Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Geometry and Dissipation Determine Temporal Depth in Delayed Stochastic Systems

Nobuchika Yamaki

nobuchika.yamaki@kcl.ac.uk

King's College London

Tenna Churiki

TNQ Tech

Article

Keywords: temporal closure, Markov blanket, propagation delay, dissipation, stochastic dynamics

Posted Date: November 27th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11776844/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License. Read Full License

Additional Declarations: No competing interests reported.

Abstract

Biological systems are spatially extended and recurrent, yet statistical boundaries are commonly formulated instantaneously. We develop a theory of temporal statistical closure in delayed stochastic systems by separating two physical determinants of history depth: propagation geometry and dissipative persistence. Using time-unrolled dependency graphs, we show that acyclic delayed systems can admit finite exact blanket-specific depth, whereas relevant recurrent cycles generically produce infinite exact depth even under asymptotic stability. At finite precision, however, residual dependence decays exponentially, yielding a temporal-depth bound consisting of a geometry-dependent propagation term and a dissipation-dependent retention term. The same structure emerges for finite-history posterior inference and variational free energy on path space. Exact Gaussian calculations, noncommuting delayed dynamics, and moving-posterior tracking confirm the predicted decay laws. In a 76-node connectome-informed network, propagation depth scaled almost exactly with inverse conduction velocity, whereas perturbation-retention depth was governed primarily by damping and increased linearly with $\log(1/\varepsilon)$. These results provide a physical basis for temporally extended statistical boundaries in recurrent delayed systems.

1. Introduction

Biological systems maintain coherent organization despite continual exchange with their surroundings. One formal approach to this problem is to partition a system into internal, external, and intervening blanket states, such that internal and external states are conditionally independent given the blanket. This statistical construction, originating in graphical models, has become central to formulations of biological self-organization, active inference, and the free-energy principle (Pearl, 1988; Friston, 2010, 2013; Kirchhoff et al., 2018; Parr et al., 2020). In these formulations, a Markov blanket specifies a statistical boundary through which internal dynamics can be related to external causes without requiring direct statistical coupling between the two domains. Related developments have extended this framework to stochastic dynamical systems, information geometry, and multiscale organization (Da Costa et al., 2021; Friston et al., 2021; Hipólito et al., 2021).

Most formal statements of blanket conditional independence, however, are expressed in terms of states at a given time. Physical and biological interactions are rarely instantaneous. Transport, conduction, biochemical reactions, and feedback introduce finite propagation times, while delayed recurrence can qualitatively alter stability, oscillation, and persistence (Mackey and Glass, 1977; Hale and Verduyn Lunel, 1993). In neural systems, anatomical distance and finite conduction velocity likewise generate heterogeneous transmission delays. Large-scale modelling has shown that these delays interact with structural coupling and local dynamics to shape collective neural activity and functional organization (Ghosh et al., 2008; Deco et al., 2009; Cabral et al., 2011; Breakspear, 2017). Whole-brain modelling therefore treats structural connectivity not only as a coupling topology but also as a spatially embedded system in which interactions require finite time to propagate (Sanz Leon et al., 2013).

Temporal structure is not absent from existing inferential formulations. Generalized-coordinate approaches represent trajectories through hierarchies of local temporal derivatives and thereby allow inference over temporally structured dynamics rather than isolated instantaneous states (Friston et al., 2008). However, this representation is conceptually distinct from explicitly retaining a finite segment of past states whose extent is imposed by heterogeneous physical propagation delays. Generalized coordinates characterize local temporal

flow around the current state, whereas delayed stochastic systems contain dependencies on states at specific earlier times that may remain dynamically relevant through recurrent pathways. Bayesian-mechanical and path-space formulations similarly admit temporally extended inference (Da Costa et al., 2021), but they do not by themselves determine how much physical history must be retained for conditional separation to hold at a specified precision. In particular, finite edge delays do not imply finite exact temporal closure when recurrent pathways are present: a directed cycle can support dependency paths extending arbitrarily far into the past even when the magnitude of those dependencies decays.

This distinction suggests that temporal closure is controlled by two physically different quantities. Propagation geometry determines when an influence can reach a state, whereas dissipation determines how long that influence remains dynamically relevant after it arrives. Spatial embedding, path length, and propagation velocity therefore determine the first temporal scale, while recurrent stability and dissipative decay determine the second. These quantities need not coincide: rapid propagation can coexist with long dynamical persistence, and slow propagation can coexist with rapid decay. Previous work has shown that temporally aligned sensory history can restore blanket-like separation in delayed sensorimotor systems and that delayed interactions can support temporally extended organization without an explicit memory substrate (Yamaki and Churiki, 2026a, 2026b). A general relationship connecting propagation geometry, recurrence, dissipation, and the history required for exact or finite-precision conditional closure has nevertheless remained unresolved.

Here, we develop such a relationship for delayed stochastic systems. We first define blanket-specific temporal depth by conditioning external–internal dependence on a finite history of blanket states and characterize exact temporal closure using time-unrolled dependency graphs. We show that acyclic delayed systems can admit finite exact depth, whereas relevant recurrent cycles generically produce infinite exact depth even under asymptotic stability. We then show that stable recurrent systems nevertheless possess finite temporal depth at every nonzero tolerance, with an upper bound separating a geometry-dependent propagation term from a dissipation-dependent retention term. The same temporal structure is derived for finite-history posterior inference and variational free energy on path space. Exact Gaussian systems are used to test the resulting analytical predictions, and connectome-informed delayed network calculations independently test the predicted effects of propagation geometry and dissipative persistence. Together, these results identify finite propagation and dissipation as separable physical determinants of temporal statistical closure.

2. Methods

2.1 Overview and notation

We developed a theoretical framework for temporal conditional separation in delayed stochastic systems and evaluated the two physical mechanisms predicted to determine its finite-precision depth: propagation through spatially embedded connections and dissipative persistence under recurrent dynamics. The theoretical analysis was formulated for systems admitting an external–blanket–internal partition and quantified residual dependence by conditional mutual information. Exact temporal closure and finite-precision temporal closure were distinguished explicitly.

The theoretical development proceeded in four stages. First, delayed stochastic systems were represented on time-unrolled dependency graphs to characterize the conditions under which a finite blanket history can exactly

separate external and internal states. Second, propagation delays were related to spatial and network geometry, and stable recurrent dynamics were used to derive exponential bounds on residual historical dependence. Third, the formulation was extended to continuous-time stochastic delay systems and to posterior inference and variational free energy on path space. Fourth, the predicted geometry and dissipation terms were tested independently in exact numerical systems and in a connectome-informed large-scale brain network.

The connectome-informed calculations were not assigned an anatomical external–blanket–internal partition. They were therefore not interpreted as direct numerical estimates of blanket-specific conditional mutual information. Instead, they tested independently the two physical quantities entering the temporal-depth bound: the history depth imposed by finite propagation delays and the persistence generated by dissipative recurrent dynamics. Separate notation was used for these mechanism-level quantities to distinguish them from blanket-specific temporal depth.

The principal notation is summarized in Table 1.

Table 1
Principal mathematical notation.

Symbol	Definition
X_t	Complete system state at discrete time t
E_t	External-state component
B_t	Blanket-state component
M_t	Internal-state component
$A^{(\ell)}$	Coupling matrix at lag ℓ
L	Maximum explicit lag
ξ_t	Stochastic innovation
H	Retained history depth
$B_t^{[H]}$	Blanket history retained through depth H
$C_B(H)$	Blanket-specific residual conditional dependence
H_B	Exact blanket-specific temporal depth
$H_{B,\varepsilon}$	Blanket-specific temporal depth at tolerance ε
ε	Finite dependence or approximation tolerance
τ_{ij}	Propagation delay from node j to node i
p	Directed network path
$T(p)$	Cumulative propagation time along path p
T_{geom}	Geometry-dependent propagation offset
d_g	Distance under spatial metric g
v	Propagation velocity
λ	Effective dissipative decay rate
$R(t)$	Continuous-time resolvent
Z	Latent trajectory over an inference interval
Y	Observed variables used for path inference
C_H	Retained recent history in path-space inference
O_H	Omitted history preceding C_H
π_H	Posterior conditioned on retained finite history
π_∞	Posterior conditioned on complete relevant history
$D_{post}(H)$	Expected posterior divergence caused by history truncation
Q	Variational probability measure

Symbol	Definition
ΠH	Target posterior measure under retained history
*F*H	Finite-history variational free energy
K	Finite system-dependent constant in an exponential bound
W_{ij}	Structural coupling weight from node *j* to node *i*
d	Positive local damping magnitude
a	Global coupling coefficient
D	Additive-noise dispersion
N	Number of network nodes
Δ*t*	Numerical integration interval
δ*f*H	Drift contribution omitted at propagation depth *H*
*J*prop(*H*,*t*)	Propagation-truncation relative-entropy rate
*H*prop,*ε*	Operational propagation depth at threshold *ε*
S[n]	Augmented current-and-delay state
*A*aug	Augmented deterministic transition matrix
ρ	Spectral radius of *A*aug
λspec	Decay rate estimated from the augmented spectrum
λimpulse	Decay rate estimated from perturbation-energy decay
*E*ret(*t*)	Relative physical-state perturbation energy
*H*ret,*ε*	Operational retention depth at energy threshold *ε*
U	Upper bound on velocity of a moving posterior target
α	Effective posterior contraction rate

2.2 Delayed stochastic systems and blanket-specific temporal closure

We first considered a discrete-time stochastic process partitioned into external, blanket, and internal variables,

$$X_t = (E_t, B_t, M_t),$$

with delayed linear dynamics

$$X_{t+1} = \sum_{l=0}^{L} A_l X_{t-l} + \xi_t. \quad (1)$$

The physical state X_t was not assumed to be Markovian. Dependence on earlier states was retained explicitly through the delayed coupling matrices.

For a candidate history depth *H*, the retained blanket history was

$$B\square^{[H]} = (B\square, B\square\square\square, \ldots, B\square\square H).$$

Residual external–internal dependence after conditioning on this history was quantified by

$$CB(H) = I(E\square; M\square \mid B\square^{[H]}). \quad (2)$$

Exact blanket-specific temporal depth was defined by

$$HB\square = inf\{H \geq 0 : CB(H) = 0\},$$

whereas finite-precision depth was

$$HB,\epsilon = inf\{H \geq 0 : CB(H) \leq \epsilon\}. \quad (3)$$

The two quantities were treated separately. *H*B□ describes exact conditional independence and can be infinite, whereas *H*B,*ε* describes operational closure at a specified nonzero tolerance.

2.3 Time-unrolled graph characterization

Delayed dependencies were represented on a time-unrolled directed graph. Each state component at each time point was represented as a separate node. A directed edge from variable *j* at time *t* − *ℓ* to variable *i* at time *t* was included whenever the corresponding element of $A^{(\ell)}$ was nonzero.

For each candidate *H*, the nodes contained in *B*□^[*H*] were conditioned upon and d-separation between *E*□ and *M*□ was evaluated. Under the graphical Markov property, d-separation implied the corresponding conditional independence. Under faithfulness, failure of d-separation implied that exact conditional independence did not hold generically.

Feed-forward and recurrent delayed structures were treated separately. In an acyclic delayed graph, all relevant directed paths terminate after finite cumulative delay, so a finite blanket history can block every external–internal path. In a recurrent graph, a relevant directed cycle can be traversed arbitrarily many times. Consequently, admissible paths can extend arbitrarily far into the past even when every individual edge has finite delay.

Dynamical stability was not used to remove such paths from the graphical analysis. Stability reduces the magnitude of influence carried by older recurrent paths but does not cause those paths to disappear after finite time. Recurrence and dissipation were therefore analyzed as distinct determinants of exact and approximate temporal closure.

2.4 Geometry-dependent propagation and dissipative temporal bounds

For a directed path *p*, cumulative propagation time was defined by

$$T(p) = \Sigma\, e \in p\, \tau\, e.$$

For spatially embedded systems, edge delay was defined by the minimum travel time over admissible spatial paths,

$$\tau \square\square = inf\gamma \int ||\dot{\gamma}(s)||g/v(s)\, ds.$$

For constant propagation velocity, this reduces to distance divided by velocity.

The geometry-dependent offset *T*geom represented the propagation scale of paths relevant to the external–blanket–internal dependency under consideration. It is a theoretical path-dependent quantity and need not correspond to a single descriptive network percentile.

For stable systems, propagation of historical perturbations was assumed to satisfy an exponential envelope after the relevant geometry-dependent offset. If $\lambda > 0$ denotes the effective dissipative rate, the blanket-specific residual dependence was bounded by

$$CB(H) \leq K\, exp[-2\lambda(H - Tgeom)\square]. \quad (4)$$

Here, $(x)_+ = \max(x,0)$. Solving the inequality for a finite tolerance gave

$$HB,\epsilon \leq Tgeom + [1/(2\lambda)]\, log(K/\epsilon). \quad (5)$$

Equation (5) was treated as an upper-bound decomposition rather than an equality model. The first term reflects the time required for physical propagation. The second reflects the persistence of historical influence after propagation and depends on both dissipation and the specified precision.

For Gaussian systems, the factor of two arises because KL divergence and weak-dependence mutual information are quadratic in the amplitude of the underlying perturbation or cross-covariance. Different observable projections can therefore share the same exponential rate while differing in the finite constant *K*.

2.5 Continuous-time delayed stochastic systems

Continuous-time dynamics were represented as stochastic delay differential equations,

$$dX(t) = [A\square X(t) + \Sigma r\, ArX(t - \tau r)]\, dt + GdW(t). \quad (6)$$

Although the instantaneous physical state is generally non-Markovian, the segment containing the trajectory over the required delay interval defines a Markov process on path space.

For asymptotically stable dynamics, the resolvent was required to satisfy

$$||R(t)|| \leq M\, exp(-\lambda t)$$

for finite *M* and positive λ. Historical covariance contributions propagated through *R*(*t*) therefore decayed exponentially, producing the same functional dependence on history depth as Eq. (4).

For matrix delay systems, stability was additionally characterized through the delay characteristic matrix

$$\Delta(z) = zI - A\square - \Sigma r\, Ar\, exp(-z\tau r). \quad (7)$$

The asymptotic decay rate was determined by the rightmost dynamically relevant characteristic root. If *s* denotes its real part and $s < 0$, the spectral decay rate was $\lambda = -s$.

2.6 Path-space inference and finite-history variational free energy

Temporal truncation was also formulated directly as an inference problem over trajectories. The latent trajectory of interest was denoted *Z*. Recent retained history was denoted *C*H and older omitted history *O*H. The finite-history and full-history posteriors were

$$\pi H \;=\; P(Z \mid Y,\, CH)$$

and

$$\pi \infty \;=\; P(Z \mid Y,\, CH,\, OH).$$

The expected divergence caused by removing *O*H obeyed

$$Dpost\,(H) \;=\; E\,[KL(\,\pi \infty \;\|\|\; \pi H\,)] \;=\; I(Z;\, OH \mid Y,\, CH). \;(8)$$

Thus, posterior truncation error was exactly equal, in expectation, to the conditional information about the latent trajectory carried uniquely by the omitted history.

For finite-dimensional Gaussian posteriors, omitted historical influence was represented by a positive semidefinite covariance contribution *V*H. If *S*Y denotes the reference observation or innovation covariance, the corresponding posterior divergence was

$$Dpost\,(H) \;=\; 1/2\; log\; det[I \;+\; SY\square\square\square\; VH\; SY\square\square\square].$$

Exponential decay of *V*H therefore generated an exponential posterior-divergence bound and the same logarithmic dependence of required history on 1/*ε*.

Finite-history variational free energy was defined by

$$FH\,[Q] \;=\; KL\,(Q\,\|\,\Pi\, H) \;-\; log\, pH(Y,\, CH). \;(9)$$

The minimizing measure was ΠH. Evaluating finite-history free energy with the full-history posterior introduced a regret equal to KL(π∞ || πH), so the expected finite-history variational penalty was governed by Eq. (8).

The same construction was extended to Gaussian path measures. Mean displacements were measured in the relevant Cameron–Martin norm, and covariance perturbations generated by omitted history were represented by positive trace-class operators.

A separate moving-target analysis examined posterior tracking when the finite-history target changed continuously in time. For effective contraction α > 0 and target velocity bounded by *U*,

$$lim\; sup\; \|\,r \;-\; mH\,\| \;\leq\; U/\alpha\;. \;(10)$$

This term was treated separately from history-truncation error.

2.7 Exact and numerical verification

2.7.1 Multimode Gaussian path-space model

The exact path-space result was first tested in a four-mode Gaussian system. The path interval was 1.5 time units. Dissipative rates were 0.35, 0.55, 0.80, and 1.10; boundary relaxation rates were 0.25, 0.40, 0.60, and 0.90; and propagation offsets were 1.0, 2.0, 3.0, and 1.5. Coupling amplitudes were 1.0, 0.8, 1.2, and 0.7. Boundary-noise amplitudes were 0.9, 1.0, 0.8, and 1.1, and observation-noise amplitudes were 0.50, 0.45, 0.55, and 0.50.

Exact posterior divergence was evaluated at history depths 4.5, 5, 6, 7, 8, 10, 12, and 14. Monte Carlo estimates used 50,000 samples per history depth and a 31-point discretization of the path interval. The random seed was fixed at 20261020.

2.7.2 Noncommuting matrix delay system

A two-dimensional system was used to test whether the asymptotic history law remained valid when delayed operators did not commute. The matrices were

$$A = [-1.1,\ 0.7;\ -0.3,\ -0.9]$$

and

$$A = [0.15,\ -0.12;\ 0.08,\ 0.05],$$

with delay 0.8.

Characteristic roots were obtained numerically from det Δ(*z*) = 0. The resolvent was independently propagated with step size 0.005 over 30 time units. Path-space posterior divergence was evaluated at *H* = 3.5, 4, 4.5, 5, 6, 7, 8, 9, and 10. The asymptotic KL slope was fitted over *H* ≥ 6 and converted to a decay rate by dividing the negative log-slope by two.

2.7.3 Geometry scaling

Geometry was varied independently of dissipative parameters by multiplying a fixed vector of propagation delays by scale factors 0.5, 1.0, 1.5, 2.0, and 2.5. Two configurations were considered. In one, the slowest dissipative mode also carried the longest propagation delay; in the other, the slowest mode and longest-delay mode differed.

Exact finite-precision depth was evaluated at *ε* = 0.01 and compared with the common-rate upper bound derived from the slowest dissipative mode and longest propagation delay.

2.7.4 Moving posterior target

Posterior tracking was tested using a sinusoidally varying target with amplitude 1.2 and angular frequency 0.4, giving *U* = 0.48. Effective contraction rates were 0.2, 0.4, 0.8, 1.5, 3.0, and 6.0. Integration used a step size of 0.002 over 40 oscillation periods. Numerical steady-state tracking amplitude was compared with the exact sinusoidal response and with the upper bound *U*/α.

2.8 Connectome-informed tests of propagation and dissipative retention

2.8.1 Structural network and common delayed dynamics

Large-scale network calculations used the 76-node demonstration structural-connectivity dataset distributed with The Virtual Brain (TVB; Sanz Leon et al., 2013; Sanz-Leon et al., 2015). The dataset represents a hybrid primate connectome incorporating CoCoMac-derived connectivity information and diffusion-imaging-derived tract lengths. It is a publicly available reference network rather than individual-level human participant data. TVB 2.10.0 and tvb-data 3.0.0 were used. The network contained 1,494 nonzero structural edges.

Self-connections were set to zero, nonfinite weights were removed, and the structural weight matrix was normalized by its maximum absolute row sum,

$$W \leftarrow W / max\Sigma \left|W\right| .$$

Tract lengths supplied with the connectivity dataset were retained. Propagation delay was obtained from tract length divided by conduction velocity.

Both connectome-informed calculations were based on the same continuous-time delayed linear drift,

$$dxi\,(t) = [-d\,xi(t) + a\,\Sigma W xj(t - \tau)]\,dt + \sqrt{(2D)}\;dWi\,(t)\,. \quad (11)$$

The global coupling coefficient was a = 0.003.

No anatomical assignment of the 76 nodes to external, blanket, and internal subsets was imposed. Consequently, the network calculations did not directly estimate $C_B(H)$ or $H_{B,\varepsilon}$. Instead, they measured the two mechanism-level quantities entering Eq. (5): the temporal extent generated by direct propagation delays and the persistence produced by dissipative recurrent dynamics. The resulting propagation and retention depths were treated as distinct observables and were not numerically equated with blanket-specific temporal depth.

2.8.2 Propagation-depth calculation

The propagation calculation examined the consequence of finite edge delays while local dynamics were held fixed. The damping magnitude was 0.006 ms⁻¹ and conduction velocity was varied across 2, 4, 8, and 16 mm/ms, numerically equivalent to 2, 4, 8, and 16 m/s.

Stochastic trajectories were generated using the TVB Linear model and Heun stochastic integrator with Δt = 0.5 ms. The additive-noise dispersion was D = 0.02, for which the TVB additive-noise convention gives diffusion amplitude $\sqrt{(2D)}$. Each trajectory was simulated for 12,000 ms, and the first 3,000 ms were discarded. Ten paired random seeds, 20261020 through 20261029, were used at each velocity.

At a candidate propagation depth H, coupling terms with $\tau_{ij} > H$ were omitted from the drift. The local term $-d\,x_i$ was identical in the complete and truncated drifts and therefore cancelled. The omitted contribution was denoted δf_H.

For two diffusion processes sharing diffusion covariance $2D\,I$ but differing in drift, the instantaneous path-space relative-entropy rate is quadratic in the drift difference. The per-node propagation-truncation rate was therefore

$$Jprop(H,t) = ||\,\delta\,fH\,(t)\,||\,/\,(4DN)\,. \quad (12)$$

The quantity has units of nats per node per millisecond. It measures loss caused by omission of direct delayed coupling terms and was not interpreted as blanket conditional mutual information.

For each *H*, *J*prop was averaged across post-burn-in samples. Ninety-five percent confidence intervals were estimated from nonoverlapping 50-ms block means using the Student *t* distribution.

Because contributions from different omitted edges can partially cancel, the raw rate need not decrease monotonically at every successive *H*. A reverse cumulative maximum was therefore applied to the mean curve and to its upper confidence limit. The operational propagation depth was defined as

$$Hprop, \epsilon \ = \ inf\{H \ : \ suph \geq \ H \ Jprop(h) \ \leq \ \epsilon \}. \ (13)$$

The conservative estimate used the reverse cumulative upper 95% confidence envelope.

The primary threshold was 1 × 10⁻⁶ nats/node/ms. Sensitivity was evaluated at 1 × 10⁻⁷, 3 × 10⁻⁷, 1 × 10⁻⁶, 3 × 10⁻⁶, and 5 × 10⁻⁶ nats/node/ms.

Because the truncation intervention operated directly on individual delayed edges, the primary geometry descriptor was the 95th percentile of discretized direct-edge delay. Mean and maximum direct-edge delays and shortest-path delay statistics were retained as secondary descriptors.

The directional 2-versus-16 m/s contrast was tested across matched seeds using a paired sign-flip permutation test with 10,000 permutations.

2.8.3 Dissipative-retention calculation

Dissipative retention was quantified independently from direct propagation truncation. The same normalized structural weights, tract lengths, conduction velocities, coupling coefficient, and continuous-time linear drift were used, but stochastic forcing was removed after initialization so that decay of an imposed perturbation could be measured directly.

To obtain the spectrum of the delayed linear operator and to propagate perturbations without stochastic contamination, the delay equation was represented as an explicit augmented discrete system. Edge delays were rounded to the nearest 0.5-ms integration step. The resulting Euler representation was

$$x[n \ + \ 1] \ = \ (1 \ - \ d\Delta \, t)x \, [n] \ + \Delta \, t \, a \, \Sigma \, \square \, W \square \square xj[n \ - \ k \square \square],$$

where k_{i} is the integer edge delay.

The augmented state was

$$S \, [n] \ = \ [x[\, n], \, x[n - 1], \, \ldots \, , \, x[n - Lmax \,]] \, ,$$

giving

$$S[n \ + \ 1] \ = \ AaugS \, [n] \, . \ (14)$$

This representation was used specifically to expose the deterministic delayed operator as a sparse matrix. It was not used to redefine the stochastic propagation-divergence calculation in Section 2.8.2.

The integration interval was Δ*t* = 0.5 ms. Damping magnitudes were 0.006, 0.012, and 0.024 ms⁻¹, and conduction velocities were 2, 4, 8, and 16 m/s. The largest local Euler decrement was therefore

$$d\Delta t \;=\; 0.024 \;\times\; 0.5 \;=\; 0.012,$$

which remained small relative to unity. The discretization consequently preserved the local continuous-time decay scale while allowing direct construction of the augmented delay operator.

The dominant spectral radius ρ of *A*aug was estimated using sparse eigensolution. Conditions with ρ < 1 were classified as spectrally stable. The corresponding decay rate was

$$\lambda\, spec \;=\; -log\,(\rho\,)\,/\Delta\, t. \;(15)$$

An independent observable decay rate was estimated by perturbation propagation. Twenty-four independent Rademacher vectors were generated in the 76-dimensional physical-state block. Each contained independent ± 1 components and was normalized to unit Euclidean norm. All delayed state blocks were initially set to zero. No additional noise was introduced after the perturbation.

The augmented system was propagated for 3,000 ms. At each time point, relative perturbation energy was calculated only from the current physical-state block,

$$Eret\,(t) \;=\; ||\,xpert\,(t)\,||\,\square\,/\;\;||\,xpert\,(0)\,||\,\square.$$

Mean energy and approximate 95% confidence intervals were calculated across the 24 perturbations. A reverse cumulative envelope was applied to prevent delayed transient revivals from producing premature estimates of decay completion.

The observable decay rate was estimated from the exponential tail,

$$log\;Eret\,(t) \;=\; c \;-\; 2\lambda\, impulset. \;(16)$$

Tail fitting began no earlier than the 95th-percentile shortest-path propagation time for the corresponding velocity. The primary fitting range included envelope energies between 10^{-10} and 0.5. When fewer than 20 observations remained, the upper limit was relaxed to 0.95. At least 10 observations were required for a valid fit.

The operational retention depth was defined as the earliest time after which the complete remaining perturbation-energy envelope stayed below threshold,

$$Hret,\epsilon \;=\; inf\{H \;:\; supt \geq\; H\; Eret(t) \;\leq\; \epsilon\,\}. \;(17)$$

Thresholds were 0.01, 0.003, 0.001, 0.0003, 0.0001, and 0.00003. The primary threshold was 0.001.

For each velocity-by-damping condition, *H*ret,*ε* was regressed on log(1/*ε*). If perturbation energy follows Eq. (16), the predicted slope is

$$1/\,(2\lambda\, impulse)\,.$$

The corresponding asymptotic prediction 1/(2λspec) was evaluated separately. Equality between λspec and λimpulse was not assumed because the spectral radius characterizes the slowest mode of the complete augmented state, whereas λimpulse reflects modes that are actually excited by the imposed perturbations and observable in the current physical-state block.

The propagation depth $H_{\mathrm{prop},\varepsilon}$ and retention depth $H_{\mathrm{ret},\varepsilon}$ were not added numerically because they were defined using different observables and different threshold units. Their role was to test separately the propagation and dissipative mechanisms appearing in the theoretical bound of Eq. (5).

2.9 Statistical analysis and reproducibility

Exact theoretical quantities were evaluated analytically where closed-form expressions were available and numerically otherwise. No predictive machine-learning model, dimensionality reduction, or learned readout was used to define any temporal-depth quantity.

For propagation-depth calculations, variability was assessed across 10 paired stochastic realizations. Confidence intervals for relative-entropy rates were calculated from nonoverlapping 50-ms blocks. The primary conduction-velocity contrast used a paired sign-flip permutation test with 10,000 permutations. Scaling with reciprocal velocity and direct-edge delay was summarized using ordinary least-squares regression.

For dissipative-retention calculations, uncertainty in perturbation energy was evaluated across 24 independent Rademacher perturbations. Exponential tail fits and regressions of $H_{\mathrm{ret},\varepsilon}$ on $\log(1/\varepsilon)$ were summarized by R^2. The empirical slopes were compared with both $1/(2\lambda_{\mathrm{impulse}})$ and $1/(2\lambda_{\mathrm{spec}})$.

3. Results

3.1 Delayed graph structure determines exact blanket-specific temporal closure

Time-unrolled graphical characterization showed that exact blanket-specific temporal closure depended on the structure of delayed dependency paths rather than on the maximum delay of any individual connection. In feed-forward or otherwise acyclic delayed systems, all directed paths connecting external and internal states terminated after a finite cumulative propagation time. Conditioning on blanket history beyond this depth d-separated the present external and internal states, yielding a finite exact blanket-specific depth, H_{B} (Fig. 1A,B).

The required depth was determined by cumulative propagation along paths relevant to the external–blanket–internal partition. Increasing delays along these paths increased the history required for exact separation, whereas delays outside the relevant dependency structure did not alter H_{B}. Finite propagation could therefore generate a nonzero blanket-specific temporal extent without requiring an explicitly defined memory variable.

3.2 Recurrence prevents finite exact blanket-specific closure

The result changed qualitatively when the delayed dependency graph contained a relevant directed cycle. A perturbation entering a recurrent path could return to the present after repeated traversals, generating admissible dependency paths with arbitrarily large cumulative delay. Consequently, no finite blanket history generically d-separated present external and internal states, giving

$$H_{\mathrm{B}} = \infty.$$

This obstruction persisted under asymptotically stable dynamics (Fig. 1C). Stability reduced the magnitude of dependence contributed by progressively older recurrent paths but did not remove those paths from the time-unrolled graph. Exact blanket-specific closure and dynamical stability were therefore distinct properties:

recurrence determined whether historical dependence terminated exactly, whereas dissipation determined how rapidly its magnitude declined.

3.3 Dissipation yields finite blanket-specific depth at finite precision

Stable recurrent systems nevertheless admitted finite blanket-specific temporal depth at any fixed nonzero tolerance. Beyond the relevant geometry-dependent propagation offset, residual conditional dependence was bounded by an exponential envelope,

$$CB(H) \leq K \exp[-2\lambda (H - Tgeom)_+].$$

The resulting finite-precision depth satisfied

$$HB,\epsilon \leq Tgeom + [1/(2\lambda)] \log(K/\epsilon). \quad (18)$$

The bound separated two contributions to blanket-specific temporal depth (Fig. 2A–C). The first was a propagation offset determined by the geometry of relevant delayed paths. The second was a dissipative-retention contribution that increased as *λ* decreased. Increasing the required precision affected the latter only logarithmically through log(1/*ε*).

Thus, a recurrent system could have infinite exact blanket-specific depth while retaining a finite operational depth for every *ε* > 0. The subsequent connectome-informed calculations tested the propagation and dissipative mechanisms entering this bound separately rather than estimating *H*B,ε directly.

3.4 Finite-history posterior error equals omitted historical information

The same temporal structure emerged for inference over trajectories. The expected divergence between the full-history posterior and the posterior conditioned only on retained recent history was

$$Dpost(H) = E[KL(\pi\infty \,||\, \pi H)] = I(Z; OH \mid Y, CH). \quad (19)$$

Finite-history posterior error was therefore exactly equal, in expectation, to the conditional information about the latent path that remained uniquely in the omitted history (Fig. 2D).

For Gaussian systems, the covariance contribution generated by omitted history decayed exponentially with increasing *H*. Posterior KL divergence inherited the same exponential tail, so the history required to achieve a specified posterior tolerance exhibited the same propagation-plus-dissipation dependence. The finite-history variational-free-energy regret was governed by the same posterior divergence.

3.5 Exact and numerical systems recover the predicted history-decay law

Exact and numerical delayed systems reproduced the predicted decay of finite-history error (Fig. 3).

In the multimode Gaussian path-space model, the exact posterior divergence at *H* = 5 was 0.996431, compared with a Monte Carlo estimate of 0.997890. At *H* = 10, the corresponding values were 0.024563 and 0.024680.

Relative differences were 0.146% and 0.476%, respectively (Fig. 3A; Table 2).

The same decay law was recovered when the delayed operators did not commute. In the two-dimensional noncommuting system, the dominant characteristic root implied a decay rate of 0.843692, whereas independent fitting of the posterior-divergence tail yielded 0.840015, differing by − 0.436% (Fig. 3B,C; Table 2).

The moving-posterior calculation showed comparable agreement. The numerical asymptotic tracking amplitude was 0.079825, compared with an exact value of 0.079823 and an analytical upper bound of 0.080000 (Fig. 3D; Table 2). These results showed that the predicted decay structure extended across Gaussian path inference, noncommuting delayed dynamics, and time-dependent posterior tracking.

Table 2
Quantitative agreement between theoretical and numerical finite-history predictions.

Quantity	Exact/theoretical value	Numerical value	Relative difference
Path-space KL at $H = 5$	0.996	0.998	0.146%
Path-space KL at $H = 10$	0.025	0.025	0.476%
Noncommuting decay rate, λ	0.844	0.840	−0.436%
Moving-posterior asymptotic amplitude	0.080	0.080	0.003%
Moving-posterior upper bound	0.080	0.080	−0.219%

3.6 Whole-brain propagation depth scales with conduction geometry

Propagation depth was examined in the 76-node structural network while local damping, coupling strength, noise intensity, and integration parameters were held fixed. Conduction velocity was varied across 2, 4, 8, and 16 m/s.

The 95th percentile of discretized direct-edge delay decreased from 52.000 ms at 2 m/s to 26.000, 13.000, and 6.500 ms at 4, 8, and 16 m/s, respectively (Fig. 4A,B). The complete propagation-truncation relative-entropy curves are shown in Figure S1.

At the primary threshold $\varepsilon = 1 \times 10^{-6}$ nats/node/ms, mean operational propagation depth $H_{\mathrm{prop},\varepsilon}$ was 42.650 ± 1.029 ms at 2 m/s, 21.450 ± 0.438 ms at 4 m/s, 10.700 ± 0.258 ms at 8 m/s, and 5.400 ± 0.211 ms at 16 m/s (Fig. 4C). Seed-level estimates are provided in Table S1.

Propagation depth scaled almost exactly with reciprocal conduction velocity. Regression of $H_{\mathrm{prop},\varepsilon}$ against $1/v$ gave $R^2 = 0.999992$, with a slope of 85.176 ms·m/s and an intercept of 0.087 ms. Regression against direct-edge p95 delay yielded a slope of 0.819 and the same R^2 (Fig. 4D; Table 3).

The prespecified paired comparison between 2 and 16 m/s yielded a mean difference of 37.250 ms in $H_{\mathrm{prop},\varepsilon}$, with a sign-flip permutation $p = 0.00230$ (Table 3).

Table 3
Whole-brain propagation-depth scaling.

Relationship	Estimate	Intercept	R^2	p
Hprop,ε versus 1/v	85.176 ms·m/s	0.087 ms	1.000	—
Hprop,ε versus direct-edge p95 delay	0.819	0.087 ms	1.000	—
2 versus 16 m/s paired difference	37.250 ms	—	—	0.002

The same velocity dependence persisted across thresholds from 1×10^{-7} to 5×10^{-6} nats/node/ms. Lowering ε increased absolute propagation depth but preserved the systematic compression produced by increasing conduction velocity (Figure S2; Table S2).

These results establish a geometry-dependent temporal scale for direct delayed coupling. Because no external–blanket–internal partition was imposed on the 76-node network, Hprop,ε was interpreted specifically as propagation depth rather than as a numerical estimate of blanket-specific HB,ε.

3.7 Dissipative dynamics determine perturbation-retention depth

All 12 combinations of conduction velocity and local damping were spectrally stable. At damping 0.006 ms^{-1}, λspec ranged from 3.937 to 4.144 s^{-1}. Increasing damping to 0.012 ms^{-1} increased λspec to 9.551–10.122 s^{-1}, whereas damping 0.024 ms^{-1} increased it further to 20.527–22.128 s^{-1}.

Decay measured directly in the current physical-state block followed the same ordering. λimpulse ranged from 4.620 to 4.887 s^{-1} at damping 0.006 ms^{-1}, from 11.342 to 11.716 s^{-1} at 0.012 ms^{-1}, and from 23.648 to 24.046 s^{-1} at 0.024 ms^{-1} (Fig. 5A; Table 4). Perturbation-energy tails were strongly exponential, with fit R^2 values between 0.994 and 1.000. Complete perturbation-energy trajectories are shown in Figure S3.

At the primary energy threshold $\varepsilon = 0.001$, conservative retention depth Hret,ε averaged 594.125 ms across conduction velocities at damping 0.006 ms^{-1}, 289.250 ms at 0.012 ms^{-1}, and 143.500 ms at 0.024 ms^{-1} (Fig. 5B,C). Stronger damping therefore produced a marked and approximately proportional reduction in retained perturbation history.

Within each damping condition, conduction velocity had only a small effect on this longer retention scale. At damping 0.006 ms^{-1}, conservative Hret,ε decreased from 596.000 to 593.000 ms as conduction velocity increased from 2 to 16 m/s. At damping 0.012 ms^{-1}, the corresponding values were 289.500 and 289.000 ms. At damping 0.024 ms^{-1}, all four velocity conditions yielded 143.500 ms.

The persistence of an already introduced perturbation was therefore governed predominantly by dissipative dynamics rather than by conduction velocity. This retention measure was distinct from Hprop,ε because it quantified decay after perturbation rather than loss caused by omitting delayed coupling terms.

3.8 Retention depth scales logarithmically with required precision

The predicted dependence of retained history on precision was tested across six perturbation-energy thresholds. In every combination of conduction velocity and damping, $H_{ret,\varepsilon}$ increased approximately linearly with $\log(1/\varepsilon)$ (Fig. 6A–C). Across all 12 conditions, R^2 ranged from 0.999179 to 0.999999.

The slope depended systematically on dissipative rate. At damping 0.006 ms^{-1}, empirical slopes ranged from 91.823 to 94.444 ms per log unit. At damping 0.012 ms^{-1}, they ranged from 42.185 to 42.617 ms, and at damping 0.024 ms^{-1} they ranged from 20.740 to 20.851 ms (Table 4).

Predictions based on directly observed perturbation decay increasingly approached the empirical slopes as damping strengthened. The ratio of empirical slope to $1/(2\lambda_{impulse})$ ranged from 0.873 to 0.898 at damping 0.006 ms^{-1}, from 0.967 to 0.988 at 0.012 ms^{-1}, and from 0.986 to 1.002 at 0.024 ms^{-1} (Fig. 6D; Table 4). Predictions based on λ_{spec} preserved the same ordering but were more conservative, particularly under the weakest damping condition (Figure S4).

Table 4
Dissipative decay and logarithmic scaling of retention depth.

Speed (m/s)	Damping (ms^{-1})	λ_{spec} (s^{-1})	$\lambda_{impulse}$ (s^{-1})	Perturbation-fit R^2	Empirical slope (ms)	$1/(2\lambda_{impulse})$ (ms)	Slope ratio	Log-scaling R^2
2	0.006	3.937	4.620	0.994	94.444	108.231	0.873	0.999
4	0.006	4.056	4.776	0.995	92.861	104.684	0.887	0.999
8	0.006	4.115	4.851	0.995	92.144	103.080	0.894	1.000
16	0.006	4.144	4.887	0.996	91.823	102.305	0.898	1.000
2	0.012	9.551	11.342	0.999	42.617	44.083	0.967	1.000
4	0.012	9.891	11.595	1.000	42.332	43.122	0.982	1.000
8	0.012	10.048	11.680	1.000	42.185	42.807	0.985	1.000
16	0.012	10.122	11.716	1.000	42.185	42.677	0.988	1.000
2	0.024	20.527	23.648	1.000	20.851	21.144	0.986	1.000
4	0.024	21.544	23.953	1.000	20.740	20.874	0.994	1.000
8	0.024	21.950	24.023	1.000	20.839	20.813	1.001	1.000
16	0.024	22.128	24.046	1.000	20.839	20.794	1.002	1.000

The complete threshold-by-condition values are provided in Table S3. Taken together, the connectome-informed results identified two distinct temporal scales predicted by the theory: finite propagation determined how much history was required to encompass delayed transmission, whereas dissipative decay determined how long an established perturbation remained dynamically detectable. These mechanism-level quantities supported the two physical contributions appearing in the blanket-specific bound without requiring an anatomical external–blanket–internal partition in the whole-brain network.

4. Discussion

4.1 Temporal closure is jointly determined by propagation and dissipation

The present study identifies two distinct physical determinants of temporal statistical closure in delayed stochastic systems: propagation geometry and dissipative persistence. Exact blanket-specific closure depended on the topology of delayed dependency paths. Acyclic delayed systems could achieve conditional separation after a finite blanket history, whereas relevant recurrent cycles generically generated dependency paths extending arbitrarily far into the past. Dynamical stability did not remove these paths and therefore did not, by itself, restore finite exact closure. Instead, stability controlled the magnitude of progressively older influences. This distinction led to a finite-precision formulation in which temporal depth was bounded by a geometry-dependent propagation term and a dissipation-dependent retention term.

The numerical results supported both components of this decomposition. Exact Gaussian calculations reproduced the predicted finite-history posterior divergence, the noncommuting delayed system recovered the decay rate independently from the dominant characteristic root, and the moving-posterior calculation approached its analytical tracking bound. In the connectome-informed system, conduction velocity strongly controlled operational propagation depth, whereas local damping primarily controlled the persistence of an established perturbation. Temporal extent therefore emerged from two separable processes rather than from a single generic notion of memory.

4.2 Exact closure and finite-precision closure are fundamentally different

A central consequence is that finite transmission delays do not imply finite exact temporal depth. If a delayed interaction belongs to a relevant recurrent cycle, repeated traversal of that cycle generates dependency paths of arbitrarily large cumulative duration. Even when every individual edge has finite delay, the corresponding time-unrolled graph therefore need not admit a finite history that blocks all external–internal paths. This is a structural consequence of recurrence rather than a consequence of dynamical instability.

Dissipation changes the operational significance of these paths without eliminating them. In an asymptotically stable system, contributions from progressively older recurrent traversals decline, potentially becoming arbitrarily small while remaining nonzero. Exact temporal depth can consequently remain infinite while finite-precision depth is finite for every nonzero tolerance. This distinction is important for biological systems, in which exact zero dependence is generally a stronger requirement than is physically or empirically meaningful. Similar distinctions between the persistence of dynamical pathways and the decay of their influence arise broadly in delay systems and recurrent neural dynamics (Hale and Verduyn Lunel, 1993; Breakspear, 2017).

The natural temporal boundary of a recurrent system is therefore precision dependent. Rather than asking only whether historical dependence exists, the relevant question becomes how far into the past one must condition before the residual dependence falls below a specified tolerance. This interpretation converts temporal closure from a binary property into a quantitatively controllable scale.

4.3 Geometry determines when influence arrives; dissipation determines how long it remains

The theoretical decomposition assigns propagation and dissipation different physical roles. Propagation geometry determines when influence can reach a state, whereas dissipation determines how long that influence remains dynamically relevant after it has arrived. These mechanisms need not covary.

The propagation results showed nearly exact inverse scaling with conduction velocity. At the primary threshold, increasing velocity from 2 to 16 m/s reduced operational propagation depth from 42.650 to 5.400 ms, closely tracking the reduction in structural edge delays. This behavior is consistent with the expected dependence of transmission time on path length and conduction velocity and with large-scale neural models in which heterogeneous anatomical delays materially influence collective dynamics (Ghosh et al., 2008; Deco et al., 2009; Cabral et al., 2011; Sanz Leon et al., 2013).

A different pattern was observed for perturbation retention. Increasing local damping from 0.006 to 0.024 ms^{-1} reduced conservative retention depth from approximately 594 to 144 ms, whereas varying conduction velocity within a fixed damping level produced comparatively small changes. The characteristic decay rate and directly observed perturbation-decay rate changed systematically with damping, and the empirical dependence of retention depth on $\log(1/\varepsilon)$ closely approached the predicted slope $1/(2\lambda_{impulse})$.

The physical distinction is therefore straightforward: spatial structure controls the latency of historical influence, whereas dissipative dynamics control its persistence. A system can transmit influence rapidly but retain its consequences for a long interval, or transmit it slowly while dissipating those consequences rapidly. Treating both effects as a single memory timescale would obscure this distinction.

4.4 Path-space inference links historical dependence to posterior error

The same temporal structure appeared at the level of inference. The expected divergence between the full-history posterior and the finite-history posterior was exactly equal to the conditional information about the latent trajectory contained in the omitted past. Finite-history approximation error was therefore not an independent computational artifact; it measured the inferential consequence of discarding dynamically relevant history.

This result provides a direct bridge between temporal conditional dependence and variational inference. When omitted historical covariance decays exponentially, posterior divergence inherits the same exponential dependence on retained history. The finite-history variational-free-energy penalty consequently follows the same geometry-plus-dissipation structure. Temporal depth can therefore be interpreted both dynamically, as the persistence of causal influence, and inferentially, as the amount of history required to approximate the full posterior within a prescribed error.

This framing also clarifies the relationship to generalized-coordinate formulations of the free-energy principle. Generalized coordinates represent local temporal flow through successive temporal derivatives and provide a powerful representation of smooth dynamics around the current state (Friston et al., 2008). The present formulation addresses a different question: how much explicit past history must be retained when dependencies arise from heterogeneous finite delays and recurrent paths. The two descriptions are therefore complementary rather than mutually exclusive. Generalized coordinates characterize local temporal structure, whereas finite-history depth quantifies the extent of explicit delayed dependence.

4.5 Implications for Markov blankets in spatially extended biological systems

Markov blankets are commonly formulated as conditional-independence structures separating internal and external states through an intervening set of blanket states (Pearl, 1988; Friston, 2010; Kirchhoff et al., 2018; Parr et al., 2020). In physically extended systems, however, the relevant separation may not be adequately characterized by an instantaneous partition alone. If interactions require finite propagation time, then the statistical boundary can inherit a temporal depth from the underlying physical geometry.

The present results suggest that a blanket may therefore possess both a spatial structure and a temporal extent. The spatial component identifies which states mediate external–internal dependencies, whereas the temporal component identifies how much of the mediating history must be retained before conditional separation becomes exact or sufficiently accurate. Recurrent dynamics make this distinction especially important because the exact history requirement can become infinite even while the finite-precision requirement remains modest.

This interpretation is consistent with previous demonstrations that temporally aligned sensory history can restore blanket-like conditional separation in delayed sensorimotor systems and that temporally extended organization can persist without an explicit memory substrate (Yamaki and Churiki, 2026a, 2026b). The present framework generalizes those observations by identifying propagation geometry and dissipation as the two quantities controlling the depth of the required history.

4.6 The connectome-informed results test mechanisms rather than anatomical blankets

The whole-brain calculations were designed to isolate the two physical mechanisms predicted by the theory rather than to assign specific cortical regions to external, blanket, and internal classes. The propagation depth quantified the loss produced by truncating delayed coupling terms, whereas the retention depth quantified persistence of an imposed perturbation under recurrent dynamics. These quantities are therefore not themselves estimates of blanket-specific conditional mutual information.

This distinction is important because an anatomical external–blanket–internal partition would require an additional hypothesis about which regions constitute the statistical boundary. Imposing such a partition solely to reproduce the theoretical notation would make the result depend on an arbitrary anatomical assignment. Instead, the large-scale network calculations tested the mechanism-level predictions that finite propagation creates a geometry-dependent temporal scale and that dissipation creates a separate retention scale.

The near-exact reciprocal-velocity scaling and the strong damping dependence support these mechanisms independently. They do not establish the existence of a particular anatomical Markov blanket in the brain, nor is such a claim required for the theoretical result.

4.7 Limitations and future directions

Several limitations should be considered. First, the principal analytical results were derived for linear Gaussian delayed systems or for regimes in which historical influence admits an exponential envelope. Nonlinear systems with state-dependent delays, switching dynamics, heavy-tailed noise, or critical slowing may exhibit non-exponential history decay and consequently different precision-depth relationships.

Second, the theoretical blanket-specific quantity and the connectome-informed mechanism-level quantities are deliberately distinct. The latter test propagation and retention without imposing an anatomical blanket partition. A future empirical test could define biologically motivated external, blanket, and internal subsets and measure conditional dependence directly as a function of retained blanket history.

Third, the connectome-informed calculations used a fixed structural network and simplified linear dynamics. Real neural systems contain heterogeneous local time constants, nonlinear population dynamics, conduction-velocity variation, plasticity, and state-dependent effective connectivity. These factors could alter both T_{geom} and the effective decay rate λ. The current framework nevertheless predicts how such alterations should enter: changes in propagation geometry should primarily shift the propagation component, whereas changes in dynamical persistence should alter the retention component and the slope with $\log(1/\varepsilon)$.

Finally, temporal depth depends explicitly on tolerance. Different experimental observables and measurement noise levels will therefore yield different operational depths even for the same underlying dynamics. This dependence is not a defect of the framework but a direct consequence of distinguishing exact from finite-precision closure.

4.8 Conclusion

Delayed recurrent systems need not possess a finite exact temporal boundary even when they are dynamically stable. Finite propagation determines when historical influence can reach a state, whereas dissipation determines how long that influence remains relevant. Their combination yields a finite operational temporal depth whose propagation component depends on geometry and whose retention component scales inversely with decay rate and logarithmically with required precision. This provides a physical basis for temporally extended statistical boundaries and connects delayed dynamics, path-space inference, and blanket-specific conditional separation within a common framework.

Declarations

Competing Interests

The authors declare no competing interests.

Ethics Statement

This study involved no new human or animal participants and no primary data collection. The connectome-informed analyses used the publicly distributed 76-region structural-connectivity and tract-length dataset provided with The Virtual Brain software. No individually identifiable participant information was accessed or analyzed; therefore, no additional ethical approval or informed consent was required for the present study.

Funding

This study received no external funding.

Author Contribution

N.Y. conceived and designed the study, developed the theoretical framework, performed the analyses, interpreted the results, prepared the figures, and drafted and revised the manuscript. T.C. developed the simulation code and contributed to implementation and validation of the computational framework. Both authors approved the final manuscript.

Data Availability

The structural-connectivity weights and tract-length data used for the connectome-informed analyses were obtained from the publicly available 76-node demonstration dataset (connectivity_76.zip) distributed with The Virtual Brain (TVB). The dataset is available from the TVB data repository (https://zenodo.org/records/10128131; https://doi.org/10.5281/zenodo.10128131). The individual connectivity file can also be accessed at https://github.com/the-virtual-brain/tvb-data/blob/master/tvb_data/connectivity/connectivity_76.zip. All numerical results reported in the manuscript were generated from this reference network or from fully specified synthetic delayed stochastic systems. The complete analysis and simulation code required to reproduce the results, tables, and figures is available at Figshare: https://doi.org/10.6084/m9.figshare.34713675.

References

1. Breakspear, M. Dynamic models of large-scale brain activity. *Nat. Neurosci.* **20**, 340–352. 10.1038/nn.4497 (2017).
2. Cabral, J., Hugues, E., Sporns, O. & Deco, G. Role of local network oscillations in resting-state functional connectivity. *NeuroImage* **57** (1), 130–139. 10.1016/j.neuroimage.2011.04.010 (2011).
3. Da Costa, L., Friston, K., Heins, C. & Pavliotis, G. A. Bayesian mechanics for stationary processes. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences, 477*(2256), 20210518. (2021). 10.1098/rspa.2021.0518
4. Deco, G., Jirsa, V., McIntosh, A. R., Sporns, O. & Kötter, R. Key role of coupling, delay, and noise in resting brain fluctuations. *Proc. Natl. Acad. Sci. U.S.A.* **106** (25), 10302–10307. 10.1073/pnas.0901831106 (2009).
5. Friston, K. The free-energy principle: A unified brain theory? *Nat. Rev. Neurosci.* **11** (2), 127–138. 10.1038/nrn2787 (2010).
6. Friston, K. Life as we know it. *J. Royal Soc. Interface*. **10** (86), 20130475. 10.1098/rsif.2013.0475 (2013).
7. Friston, K. J. et al. Parcels and particles: Markov blankets in the brain. *Netw. Neurosci.* **5** (1), 211–251. 10.1162/netn_a_00175 (2021).
8. Friston, K. J., Trujillo-Barreto, N. & Daunizeau, J. DEM: A variational treatment of dynamic systems. *NeuroImage* **41** (3), 849–885. 10.1016/j.neuroimage.2008.02.054 (2008).
9. Ghosh, A., Rho, Y., McIntosh, A. R., Kötter, R. & Jirsa, V. K. Noise during rest enables the exploration of the brain's dynamic repertoire. *PLoS Comput. Biol.* **4** (10), e1000196. 10.1371/journal.pcbi.1000196 (2008).
10. Hale, J. K. & Verduyn Lunel, S. M. *Introduction to Functional Differential Equations* (Springer-, 1993). 10.1007/978-1-4612-4342-7

11. Hipólito, I. et al. Markov blankets in the brain. *Neurosci. Biobehavioral Reviews*. **125**, 88–97. 10.1016/j.neubiorev.2021.02.003 (2021).
12. Kirchhoff, M., Parr, T., Palacios, E., Friston, K. & Kiverstein, J. The Markov blankets of life: Autonomy, active inference and the free energy principle. *J. Royal Soc. Interface*. **15** (138), 20170792. 10.1098/rsif.2017.0792 (2018).
13. Mackey, M. C. & Glass, L. Oscillation and chaos in physiological control systems. *Science* **197** (4300), 287–289. 10.1126/science.267326 (1977).
14. Parr, T., Da Costa, L. & Friston, K. Markov blankets, information geometry and stochastic thermodynamics. *Philosophical Trans. Royal Soc. A: Math. Phys. Eng. Sci.* **378** (2164), 20190159. 10.1098/rsta.2019.0159 (2020).
15. Pearl, J. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference* (Morgan Kaufmann, 1988).
16. Sanz Leon, P. et al. The Virtual Brain: A simulator of primate brain network dynamics. *Front. Neuroinformatics*. **7**, 10. 10.3389/fninf.2013.00010 (2013).
17. Yamaki, N. & Churiki, T. Temporal membranes: Delay-induced self-retention without clocks or memory substrates. *Artif. Life*. 1–14. 10.1162/ARTL.a.474 (2026a).
18. Yamaki, N. & Churiki, T. Temporally aligned sensory history restores Markov-blanket function in delayed sensorimotor systems. *BioSystems* **268**, 105939. 10.1016/j.biosystems.2026.105939 (2026b).

Figures

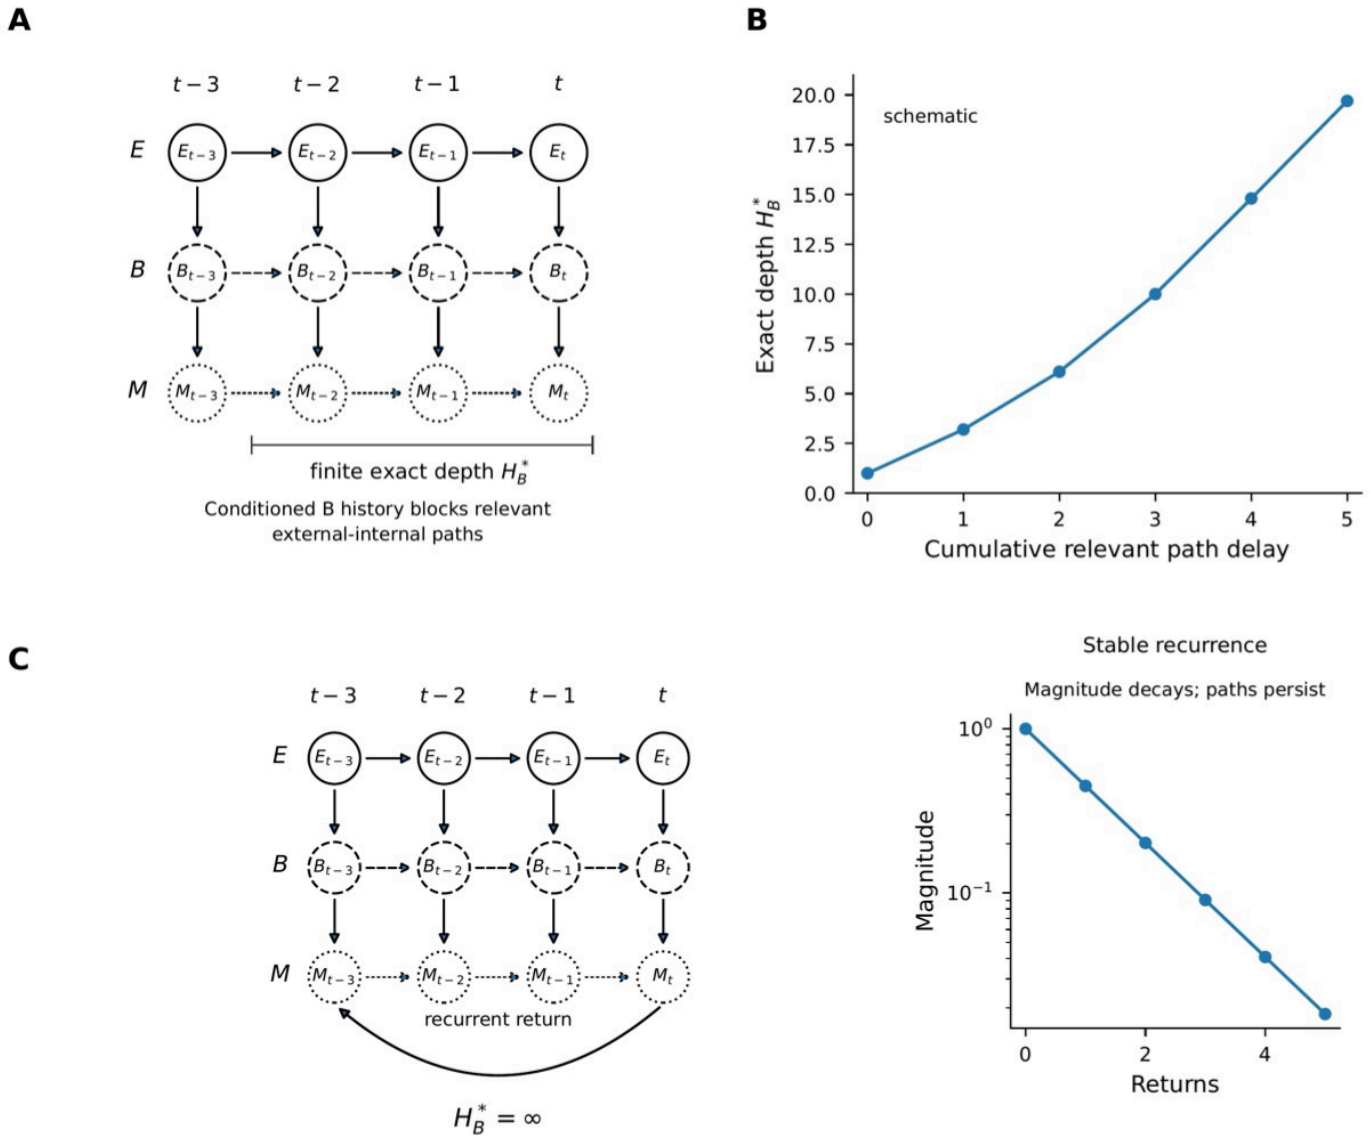

Figure 1

Exact blanket-specific temporal closure and recurrence obstruction.
(A) Time-unrolled representation of a delayed system with external (*E*), blanket (*B*), and internal (*M*) states. In an acyclic dependency structure, conditioning on a sufficiently long blanket history blocks all relevant external–internal paths and yields a finite exact blanket-specific temporal depth, *H*B*. (B) Schematic relationship between cumulative delay along relevant dependency paths and exact temporal depth. Delays outside the relevant dependency structure do not alter *H*B*. (C) In the presence of a relevant delayed recurrent cycle, admissible dependency paths can extend arbitrarily far into the past, giving *H*B* = ∞. Dynamical stability decreases the magnitude of progressively older recurrent contributions but does not eliminate the corresponding paths.

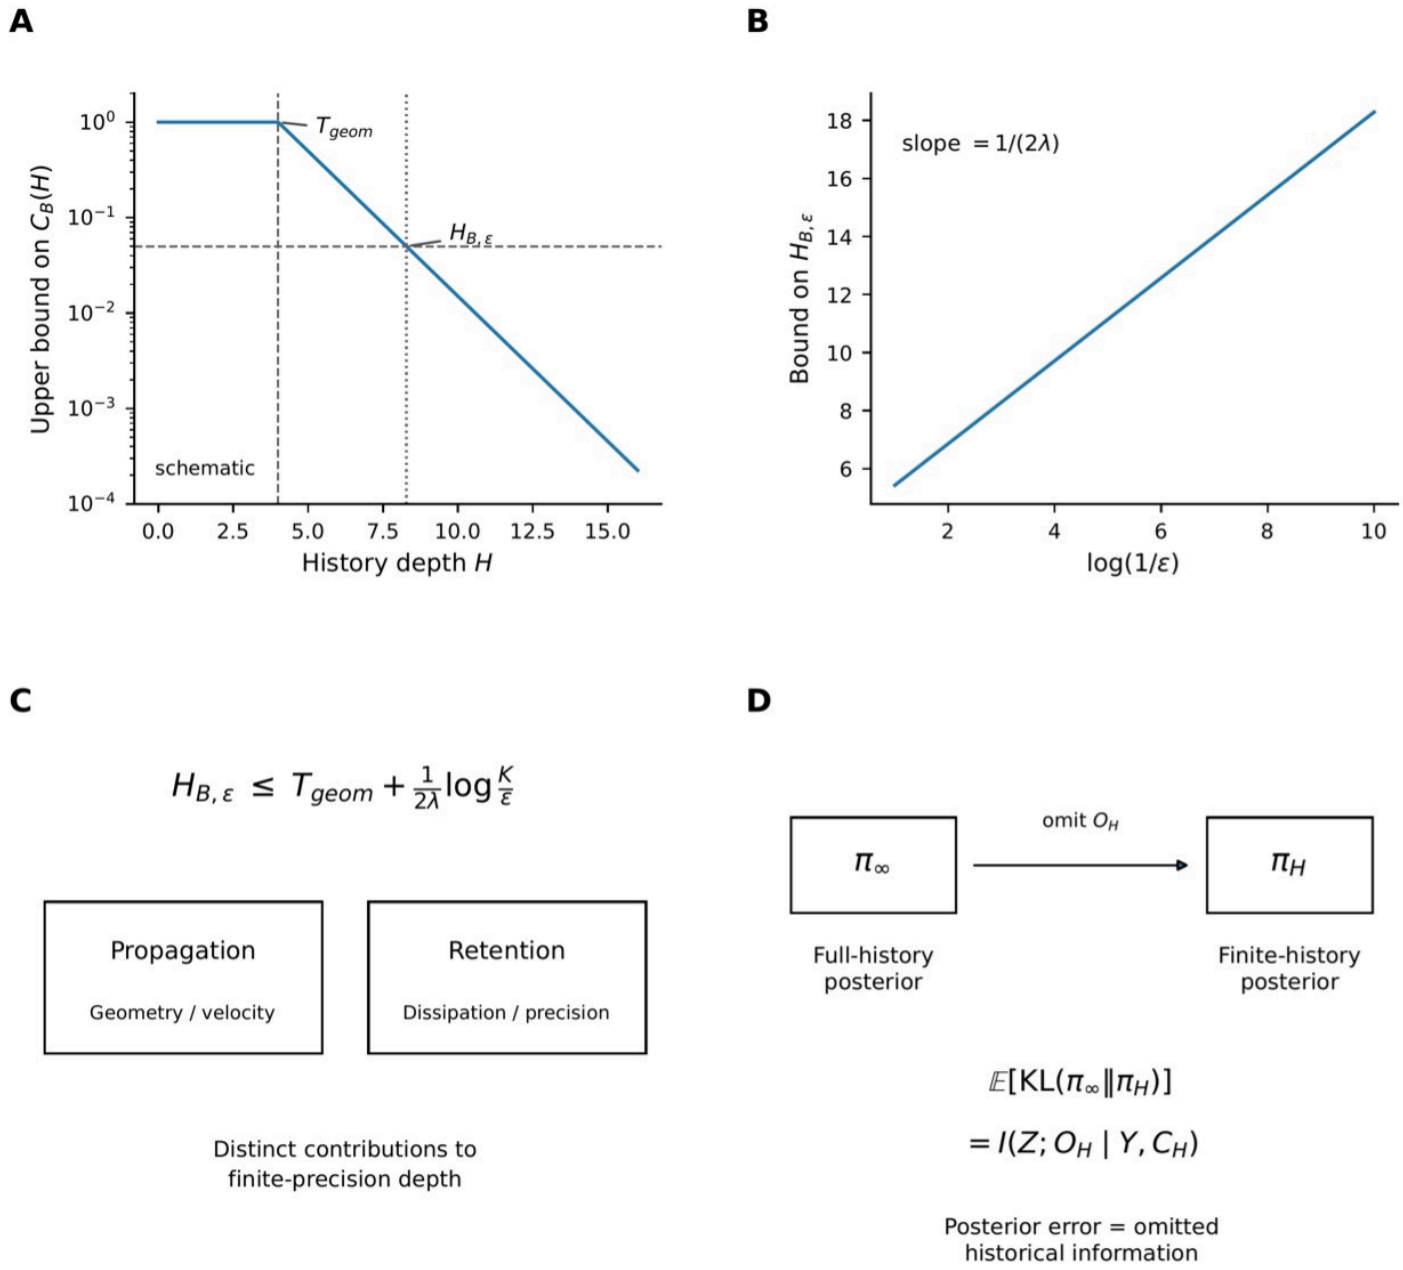

Figure 2

Finite-precision temporal depth and path-space interpretation.
(A) Exponential upper bound on residual blanket-specific dependence as retained history depth *H* increases. *T*geom denotes the geometry-dependent propagation offset, and *H*B,ε is the history depth at which the residual dependence falls below tolerance *ε*. (B) The resulting bound predicts a linear increase of *H*B,ε with log(1/*ε*), with slope 1/(2*λ*), where *λ* is the effective dissipative decay rate. (C) Decomposition of finite-precision temporal depth into a propagation contribution determined by geometry and velocity and a retention contribution determined by dissipation and required precision. (D) Path-space interpretation of history truncation. Removing the older history *O*H changes the full-history posterior π∞ to the finite-history posterior πH, and the expected posterior divergence equals the conditional information carried by the omitted history, E[KL(π∞ || πH)] = I(*Z*; *O*H | *Y*, *C*H).

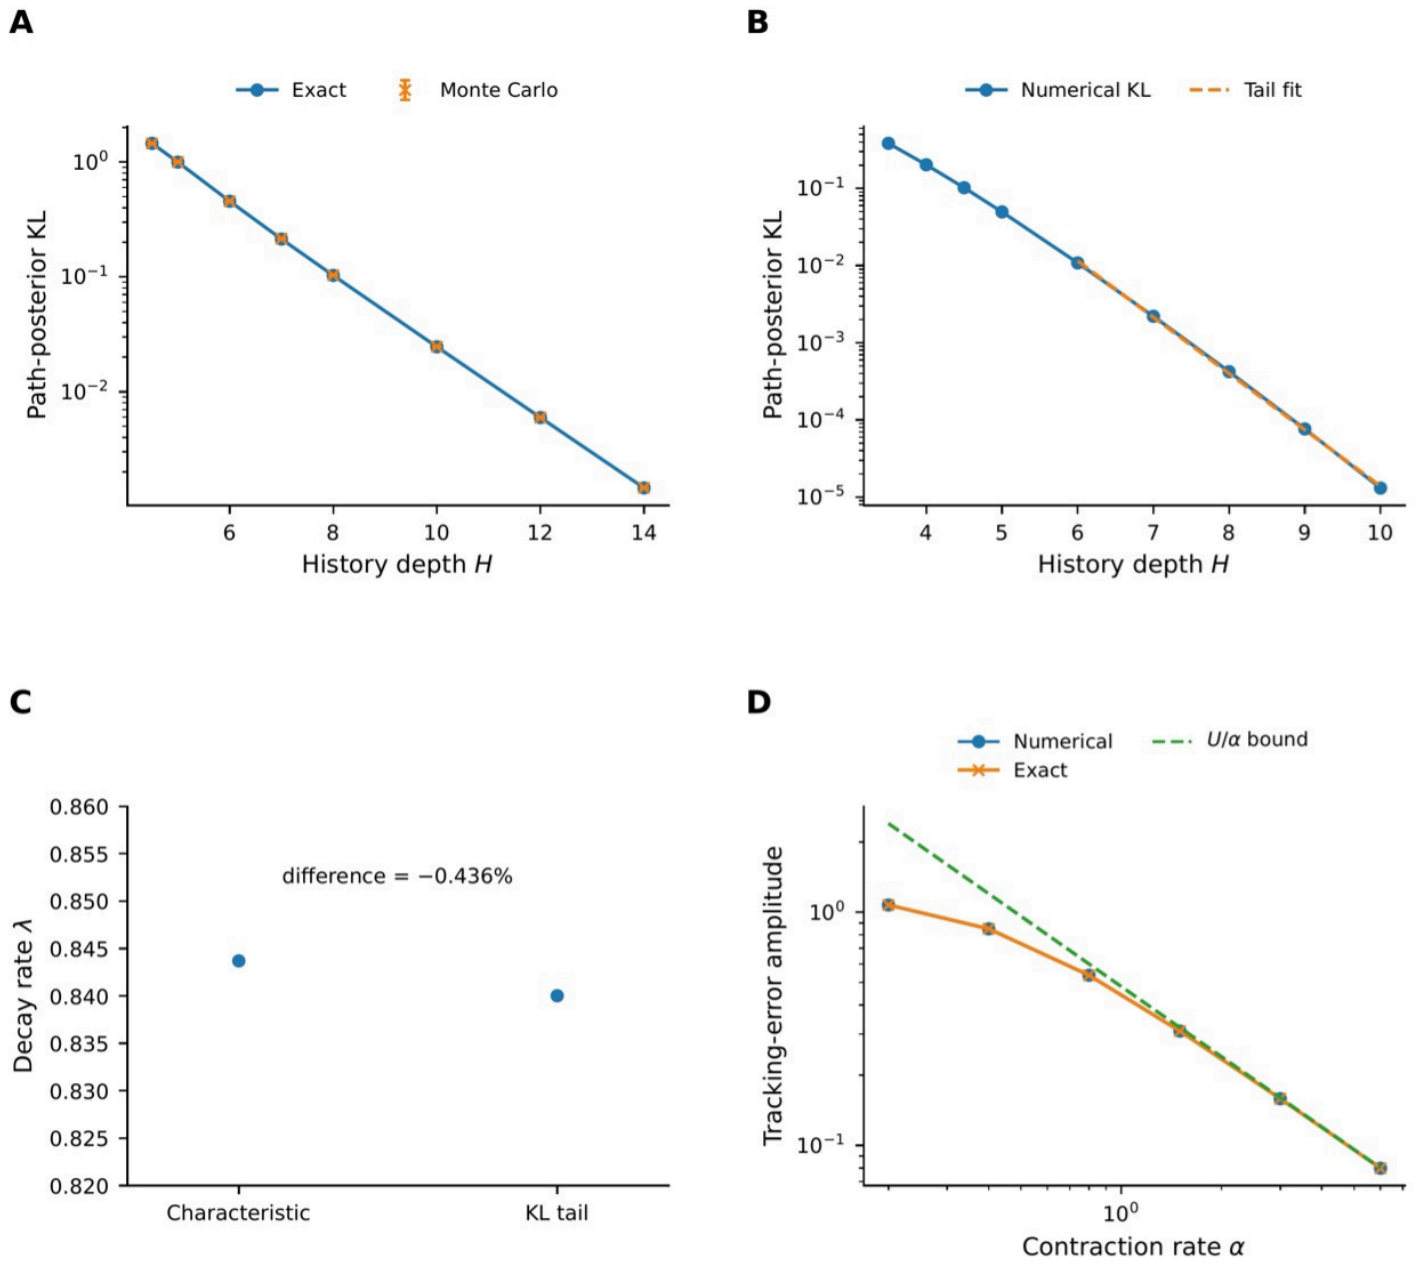

Figure 3

Exact and numerical verification of finite-history predictions.
(A) Exact path-posterior KL divergence and Monte Carlo estimates as a function of retained history depth *H* in the multimode Gaussian system. Monte Carlo error bars denote the corresponding standard errors. (B) Posterior KL divergence in the noncommuting delayed matrix system. The dashed line shows the exponential tail fit over the asymptotic regime. (C) Comparison between the decay rate obtained from the dominant characteristic root and that obtained independently from the posterior-KL tail; the estimates differed by −0.436%. (D) Moving-posterior tracking error as a function of contraction rate α. Numerical amplitudes are compared with the exact sinusoidal solution and the analytical upper bound *U*/α.

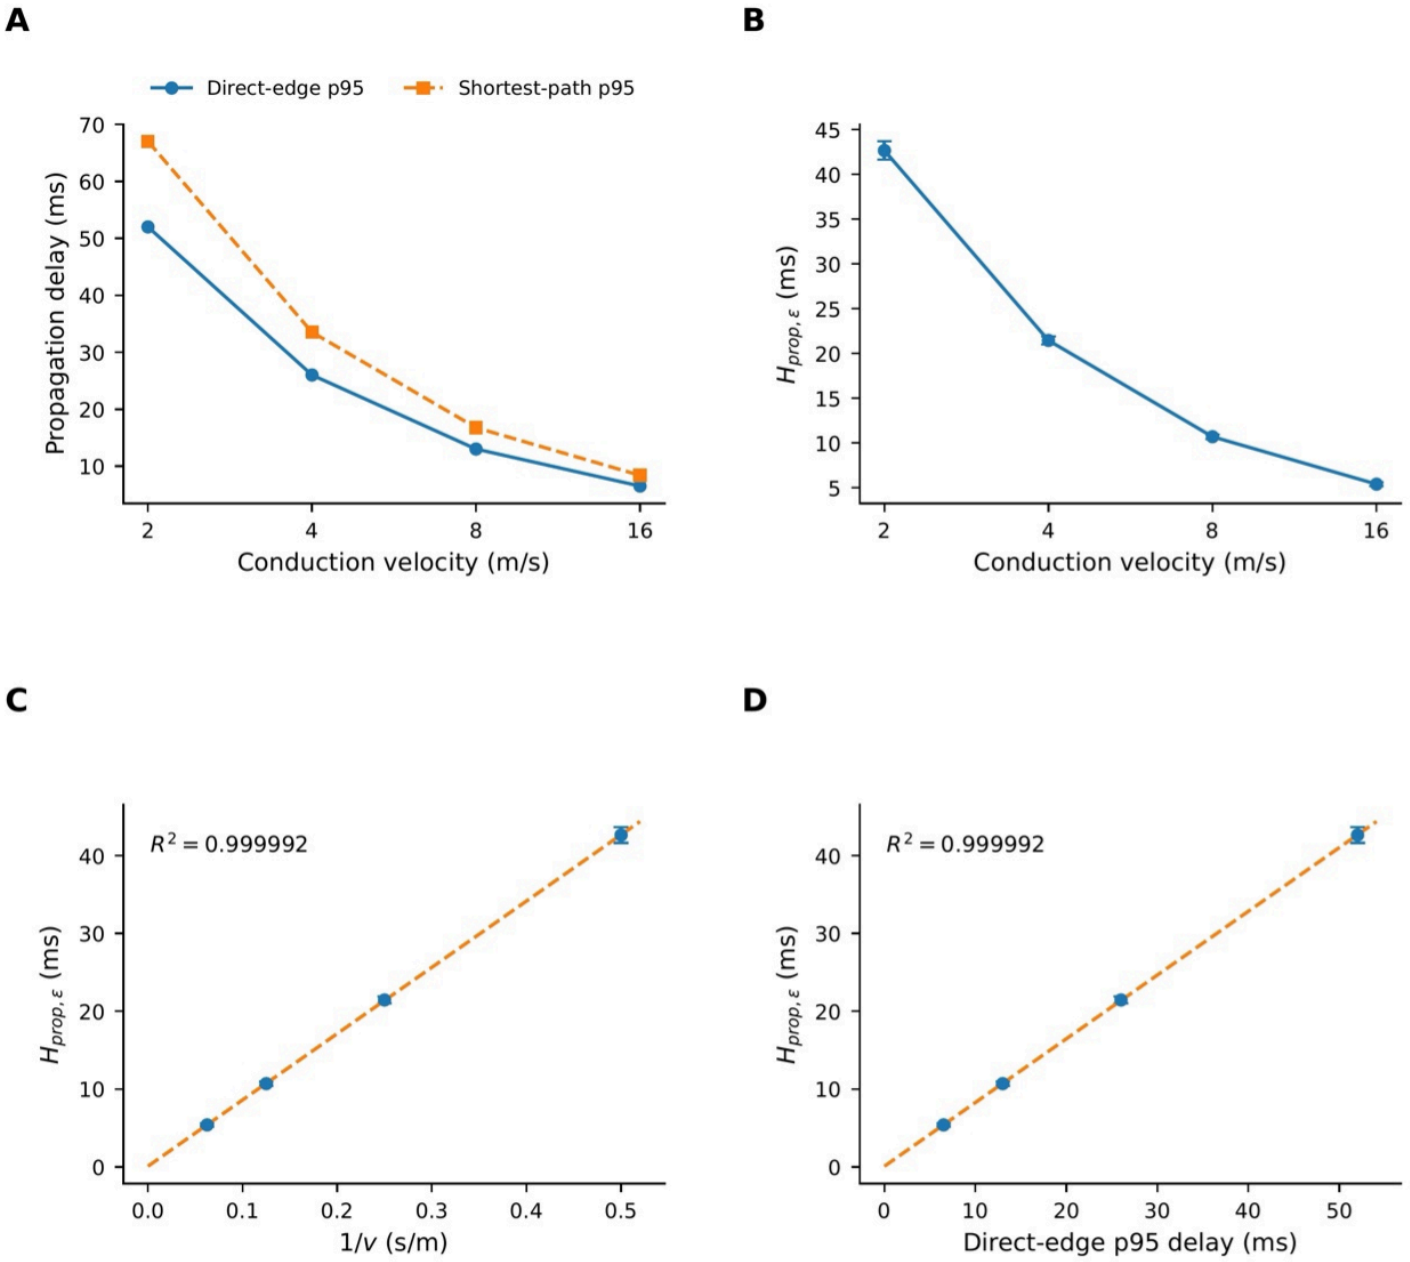

Figure 4

Whole-brain propagation depth scales with conduction geometry.
(A) Direct-edge and shortest-path 95th-percentile propagation delays as functions of conduction velocity in the 76-node structural network. (B) Operational propagation depth *H*prop,ε at the primary threshold $\varepsilon = 1 \times 10^{-6}$ nats/node/ms for conduction velocities of 2, 4, 8, and 16 m/s. Points show means and error bars show SD across 10 paired stochastic realizations. (C) *H*prop,ε as a function of reciprocal conduction velocity. The dashed line shows the linear fit. (D) *H*prop,ε as a function of the 95th percentile of discretized direct-edge delay. The dashed line shows the linear fit. Both relationships were nearly linear, with $R^2 = 0.999992$.

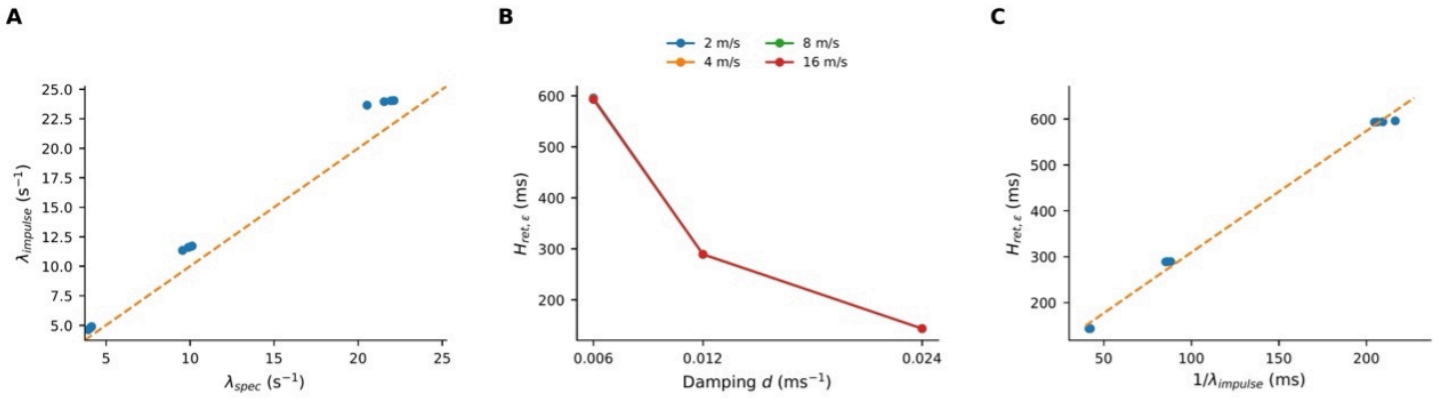

Figure 5

Dissipative dynamics determine perturbation-retention depth.
(A) Relationship between the asymptotic spectral decay rate λspec obtained from the augmented delayed transition matrix and the observable perturbation-decay rate λimpulse obtained from physical-state impulse responses. The dashed line denotes equality. (B) Conservative retention depth *H*ret,ε at the primary energy threshold $\varepsilon = 0.001$ as a function of local damping *d*for conduction velocities of 2, 4, 8, and 16 m/s. Increasing damping markedly reduced retained-history depth, whereas conduction velocity had comparatively little effect within a fixed damping condition. (C) Relationship between *H*ret,ε and the observable dissipative timescale 1/ λimpulse. The dashed line shows the linear fit.

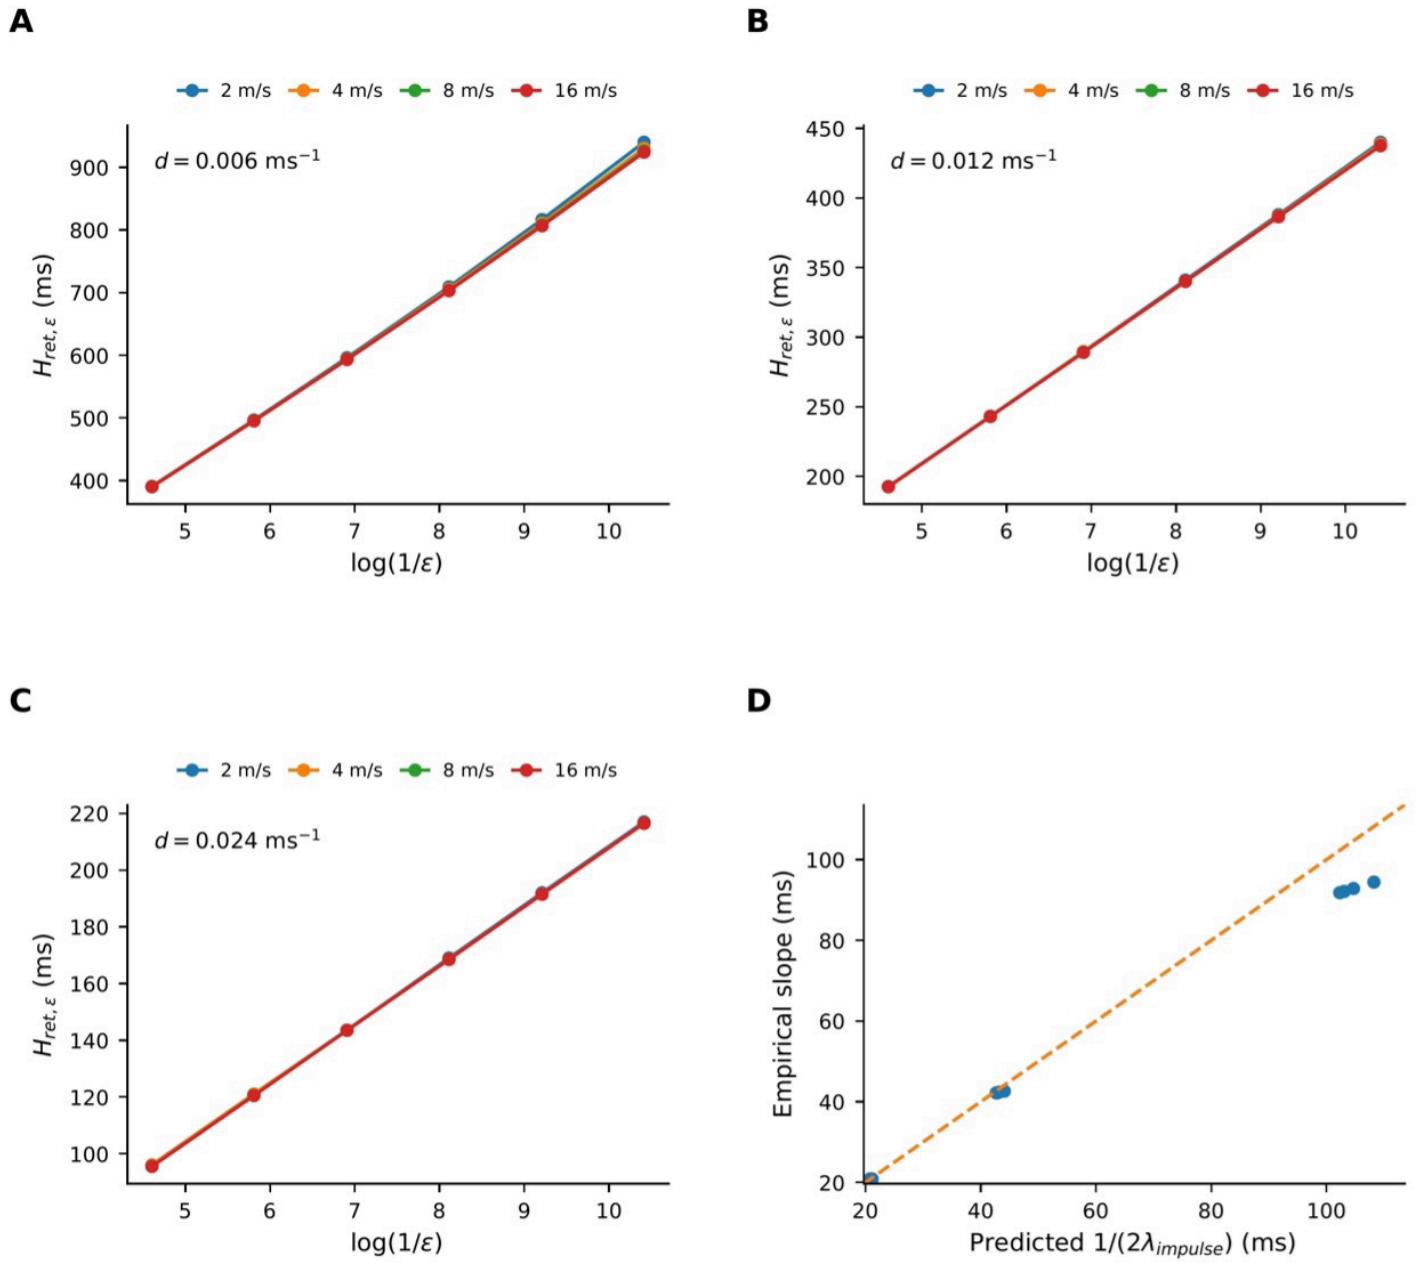

Figure 6

Retention depth scales logarithmically with required precision.
(A–C) Conservative retention depth *H*ret,ε as a function of log(1/ε) at local damping values of 0.006, 0.012, and 0.024 ms⁻¹, respectively. Each panel contains conduction velocities of 2, 4, 8, and 16 m/s. Across all conditions, retention depth increased approximately linearly with log(1/ε). (D) Empirical slopes of the *H*ret,ε–log(1/ε) relationships plotted against the observable prediction 1/(2λimpulse). The dashed line denotes equality.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- GeometrySupplementalInformation.docx