Supplementary Material 2: Supplementary Proofs

Predictive Structure and Realized Forecast Gain in Nonlinear Fields

International Journal of Dynamics and Control

This supplement gives the analytic proofs and computational contracts for the fixed 257-state query-capacity proposition and the frozen-model mean-error propositions on 257 and fresh 360 states. Supplementary Evidence Archive SA1 contains the exact rational endpoints, pointwise bounds, codebook tables, and a summary-verification program. All displayed decimal intervals are rounded outward.

The target throughout is the designated 64-node ordinary differential equation (ODE). A certificate for this target does not certify the continuous reaction-diffusion partial differential equation. The capacity result concerns a fixed-width message transmitted for each query. The model comparison concerns specified saved codes and readers. These are separate propositions.

Role in the combined manuscript. These complete arguments are retained from the field study for separate coded-model tasks. They do not certify the new unquantized physical readers or any lattice information endpoint. The latter proof is in Supplement B8.

SP0. Physical target, arithmetic contract, and target enclosures

SP0.1. Designated ODE and initial states

There are 64 spatial nodes, indexed by j = 0,...,63. The Neumann cosine basis is C_n(j) = cos($n\pi j/63$), n = 0,...,63. The designated diffusion operator acts diagonally on this basis with eigenvalues -n^2/4 for X and -16 n^2/4 for Y. Reaction terms are evaluated at all 64 nodes. No high cosine mode or nodal alias is removed. At B = 25/16, write u = X - 1 and v = Y - B. The deviation vector field is F = L + N, with species blocks and nonlinear term

$$M_n = \begin{bmatrix} 9/16 - q_n & 1 \\ -25/16 & -1 - 16q_n \end{bmatrix}, \qquad q_n = n^2/4. \tag{1}$$

$$\mathcal{N}(u, v) = (Q, -Q), \qquad Q = Bu^2 + 2uv + u^2v. \tag{2}$$

The physical target Y_i is the 128-component deviation vector at T = 10. The numerical target y_i is the saved float32 future. Every saved float32 or float64 value denotes its exact binary rational value. The risk regularizer is

$$\eta = \frac{7737125245533627}{773712524553362671811952 64} > 0. \tag{3}$$

This is the exact stored binary64 constant used for 10^{-10}, rather than the distinct decimal rational 1/10^{10}.

The 257-state set consists of the first 257 source indices in ascending order from the fixed 273-state separation group. It includes source indices 3, 7, and 12. Its physical initial-state contract is the nearest-cell reconstruction from the saved float32 current, including the declared outward binary64 arithmetic allowances. For a finite float32 component x, its nearest-rounding preimage is enclosed by

$$r(x) = \frac{1}{2}\max\{x - \text{prev}(x), \text{next}(x) - x\}. \tag{4}$$

Ties and unequal spacings are included. The auxiliary initial error bounds enclose these component cells and reconstruction allowances; they do not recover an unavailable original Float64 initial array. The stated ODE target certificate is conditional on this declared cell and arithmetic contract.

The fresh 360 states instead use the exact original binary64 X0/Y0 values preserved with that dataset. Their recorded curve agrees bitwise with the original physical initial point, so its true initial discrepancy is zero. The integer checker uses a positive component upper bound 2^{-80}, which encloses zero without perturbing the initial point. Neural inputs remain the saved float32 current fields. This initial-state contract is separate from the 257-state contract.

SP0.2. Auxiliary energy and linear rates

Use trapezoidal weights w_0 = w_{63} = 1/2 and w_j = 1 otherwise. Let h_0 = h_{63} = 63 and h_n = 63/2 for the remaining modes. The species energy matrices are P_0 = (25/16, 25/32; 25/32, 1), P_1 = (307/225, 697/1125; 697/1125, 6937/5625), and P_n = I for n $\geq$ 2. The first-mode coordinate matrix T has rows (16/15, 16/75) and (-1/15, -16/75), and

$$P_1 = T^T \,\mathrm{diag}(281/256, 26)T. \tag{5}$$

Positive principal minors give m I < P_n < M I with m = 9/20 and M = 17/8. If c_n are the cosine coefficients, define

$$\|z\|_*^2 = \sum_{n=0}^{63} h_n c_n^T P_n c_n. \tag{6}$$

Discrete cosine orthogonality gives m $\|z\|_W^2 \leq \|z\|_*^2 \leq$ M $\|z\|_W^2$, while $\|z\|_2 \leq \sqrt{2}\,\|z\|_W$. Consequently

$$\|z\|_2 \leq \sqrt{40/9}\|z\|_*. \tag{7}$$

The uniform block satisfies M_0^T P_0 + P_0 M_0 = -(7/16) P_0. The first-mode transformation gives T M_1 = diag$(0, -75/16)$ T; its critical and fast directions are orthogonal in P_1. For n $\geq$ 2, sym$(M_n) \leq$ -(q_n-5/8) I: the positive matrix (5/8) I - sym(M_0) has determinant 23/1024. Energy differentiation therefore gives 65 mutually orthogonal channels, ordered as uniform, first critical, first fast, and n = 2,...,63, with norm-decay rates

$$\alpha = (7/32, 0, 75/16, 3/8, 13/8, \ldots, 7933/8). \tag{8}$$

Their total dimension is 128. This calculation retains stable-mode decay and the nonnormal species geometry; it does not supply physical labels to the trained encoders.

SP0.3. Whole-interval residual and nonlinear coupling

A C^1 cubic Hermite curve is constructed from saved binary64 physical values and native slopes at 4, 001 canonical boundaries. Its mathematical times are t_j = j/400, and h = 1/400. The slopes define the curve; they are not assumed to equal the exact ODE right-hand side. Thus residuals also account for the discrepancy between the native numerical time arithmetic and this exact time grid. On one interval, with θ in [0,1],

$$\widehat{U}_j = H_{00}C_j + hH_{10}F_j + H_{01}C_{j+1} + hH_{11}F_{j+1}. \tag{9}$$

The Hermite basis is H_{00} = 2 θ^3 - 3 θ^2 + 1, H_{10} = θ^3 - 2 θ^2 + θ, H_{01} = -2 θ^3 + 3 θ^2, and H_{11} = θ^3 - θ^2. Subtract the equilibrium to obtain the deviation curve V(t). The residual r = V' - F(V) has degree at most 9 in θ. The nonlinear gradient has degree at most 6:

$$g_u = 2Bu + 2v + 2uv = 2XY - 2B, \qquad g_v = 2u + u^2 = X^2 - 1. \tag{10}$$

For a degree-d polynomial, its power coefficients a_k give Bernstein coefficients

$$b_i = \sum_{k=0}^{i} a_k \frac{\binom{i}{k}}{\binom{d}{k}}. \tag{11}$$

The Bernstein basis is nonnegative and sums to one on [0,1], so the convex hull of coefficient intervals encloses the entire interval. This produces whole-interval residual bounds $R_a \geq \|\Pi_a r\|_*$ and coupling bounds $K_{ab} \geq$ ‖Π_a DN(V) Π_b‖, rather than sampled maxima.

Cosine multiplication retains the highest-mode alias. With ell_- = |n-m| and ell_+ = min(n+m,126-n-m), the modal multiplication coefficient is

$$H_{nm}(g) = \frac{h_{\ell_-}\widehat{g}_{\ell_-} + h_{\ell_+}\widehat{g}_{\ell_+}}{2h_n}. \tag{12}$$

The reaction Jacobian is the rank-one species block (1,-1)^T (g_u,g_v). The input dual and output energy quadratic forms, maximized over coefficient rectangles, give K_{ab}; for a positive quadratic form the maximum on a rectangle occurs at a vertex. Critical and fast directions are combined before taking their hull. No high-mode-to-critical coupling is set to zero.

For a candidate energy tube $\|z\|_* \le \rho$, each nodal perturbation has Euclidean norm below 2.11 ρ. Let d = 2.11 ρ and let r_u,r_v bound |u|,|v| on the curve interval. Then

$$d_u = (2B + 2 + 2r_u + 2r_v)d + 2d^2, \qquad d_v = (2 + 2r_u)d + d^2. \tag{13}$$

The species and energy norm bounds imply

$$\|DN(V + z) - DN(V)\|_{*\to *} \le 3.1\sqrt{d_u^2 + d_v^2} = \sigma(\rho), \tag{14}$$

because $\sqrt{2M/m}$ = $\sqrt{85/9}$ < 3.1. The comparison matrix uses C_{ab} = K_{ab} + σ for every block. This is a conservative blockwise bound and includes all cross-couplings.

SP0.4. Certified propagation and existence

Let Z_a bound the error in channel a at the interval start. For α_a > 0 put κ_a = min(h,1/α_a); for α_a = 0 put κ_a = h. Put d_a = 1/(1+α_a h). The integral semigroup kernel is bounded by κ throughout the interval, and $\exp(-\alpha_a h) \le d_a$ follows from exp(x) ≥ 1+x. These are separate bounds; κ is not replaced by (1-d_a)/α_a.

Variation of constants, while the solution remains inside the tube, gives

$$E_a(t_j + s) \le e^{-\alpha_a s} Z_a + \int_0^s e^{-\alpha_a(s-t)} \left[\sum_b C_{ab} E_b(t_j + t) + R_a \right] dt. \tag{15}$$

Define b = Z + diag(κ) R and A = diag(κ) C. A strictly positive vector W satisfying

$$W \ge b + AW, \qquad \sum_a W_a^2 < \rho^2 \tag{16}$$

controls the entire interval. Indeed b is positive, so AW < W, making A a contraction in the W-weighted supremum norm. The intervalwise error supremum M satisfies M ≤ b + AM, whence M ≤ (I-A)^(-1)b ≤ W. A first exit from the candidate tube would contradict the strict radius inequality. The polynomial vector field is locally Lipschitz; the bounded tube also permits continuation of the designated ODE solution. Its endpoint is enclosed by

$$Z_{\text{next}} \ge \text{diag}(d)Z + \text{diag}(\kappa)(CW + R). \tag{17}$$

Induction across all 4,000 intervals gives a full-time bound. Construction iterations alone are not a certificate: the supersolution, radius, endpoint, and interval-connection inequalities must all pass.

The computation assumes IEEE binary64 nearest rounding for elementary operations and correct adjacent-value expansion. Additions, products, divisions, projections, and sums are enclosed outward. Cosine values use rational Machin/Taylor bounds. Square-root uppers are checked by a downward square comparison. A separate checker maps exact binary64 inputs outward to a 2^{-96} integer grid and checks the propagation inequalities by arbitrary-precision integer operations. Selected physical-formula intervals use an independent 110-decimal-digit path. These explicit arithmetic assumptions are the computation's trust boundary; they are not a formal verification of every processor or compiler.

The 257-state certificate combines the three fully certified states with 254 additional certified trajectories; the latter have 1,016,000 full-interval integer checks and 2,032 fixed physical-formula checks. The

fresh 360 certificate has 1,440,000 full-interval checks and 2,880 fixed physical checks at segment IDs [3,4,5,500,1000,2000,3000,3996]. High-precision physical checks are selected independent checks, not high-precision re-evaluation of every interval.

SP0.5. Endpoint transfer to saved targets

If w_i is the curve endpoint, G_i bounds $\|w_i - y_i\|_2$, and Z_K is the final channel bound, then

$$\|Y_i - y_i\|_2 \le \sqrt{40/9}\left(\sum_a Z_{K,a}^2\right)^{1/2} + G_i. \tag{18}$$

The endpoint gap includes the conversion to the saved float32 target and is not counted a second time. Rational norm lower bounds give the recorded endpoint certificates

$$\|Y_i - y_i\|_2 \le \beta_i(\|y_i\|_2 + \eta). \tag{19}$$

The maximum certified β is below 0.278222% for the 257 set and below 0.384291% for the fresh 360 set. These quantify target uncertainty, not predictor error or sampling uncertainty. The following proofs use the exact rational β values, rather than those decimal displays.

SA1 contains the small endpoint and comparison evidence. It does not contain the large full trajectory and interval-kernel payloads. Its summary checker cannot replace from-zero verification of all physical interval enclosures. The analytic implication above, the declared computation contract, and the recorded completed checks specify precisely the premises of the ensuing computer-assisted propositions.

SP1. Exact nine-query-bit capacity on the fixed 257-state task

SP1.1. Query model and proposition

The encoder receives one already-sensed 128-component current x_i. The common deterministic reader receives only a fixed-width B-bit word and pre-fixed common rules. It receives no further current vector, source index, query history, variable-length message, timing, or routing information. There are at most 2^B reader outputs. Common encoder/reader programs, dictionaries, parameters, computation, and query workspace are unrestricted in this proposition and are counted separately in the resource record.

Define the true relative loss

$$\ell_i(p) = \frac{\|Y_i - p\|_2}{\|Y_i\|_2 + \eta}. \tag{20}$$

Proposition SP1. Under SP0's 257-state physical and arithmetic contract, the smallest effective fixed-width query budget permitting loss at most 1/20 at every state is exactly nine bits.

SP1.2. Necessity

Let S_i^+ be the certified upper bound for $\|y_i\|_2$ + η obtained with the common dyadic norm-ceiling grid. The 32,896 exact pair checks give

$$\|y_i - y_j\|_2 > \frac{71}{1000}(S_i^+ + S_j^+), \qquad i \ne j. \tag{21}$$

Write Δ_i = β_i ($\|y_i\|_2$+η). All 257 β_i are at most 1/50. For ε = 1/20,

$$\varepsilon + (1+\varepsilon)/50 = 71/1000. \tag{22}$$

Consequently the saved distance is strictly larger than $\varepsilon(S_i^+ + S_j^+)$ + (1+ε)(Δ_i+Δ_j). Triangle and reverse-triangle inequalities now give

$$\|Y_i - Y_j\|_2 > \varepsilon \left[(\|Y_i\|_2 + \eta) + (\|Y_j\|_2 + \eta) \right]. \tag{23}$$

If one reader output p attained loss at most ε for both states, the ordinary triangle inequality would give the opposite weak inequality. Therefore one codeword can serve at most one successful state, including success exactly at the threshold. Since

$$2^8 = 256 < 257 \leq 512 = 2^9, \tag{24}$$

every qualifying protocol has B $\geq$ 9.

SP1.3. Sufficiency and cost

The 257 saved current rows are numerically distinct. A fixed input table maps them to integers 0,...,256, transmitted as nine-bit words. The common reader returns the corresponding saved future from a fixed target table; the other 255 words are unused. This is a single common lookup rule, not a query-dependent index supplied through a second channel.

Reverse triangle gives

$$\ell_i(y_i) \leq \frac{\beta_i}{1 - \beta_i}. \tag{25}$$

Every exact upper is below 1/20; the maximum is below 0.279%. This protocol therefore attains the required accuracy on every state, completing the proof.

Each float32 table has 257 x 128 x 4 = 131,584 bytes. Together they occupy 263,168 bytes, or 2,105,344 bits, before program and runtime costs. A nine-character bit-string implementation also has host-object storage exceeding nine bits. Sufficiency is an oracle finite-task existence result using the stored futures, rather than a train-only learned baseline. A real-valued scalar can encode 257 distinct labels, so this result gives no lower bound on real latent dimension and no nine-bit bound on total persistent memory.

SP1.4. Uniform-task corollaries

For a uniformly selected state from this fixed task and any B-bit deterministic protocol, at most min(2^B,257) states can meet the 5% threshold. Thus

$$\Pr\{\ell_I \geq \varepsilon\} \geq \frac{(257 - 2^B)_+}{257}, \qquad \mathbb{E}\ell_I \geq \varepsilon \frac{(257 - 2^B)_+}{257}. \tag{26}$$

Query bits	Failed states at least	Failure fraction lower	Mean risk lower
2	253	253/257	253/5140
4	241	241/257	241/5140
8	1	1/257	1/5140
9	0	0	0 from counting alone

The same lower bounds hold after averaging query-independent random seeds, by conditioning on each seed. They do not assert a specified state's expected loss or identify the optimum average risk. Uniformity refers to a defined finite-task query distribution, rather than the physical initial-state population.

The same stored pair threshold and exact maximum β satisfy $\varepsilon + (1+\varepsilon)\,\beta_{max} < 71/1000$ at $\varepsilon = 17/250$. If u_{max} is the exact maximum oracle upper, the minimum query budget is therefore nine for each $u_{max} < \varepsilon \leq 17/250$. The exact positive margin is in SA1. This post-certificate algebraic corollary concerns the task's accuracy requirement; it is not a new primary experiment, a sensing-effect range, or an E3 parameter regime. No conclusion about the other 103 original test states follows from this packing witness, and the fresh 360-state comparison below is a separate task.

SP2. Frozen PRED/RECON mean ordering on 257 states

SP2.1. Prediction identities

Fix the six coded pipelines: prediction-trained PRED and reconstruction-trained RECON at seeds 20261004, 20261005, and 20261006. They use dimension d = 4, four quantization bits per component, and the frozen train-fitted quantizer and ridge reader. Each model receives 16 effective query bits. The unfused inference-parameter payload is 106,416 bytes per model; code storage is not the model's total memory.

Let p_{is}^{M} be the archived binary64 forecast for method M in {P,R}, state i, and seed s. A second prediction identity is the exact rational evaluation of the same saved code and reader parameters. The statements below hold separately for both identities, conditional on these recorded codes. They do not formally verify that every neural-encoder implementation produces those codes.

SP2.2. Exact square roots and positive denominators

For a nonnegative rational a, with G = 2^{96}, integer square-root and cross multiplication find n such that

$$n^2 \le aG^2 < (n+1)^2. \tag{27}$$

If the left equality holds, $\sqrt{a}$ = n/G. Otherwise [n/G,(n+1)/G] encloses $\sqrt{a}$. This gives exact rational norm intervals without relying on an unbounded floating-point square root.

Let $L_i \le \|y_i\|_2 + \eta \le U_i$ and let $\delta_i = \beta_i\, U_i$. SP0 then gives $\|Y_i - y_i\|_2 \le \delta_i$. Define

$$D_i^- = \max(\eta, L_i - \delta_i), \qquad D_i^+ = U_i + \delta_i. \tag{28}$$

Then $0 < D_i^- \le \|Y_i\|_2 + \eta \le D_i^+$. For this 257-state certificate L_i-δ_i is already positive, so the maximum leaves its endpoints unchanged. The η floor also supplies a valid positive bound in the fresh case. No small-norm state is discarded.

SP2.3. Individual risks and paired differences

Enclose each saved forecast distance by $n_{is}^- \le \|y_i - p_{is}\|_2 \le n_{is}^+$. The true risk lies between

$$\frac{\max(0, n_{is}^- - \delta_i)}{D_i^+} \le \ell_i(p_{is}) \le \frac{n_{is}^+ + \delta_i}{D_i^-}. \tag{29}$$

For the paired difference g_{is} = $\ell_i(p_{is}^R)$ - $\ell_i(p_{is}^P)$, its two risks share the same true denominator. A valid numerator interval is

$$a_{is} = n_{is}^{R,-} - n_{is}^{P,+} - 2\delta_i, \qquad b_{is} = n_{is}^{R,+} - n_{is}^{P,-} + 2\delta_i. \tag{30}$$

Take the minimum and maximum of the four ratios a_{is}/D_i^-, a_{is}/D_i^+, b_{is}/D_i^-, and b_{is}/D_i^+. Equivalently, a nonnegative lower numerator uses D_i^+, while a negative lower numerator uses D_i^-; the upper endpoint has the opposite sign-dependent choice. Negative pointwise lower bounds are retained.

For each state, average the three paired numerators before applying the common denominator interval, keeping the target-uncertainty term $2\,\delta_i$. This yields a pointwise three-seed loss-difference interval $[\alpha_i,\omega_i]$. Final endpoints are rounded outward to a 2^{-96} grid. Exact rational summation gives

$$\frac{1}{257}\sum_i \alpha_i \le \overline{g} \le \frac{1}{257}\sum_i \omega_i, \qquad \overline{g} = \frac{1}{771}\sum_i\sum_s g_{is}. \tag{31}$$

SP2.4. Extension to exact code readers

Decode each component c_j in {0,...,15} from the archived little-endian nibble stream. With the frozen range endpoints l_j,h_j, quantized mean μ_j, standard deviation $\sigma_j > 0$, and ridge weights w, define

$$z_j(c) = \frac{l_j + (c_j/15)(h_j - l_j) - \mu_j}{\sigma_j}, \qquad \widetilde{p}_k(c) = w_{4k} + \sum_{j=0}^{3} z_j(c) w_{jk}. \tag{32}$$

All parameters are interpreted as their exact IEEE rationals. The mapping from source indices to test slots is the recorded one, rather than an assumption that they coincide. Exact component reconstruction gives an upper ρ_{is}^M for $\|\widetilde{p}_{is}^M - p_{is}^M\|_2$. Replace δ_i by $\delta_i + \rho_{is}^M$ in individual-risk bounds and replace 2 δ_i by 2 $\delta_i + \rho_{is}^P + \rho_{is}^R$ in the paired numerator. For a three-seed comparison, average the reader margins. The same positive-denominator and four-corner arguments apply.

SP2.5. Strict mean result and pointwise counterexamples

Proposition SP2. For both prediction identities and each of the three seeds, PRED has strictly lower mean true-ODE relative L2 loss on this fixed 257-state task. The three-seed loss average also has the following outward bounds:

Quantity	Certified enclosure
PRED mean relative L2	[0.209014, 0.210791]
RECON mean relative L2	[0.881085, 0.883772]
Paired mean RECON-PRED	[0.670678, 0.674376]

The exact mean interval, for both identities, satisfies

$$\frac{335339}{500000} < \overline{g} < \frac{84297}{125000}. \tag{33}$$

Each per-seed paired mean lower is also strictly positive, and each PRED mean upper is below its paired RECON mean lower. These are exact inequalities obtained by summing all saved rational point bounds, rather than bootstrap confidence intervals or a floating replay tolerance. They apply to every target collection satisfying the certified balls, including the designated true ODE futures.

The three-seed pointwise loss average favors PRED at 253 states and RECON at four, with none unresolved. Reverse source indices are 212, 809, 1012, and 1376. Per-seed pointwise unresolved states are retained. A mean result therefore does not imply pointwise dominance. Averaging losses from three models does not construct a single 16-bit ensemble prediction.

SP3. Separate original-Float64 certification on fresh 360 states

The fresh set was generated with seed 202610110801 from the same controlled initial-state distribution after all six models, quantizers, scalers, and readers were frozen. It is disjoint from the original 1,800 states. The d4/b4 setting had already been chosen using the earlier dimension/bit grid; it is not an unselected optimum. No fresh-state fitting or re-selection is performed.

Use the fresh initial-state contract in SP0.1 and its own 360 target enclosures. The risk regularizer, forecast and reader identities, square-root enclosure, four-corner division, and exact summation are unchanged. Use $D_i^- = \max(\eta, L_i - \delta_i)$. The independent physical and endpoint checks belong to this set, rather than being inferred from the old 257 certificate.

Proposition SP3. PRED has strictly lower mean true-ODE relative L2 loss on these fixed fresh 360 states, for both identities, each seed, and their three-seed loss average.

Quantity	Certified enclosure
PRED mean relative L2	[0.201611, 0.203304]
RECON mean relative L2	[0.880296, 0.882868]
Paired mean RECON-PRED	[0.677351, 0.680901]

Proof. Apply SP2.2-SP2.4 with the fresh target balls and the exact original-Float64 initial contract. Sum all 360 point endpoints as rational numbers. The exact mean difference lower is strictly positive for both forecast identities and each seed. The resulting display intervals enclose both identities' exact endpoints, establishing the proposition.

The pointwise three-seed average favors PRED at 354 states and RECON at six, with none unresolved. Reverse local IDs are 69, 86, 256, 266, 293, and 319. Every reverse or small-norm state remains in the averages.

Task	Seed	Outward mean RECON-PRED	PRED / RECON / unresolved points
Fixed 257	20261004	[0.660191, 0.663870]	244 / 13 / 0
Fixed 257	20261005	[0.714230, 0.718028]	249 / 8 / 0
Fixed 257	20261006	[0.637595, 0.641248]	248 / 8 / 1
Fresh 360	20261004	[0.679436, 0.683028]	346 / 13 / 1
Fresh 360	20261005	[0.698641, 0.702256]	352 / 6 / 2
Fresh 360	20261006	[0.653944, 0.657450]	343 / 15 / 2

Zero unresolved points in a three-seed loss average does not imply zero unresolved points for each seed. The two sets are reported separately; they do not form a pooled 617-state population experiment.

SP4. Evidence, verification scope, and relation to sensing results

Supplementary Evidence Archive SA1 supplies the exact capacity certificate, all 257 oracle risk uppers, both query tables, both comparison certificates, all point/seed rational risk bounds, target enclosures, endpoint β values, and the rational energy/rate certificates. SHA256 manifests bind each published member and its source object. Scientific numeric fields are preserved; computer-specific paths and execution-administration fields are excluded from the reader-facing derivatives.

The included summary checker verifies member hashes, exact point-to-mean summation, positive per-seed mean differences, pointwise counts, and target-denominator consistency. It performs no physical integration, training, encoder inference, or fitting. It is a checker for the supplied certified evidence, not a from-zero reproduction of the full numerical study. Full physical trajectory payloads and historical training environments are outside this small archive. Hash agreement establishes file identity, not a replacement for the analytic or numerical validation assumptions in SP0. Numerical-certification source snapshots, their identity map and the large-payload reconstruction contract are supplied in the repository archive linked by Supplementary material 3, under certification/README.md. They expose the retained physical-validation implementation without changing the summary checker into an independent physical-kernel audit.

Both mean propositions remain restricted to the designated ODE, declared initial-state/arithmetic contracts, fixed sets, and frozen code/reader identities. They do not imply population mean ordering, continuous-PDE accuracy, arbitrary PRED superiority, encoder optimality, equal total computational cost, or optimal real latent dimension. The 16-bit models do not meet the nine-bit oracle's 5% all-state criterion merely because their mean ordering is certified. Related finite-persistent-memory results for independent Gaussian streams are not applied to these controlled physical states or correlated KS snapshots.

The certified representation propositions do not alter sensing-allocation evidence. Brusselator E3 after the solver correction remains directionally favorable but its prespecified 25% primary comparison is not statistically confirmed (relative improvement approximately 1.68%, one-sided randomization p approximately 0.085). KS F1b passes the archived primary allocation comparison on its saved pool; independent physical-data repetition is unavailable. The regenerated Brusselator follow-ups show a near-zero

block mean and short-horizon negative relative-error effects. The full sensing evidence is system- and task-dependent, rather than directionally consistent across systems. The proposed old-solver mechanism remains supported or plausible; it is not established by these representation certificates.