \chi(z)^n)$

The companion analysis evaluates fcoll(z) numerically for a chosen mass function and Mmin, then computes E(z), the comoving distance DM(z), and the Hubble distance DH(z) = c/H(z) for direct comparison with BAO and SNe distance constraints. In the companion analysis, dn/dM is computed using the background expansion H(z) and growth function corresponding to each parameter point {H_0, $\Omega_{m,0}$, n, Mmin, ensuring internal consistency between halo statistics and the modified Friedmann dynamics. This prevents the implicit ΛCDM circularity that would arise from using a fixed ΛCDM halo mass function at non-ΛCDM parameter points.

C.4 Testable Consequence

Model 1 predicts that the baseline deviations in Section 6.5 are moderated at $z \lesssim 1$ by structure growth through fcoll(z): as halos form and fcoll rises, $\chi(a)$ is held up relative to the a^{-3} baseline, reducing the closure term activation and suppressing the H(z) deficit. Quantitatively, viability requires that the resulting E(z) and distance ratios fit BAO + SNe within current uncertainties for some parameter point {H_0, $\Omega_{m,0}$, n, Mmin with Mmin in a physically motivated halo mass range. Failure to find such a region falsifies the closure functional in its present form.

References

1. Planck Collaboration (2018). *Planck 2018 results VI: Cosmological parameters.* A&A, 641, A6.
2. Planck Collaboration (2020). *Planck 2018 results VIII: Gravitational lensing.* A&A, 641, A8.
3. Riess, A.G., et al. (1998). *Observational evidence from supernovae for an accelerating universe.* AJ, 116, 1009.
4. Perlmutter, S., et al. (1999). *Measurements of Ω and Λ from 42 high-redshift supernovae.* ApJ, 517, 565.
5. DESI Collaboration (2024). *DESI 2024 VI: Cosmological constraints from BAO.* arXiv:2404.03002.
6. Weinberg, S. (1989). *The cosmological constant problem.* Rev. Mod. Phys., 61, 1.
7. Bekenstein, J.D. (1973). *Black holes and entropy.* PRD, 7, 2333.
8. Eisenstein, D.J. & Hu, W. (1998). *Baryonic features in the matter transfer function.* ApJ, 496, 605.
9. Carroll, S.M. (2001). *The cosmological constant.* Living Rev. Relativ., 4, 1.
10. Chevallier, M. & Polarski, D. (2001). *Accelerating universes with scaling dark matter.* Int. J. Mod. Phys. D, 10, 213.
11. Linder, E.V. (2003). *Exploring the expansion history of the universe.* Phys. Rev. Lett., 90, 091301.
12. Press, W.H. & Schechter, P. (1974). *Formation of galaxies and clusters of galaxies by self-similar gravitational condensation.* ApJ, 187, 425.
13. Sheth, R.K. & Tormen, G. (1999). *Large-scale bias and the peak-background split.* MNRAS, 308, 119.
14. Hill, J.C., et al. (2022). *Atacama Cosmology Telescope: Constraints on pre-recombination early dark energy.* Phys. Rev. D, 105, 123536.