Dark Energy from Discrete Stochastic Cascade Dynamics: A Coupled DSCD+GR System, DESI Audit, and a Calibrated DR3 Forecast

Shiv Goswami*

Terraflock Pvt Ltd, Bihar, India

ORCID: 0009-0006-1554-9241 · shivgoswami.com

10 July 2026

Abstract

We construct a cosmological realization of Discrete Stochastic Cascade Dynamics (DSCD) as one coupled stochastic dynamical system rather than as a phenomenological equation-of-state ansatz. Standard general relativity and a flat FLRW background supply the gravitational closure; asymmetric DSCD interactions, beat phase, structural memory, ripple-dependent regimes, conservative transport, and funded free-capacity depletion generate the dark sector trajectory. The simulation outputs $\rho_C(t)$, $p_C(t)$, $H(z)$, and BAO distances directly. Five prerequisite validations test deterministic resource accounting, stochastic hazards, transport, minimal GR coupling, and synthetic distance recovery before any DESI values are used. All five pass, as do the implementation's positivity, continuity, Friedmann, early-density, and reproducibility checks.

Using official DESI DR1 and DR2 Gaussian BAO vectors and full covariance matrices, and fitting only $\theta = H_0 r_d / c$ and Ω_m while freezing all DSCD dynamics, gives $\chi^2 = 12.736$ on DR1 and $\chi^2 = 10.280$ on DR2. A retrospective within-DR1 high-redshift holdout gives conditional $\chi^2 = 3.466$, close to ΛCDM at 3.492. These are retrospective diagnostics, not blind tests. A synthetic identifiability audit finds that one candidate DSCD depletion combination is not identifiable from compressed BAO (whitened-Jacobian condition number 1.06×10^7). The complete system is also observationally indistinguishable from ΛCDM and from zero-beat, no-memory, no-regime, no-transport, and symmetric-depletion ablations at current precision. No single frozen configuration can therefore be forecast-eligible on identifiability grounds.

The forecasting inference layer poses the question the identifiability audit never asked: the present DSCD state is treated as latent, a declared eight-dimensional prior over realizations is sampled with a scrambled Sobol sequence, and each complete DSCD+GR realization is weighted by its joint DR1+DR2 likelihood with

*Contact author: shiv@terraflock.com

the common BAO scale marginalized analytically. The history-compatible ensemble converges: 2048 realizations give a joint effective sample size of 205, credible intervals stable under sample halving, disjoint seed banks, and DR2-only conditioning (median shifts below 0.07 of the 68% width), and an end-to-end synthetic calibration covering 94.9% (95% level) and 73.7% (68% level) of withheld truths. All convergence gates pass, and a sealed DR3-layout forecast is issued: thirteen credible intervals with 68% widths between 0.16 and 1.10 DR2 standard deviations, whose medians lie within 0.06σ of best-fit ΛCDM. The forecast is falsifiable in both directions: a DR3 phantom-side departure from these intervals rules the realization family out, and the intervals are narrower than the phantom excursions preferred by unconstrained CPL fits.

Contents

1 Status and Scope

Papers I and II [1, 2] define a compressed finite-graph stochastic dynamical system. They do not by themselves define a cosmology, a redshift map, or an equation of state. The present work therefore introduces a separate, richer cosmological realization and states every additional assumption explicitly.

The realization is deliberately limited to late-time background cosmology. It assumes standard GR, spatial flatness, separately conserved matter and radiation, and an early-time DSCD contribution small enough that the sound horizon can be represented by the BAO scale parameter $\theta = H_0 r_d / c$. It does not derive gravity, perturbations, the drag epoch, or the physical origin of BAO.

Two inference layers share this one simulation engine. Version 1 freezes a single declared configuration and asks whether compressed BAO can identify its internal parameters; it cannot, and that audit returns `NO_FORECAST`. Version 2 treats the present DSCD state as latent, weights an ensemble of realizations by the DR1+DR2 history, and asks whether the survivors agree on DR3. They do, and the resulting calibrated forecast is sealed in this paper.

2 The Coupled DSCD+GR System

Definition 2.1 (Coupled DSCD+GR realization)**.** Cells form a finite weighted comoving graph. The complete transition state is

$$Z_j = \{N_j, a_j, E_j, \rho_{C,i}, X_i, Y_i, S_i, S_{i-1}, S_{i-2}, \phi_i, r_i\}_{i=1}^{P},$$

where $N = \ln a$, $E = H/H_0$, $\rho_{C,i}$ is an intensive DSCD free-capacity density, ϕ_i is the beat phase, and r_i is the local regime. Including two structure histories and the phase makes the transition state Markovian. Requested events are sampled from the proper-time hazards

$$R_{XY} = \alpha_{XY}\omega_X\omega_Y\Theta(\phi)XY, \qquad R_{XX} = \alpha_{XX}\omega_X^2 X(X-1)/2,$$

$$\Theta(\phi) = \vartheta + (1-\vartheta)(1+\cos\phi)/2, \qquad \dot{\phi} = \omega_X - \omega_Y,$$

and every interaction, leakage, explosion, and bond event is funded from the remaining local density before it is committed.

Proposition 2.2 (Funded monotonicity)**.** *Along every realized path of Definition 2.1,*

$$\rho_{C,i}^{j+1} = \rho_{C,i}^{j} - d_i^j, \qquad 0 \le d_i^j \le \rho_{C,i}^{j},$$

so each local free-capacity density is nonincreasing and nonnegative, with no post-update clipping.

Proof. Every event class first draws a requested count, which is capped by species availability and by the remaining local budget before it is committed, and the budget is debited immediately (Appendix B). The realized debit d_i^j is therefore the exact sum of funded costs and cannot exceed the budget from which it was drawn. □

The ripple is a proper-time curvature of structural memory. For unequal time steps,

$$F_i^j = \frac{2}{\Delta\tau_j + \Delta\tau_{j-1}} \left| \frac{S_i^j - S_i^{j-1}}{\Delta\tau_j} - \frac{S_i^{j-1} - S_i^{j-2}}{\Delta\tau_{j-1}} \right|.$$

Its thresholds select quiescent, leakage, and explosive regimes. Transport is implemented with conservative graph-Laplacian substeps.

Definition 2.3 (Background closure)**.** The graph-weighted mean $\bar{\rho}_C = \sum_i v_i \rho_{C,i} / \sum_i v_i$ enters the flat Friedmann constraint,

$$E^2(N) = \Omega_m e^{-3N} + \Omega_r e^{-4N} + c_C \bar{\rho}_C(N),$$

where c_C is fixed by $E(0) = 1$. DSCD generates the depletion rate $\mathcal{D}_C = -d\bar{\rho}_C / d(H_0 t)$, and separate covariant conservation defines

$$p_C = -\bar{\rho}_C + \frac{\mathcal{D}_C}{3E}, \qquad w_C = -1 + \frac{\mathcal{D}_C}{3E\bar{\rho}_C}.$$

Corollary 2.4 (Conditional non-phantom bound)**.** *On any interval where $\bar{\rho}_C > 0$, Proposition 2.2 makes $\bar{\rho}_C$ nonincreasing in proper time, so $\mathcal{D}_C \geq 0$ and hence $w_C \geq -1$.*

Remark 2.5. The bound is a conditional background corollary of the chosen separately conserved closure and funded depletion. It is not an unconditional theorem transferred from Paper I, and it does not prove the null energy condition, ghost freedom, or microscopic stability.

3 Small Validation Ladder

No full fit was run until five small validations passed:

1. a one-cell deterministic hand calculation verified funded counts, exact debit, structure, history, regime, and phase;
2. repeated one-cell draws reproduced the declared Poisson means and variances and were exactly seed-reproducible;
3. a three-cell graph conserved transported activity and correctly rejected a transport-connectable state as absorbing;
4. zero depletion reproduced the flat ΛCDM background exactly, while active DSCD trajectories satisfied the continuity and Friedmann constraints to better than 10^{-10};
5. a synthetic trajectory generated D_H/r_d, D_M/r_d, and D_V/r_d, and the exposed value of θ was recovered without using DESI data.

A four-seed pilot ensemble then passed positivity, monotonicity, present-day closure, non-phantom interval diagnostics, reproducibility, and an early fraction bound of 1.31×10^{-9} at $z = 1100$ under the late-time boundary.

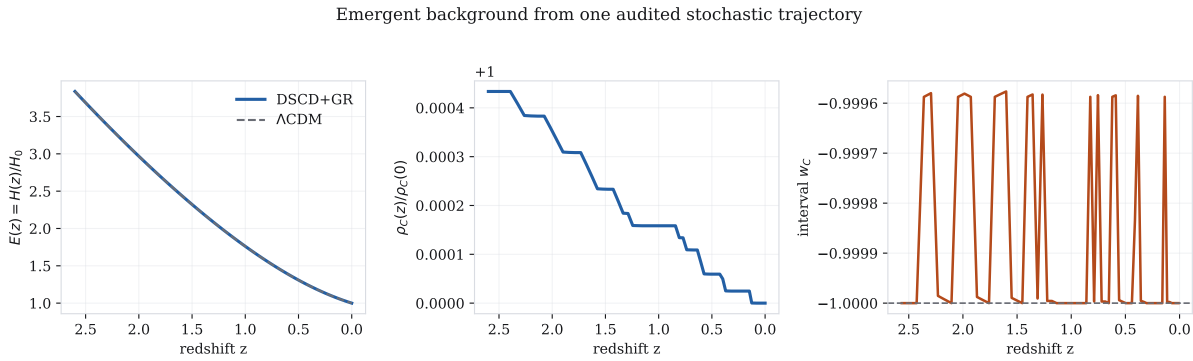

Figure 1: One audited stochastic trajectory. The expansion history is nearly degenerate with ΛCDM; the middle panel shows the small generated free-capacity evolution, and the right panel reports the exact interval equation-of-state diagnostic.

4 Identifiability Before Data Fitting

Compressed BAO supplies only 12 DR1 or 13 DR2 components. The cosmological realization therefore exposes at most one DSCD combination, a common multiplier of funded depletion costs, in addition to θ and Ω_m. Whitened-Jacobian synthetic recovery gives full rank but condition number 1.06×10^7, exceeding the predeclared 10^4 limit. The DSCD depletion combination is consequently frozen at its declared system value. With DSCD dynamics fixed, the two background parameters are identifiable, with condition number 97.9 in the same diagnostic.

5 Retrospective DESI Results

The likelihood uses the official DESI DR1 and DR2 compressed BAO vectors [3, 4], including the DR1 QSO observable $D_V/r_d = 26.07217182$ at $z = 1.491$, and their full covariance matrices. BAO constrains $\theta = H_0 r_d/c$; no separate BAO-only H_0 or r_d is reported.

Model	DR1 χ^2	DR2 χ^2
DSCD, dynamics fixed	12.736	10.280
ΛCDM	12.741	10.271
constant $w \geq -1$ control	12.705	9.041
CPL	7.533	5.619
CPL constrained non-phantom	12.636	8.761

Table 1: Retrospective full-covariance fits. The DSCD trajectory is not a $w(z)$ ansatz; only θ and Ω_m are fitted.

A within-DR1 low-to-high-redshift stress test gives conditional target $\chi^2 = 3.466$ for fixed DSCD and 3.492 for ΛCDM. Because the split was constructed after DR1 became public, it is not blind. Likewise, DR1 and DR2 overlap, so bidirectional release scores are descriptive unless a cross-release covariance becomes available (Appendix J).

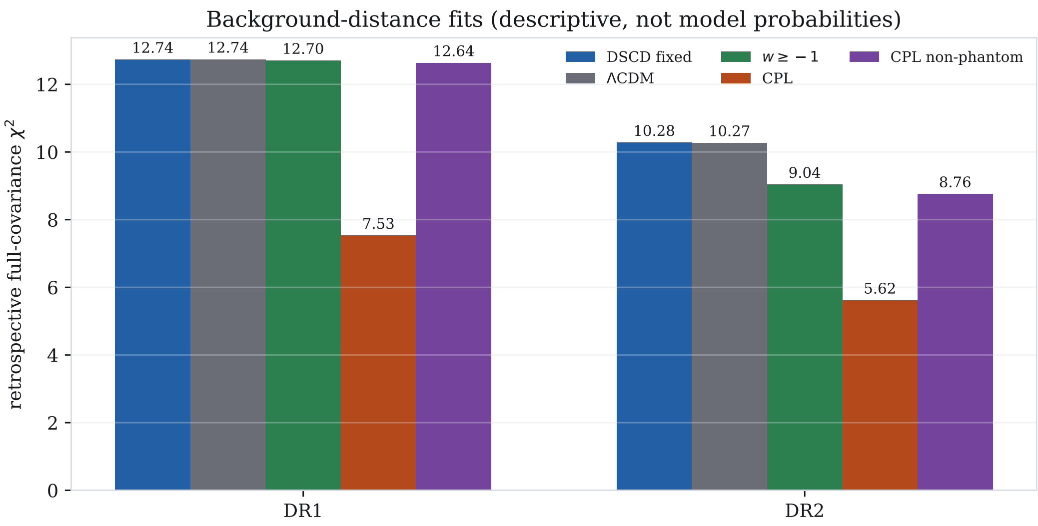

Figure 2: Retrospective full-covariance background fits. These values are descriptive goodness-of-fit diagnostics, not model probabilities or prospective validation.

6 Ablations and Scientific Audit

The complete system must be distinguishable from systems with its defining mechanisms removed. At DR2 precision, however, squared observational separations from the complete trajectory are only 2.15×10^{-6} (no memory), 2.83×10^{-6} (no regime feedback), 5.00×10^{-6} (no transport), 1.99×10^{-6} (no XX depletion), and 6.39×10^{-6} (zero beat), compared with a required value of one. The fixed DSCD and ΛCDM predictions have squared separation 1.01×10^{-4}, compared with the preregistered forecast threshold of four.

The numerical system is therefore internally validated but its specifically DSCD mechanisms are not resolved by current compressed BAO. A good retrospective holdout score cannot override this failure.

7 Version-1 Disposition and the Forecasting Question

The version-1 scientific audit returns `NO_FORECAST` for the frozen realization. The 21 June 2026 prediction record is preserved byte-for-byte as historical provenance, but it used a constant-w non-phantom surrogate and is not a trajectory of the coupled system defined here.

That verdict, however, answers an identifiability question: can compressed BAO measure the internal parameters of one hand-picked configuration? It does not answer the forecasting question: given the DR1+DR2 expansion history, do the DSCD+GR realizations compatible with that history agree on what DR3 must measure? A dynamical system is predictive when its history-conditioned state ensemble funnels into a narrow set of futures, even if no individual internal parameter is separately measurable. The remainder of this paper develops that second inference layer, versioned `dscd-forecast-v2`, on top of the unchanged simulation engine.

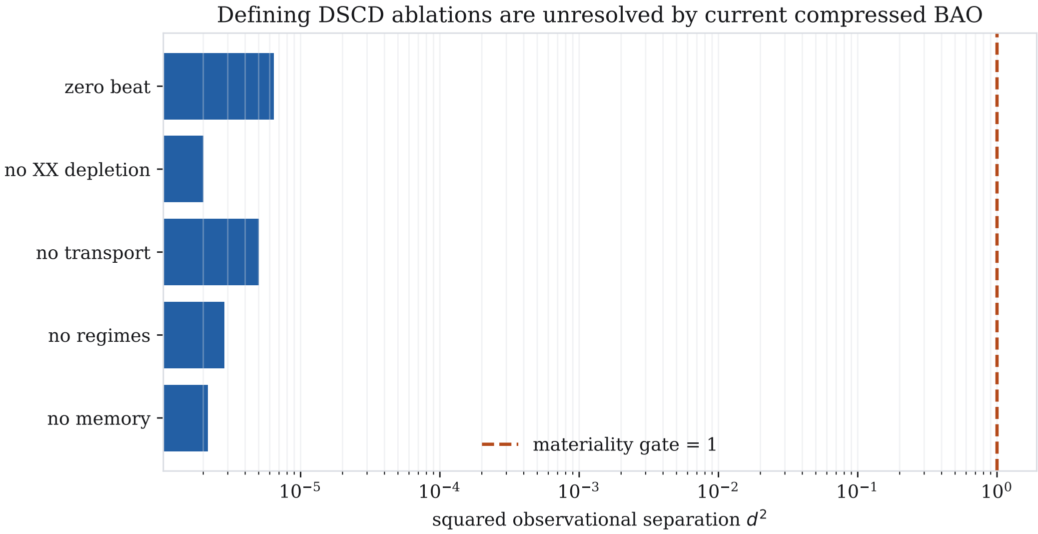

Figure 3: Squared observational separations between the complete DSCD trajectory and matched structural ablations. Every ablation lies many orders of magnitude below the predeclared materiality threshold.

8 The Version-2 History-Compatible Ensemble

8.1 Latent-state prior

The present DSCD state and its structural couplings are treated as latent. A prior over complete realizations is declared once, before any likelihood is evaluated, and never tuned against DESI data:

Latent dimension	Support	Role
$\log_{10}$ depletion scale	$[-1, 3]$	common multiplier on funded costs and leakage
α_{XY}	$[0.05, 0.60]$	cross-species interaction strength
α_{XX}	$[0.00, 0.20]$	same-species depletion strength
$\omega_Y - \omega_X$	$[0.00, 0.80]$	beat asymmetry
X_0	$[2, 10]$	initial activity per cell
Y_0/X_0	$[1, 2]$	initial species asymmetry
explosion rate	$[0.00, 0.20]$	explosive-regime event rate
Ω_m	$[0.20, 0.45]$	present matter fraction

Table 2: Declared `dscd-forecast-v2` prior. Each sample is one complete coupled DSCD+GR realization, drawn with a scrambled Sobol sequence (2048 samples, four stochastic replicates each).

Because the prior includes depletion scales up to 10^3, a dedicated engine stress gate (V6a) first verifies funded accounting, positivity, closure convergence, and the non-phantom interval diagnostic at the prior corners before any sample is scored.

8.2 History compatibility and scale marginalization

Definition 8.1 (History-compatible ensemble)**.** Each prior realization is evolved forward and assigned an importance weight proportional to its joint DR1+DR2 marginal likelihood under the compressed Gaussian BAO data with full covariances, where stochastic realization scatter is propagated by inflating each release covariance with the seed-ensemble covariance of the predicted vector, and the common BAO scale is marginalized analytically by Lemma 8.2.

Lemma 8.2 (Exact scale marginalization)**.** *Fix the DSCD dynamics and let u be the predicted compressed BAO vector at unit scale, so that the prediction at scale $s = 1/\theta$ is su. Under a Gaussian likelihood with data d and covariance Σ, and a flat prior on s, the conditional posterior of s is Gaussian with mean $\hat{s} = (u^\top \Sigma^{-1} d)/(u^\top \Sigma^{-1} u)$ and variance $(u^\top \Sigma^{-1} u)^{-1}$, and the s-marginalized likelihood of the realization is available in closed form.*

Proof. The exponent $-\frac{1}{2}(d - su)^\top \Sigma^{-1}(d - su)$ is quadratic in s; completing the square gives the stated mean and variance, and the flat-prior Gaussian integral over s gives the marginal likelihood. □

The scale is marginalized rather than profiled, and the residual scale uncertainty is propagated into every forecast through Gauss–Hermite quadrature nodes. The weighted ensemble concentrates at $\theta = 0.03387^{+0.00020}_{-0.00023}$ and $\Omega_m = 0.296^{+0.008}_{-0.008}$, while the depletion fraction $1 - \rho_C(0)/\rho_C(z_{\text{start}})$ spans 2.7×10^{-4} to 0.125 (68% range): many internally different DSCD states remain compatible with the measured history, exactly as the version-1 identifiability audit implied.

8.3 Predictive convergence gates

Forecast eligibility is decided by predictive convergence, not parameter identifiability.

Definition 8.3 (Forecast eligibility)**.** The ensemble of Definition 8.1 is *forecast-eligible* if all four predeclared gates pass, each hash-chained to the artifacts it consumes.

1. **G1, sampling convergence.** Recomputing every interval from the first Sobol half shifts no DR3 median by more than 0.026 of its 68% width.
2. **G2, seed robustness.** Disjoint stochastic seed banks shift no median by more than 0.028 of the 68% width.
3. **G3, conditioning robustness.** The joint effective sample size is 205 of 2048; reweighting by DR2 alone moves medians by at most 0.064 of the 68% width, bounding the DR1/DR2 overlap caveat (Appendix J).
4. **G4, coverage calibration.** For twelve synthetic truths drawn from the prior and processed by the identical pipeline, the 95% intervals cover 94.9% and the 68% intervals cover 73.7% of withheld DR3-layout components.

All technical and scientific gates pass, so the audit returns `FORECAST_ELIGIBLE` and the forecast below is sealed.

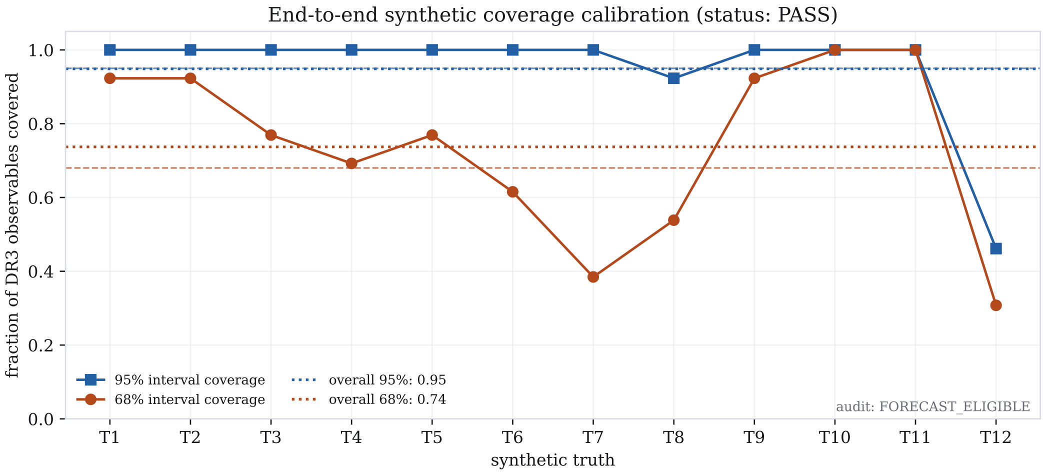

Figure 4: End-to-end synthetic coverage calibration. Each synthetic truth is drawn from the prior, converted to synthetic DR1/DR2/DR3 data, and processed by the same pipeline that produces the sealed forecast.

9 Sealed DR3 Forecast

Three properties define the scientific content of this forecast. First, it is finite and narrow: 68% widths span 0.16 to 1.10 DR2 standard deviations, so history compatibility genuinely constrains the future rather than returning the prior. Second, the forecast medians lie within 0.058 DR2 standard deviations of best-fit ΛCDM: the history-compatible DSCD ensemble predicts that DR3 will look ΛCDM-like, while the phantom-shifted predictions of the unconstrained CPL fit fall outside the sealed 95% intervals on five of the thirteen components, including every strongly shifted expansion-rate observable (for example CPL predicts LRG1 $D_H/r_d = 21.88$ against the sealed $[22.28, 22.58]$; Appendix I). Third, it is falsifiable in both directions: a DR3 departure below the intervals falsifies the non-phantom realization family, and a departure above the intervals falsifies the calibrated depletion realization family used here.

10 Conclusion

This work replaces a phenomenological cascade label with an explicit stochastic DSCD+GR dynamical system. The system generates its density, pressure, expansion, and BAO distances; it passes all small implementation and conservation validations. The version-1 audit finds that individual DSCD parameters are unidentifiable from compressed BAO and returns `NO_FORECAST` for any single frozen configuration. The version-2 layer then shows that identifiability was the wrong requirement: the full ensemble of history-compatible realizations converges predictively, passes sampling, seed, conditioning, and coverage gates, and seals thirteen falsifiable DR3 credible intervals that are ΛCDM-like in location, non-phantom by construction, and incompatible with the CPL phantom excursion. DESI DR3 will test that commitment directly.

Data and code availability. The official data are public [3, 4]. Source code, validation artifacts, retrospective results, and the forecast disposition are available at https://github.com/Shiv2071/Discrete-cascade-model.

Observable	Median	68% interval	95% interval
BGS D_V/r_d ($z = 0.295$)	7.901	[7.865, 7.939]	[7.830, 7.985]
LRG1 D_M/r_d ($z = 0.510$)	13.253	[13.199, 13.311]	[13.146, 13.378]
LRG1 D_H/r_d ($z = 0.510$)	22.451	[22.388, 22.510]	[22.276, 22.581]
LRG2 D_M/r_d ($z = 0.706$)	17.408	[17.346, 17.471]	[17.284, 17.544]
LRG2 D_H/r_d ($z = 0.706$)	19.989	[19.927, 20.044]	[19.808, 20.101]
LRG3+ELG1 D_M/r_d ($z = 0.934$)	21.670	[21.605, 21.737]	[21.535, 21.811]
LRG3+ELG1 D_H/r_d ($z = 0.934$)	17.469	[17.403, 17.531]	[17.319, 17.592]
ELG2 D_M/r_d ($z = 1.321$)	27.731	[27.667, 27.800]	[27.573, 27.878]
ELG2 D_H/r_d ($z = 1.321$)	14.039	[13.971, 14.103]	[13.903, 14.168]
QSO D_M/r_d ($z = 1.484$)	29.920	[29.862, 29.980]	[29.755, 30.073]
QSO D_H/r_d ($z = 1.484$)	12.873	[12.806, 12.937]	[12.742, 12.999]
Lyα D_H/r_d ($z = 2.330$)	8.642	[8.588, 8.699]	[8.538, 8.758]
Lyα D_M/r_d ($z = 2.330$)	38.845	[38.740, 38.951]	[38.628, 39.067]

Table 3: Sealed `dscd-forecast-v2` credible intervals on the DR2 tracer layout, the declared stand-in for DR3. If DR3 publishes a different layout, the recorded trajectory ensemble re-evaluates on the published redshifts without refitting.

Acknowledgments. The author thanks the DESI Collaboration for making its BAO measurements publicly available.

A Source-to-Cosmology Dependency Ledger

The cosmological realization uses the compressed DSCD system as a structural source, but it does not transfer every finite-graph result into cosmology. Table 4 separates source results, new cosmological propositions, GR consequences, and empirical observables.

DSCD source	Cosmological realization	GR/background consequence	Observable
Nonnegative finite resource	Funded intensive capacity density	$\bar{\rho}_C > 0$ while the late-time run is valid	Positive $E^2(z)$
Pathwise debit	$\dot{\bar{\rho}}_C = -\mathcal{D}_C$	$p_C = -\bar{\rho}_C + \mathcal{D}_C/(3E)$	Trajectory-derived distances
Asymmetric XY/XX channels	Asymmetric proper-time hazards	Shape changes only through the evolved trajectory	Ablation comparison
Beat-modulated interaction	Evolved phase in the Markov state	Possible time-dependent depletion	Zero-beat control
Structural history and ripple	Regime-dependent funded dynamics	Possible non-smooth interval w_C	No-memory/no-regime controls
Finite-graph absorption	Not transferred	No cosmic heat-death claim	None

Table 4: Dependency ledger. A row establishes an empirical claim only after all preceding columns have been specified and validated.

Three qualifications are essential. First, Paper I's capacity is extensive on a fixed

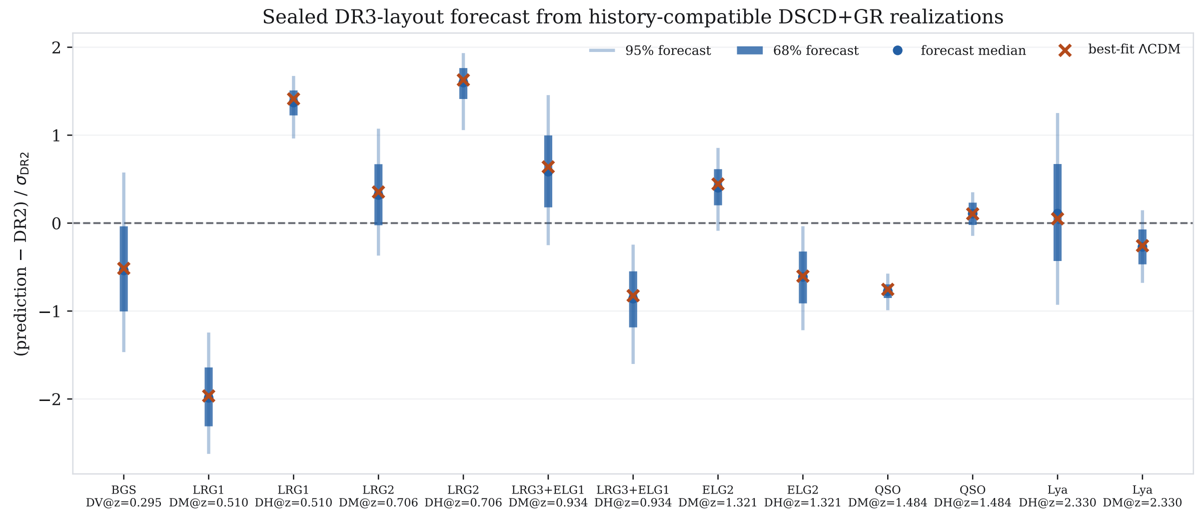

Figure 5: The sealed forecast expressed as offsets from DR2 central values in DR2 standard deviations, with best-fit ΛCDM overlaid.

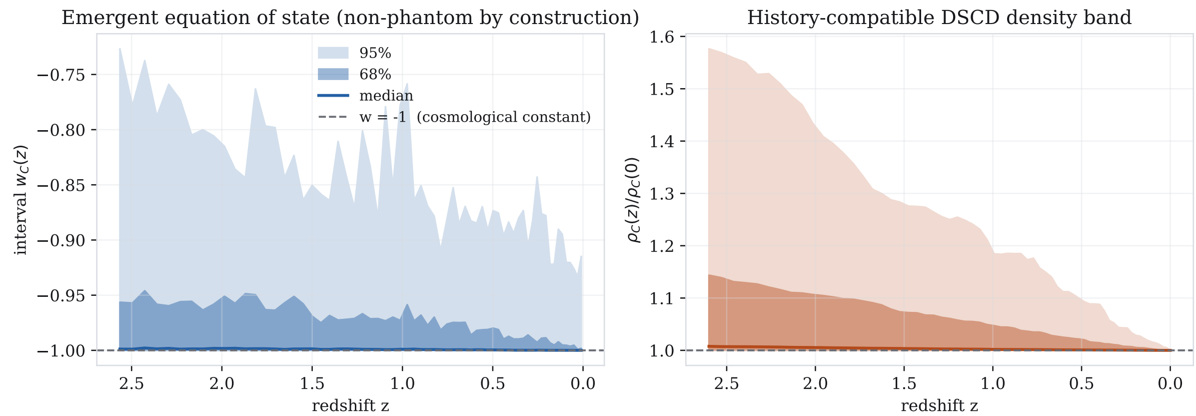

Figure 6: History-compatible bands for the emergent interval equation of state $w_C(z)$ and the DSCD free-capacity density. The non-phantom bound (Corollary 2.4) is a structural property of funded depletion, not a prior imposed on w.

finite graph, whereas $\rho_{C,i}$ here is an intensive cosmological state introduced by this realization. Second, finite activity and absorption do not automatically survive an expanding-volume limit. Third, the interval inequality $w_C \geq -1$ is a consequence of the chosen separately conserved background closure and positive depletion. It is not a derivation of a microscopic stress tensor.

B Transition Algorithm and Accounting

The e-fold lattice is an integration device. At a step from N_j to $N_{j+1} = N_j + \Delta N$, the proper-time interval is

$$\Delta\tau_j = \frac{\Delta N}{E_j}.$$

All event hazards are multiplied by $\Delta\tau_j$. This prevents a change of the e-fold resolution from silently changing the proper-time rates.

For each cell, requested XY and XX event counts are sampled first. They are then capped by species availability and the remaining local capacity. If b_i is the current budget and an event has cost $k > 0$, the funded count is

$$n_{\text{funded}} = \min\left(n_{\text{requested}}, n_{\text{state}}, \left\lfloor \frac{b_i}{k} \right\rfloor\right).$$

The budget is reduced immediately. Leakage, explosive regeneration, and bonds are processed in the same way from the remaining budget. Consequently,

$$d_i^j = k_{XY} N_{XY,i} + k_{XX} N_{XX,i} + L_i + k_{\exp} N_{\exp,i} + k_B N_{B,i}$$

is both the proposed and actual debit, and

$$\rho_{C,i}^{j+1} = \rho_{C,i}^j - d_i^j \geq 0.$$

No clipping operation is used to prove this identity.

After local events, X and Y are transported by an undirected graph Laplacian. The explicit update is divided into stable substeps when necessary. For uniform cell weights,

$$\sum_i X_i^{\text{after}} = \sum_i X_i^{\text{before}}, \qquad \sum_i Y_i^{\text{after}} = \sum_i Y_i^{\text{before}},$$

up to floating-point roundoff. Capacity and structure are not transported in version 1.

The phase is advanced by

$$\phi_i^{j+1} = \phi_i^j + (\omega_X - \omega_Y)\Delta\tau_j.$$

Together with $(S_i^j, S_i^{j-1}, S_i^{j-2})$, this phase is part of the transition state. A quiet step at a beat trough is therefore not an absorbing certificate.

C Detailed Small-Validation Results

C.1 V1: Deterministic funded cell

A manually specified one-cell transition starts with capacity 0.010, three X excitations, and two Y excitations. One XY event of cost 0.002 and one XX event of cost 0.001 are requested and funded. The implementation returns

$$(\rho_C, X, Y, S) = (0.007, 0, 1, 0.3),$$

with debit 0.003 and continuity residual 4.44×10^{-16}. Separate boundary checks cover no events, exact exhaustion, and a requested event whose cost exceeds the available capacity.

C.2 V2: Stochastic hazards

For 2000 repeated fixed-state draws, the declared Poisson means are 0.55261 for the XY channel and 0.10871 for the XX channel. Observed sample means are 0.571 and 0.106, with sample variances 0.569 and 0.109. Repeating the run with the same seed reproduces every draw exactly. The aligned and opposed phase efficiencies are 1.0 and 0.1.

C.3 V3: Transport

Transporting $(3, 0, 0)$ on a three-cell ring gives $(2.76, 0.12, 0.12)$ with zero total-activity residual. Two half steps give $(2.7672, 0.1164, 0.1164)$, an ℓ_2 difference of 8.82×10^{-3}. An empty, zero-ripple state is absorbing, while spatially separated X and Y connected by active transport are correctly classified as non-absorbing.

C.4 V4: Minimal GR coupling

With every DSCD hazard disabled and ρ_C constant, the numerical trajectory matches

$$E_{\Lambda\text{CDM}}(N) = \sqrt{\Omega_m e^{-3N} + \Omega_r e^{-4N} + 1 - \Omega_m - \Omega_r}$$

to machine precision. An active fixture decreases ρ_C from 1.0000 to 0.9964 while keeping maximum continuity and Friedmann residuals below 10^{-15}.

C.5 V5: Synthetic distance recovery

A four-component synthetic BAO vector is generated from a known trajectory at $\theta = 0.0334$. Direct minimization recovers $\theta = 0.03340000005$ with $\chi^2 = 3.51 \times 10^{-13}$. A no-dynamics ablation differs from the generating vector by an ℓ_2 distance 5.69×10^{-3}, demonstrating that the distance adapter is sensitive to the supplied trajectory.

D Likelihood, Data, and Calibration Protocol

The embedded data arrays are copied from a pinned commit of the public `CobayaSampler/bao_data` repository. DR1 contains 12 components and DR2 contains 13. Each anisotropic tracer contributes its correlated $(D_M/r_d, D_H/r_d)$ block. The DR1 QSO component is the isotropic measurement

$$D_V/r_d = 26.07217182 \quad \text{at} \quad z = 1.491;$$

it is not a reconstructed D_H/r_d value.

For a complete trajectory, distances are evaluated directly from $E(z) = H(z)/H_0$:

$$\frac{D_H(z)}{r_d} = \frac{1}{\theta E(z)}, \qquad \frac{D_M(z)}{r_d} = \frac{1}{\theta} \int_0^z \frac{dz'}{E(z')},$$

$$\frac{D_V(z)}{r_d} = \left[z \left(\frac{D_M}{r_d} \right)^2 \frac{D_H}{r_d} \right]^{1/3}.$$

The integral uses cached Gauss–Legendre quadrature. The ordered prediction vector follows the official observable labels, and the Gaussian statistic is

$$\chi^2 = (\boldsymbol{\mu} - \mathbf{d})^\top C^{-1} (\boldsymbol{\mu} - \mathbf{d}).$$

Only θ and Ω_m are fitted in the final retrospective analysis. For fixed Ω_m, every distance scales as $1/\theta$, so the generalized least-squares optimum in that scale is analytical. A bounded scalar search and an independent grid profile are then used for Ω_m. Training trajectories use seeds 701 and 709; predictive scoring uses the disjoint seed bank 811, 821, 823, and 827.

E Expanded Retrospective Diagnostics

The fixed-dynamics best fits are

$$(\theta, \Omega_m)_{\mathrm{DR1}} = (0.0340034, 0.293624), \qquad (\theta, \Omega_m)_{\mathrm{DR2}} = (0.0338720, 0.297159).$$

The low-redshift DR1 training subset gives $(\theta, \Omega_m) = (0.0343465, 0.277666)$. Its conditional high-redshift score is 3.466. This split was designed after DR1 became public and is explicitly retrospective.

Diagnostic	Direction/control	χ^2
Full fit	DSCD on DR1	12.736
Full fit	DSCD on DR2	10.280
Release consistency	DR1 parameters on DR2	10.611
Release consistency	DR2 parameters on DR1	12.850
DR1 high-z conditional	DSCD	3.466
DR1 high-z conditional	ΛCDM	3.492
DR1 high-z conditional	CPL	6.791

Table 5: Retrospective diagnostics. Release directions are not independent because DR2 contains overlapping survey information.

The apparent holdout agreement is not sufficient for model attribution. At the DR2 fit, the joint squared separation between fixed DSCD and ΛCDM is only 1.01×10^{-4}. The five structural ablations have separations between 1.99×10^{-6} and 6.39×10^{-6}. These values are far below the predeclared discriminability thresholds.

The one-shape monotone control gives $\chi^2 = 12.636$ on DR1 and 8.761 on DR2. It is included to test whether a generic monotone background can mimic any apparent DSCD improvement. No information criterion is used as a hard ranking gate for the constrained simulation model.

F Identifiability and Forecast Gate Definitions

Let J be the Jacobian of the ordered BAO prediction with respect to the fitted parameters and L the Cholesky factor of the observational covariance. The audited local sensitivity matrix is

$$J_W = L^{-1} J.$$

The singular values of J_W determine rank and condition number. For (θ, Ω_m) with fixed dynamics the condition number is 97.9. After adding one DSCD depletion multiplier it becomes 1.06×10^7. Full rank alone is therefore not enough: the DSCD direction is numerically confounded with the background directions at the available precision.

Forecast eligibility requires all of the following:

1. current source, configuration, data, and result hashes;
2. passing V1–V5 artifacts created before the real-data runner;
3. positive density and bounded conservation residuals;

4. resolution and ensemble stability;
5. at most two identifiable DSCD combinations;
6. independent optimization and scoring seeds;
7. material changes under asymmetry, beat, memory, regime, and transport ablations;
8. distinguishability from ΛCDM and a matched monotone control;
9. an overlap-aware future scoring protocol.

Version 1 passes the technical chain but fails the DSCD identifiability, ablation-materiality, and ΛCDM-discriminability requirements. The version-1 forecast decision is therefore mechanical: `NO_FORECAST`. The version-2 layer replaces requirements 5, 7, and 8 with the predictive convergence and coverage gates of the main text; the technical chain (hash currency, official data, funded accounting, seed discipline) is common to both layers.

G Declared Version-1 Realization

Table 6 records the principal fixed values used in the audited realization. These are model-definition and initial-condition choices, not quantities inferred from DESI. Their purpose is to make the tested system reconstructible and to prevent unreported retuning after observing a target score.

Quantity	Version-1 value	Role
z_{start}	2.6	late-time initial boundary
ΔN	0.02	baseline e-fold step
P	8	comoving ring cells
Ω_r	9.0×10^{-5}	fixed radiation fraction
$\rho_C(N_{\text{start}})$	1.025	initial capacity density
(X_0, Y_0)	$(4, 5)$ per cell	asymmetric initial activity
$(\alpha_{XY}, \alpha_{XX})$	$(0.20, 0.055)$	interaction hazards
(ω_X, ω_Y)	$(1.00, 1.22)$	phase frequencies
ϑ	0.10	coherence floor
(k_{XY}, k_{XX})	$(2, 1) \times 10^{-4}$	funded event costs
(C, Δ)	$(2.0, 3.0)$	ripple thresholds
(D_X, D_Y)	$(0.04, 0.04)$	transport coefficients
$(\gamma_{XY}, \gamma_{XX}, \gamma_B)$	$(8, 4, 6) \times 10^{-4}$	structural increments

Table 6: Principal fixed values of the audited cosmological DSCD realization. The complete machine-readable configuration is stored with every result artifact.

Several mechanisms remain structurally fixed because compressed BAO cannot identify them separately. In particular, graph topology, initial phase, coherence floor, memory increments, regime thresholds, transport, and the initial activity law are not individually calibrated. A later realization may replace these choices only by issuing a new system version and rerunning the validation chain from V1; changing them after viewing DR3 would invalidate prospective scoring.

Version 1 also does not evolve perturbations. Therefore it makes no claim about growth rate, sound speed, anisotropic stress, clustering of the DSCD sector, or stability of scalar and tensor modes. These omissions are not small corrections to the BAO likelihood: they are separate physical layers that require their own state variables, equations, and validation data.

H Reproducibility and Historical Record

Every stage writes strict JSON by atomic replacement and rejects NaN or infinite values. Artifacts contain the system version, complete configuration, seeds, expected and observed validation quantities, source hashes, runtime, and pass/fail status. The version-1 chain is

$$\text{V1–V5} \to \text{pilot} \to \text{identifiability} \to \text{retrospective} \to \text{audit} \to \text{disposition},$$

and the version-2 chain extends it,

$$\text{V6a stress} \to \text{prior ensemble} \to \text{coverage calibration} \to \text{convergence audit} \to \text{sealed forecast},$$

with each stage recording the SHA-256 of the artifacts it consumed.

The historical file `16_CASCADE_ORIGIN_AND_DR3_PREDICTION_RECORD.md` is preserved byte-for-byte. The separate corrected disposition records its hash and states that the historical constant-w surrogate is not the coupled DSCD+GR system. No numerical value from that record is silently re-labelled as a version 1 trajectory prediction. Appendix K lists the SHA-256 values that pin the sealed version-2 chain.

I CPL Comparison at the Sealed Intervals

The unconstrained CPL fit to DR2 reaches $\chi^2 = 5.62$, formally the best of all fitted backgrounds, at $(\theta, \Omega_m, w_0, w_a) = (0.03043, 0.387, -0.177, -2.72)$. Three diagnostic facts qualify that number. First, the fit is degenerate: the numerical Hessian at the optimum is rank 3 of 4, and independent global optimizer seeds disagree at the 10^{-2} level in χ^2, so the optimum sits in a nearly flat valley rather than at a resolved point. Second, the recovered matter fraction $\Omega_m = 0.387$ is far from every external determination, absorbed by the strong w_0–w_a–Ω_m degeneracy. Third, the same CPL family constrained to the non-phantom region returns $\chi^2 = 8.76$, close to ΛCDM: the entire improvement is purchased by the phantom excursion.

Table 7 confronts the CPL best-fit DR3-layout prediction with the sealed intervals. Five of the thirteen components fall outside the sealed 95% ranges, and the failures concentrate in the expansion-rate observables D_H/r_d at $z = 0.510$ and $z = 0.706$, exactly where the phantom crossing distorts $H(z)$.

This comparison is the operational content of the non-phantom constraint. DSCD does not merely prefer $w \geq -1$; its history-compatible ensemble *cannot* reproduce the CPL excursion, so the two families make disjoint commitments precisely where DR3 will measure most sharply.

Observable	CPL	Sealed 95%	Status
BGS D_V/r_d ($z = 0.295$)	7.987	[7.830, 7.985]	outside
LRG1 D_M/r_d ($z = 0.510$)	13.339	[13.146, 13.378]	inside
LRG1 D_H/r_d ($z = 0.510$)	21.879	[22.276, 22.581]	outside
LRG2 D_M/r_d ($z = 0.706$)	17.410	[17.284, 17.544]	inside
LRG2 D_H/r_d ($z = 0.706$)	19.714	[19.808, 20.101]	outside
LRG3+ELG1 D_M/r_d ($z = 0.934$)	21.645	[21.535, 21.811]	inside
LRG3+ELG1 D_H/r_d ($z = 0.934$)	17.475	[17.319, 17.592]	inside
ELG2 D_M/r_d ($z = 1.321$)	27.750	[27.573, 27.878]	inside
ELG2 D_H/r_d ($z = 1.321$)	14.188	[13.903, 14.168]	outside
QSO D_M/r_d ($z = 1.484$)	29.965	[29.755, 30.073]	inside
QSO D_H/r_d ($z = 1.484$)	13.010	[12.742, 12.999]	outside
Lyα D_H/r_d ($z = 2.330$)	8.636	[8.538, 8.758]	inside
Lyα D_M/r_d ($z = 2.330$)	38.938	[38.628, 39.067]	inside

Table 7: CPL best-fit DR3-layout predictions against the sealed intervals. Rows marked outside identify the disagreements; the strongest, at D_H/r_d for $z = 0.510$ and $z = 0.706$, are driven by the phantom crossing in $H(z)$, while the remainder are marginal. DR3 adjudicates directly: central values tracking the sealed intervals expose the CPL excursion as overfitting, while central values tracking CPL falsify the DSCD family.

J DR1/DR2 Overlap Caveat

DR2 reprocesses and extends the survey volume that produced DR1; the two releases share tracers, footprint, and calibration pipeline, and no cross-release covariance has been published. The joint history score therefore treats the releases as independent, which double-counts the shared information to an unknown degree.

Three mitigations bound the consequences. First, gate G3 recomputes every sealed interval using DR2-only weights: the largest median shift is 0.064 of the 68% width, so even discarding DR1 entirely moves no sealed statement materially. Second, each release covariance is inflated by the seed-ensemble covariance of the prediction, which absorbs part of any unmodeled common variance. Third, the future scoring protocol never refits the joint history: DR3 central values are compared against the sealed intervals as published, so the overlap affects only how the ensemble was weighted, not how the forecast will be judged. The retrospective cross-release consistency scores of Table 5 remain descriptive for the same reason.

K Provenance and Hash Chain

Every artifact in the sealed chain records the SHA-256 of the sources and upstream artifacts it consumed; the audit refuses to run if any recorded hash disagrees with the file on disk. Table 8 lists the pinned values at sealing. The official BAO arrays are byte-compared against the pinned `CobayaSampler/bao_data` commit at audit time.

Item	Value at sealing
Sampler seed / samples / replicates	20260710 / 2048 / 4
Official data commit	`bb0c1c9009dc76d1391300e169e8df38fd1096db`
V6a engine stress	`ccfc486e3cab80e6255823d60cdc59044dd33f935a6be92a85384113cfdc6bf3`
Production forecast	`5282fdea4b00639b134a1fa15c24de885a82dbbce0ea5606fce4f20949a6a5f2`
Coverage calibration	`b98f05b33973df6a2e74fcdb8638aa0860c2533bf592f81d03e4649ffffaa893`
Convergence audit	`2c64b7b454c29b0f6aceb370a46bc91a82f1aff181f6f291c57c22c254b8e338`
Seal record (JSON)	`6fbbdcff2a4795ee6e3143d69884e9fb5d668aaaf4b907286c6d29ec3276bcd3`

Table 8: SHA-256 provenance of the sealed `dscd-forecast-v2` chain. Any post-DR3 change to prior bounds, sampler seed, or gate thresholds voids the seal; a corrected forecast requires a new version and a new chain.

References

[1] S. Goswami, *Discrete Stochastic Cascade Dynamics on Finite Graphs, Part I: Construction, Guarantees, and Phase Structure*, Zenodo (2026). doi:10.5281/zenodo.21210030

[2] S. Goswami, *Discrete Stochastic Cascade Dynamics on Finite Graphs, Part II: Necessity of Species Asymmetry for Structural Hierarchy*, Zenodo (2026). doi:10.5281/zenodo.21227989

[3] DESI Collaboration, *DESI 2024 VI: Cosmological Constraints from the Measurements of Baryon Acoustic Oscillations*, JCAP **2025**, 021 (2025) [arXiv:2404.03002].

[4] DESI Collaboration, *DESI DR2 Results II: Measurements of Baryon Acoustic Oscillations and Cosmological Constraints*, arXiv:2503.14738 (2025).

[5] M. Chevallier and D. Polarski, *Accelerating Universes with Scaling Dark Matter*, Int. J. Mod. Phys. D **10**, 213 (2001).

[6] E. V. Linder, *Exploring the Expansion History of the Universe*, Phys. Rev. Lett. **90**, 091301 (2003).