

RESEARCH ARTICLE | OCTOBER 06 2026

Finite-time scalar retention and slow mixing tails in uncertain passive-scalar transport

Srimanta Santra 



Physics of Fluids 38, 103606 (2026)

<https://doi.org/10.1063/5.0345598>



Articles You May Be Interested In

Reinforcement learning of fixed-budget moving-rod stirring for passive scalar mixing in Navier–Stokes–Brinkman flows

Physics of Fluids (August 2026)

Transport–reaction coupling under hydrodynamic mixing in batch systems

Physics of Fluids (July 2026)

Continuum model for extraction and retention in porous media

Physics of Fluids (December 2023)



Physics of Fluids

Special Topics Open
for Submissions

[Learn More](#)



Finite-time scalar retention and slow mixing tails in uncertain passive-scalar transport

Cite as: Phys. Fluids **38**, 103606 (2026); doi: [10.1063/5.0345598](https://doi.org/10.1063/5.0345598)

Submitted: 28 May 2026 · Accepted: 20 August 2026 ·

Published Online: 6 October 2026



View Online



Export Citation



CrossMark

Srimanta Santra^{1,2,a)}

AFFILIATIONS

¹Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA

²Institute of Control Systems, Hamburg University of Technology, Harburger Schloßstr. 22A, 21079 Hamburg, Germany

^{a)}Author to whom correspondence should be addressed: santrasrimanta7@gmail.com

ABSTRACT

Mixing of a passive scalar in a time-dependent flow is usually judged from the variance decay of a nominal flow. In many applications, though, stirring protocols and flow parameters are uncertain, and the more relevant question is whether the required level of homogenization is reached within a finite time. We study this reliability question for incompressible advection–diffusion with uncertain forcing parameters, using a time-periodic double-gyre flow as the main benchmark and a periodically stirred cavity as a secondary test. Direct numerical simulations show that delayed mixing is not a fluctuation around the mean variance decay curve. Instead, slow realizations retain coherent scalar remnants near weakly exchanging regions of the flow, which produces a persistent upper tail in the mixing time distribution. To identify these slow events without running a full simulation for every parameter sample, we build a parameter- and phase-dependent reduced transport operator by combining a proper orthogonal decomposition-based Koopman representation with a polynomial-chaos expansion in the uncertain parameters. From this operator, we compute a finite-time scalar-retention score based on the leading singular direction of the learned propagator in a scalar-energy metric. The diagnostic ranks parameter realizations by their tendency to retain scalar energy, identifies the realizations contributing to the upper tail of the mixing time distribution, and recovers the scalar structures responsible for delayed homogenization. Reliable mixing under uncertainty is therefore controlled not only by the mean variance decay, but also by the least-damped scalar directions of the uncertain transport operator, which suggests that designs targeting reliable mixing should act on these directions directly.

© 2026 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/5.0345598>

I. INTRODUCTION

Mixing of a passive scalar by advection and diffusion controls how quickly scalar inhomogeneity is removed in a wide range of natural and engineered flows, including chemical and biomanufacturing processes, stirred reactors, microfluidic devices, environmental transport, and combustion. In laminar and weakly turbulent flows, mixing is often produced by repeated stretching and folding of material elements, which creates thin scalar filaments that are smoothed by molecular diffusion.^{1–3} Classical diagnostics therefore focus on stretching rates, finite-time Lyapunov exponents, scalar variance decay, transfer operators, and spectral properties of advection–diffusion dynamics.^{4–7}

A complementary physical interpretation is provided by the classical lamellar description of mixing. Incompressible advection redistributes and deforms the scalar field but does not directly reduce its L^2 energy. Stretching and folding compress scalar patches into thin lamellae, and molecular diffusion then acts across their decreasing

transverse scale.⁸ This description has also been extended from two-dimensional strips to three-dimensional diffusive sheets.⁹ More generally, scalar decay depends on the accumulated deformation history, because this history determines how quickly small scalar scales and strong gradients are formed.¹⁰ The diffuselet formulation provides a related reduced description of the interaction between deformation and diffusion.¹¹ Recent hyperbolic formulations give a complementary deterministic Lagrangian description of advection–diffusion transport.¹² In the uncertain setting considered here, slow realizations correspond to parameter values for which the flow delays the formation of thin scalar structures and therefore delays scalar dissipation and threshold crossing.

Many mixing problems, however, are not described by a single known flow. Stirring amplitude, actuation frequency, vortex strength, and similar flow parameters may vary between realizations. In that setting, the relevant question is not only whether a nominal flow

mixes rapidly, but whether the required level of mixing is reached reliably. A flow can have a favorable mean decay rate while still producing a small but important set of slow realizations. These realizations control the probability that homogenization has not been achieved by a specified operating time. In this work, we study this question for two canonical time-dependent flows, a time-periodic double-gyre (DG) flow and a periodically stirred cavity (SC), both with uncertain forcing parameters.

To make this reliability question precise, we treat the uncertain parameters as a random vector ξ with a prescribed probability density, so that each realization of ξ corresponds to one possible flow. For each realization, we define the mixing time $T_m(\xi)$ as the first time at which a scalar variance measure falls below a prescribed tolerance, so that reaching the tolerance later means slower mixing for that particular flow. The quantity of interest is then the exceedance probability

$$\mathbb{P}\{T_m(\xi) > \tau\}, \quad (1)$$

which measures the probability that the scalar has not yet reached the tolerance by time τ . This tail probability is not determined by the mean decay rate alone. Two uncertain flows can have similar average scalar variance decay while having very different probabilities of remaining unmixed at time τ . The tail of the mixing time distribution therefore contains information that is missed by a purely nominal or mean-flow analysis.

A useful way to understand these delayed events is through the geometry of transport. In deterministic flows, coherent structures organize scalar transport by separating regions that exchange material slowly over finite times. Finite-time Lyapunov exponent ridges identify material structures that shape transport.^{13,14} Transfer-operator and coherent-set approaches identify regions with weak finite-time exchange.^{15–17} Related passive-scalar studies have also connected persistent concentration patterns and mixing quality to strange-eigenmode behavior in chaotic open flows.^{18,19} Koopman operator methods provide a complementary description through the evolution of observables of the flow and scalar field.^{20,21} Each of these approaches is designed for a single, fully specified flow. They have clarified transport organization for fixed velocity fields, but they do not, by themselves, identify which members of an uncertain family of flows are responsible for the slow tail of a mixing time distribution.

The central point of this paper is that slow mixing under parametric uncertainty is governed by parameter-conditioned coherent scalar structures. These structures are not simply the transport barriers of the nominal or mean flow. They are scalar patterns selected by particular members of the uncertain transport-operator family. For these parameter realizations, scalar gradients remain organized near weakly exchanging regions for long enough to delay variance decay. Thus, the object relevant to reliability is not only a coherent structure in physical space, but a coherent scalar structure conditioned on the uncertain parameters.

This viewpoint can be stated in terms of finite-time scalar retention. For the mean-zero passive scalar, diffusion makes the physical L^2 variance non-increasing for each fixed parameter realization. Slow mixing therefore does not require growth of total scalar variance. It occurs when the finite-time transport operator has directions that are only weakly damped over the time interval of interest. Even when all directions eventually decay, some directions can decay much more slowly than others over the horizon of interest. This is related to

non-normal finite-time dynamics: eigenvalues or asymptotic decay rates alone may not identify the directions that dominate finite-time response.^{22,23} Similar operator and eigenmode viewpoints have been used to analyze mixing and transient distributive transport in recirculating or periodically driven flows.^{24,25} In the present problem, the important directions are scalar structures that remain least damped for particular parameter values.

To study this mechanism, we build a reduced stochastic operator model from high-fidelity scalar simulations. Mean-removed scalar fields are first compressed into low-dimensional proper orthogonal decomposition (POD) coordinates. These reduced coordinates are advanced by a parameter- and phase-dependent Koopman approximation. The dependence of this reduced transport operator on uncertain parameters is represented by a polynomial-chaos expansion.^{26,27} The purpose is not only to obtain a fast surrogate for scalar snapshots. The purpose is to approximate the uncertain family of transport operators well enough to evaluate finite-time scalar retention and to identify the parameter values responsible for delayed mixing.

The novelty is not the use of Koopman approximation or polynomial chaos alone. Rather, the novelty is to use a parameter-dependent transport-operator representation to identify least-damped scalar structures that control finite-time mixing failure under uncertainty. The paper makes three contributions. First, it formulates mixing reliability as a tail problem for uncertain scalar transport, rather than as a mean variance decay problem. Second, it shows that slow mixing events are associated with coherent scalar structures that are strengthened or sustained by particular parameter realizations. Third, it introduces a finite-time retention diagnostic, computed from a Koopman polynomial-chaos representation of the uncertain transport operator, that identifies parameter regimes responsible for mixing failure.

The double-gyre flow serves as the primary benchmark, and the periodically stirred cavity provides an independent test of whether the same workflow extends beyond that geometry. In the primary results, the slowest realizations are not random deviations from the mean decay curve. They correspond to identifiable parameter regimes in which coherent unmixed scalar structures persist over finite times. The finite-time retention diagnostic separates these slow realizations from the rapidly mixing bulk and predicts the tail of the mixing time distribution. These results suggest that reliability-aware mixing design should not only improve nominal mixing rates, but should also suppress weak finite-time contraction in the uncertain scalar transport operator.

II. PROBLEM FORMULATION: UNCERTAIN PASSIVE SCALAR MIXING

We consider the transport of a passive scalar in an incompressible flow whose dynamics depend on uncertain parameters. Let $\Omega \subset \mathbb{R}^2$ be the spatial domain and let $c(\mathbf{x}, t; \xi)$ denote the scalar concentration at position $\mathbf{x} \in \Omega$ and time $t \geq 0$. The scalar evolves according to the advection–diffusion equation

$$\frac{\partial c}{\partial t} + \mathbf{u}(\mathbf{x}, t; \xi) \cdot \nabla c = \text{Pe}^{-1} \Delta c, \quad \nabla \cdot \mathbf{u} = 0. \quad (2)$$

Here, $\mathbf{u}(\mathbf{x}, t; \xi)$ is a prescribed incompressible velocity field, Pe is the Peclet number, and $\xi \in \Xi \subset \mathbb{R}^{d_\xi}$ is a vector of uncertain but time-independent parameters. The probability density of ξ is denoted by

$\rho(\xi)$. In the examples considered later, ξ represents quantities such as forcing amplitude, oscillation frequency, vortex strength, and switching period.

Throughout this work, the initial scalar field c_0 is kept fixed. Randomness enters only through the parameter vector ξ , so the variability of $T_m(\xi)$ reflects uncertainty in the transport dynamics rather than uncertainty in the initial condition.

We assume that the velocity field is sufficiently smooth for (2) to be well posed. Both benchmark problems are treated as closed domains. The prescribed velocity satisfies the no-penetration condition

$$\mathbf{u} \cdot \mathbf{n} = 0 \quad \text{on } \partial\Omega, \quad (3)$$

and the passive scalar satisfies the homogeneous no-flux condition

$$\nabla c \cdot \mathbf{n} = 0 \quad \text{on } \partial\Omega. \quad (4)$$

The spatial mean of the scalar is

$$\bar{c} = \frac{1}{|\Omega|} \int_{\Omega} c(\mathbf{x}, t; \xi) d\mathbf{x}. \quad (5)$$

Integrating (2) over Ω and using incompressibility together with the no-penetration and no-flux boundary conditions shows that $\bar{c}$ is conserved in time. It therefore provides a well-defined reference value against which the degree of homogenization can be measured.

To measure the degree of mixing, we use the normalized scalar variance

$$M(t, \xi) = \frac{\|c(\cdot, t; \xi) - \bar{c}\|_{L^2(\Omega)}^2}{\|c(\cdot, 0; \xi) - \bar{c}\|_{L^2(\Omega)}^2}. \quad (6)$$

Variance-based measures are widely used in passive-scalar mixing because they quantify the decay of large-scale scalar inhomogeneity under the combined action of advection and diffusion.^{4,5} At $t = 0$, $M(0, \xi) = 1$. As mixing proceeds, advection creates fine scalar gradients and diffusion reduces the scalar variance, so that $M(t, \xi)$ decreases toward zero.

For a prescribed tolerance $\epsilon_{\text{mix}} \in (0, 1)$, we define the mixing time

$$T_m(\xi) = \inf\{t > 0 : M(t, \xi) \leq \epsilon_{\text{mix}}\}. \quad (7)$$

Thus, $T_m(\xi)$ is the first time at which the scalar field reaches the required level of homogeneity for a given realization of the uncertain parameters. Since ξ is random, the mixing time is also a random variable. This definition is consistent with finite-time descriptions of mixing, where the objective is not only asymptotic decay but the time required to reach a prescribed level of scalar uniformity.^{4,5}

The quantity of interest in this work is not the average value of T_m alone, but the probability that T_m exceeds a prescribed deadline. For a deadline $\tau > 0$, we define the mixing-failure probability

$$P_f(\tau) = \mathbb{P}\{T_m(\xi) > \tau\} = \int_{\{\xi \in \Xi : T_m(\xi) > \tau\}} \rho(\xi) d\xi. \quad (8)$$

This probability measures the chance that the scalar field has not mixed sufficiently by time τ . The distinction between mean mixing and reliable mixing is important. A flow may have a small expected mixing time and still produce a non-negligible probability of slow

events. Conversely, a rapid initial decay of scalar variance does not necessarily imply a small tail probability. The tail of $T_m(\xi)$ can be controlled by parameter values for which scalar gradients remain organized near weakly exchanging regions of the flow for finite times.

It is useful to express this problem in operator form. For each parameter realization ξ , let $\mathcal{S}_t(\xi)$ denote the solution operator of (2) so that

$$c(\cdot, t; \xi) = \mathcal{S}_t(\xi)c_0. \quad (9)$$

For fixed ξ , classical mixing diagnostics describe the decay properties of $\mathcal{S}_t(\xi)$ through scalar variance decay and stretching rates.^{4,5} Coherent-structure and transfer-operator methods describe weak exchange and transport barriers.^{15,16} Koopman methods provide a complementary spectral description of observable evolution.^{20,21} In the present setting, however, the object of interest is the uncertain family of operators

$$\{\mathcal{S}_t(\xi) : \xi \in \Xi\}. \quad (10)$$

Slow mixing occurs when particular members of this family sustain scalar structures that decay more slowly than expected from the nominal flow or from the ensemble-averaged response.

This formulation leads to two questions. First, which parameter realizations contribute most strongly to the tail probability $P_f(\tau)$? Second, can these realizations be detected from a reduced representation of the uncertain transport operator, without performing a full high-fidelity simulation for every sample of ξ ? Sections III and IV address these questions by introducing a finite-time scalar-retention diagnostic and a Koopman polynomial-chaos reduced model for the uncertain transport dynamics.

III. FINITE-TIME SCALAR RETENTION AND SLOW MIXING

We now introduce the diagnostic used to identify parameter realizations that produce delayed mixing. Let

$$q(\mathbf{x}, t; \xi) = c(\mathbf{x}, t; \xi) - \bar{c} \quad (11)$$

denote the mean-removed scalar field. Because $\bar{c}$ is independent of both space and time (Sec. II), the mean-removed field q satisfies the same advection-diffusion equation as c , i.e.,

$$\frac{\partial q}{\partial t} + \mathbf{u}(\mathbf{x}, t; \xi) \cdot \nabla q = \text{Pe}^{-1} \Delta q. \quad (12)$$

Multiplying (12) by q and integrating over Ω , together with incompressibility and the no-penetration and no-flux boundary conditions, gives the scalar-energy identity

$$\frac{1}{2} \frac{d}{dt} \|q(\cdot, t; \xi)\|_{L^2(\Omega)}^2 = -\text{Pe}^{-1} \|\nabla q(\cdot, t; \xi)\|_{L^2(\Omega)}^2 \leq 0. \quad (13)$$

Thus, diffusion makes the physical scalar variance non-increasing for each fixed parameter realization. Slow mixing is therefore not caused by growth of the total scalar variance above its initial value. It is caused by weak finite-time contraction: for some parameter values, coherent scalar structures remain organized long enough that diffusion removes variance much more slowly than in typical realizations.

Let $S_t(\xi)$ denote the solution operator of (12) on mean-zero scalar fields. For a fixed initial condition q_0 , the mixing index introduced in (6) can be written as

$$M(t, \xi) = \frac{\|S_t(\xi)q_0\|_{L^2(\Omega)}^2}{\|q_0\|_{L^2(\Omega)}^2}. \quad (14)$$

This quantity measures the decay of one prescribed scalar field. To identify the structures most strongly retained by a given parameter realization, we consider the finite-time operator gain

$$\mathcal{G}_t(\xi) = \sup_{q \neq 0} \frac{\|S_t(\xi)q\|_{L^2(\Omega)}}{\|q\|_{L^2(\Omega)}}. \quad (15)$$

The maximizing direction in (15) when it exists is the scalar structure that is least damped over the time interval $[0, t]$. In a passive-scalar problem with diffusion, $\mathcal{G}_t(\xi)$ is bounded by one in the exact physical L^2 energy norm. Large values of $\mathcal{G}_t(\xi)$ therefore indicate weak contraction, not indefinite growth. The physical interpretation of (15) is straightforward. If $\mathcal{G}_t(\xi)$ is small, all scalar structures are rapidly damped over the time interval. If $\mathcal{G}_t(\xi)$ is close to one, at least one scalar structure is retained over that interval. Such a structure can delay the crossing of the mixing threshold in (7). The associated parameter value then contributes to the slow tail of the mixing time distribution.

This diagnostic is related to coherent-structure ideas, but it is not the same object as a deterministic transport barrier. A deterministic barrier describes transport organization for one fixed flow. Here, the relevant structure depends on ξ . It is selected by the parameter-dependent transport operator $S_t(\xi)$ and is important only if it delays scalar decay over the finite time horizon of interest. This parameter-conditioned viewpoint is consistent with finite-time transport analysis and coherent-set descriptions of weak exchange.^{14,16}

In practice, the explicitly time-dependent flow is represented by a phase-conditioned family of reduced operators. The reduced coordinates and scalar-energy metric used in this section are constructed formally in Sec. IV; here, we introduce them only to the extent needed to define the retention diagnostic. The reduced state is taken directly as the POD coordinate, so that $y_k = z_k \in \mathbb{R}^r$ and $p = r$. For a fixed parameter realization,

$$y_{k+1} = K(\xi, \theta_k)y_k, \quad (16)$$

where $t_k = k\Delta t_s$ and θ_k is the known forcing or switching phase. For the double-gyre flow, $\theta_k = \omega t_k \bmod 2\pi$. For the stirred cavity, $\theta_k = 2\pi t_k / T_s \bmod 2\pi$, and θ_k selects the currently active vortex. The finite-time reduced propagator over N snapshot intervals is the ordered product

$$K_{N:0}(\xi) = K(\xi, \theta_{N-1}) \cdots K(\xi, \theta_1)K(\xi, \theta_0). \quad (17)$$

The associated physical retention horizon is $T_{\text{ret}} = N\Delta t_s$. Let $H \succ 0$ denote the scalar-energy metric in the POD coordinates, as constructed in Sec. IV. The metric-weighted finite-time propagator is

$$A_N(\xi) = H^{1/2}K_{N:0}(\xi)H^{-1/2}. \quad (18)$$

The largest finite-time retention is

$$\Gamma_N(\xi) = \sigma_{\max}(A_N(\xi)), \quad (19)$$

and the normalized retention score is

$$G_N(\xi) = \frac{\Gamma_N(\xi)}{\text{median}_{\xi \in \mathcal{E}_{\text{dev}}} \Gamma_N(\xi)}. \quad (20)$$

The normalization is fixed from the development ensemble before any held-out or Monte Carlo evaluation. Thus, $G_N(\xi) > 1$ indicates stronger-than-typical retention in the resolved POD subspace and does not imply physical growth of scalar variance.

The operator score and the mixing time measure related but distinct quantities. The mixing time $T_m(\xi)$ is computed for the prescribed initial scalar field q_0 , whereas $\Gamma_N(\xi)$ is maximized over all resolved scalar-energy directions. To connect the two, we also compute the fixed-initial-condition retention

$$g_N^{\text{fix}}(\xi) = \frac{\|K_{N:0}(\xi)y_0\|_H}{\|y_0\|_H}, \quad \|y\|_H = \sqrt{y^T H y}, \quad (21)$$

where $y_0 = z_0$ is the POD representation of the prescribed initial scalar field. By construction,

$$g_N^{\text{fix}}(\xi) \leq \Gamma_N(\xi). \quad (22)$$

Let $v_1(\xi)$ denote the leading right singular vector of $A_N(\xi)$. We additionally report the alignment

$$a_0(\xi) = \frac{|v_1(\xi)^T H^{1/2} y_0|}{\|H^{1/2} y_0\|_2}. \quad (23)$$

The fixed-initial-condition retention is the direct connection to the initial field used in the mixing time calculation; the single-vector alignment is reported only as a secondary diagnostic. The working hypothesis tested numerically is that parameter realizations with large G_N are more likely to produce delayed mixing, while g_N^{fix} quantifies the corresponding retention of the prescribed initial scalar field.

Section IV describes the POD coordinates, scalar-energy metric, and phase-conditioned Koopman polynomial-chaos operator.

IV. REDUCED OPERATOR MODEL

We now describe how the phase-conditioned, parameter-dependent operator $K(\xi, \theta)$ in (16) is obtained from scalar-field data. The reduced representation is used both to predict scalar snapshots and, more importantly, to evaluate the mixing time and finite-time retention over many parameter realizations at low cost.

Let

$$\mathcal{D} = \left\{ c_k^{(j)}, \xi_j \right\}_{j=1, k=0}^{N_s, N_t} \quad (24)$$

denote a set of high-fidelity scalar snapshots, where $c_k^{(j)} = c(\cdot, t_k; \xi_j)$ is the scalar field at time t_k for the parameter realization ξ_j . The parameter samples ξ_j are drawn from the probability density $\rho(\xi)$. The snapshots are generated by solving the advection-diffusion equation (2) for a finite training set of parameter values.

A. Mean-removed POD representation of scalar fields

Passive-scalar snapshots typically lie close to a low-dimensional linear subspace, because the dominant scalar structures in a stirred flow are much lower in dimension than the full grid. POD is the standard tool for extracting such a subspace from snapshot data,^{28,29} and we use it here to obtain a compressed coordinate system for the mean-removed scalar field. For each scalar snapshot, the spatial mean is first removed,

$$\mathbf{q}_k^{(j)} = \mathbf{c}_k^{(j)} - \bar{\mathbf{c}}_k^{(j)} \mathbf{1}. \quad (25)$$

The mean-removed training snapshots are arranged as row vectors in a snapshot matrix. A rank- r proper orthogonal decomposition basis $\mathbf{V}_r \in \mathbb{R}^{n_s \times r}$ is then obtained from its singular-value decomposition. The implementation uses Euclidean-orthonormal POD modes,

$$\mathbf{V}_r^\top \mathbf{V}_r = \mathbf{I}_r. \quad (26)$$

The POD coordinates and reconstructed mean-removed field are

$$\mathbf{z}_k^{(j)} = \mathbf{V}_r^\top \mathbf{q}_k^{(j)}, \quad (27)$$

$$\hat{\mathbf{q}}_k^{(j)} = \mathbf{V}_r \mathbf{z}_k^{(j)}. \quad (28)$$

The POD rank is selected using only the training trajectories and is then fixed for validation and testing. The POD coordinates are used as identity observables without an additional nonlinear lifting. We therefore set $\mathbf{y}_k = \mathbf{z}_k$, $p = r$.

B. Scalar-energy metric

The retention diagnostic in Sec. III is defined in the physical L^2 scalar-energy norm. For the diagnostic to remain physically meaningful in the reduced coordinates, we need a metric $\mathbf{H}$ on those coordinates that agrees with the discrete L^2 energy of the reconstructed field. This is the standard energy-consistent inner product used in energy-based reduced-order modeling.³⁰ On the uniform cell-centered grid, let $\Delta V = \Delta x \Delta y$. Because the POD modes are Euclidean orthonormal [Eq. (26)], the discrete L^2 energy of the POD reconstruction is

$$\begin{aligned} \|\hat{\mathbf{q}}_k\|_{L_h^2}^2 &= \Delta V \hat{\mathbf{q}}_k^\top \hat{\mathbf{q}}_k \\ &= \Delta V \mathbf{z}_k^\top \mathbf{z}_k. \end{aligned} \quad (29)$$

Accordingly, the scalar-energy metric in the reduced coordinates is

$$\mathbf{H} = \Delta x \Delta y \mathbf{I}_r \quad (30)$$

and

$$\mathbf{y}_k^\top \mathbf{H} \mathbf{y}_k = \|\hat{\mathbf{q}}_k\|_{L_h^2}^2. \quad (31)$$

The matrix $\mathbf{H}$ is positive definite. The difference between $\mathbf{y}_k^\top \mathbf{H} \mathbf{y}_k$ and the energy of the full high-fidelity field is therefore caused only by POD truncation and is evaluated on independent held-out snapshots in Sec. V D.

C. Phase-conditioned Koopman approximation

The scalar-field dynamics under a periodically forced or switched flow are not autonomous, so a single time-invariant linear operator cannot represent the transition from $\mathbf{y}_k$ to $\mathbf{y}_{k+1}$. We therefore approximate the reduced dynamics by a phase-conditioned Koopman operator, following the extended dynamic mode decomposition and Koopman-operator viewpoint for non-autonomous systems.^{20,21} Because the flow is explicitly time dependent, the transition from $\mathbf{y}_k$ to $\mathbf{y}_{k+1}$ depends on the forcing or switching phase θ_k as well as on ξ . The reduced dynamics are therefore approximated by

$$\mathbf{y}_{k+1}^{(j)} \approx \mathbf{K}(\xi_j, \theta_k) \mathbf{y}_k^{(j)}. \quad (32)$$

The deterministic phase dependence is represented separately from the random-parameter dependence. For the double gyre, the phase basis is

$$\mathbf{b}_{\text{DG}}(\theta) = [1, \cos \theta, \sin \theta, \cos(2\theta), \sin(2\theta)]^\top. \quad (33)$$

For the stirred cavity, the velocity is constant within each half-cycle, and the phase basis is

$$\mathbf{b}_C(\theta) = [1, s(\theta)]^\top, \quad s(\theta) = \begin{cases} 1, & 0 \leq \theta < \pi, \\ -1, & \pi \leq \theta < 2\pi. \end{cases} \quad (34)$$

This representation does not treat the deterministic phase variables as additional independent random parameters.

D. Polynomial-chaos representation of the operator

The dependence of the reduced operator on the uncertain parameters ξ needs to be represented in a way that admits inexpensive evaluation at any new sample ξ . When ξ has a smooth influence on the dynamics, generalized polynomial chaos provides a compact spectral representation for such parametric dependence.^{26,27} For the three uniformly distributed uncertain parameters used here, the corresponding orthogonal basis is the total-degree Legendre basis.³¹

Let $\{\Phi_\alpha\}_{\alpha \in \mathcal{I}_p}$ be the total-degree Legendre basis associated with the three independent uniform uncertain parameters. The phase-conditioned reduced operator is represented as

$$\mathbf{K}(\xi, \theta) = \sum_{\alpha \in \mathcal{I}_p} \sum_{\ell=1}^{n_\theta} \mathbf{K}_{\alpha\ell} \Phi_\alpha(\xi) \mathbf{b}_\ell(\theta), \quad (35)$$

where $n_\theta = 5$ for the double gyre and $n_\theta = 2$ for the cavity. This is a polynomial-chaos representation in the uncertain parameters combined with a deterministic Fourier or switching basis in phase. It is not an intrusive stochastic Galerkin discretization of the full advection-diffusion equation. Each uncertain parameter ξ_i is affinely mapped to $[-1, 1]$ using the endpoints of its uniform support. The Legendre basis $\{\Phi_\alpha\}$ is normalized so that $\int_{[-1,1]^3} \Phi_\alpha \Phi_\beta d\xi / 2^3 = \delta_{\alpha\beta}$ (orthonormal). Before the ridge regression in Eq. (36), each parameter-phase feature $\Phi_\alpha(\xi_j) \mathbf{b}_\ell(\theta_k^{(j)})$ is standardized to zero mean and unit variance across the development set; the regularization $\lambda = 10^{-4}$ therefore acts on standardized features.

The coefficient matrices $\mathbf{K}_{\alpha\ell}$ are identified jointly from all reduced snapshot pairs $(\mathbf{y}_k^{(j)}, \mathbf{y}_{k+1}^{(j)})$ in the development set by solving

$$\begin{aligned} \min_{\{\mathbf{K}_{\alpha\ell}\}} \sum_{j=1}^{N_s} \sum_{k=0}^{N_t-1} \left\| \mathbf{y}_{k+1}^{(j)} - \left(\sum_{\alpha \in \mathcal{I}_p} \sum_{\ell=1}^{n_\theta} \mathbf{K}_{\alpha\ell} \Phi_\alpha(\xi_j) \mathbf{b}_\ell(\theta_k^{(j)}) \right) \mathbf{y}_k^{(j)} \right\|_2^2 \\ + \lambda \sum_{\alpha \in \mathcal{I}_p} \sum_{\ell=1}^{n_\theta} \|\mathbf{K}_{\alpha\ell}\|_F^2. \end{aligned} \quad (36)$$

The regularization parameter $\lambda \geq 0$ improves conditioning and limits overfitting. For a new parameter realization, the operator is evaluated at each known forcing or switching phase and the reduced trajectory is advanced using the ordered product in (17).

E. Prediction of scalar variance and mixing time

Once the operator $\mathbf{K}(\xi, \theta)$ is available, a full scalar trajectory for any new parameter realization can be obtained by propagation in the

POD coordinates and reconstruction in physical space. This is the surrogate step that replaces a high-fidelity advection–diffusion solve when evaluating the mixing time distribution and the retention score over the Monte Carlo ensemble. For a new parameter realization ξ , the reduced model is initialized from the prescribed mean-removed scalar field, $\mathbf{z}_0 = \mathbf{V}_r^T \mathbf{q}_0$, and propagated according to

$$\widehat{\mathbf{z}}_{k+1} = \mathbf{K}(\xi, \theta_k) \widehat{\mathbf{z}}_k, \quad \widehat{\mathbf{z}}_0 = \mathbf{z}_0. \quad (37)$$

The reconstructed mean-removed and full scalar fields are $\widehat{\mathbf{q}}_k = \mathbf{V}_r \widehat{\mathbf{z}}_k$, and $\widehat{\mathbf{c}}_k = \widehat{\mathbf{q}}_k + \bar{c} \mathbf{1}$. The predicted mixing index is

$$\widehat{M}(t_k, \xi) = \frac{\|\widehat{\mathbf{q}}_k(\cdot; \xi)\|_{L^2(\Omega)}^2}{\|\mathbf{q}_0\|_{L^2(\Omega)}^2}. \quad (38)$$

The surrogate mixing time and failure probability are

$$\widehat{T}_m(\xi) = \inf \left\{ t_k > 0 : \widehat{M}(t_k, \xi) \leq \varepsilon_{\text{mix}} \right\}, \quad (39)$$

$$\widehat{P}_f(\tau) = \mathbb{P} \left\{ \widehat{T}_m(\xi) > \tau \right\}. \quad (40)$$

The reduced model is validated using (i) scalar variance decay, (ii) mixing time prediction, and (iii) preservation of the upper tail of the mixing time distribution.

V. NUMERICAL SETUP AND VALIDATION PROTOCOL

This section gives the numerical details needed to reproduce the results. The definitions of the mixing index, mixing time, finite-time retention score and Koopman polynomial-chaos operator were introduced in Secs. II–IV. Here, we specify the benchmark flows, discretization, data generation, model-training protocol, validation metrics, and consistency checks.

A. Benchmark flows and uncertain parameters

We use two canonical two-dimensional time-dependent flows. The main benchmark is the time-periodic double-gyre flow, which is a standard testbed for finite-time transport and coherent-structure analysis.^{13,14} The secondary benchmark is a periodically stirred cavity of the alternating-vortex type, which has long been used to study chaotic advection and mixing in bounded domains.^{1,2} Both flows are prescribed analytically, so all uncertainty enters only through their parameters.

The double-gyre flow is defined on $\Omega = [0, 2] \times [0, 1]$. The incompressible velocity field is generated from the streamfunction

$$\psi(x, y, t; \xi) = -A \sin(\pi f(x, t)) \sin(\pi y), \quad (41)$$

with $f(x, t) = a(t)x^2 + b(t)x$, $a(t) = \epsilon \sin(\omega t)$, $b(t) = 1 - 2\epsilon \sin(\omega t)$. The velocity is $\mathbf{u} = (u_x, u_y) = (\partial\psi/\partial y, -\partial\psi/\partial x)$, which is divergence-free by construction. The uncertain parameter vector is $\xi = (A, \epsilon, \omega)$, with independent uniform distributions $A \sim \mathcal{U}(0.15, 0.35)$, $\epsilon \sim \mathcal{U}(0.15, 0.35)$, $\omega \sim \mathcal{U}(0.7 \times 2\pi, 1.3 \times 2\pi)$. These ranges keep the flow in the same qualitative transport regime while introducing enough variation to produce different finite-time mixing behavior.

As a secondary benchmark, we consider a periodically stirred cavity on $\Omega = [0, 1] \times [0, 1]$. The flow is generated by two smooth vortices that alternate every half-period $T_s/2$. For an active vortex centered at $\mathbf{x}_m = (x_m, y_m)$, the streamfunction is

$$\psi_m(x, y) = \frac{\Gamma}{2\pi} \exp \left[-\frac{(x - x_m)^2 + (y - y_m)^2}{2\delta_v^2} \right] \sin^2(\pi x) \sin^2(\pi y), \quad (42)$$

and the corresponding incompressible velocity field is

$$\mathbf{u}_m(x, y) = \left(\frac{\partial\psi_m}{\partial y}, -\frac{\partial\psi_m}{\partial x} \right). \quad (43)$$

The factor $\sin^2(\pi x) \sin^2(\pi y)$ makes the streamfunction vanish on all four walls, so that $\mathbf{u}_m \cdot \mathbf{n} = 0$ on $\partial\Omega$. The vortex centers are

$$\mathbf{x}_1 = (0.35 + \Delta x_1, 0.50), \quad (44)$$

$$\mathbf{x}_2 = (0.65, 0.50), \quad (45)$$

and the fixed vortex-core parameter is $\delta_v = 0.20$. The first vortex is active for $0 \leq t \bmod T_s < T_s/2$ and the second is active for $T_s/2 \leq t \bmod T_s < T_s$. The uncertain parameter vector is $\xi = (\Gamma, T_s, \Delta x_1)$, with independent distributions $\Gamma \sim \mathcal{U}(3, 7)$, $T_s \sim \mathcal{U}(0.6, 1.4)$, $\Delta x_1 \sim \mathcal{U}(-0.1, 0.1)$.

B. High-fidelity scalar solver

The high-fidelity solver is used to generate both the training snapshots for the reduced model and the ground-truth mixing time distributions against which the reduced model is validated. It must therefore be accurate enough to resolve the fine scalar filaments produced at $\text{Pe} = 1000$ and stable enough to run reliably over many parameter realizations.

For each parameter realization, the passive-scalar equation (2) is solved from the same smooth initial condition, so that all variability in the mixing time comes from the uncertain transport dynamics. In the double gyre, the initial scalar field is

$$c_0(x, y) = \frac{1}{2} \left[1 - \tanh \left(\frac{x - x_c}{\ell_c} \right) \right]. \quad (46)$$

For the double-gyre benchmark, the interface is centered at $x_c = 1.0$ with thickness $\ell_c = 0.05$, so that it is initially centered between the two gyres. The stirred-cavity benchmark uses the same smooth-interface form,

$$c_{0,c}(x, y) = \frac{1}{2} \left[1 - \tanh \left(\frac{x - 0.50}{0.03} \right) \right], \quad (47)$$

with $x_c = 0.50$ and $\ell_c = 0.03$. A smooth interface is used in both benchmarks to avoid early-time numerical oscillations associated with a discontinuous scalar jump. The initial scalar field is fixed across all parameter realizations within each benchmark.

The Peclet number is fixed at $\text{Pe} = 1000$. The double-gyre simulations use a uniform grid with $N_x \times N_y = 256 \times 256$, time step $\Delta t = 0.005$, and final time $T_f = 20$. The stirred-cavity simulations use $N_x \times N_y = 128 \times 128$ with the same time step and final time. The mixing threshold used in all reported results is $\varepsilon_{\text{mix}} = 5 \times 10^{-2}$.

For each parameter realization, we record the observed time

$$\tilde{T}_m(\xi) = \min \{ T_m(\xi), T_f \} \quad (48)$$

and the censoring indicator $\delta(\xi) = \mathbf{1}_{\{T_m(\xi) \leq T_f\}}$. If $\delta(\xi) = 0$, the realization has not crossed the mixing threshold by the final simulation time. Such a realization is recorded as $T_m(\xi) > T_f$ and is not treated as an exact observation at $T_m = T_f$. For any deadline $\tau \leq T_f$, the

exceedance event is observed directly from the computed mixing index: $T_m(\xi) > \tau \iff M(\tau, \xi) > \epsilon_{\text{mix}}$. The failure probability over the simulated time interval can therefore be estimated without assigning an artificial mixing time to censored realizations.

The passive-scalar equation is discretized on a uniform cell-centered finite-volume grid. The advective term is written in conservative flux form. Face values are obtained using MUSCL reconstruction with a minmod slope limiter,³² and advection is advanced using the second-order SSP–RK2 scheme.³³ The wall-normal advective flux is set to zero at every physical boundary. Diffusion is advanced using the Crank–Nicolson method with a symmetric cell-centered discrete Laplacian satisfying the homogeneous Neumann condition $\nabla c \cdot \mathbf{n} = 0$, on $\partial\Omega$. This treatment preserves the discrete spatial mean. At each time step, we monitor the spatial mean and normalized scalar variance. The maximum relative scalar-mean drift and the maximum positive increment of the normalized scalar variance are recorded as numerical consistency checks. The normalized mixing index and the corresponding censored or uncensored mixing time are computed from Eqs. (6) and (7).

C. Reference ensemble, training data, and model identification

To distinguish uncertainty in the mixing response from any sensitivity of the reduced model to its own training data, all reference statistics are computed on ensembles that are disjoint from the trajectories used to fit or select the reduced model.

The double-gyre reference statistics are obtained from an independent high-fidelity Monte Carlo ensemble of $N_{\text{MC}} = 1000$ parameter realizations. This ensemble is disjoint from all trajectories used for reduced-model identification and model selection. It is used for the high-fidelity statistics in Figs. 1 and 2. After the double-gyre reduced model has been identified, it is also evaluated on these same parameter values for the distributional and diagnostic comparisons in Figs. 3 and 4. The stirred-cavity model is evaluated on the separate 30-trajectory held-out ensemble described below.

For a deadline $\tau \leq T_f$, the empirical exceedance probability is computed directly from the observed mixing-index event,

$$\hat{P}_f(\tau) = \frac{1}{N_{\text{MC}}} \sum_{j=1}^{N_{\text{MC}}} \mathbf{1}_{\{M(\tau, \xi_j) > \epsilon_{\text{mix}}\}}. \quad (49)$$

The corresponding cumulative probability of mixing by time τ is

$$\hat{F}_{T_m}(\tau) = 1 - \hat{P}_f(\tau), \quad 0 \leq \tau \leq T_f. \quad (50)$$

This definition does not assign $T_m = T_f$ to a realization that remains unmixed at the end of the simulation. Consequently, $\hat{F}_{T_m}(T_f)$ need not equal one and $\hat{P}_f(T_f)$ need not equal zero.

Pointwise uncertainty in the exceedance probability is reported with a 95% Wilson interval.^{34,35} Let $\hat{p} = \hat{P}_f(\tau)$, $n = N_{\text{MC}}$, and $z = 1.96$. The lower and upper limits are




FIG. 1. Uncertainty creates a slow tail in the double-gyre response. Panels (a)–(c) show physical concentration on the common scale $c \in [0, 1]$: (a) the fixed initial field, (b) a fast realization at the deadline $\tau_S = 15.20$, and (c) a slow realization at the same deadline. (d) Mixing-index trajectories from the independent high-fidelity ensemble of $N_{\text{MC}} = 1000$ realizations; the shaded bands are the 5%–95% and 10%–90% percentile envelopes and the solid line is the median trajectory. (e) Histogram of the exact mixing times $T_m(\xi)$ for the uncensored subset only. The horizontal axis is truncated at the largest observed uncensored mixing time, strictly less than $T_f = 20$. The number and fraction of right-censored realizations ($T_m > T_f$) are printed in the panel and are not plotted as an accumulation at $T_m = T_f$. (f) Direct estimate of the exceedance probability $\hat{P}_f(\tau) = N_{\text{MC}}^{-1} \sum_j \mathbf{1}_{\{M(\tau, \xi_j) > \epsilon_{\text{mix}}\}}$ with a pointwise 95% Wilson interval. The value at $\tau = T_f$ is the observed remaining failure fraction and is not forced to zero.




FIG. 2. Parameter conditioning of slow mixing for the independent high-fidelity ensemble of $N_{MC} = 1000$ realizations. (a) Samples in the $(\epsilon, \omega/(2\pi))$ plane colored by the observed response $\bar{T}_m = \min(T_m, T_f)$. (b) Samples in the (A, ϵ) plane using the same response scale. (c) Overlaid marginal distributions of the amplitude A and the separatrix parameter ϵ for the full ensemble and for the slow set $T_m > \tau_{\delta}$; a common horizontal axis is used because A and ϵ share the same sampled range $[0.15, 0.35]$. (d) Mean observed response $\bar{T}_m = \min(T_m, T_f)$ on a uniform 25×25 bin grid in the $(\epsilon, \omega/(2\pi))$ plane. Only bins containing at least five realizations are shown, and no smoothing is applied.

$$P_{\pm}(\tau) = \frac{\hat{p} + \frac{z^2}{2n} \pm z \sqrt{\frac{\hat{p}(1-\hat{p})}{n} + \frac{z^2}{4n^2}}}{1 + \frac{z^2}{n}}. \quad (51)$$

The reduced-model database contains 200 independently sampled trajectories for each benchmark. The data are split by parameter realization, not by snapshot: 170 trajectories form the development set and 30 trajectories form an independent held-out set. For the double gyre, the 170 development trajectories are used directly for fitting. For the cavity, they are further split into fitting and internal model-selection subsets, as described below (Table I). Scalar snapshots are stored at $t_k = k\Delta t_s$, $\Delta t_s = 0.20$. Each complete trajectory therefore contains $N_t + 1 = 101$ stored snapshots at $t_k = 0, 0.2, \dots, 20$. The default retention calculation uses $N = 50$ snapshot intervals, giving $T_{ret} = N\Delta t_s = 10$. For the double gyre, development-set reconstruction and rollout checks are used to fix $r = p = 16$, $P = 3\lambda = 10^{-4}$. Neither the independent held-out set nor the $N_{MC} = 1000$ evaluation ensemble is used for model selection or score normalization. The total-degree Legendre basis in three uncertain parameters contains $|\mathcal{I}_3| = 20$ terms. Together with the five phase functions, this gives 100 parameter-phase terms. The independent 30-trajectory set is used for field-level, rollout, and scalar-energy checks. Distributional tail statistics are then evaluated on the independent $N_{MC} = 1000$ ensemble defined above.

For the stirred cavity, the 170 development trajectories are divided into 140 fitting trajectories and 30 internal model-selection trajectories. Stable models are compared first by deadline-status agreement at $\tau_C = 12$, then by the median absolute mixing time error and the 90th percentile of the trajectory-averaged mixing-index error. This procedure selects $r = p = 6$, $P = 2$. The corresponding

Legendre basis contains $|\mathcal{I}_2| = 10$ terms and, with the two switching functions, 20 parameter-phase terms. The selected model is refitted on all 170 development trajectories and evaluated on the independent 30-trajectory held-out set. The regularization parameter is $\lambda = 10^{-4}$ for both benchmarks.

D. Validation metrics and numerical consistency checks

Because the reduced model is meant to predict both individual mixing trajectories and the distribution of mixing times across the uncertain family, it is validated at three levels: pointwise scalar fields, per-realization mixing indices and mixing times, and distributional agreement with the high-fidelity mixing time law. Numerical consistency of the high-fidelity solver itself is verified separately through grid and time step refinement.

For a held-out realization, the relative field error is

$$e_c(t_k, \xi) = \frac{\|c(\cdot, t_k; \xi) - \hat{c}(\cdot, t_k; \xi)\|_2}{\|c(\cdot, t_k; \xi)\|_2}. \quad (52)$$

The pointwise and trajectory-averaged mixing-index errors are

$$e_M(t_k, \xi) = |M(t_k, \xi) - \hat{M}(t_k, \xi)|, \quad (53)$$

$$\bar{e}_M(\xi) = \frac{1}{N_t + 1} \sum_{k=0}^{N_t} e_M(t_k, \xi). \quad (54)$$

The mixing time error is evaluated only when both calculations are uncensored,




FIG. 3. Validation of the reduced Koopman polynomial-chaos model on the independent double-gyre evaluation ensemble of $N_{MC} = 1000$ realizations. (a) High-fidelity and reduced mixing-index trajectories for representative fast, median-response, and slow realizations. (b) Mixing time parity for jointly uncensored realizations only; the four HFM/ROM censoring status counts are reported separately in the panel. (c) Cumulative probability of mixing computed as $1 - \hat{P}_f(\tau)$ from the directly observed deadline event. (d) High-fidelity and reduced exceedance probabilities with pointwise 95% Wilson intervals. No end point is appended at T_f in either panel. (e) Absolute mixing time error for jointly uncensored realizations. (f) Trajectory-averaged mixing-index error vs the high-fidelity observed response $\bar{T}_m$, with censored cases identified separately.

$$e_T(\xi) = |T_m(\xi) - \hat{T}_m(\xi)|. \quad (55)$$

Censored realizations are instead compared through their deadline-status indicators. Distributional agreement is measured by

$$e_{CDF} = \sup_{0 \leq \tau \leq T_f} |\hat{F}_{T_m}^{HFM}(\tau) - \hat{F}_{T_m}^{ROM}(\tau)|. \quad (56)$$

The scalar-energy approximation is checked on independent held-out snapshots. For each snapshot,

$$e_H = \frac{|||\mathbf{q}|||_{L_h^2}^2 - \mathbf{z}^T \mathbf{H} \mathbf{z}}{\max\{|||\mathbf{q}|||_{L_h^2}^2, \varepsilon_H\}}, \quad (57)$$

where

$$\varepsilon_H = \max\{10^{-6} Q_{0.01}(||\mathbf{q}|||_{L_h^2}^2), 10^{-15}\}. \quad (58)$$

Because $\mathbf{z}^T \mathbf{H} \mathbf{z}$ is exactly the discrete energy of the POD reconstruction, e_H measures POD truncation error. Its median, 95th percentile, and maximum are reported for both benchmarks in Fig. S1 and Table S1 of the [supplementary material](#).

Grid convergence was evaluated for representative low-, median-, and high-retention double-gyre screening cases using 128×128 ,

192×192 , and 256×256 cells at $\Delta t = 0.005$. The cases are labeled according to their terminal mixing-index values because all three remain above the threshold at $T_f = 20$. Relative to the 256×256 calculations, the maximum differences in the complete mixing-index curves on the 192×192 grid were 4.06×10^{-3} , 5.88×10^{-3} , and 3.67×10^{-3} . The corresponding differences on the 128×128 grid were 1.40×10^{-2} , 2.00×10^{-2} , and 1.33×10^{-2} . Time convergence was evaluated on the 256×256 grid using $\Delta t \in \{0.01, 0.005, 0.0025\}$. Relative to $\Delta t = 0.0025$, the maximum curve differences at the adopted step $\Delta t = 0.005$ were 1.55×10^{-4} , 6.98×10^{-5} , and 2.48×10^{-5} . The maximum relative scalar drift was of order 10^{-12} , and no positive increment of the normalized scalar variance was detected. Figure S3 and Table S2 of the [supplementary material](#) report the complete refinement results. Figure S4 and Table S3 of the [supplementary material](#) report the reduced-model sensitivity study.

VI. RESULTS

The results test the following mechanism: under parametric uncertainty, slow mixing is produced by scalar structures that remain weakly contracted over a finite interval for particular parameter realizations. We first establish the slow tail and its localization in parameter space. We then validate the reduced operator, evaluate the finite-time retention score, and identify the associated retained scalar




FIG. 4. Censoring-consistent comparison of the finite-time retention score G_N with three baseline diagnostics: an early-decay diagnostic, the forcing amplitude A , and a stretching proxy on the independent double-gyre ensemble of $N_{MC} = 1000$ realizations. (a)–(c) Each diagnostic, standardized as $\bar{D} = (D - \bar{D})/\sigma_D$ so that all three share a common horizontal scale, vs the observed response $\bar{T}_m = \min(T_m, T_f)$: (a) retention score G_N , (b) early-decay diagnostic, (c) stretching proxy. Realizations satisfying $T_m > \tau_s$ are marked separately ($\tau_s = 15.20$). (d) ROC curves for the fully observed deadline event $T_m > \tau_s$ for all four diagnostics. (e) Tail-enrichment factor $\mathcal{E}_\alpha$ vs the selected upper fraction α , for the same event. (f) Numerical summary of the deadline event, the failure probability at τ_s , and the diagnostic ranking (Spearman correlation ρ_s , area under the ROC curve, and tail enrichment $\mathcal{E}_{0.1}$) for each of the four diagnostics. Panels (d)–(f) and Table III use each diagnostic in its raw scale; standardization affects only the horizontal scale of panels (a)–(c) and leaves ρ_s , AUC, and $\mathcal{E}_\alpha$ unchanged. No censored realization is treated as an exact mixing time at T_f .

TABLE I. Numerical and reduced-model settings for the double-gyre (DG) and stirred-cavity (SC) benchmarks. Common settings are repeated only where needed for clarity.

Physical and numerical setup			Reduced model and data		
Quantity	DG	SC	Quantity	DG	SC
Domain	$[0, 2] \times [0, 1]$	$[0, 1] \times [0, 1]$	Reduced dimensions $r = p$	16	6
Grid	256×256	128×128	PCE degree P	3	2
Interface (x_c, ℓ_c)	$(1.0, 0.05)$	$(0.50, 0.03)$	Polynomial terms $ \mathcal{I}_P $	20	10
Vortex centers $(\mathbf{x}_1, \mathbf{x}_2)$	$\dots$	$(0.35 + \Delta x_1, 0.50); (0.65, 0.50)$	Phase terms n_θ	5	2
Core/switch $(\delta_v, T_s/2)$	$\dots$	$(0.20, T_s/2)$	Combined terms $ \mathcal{I}_P n_\theta$	100	20
Boundary conditions	$\mathbf{u} \cdot \mathbf{n} = 0;$ $\nabla c \cdot \mathbf{n} = 0$	$\mathbf{u} \cdot \mathbf{n} = 0;$ $\nabla c \cdot \mathbf{n} = 0$	Regularization λ	10^{-4}	10^{-4}
Peclet number Pe	1000	1000	Retention steps N	50	50
Solver step Δt	0.005	0.005	Retention horizon $N\Delta t_s$	10	10
Snapshot step Δt_s	0.20	0.20	Development traj.	170	170
Final time T_f	20	20	Fit/selection split	$\dots$	140/30
Mixing threshold ε_{mix}	5×10^{-2}	5×10^{-2}	Held-out trajectories	30	30
			DG evaluation set	1000	$\dots$

structures. The periodically stirred cavity provides an independent test in a different transport geometry.

Slow double-gyre realizations are defined by the deadline event

$$\mathcal{S}_{\text{slow}}(\tau_s) = \{\xi_j : M(\tau_s, \xi_j) > \varepsilon_{\text{mix}}\} = \{\xi_j : T_m(\xi_j) > \tau_s\}, \quad (59)$$

with $\tau_s = 15.20 < T_f$. The high-fidelity failure probability at this deadline is 0.15. This definition remains valid for realizations that are right-censored at T_f . Fast realizations used in representative field comparisons are selected from the lower 10% of observed, uncensored mixing times.

A. Slow mixing tails in the uncertain double-gyre flow

Figure 1 shows the high-fidelity Monte Carlo results. Most realizations mix within the central part of the observation window, while a smaller subset remains above the threshold substantially longer. The resulting distribution therefore has a persistent right tail rather than a narrow fluctuation around a nominal decay curve.

The scalar fields explain the tail physically. A fast realization rapidly forms fine filaments that are removed by diffusion. At the same deadline, a slow realization retains a large-scale scalar remnant near a weakly exchanging region. Because scalar variance is non-increasing for every realization, the slow event reflects weak finite-time contraction rather than scalar-energy growth.

The slow realizations are not distributed uniformly in parameter space. Figure 2 shows that the deadline failures occupy restricted regions of the uncertain parameter domain. The conditional parameter distributions and the binned mean observed response identify combinations of forcing amplitude, separatrix oscillation, and forcing frequency that sustain weak exchange between the two gyres. The slow tail is therefore parameter-conditioned rather than numerical scatter around the mean response.

The double-gyre reduced model is evaluated on the independent $N_{\text{MC}} = 1000$ ensemble, which is disjoint from the model-building data. Representative trajectories test the temporal decay, jointly uncensored samples test mixing time prediction, and the full evaluation ensemble tests the cumulative and exceedance probabilities. Tail-specific statistics are reported because agreement in the bulk alone does not establish reliable prediction of delayed mixing (Table II).

B. Retention score as a predictor of slow mixing

The finite-time retention score $G_N(\xi)$ ranks parameter realizations by the least contraction of resolved scalar-energy directions over $T_{\text{ret}} = 10$. Because some realizations are right-censored at T_f , rank

TABLE II. Tail-aware validation of the double-gyre reduced model on the independent evaluation ensemble. Mixing time errors use jointly uncensored realizations only. The upper-tail column uses the upper 10% of observed, uncensored high-fidelity mixing times. Mixing times are recorded at the snapshot resolution $\Delta t_s = 0.20$, so exact ties between the high-fidelity and reduced predictions are common in the upper-tail subset.

Quantity	Full evaluation set	Upper 10% of T_m
Median $\bar{e}_M$	0.007	0.007
Median e_T	0.321	$< \Delta t_s$
CDF error e_{CDF}	0.077	...

TABLE III. Comparison of the finite-time retention score G_N with three baseline diagnostics for identifying double-gyre deadline failures on the independent $N_{\text{MC}} = 1000$ ensemble. The Spearman correlation ρ_s is computed with the observed response $\tilde{T}_m = \min(T_m, T_f)$. AUC and the tail-enrichment factor $\mathcal{E}_{0.1}$ use the fully observed deadline event $T_m > \tau_s$ with $\tau_s = 15.20$, and $\mathcal{E}_{0.1}$ uses the largest 10% of each diagnostic. The values agree, up to rounding, with the summary printed in Fig. 4(f).

Diagnostic	Spearman ρ_s with $\tilde{T}_m$	AUC	$\mathcal{E}_{0.1}$
G_N	0.815	0.924	4.963
Forcing amplitude A	0.653	0.870	4.000
Early-decay diagnostic	0.247	0.587	1.333
Stretching proxy	0.416	0.745	2.960

correlations are reported against the observed response $\tilde{T}_m = \min(T_m, T_f)$, whereas classification and enrichment are computed from the fully observed deadline event $T_m > \tau_s$ with $\tau_s = 15.20 < T_f$. Figure 4 compares G_N with three baseline diagnostics: an early-decay diagnostic, the forcing amplitude A (the strongest single-parameter baseline in this benchmark), and a stretching proxy. The operator score gives the strongest rank association, deadline classification, and concentration of slow events. Table III tabulates the same four-diagnostic comparison numerically, and Fig. 4(f) prints the matching in-panel summary.

The three summary statistics reported in Table III and in Fig. 4(f) are the Spearman rank correlation ρ_s between the diagnostic and the observed response $\tilde{T}_m$, the area under the receiver operating characteristic (ROC) curve (AUC) for the deadline event $T_m > \tau_s$, and the tail-enrichment factor $\mathcal{E}_{0.1}$ defined below. For visual comparison across diagnostics with different physical units, panels (a)–(c) of Fig. 4 plot each diagnostic after standardization, $\tilde{D} = (D - \bar{D})/\sigma_D$, where $\bar{D}$ and σ_D are the sample mean and standard deviation on the evaluation ensemble. Ranks, ROC curves, and enrichment factors are invariant under this monotone transformation, so panels (d)–(f) and Table III use each diagnostic in its raw scale; standardization affects only the horizontal scale of panels (a)–(c) and does not affect ρ_s , AUC, or $\mathcal{E}_\alpha$. For a diagnostic D oriented so that larger values indicate slower mixing, the tail-enrichment factor is

$$\mathcal{E}_\alpha(D) = \frac{\mathbb{P}\{T_m > \tau_s | D \geq Q_{1-\alpha}(D)\}}{\mathbb{P}\{T_m > \tau_s\}}. \quad (60)$$

A value greater than one means that selecting the largest diagnostic values concentrates deadline failures more effectively than random sampling.

C. Decoded retention modes and retained scalar structures

For each realization, let

$$\mathbf{A}_N(\xi) = \mathbf{U}(\xi)\mathbf{\Sigma}(\xi)\mathbf{V}(\xi)^\top. \quad (61)$$

The leading right singular vector defines the least-damped initial direction in the resolved POD subspace,

$$\mathbf{y}_0^* = \mathbf{H}^{-1/2}\mathbf{v}_1. \quad (62)$$

TABLE IV. Independent held-out validation of the selected stirred-cavity reduced model ($r = 6$, $P = 2$).

Quantity	Value
Median $\bar{e}_M$	0.006 43
90th percentile $\bar{e}_M$	0.008 45
Maximum $\bar{e}_M$	0.010 27
Median e_T	0.0801
$\hat{P}_{f,HFM}(12)$	0.5333
$\hat{P}_{f,ROM}(12)$	0.5000
Absolute failure-probability error	0.0333
Deadline-status disagreements	1/30
Maximum relative scalar-mean drift	1.05×10^{-12}
Maximum positive increment of M	0

The physical field compared with the high-fidelity remnant is not this initial direction. It is first propagated to

$$y_N^* = K_{N:0}(\xi)y_0^*, \quad \phi_N^* = V_r y_N^*, \quad (63)$$

at $t_N = N\Delta t_s = 10$. The high-fidelity mean-removed scalar and ϕ_N^* are individually L^2 normalized. The sign of ϕ_N^* is chosen to make

their inner product non-negative, and all field panels use a common symmetric color scale.

Their spatial overlap is

$$\eta(\xi) = \frac{|\langle q(\cdot, t_N; \xi), \phi_N^*(\cdot; \xi) \rangle_{L^2}|}{\|q(\cdot, t_N; \xi)\|_{L^2} \|\phi_N^*(\cdot; \xi)\|_{L^2}}. \quad (64)$$

The fixed-initial-condition retention g_N^{fix} and alignment a_0 defined in Eqs. (21) and (23) are reported together with the operator score so that G_N is not interpreted as the decay of the prescribed initial field.

The stirred cavity tests the same diagnostic in a flow organized by alternating vortices rather than an oscillating separatrix. The selected model uses $r = 6$ and $P = 2$ and is evaluated on 30 independent held-out trajectories (Table IV). The reduced model reproduces the mixing-index decay and deadline status, while G_N remains strongly associated with the mixing response (Fig. 5).

For a held-out cavity trajectory, the error $\bar{e}_M$ is defined by Eq. (54). Its median, 90th percentile, and maximum are 0.006 43, 0.008 45, and 0.010 27, respectively. The median absolute mixing time error is 0.0801. At $\tau_C = 12$, the high-fidelity and reduced failure probabilities are 0.5333 and 0.5000, giving an absolute error of 0.0333 and one deadline-status disagreement. The maximum relative scalar-mean drift is 1.05×10^{-12} , and no positive increment of M is detected (Fig. 6).




FIG. 5. Decoded retention modes for representative slow and fast double-gyre realizations using the final configuration $r = 16$, $P = 3$, $N = 50$, and $t_N = 10$. (a) High-fidelity mean-removed remnant for a slow realization. (b) Propagated retention mode for the same slow realization. (c) High-fidelity mean-removed remnant for a fast realization. (d) Propagated retention mode for the same fast realization. (e) Group summaries of the spatial overlap η , fixed-initial-condition retention g_N^{fix} , and initial-direction alignment a_0 for the selected fast and slow sets. (f) All and only the 16 singular values of the metric-weighted finite-time propagator for the two representative realizations. The four field panels are individually L^2 normalized, sign aligned where appropriate, and use one common symmetric color scale. Figure S2 of the [supplementary material](#) displays the initial and propagated directions separately.




FIG. 6. Verification on the periodically stirred-cavity benchmark using the independent held-out set of 30 realizations and the deadline $\tau_c = 12$. Panels (a)–(c) show physical concentration on the common scale $c \in [0, 1]$: (a) the fixed initial field with nominal vortex centers marked only to indicate the stirring geometry, (b) a fast field at the deadline, and (c) a slow field at the same deadline. (d) High-fidelity and reduced cumulative mixing probabilities with pointwise 95% Wilson intervals, computed without assigning censored cases to $T_m = 20$. (e) Retention score vs the high-fidelity observed response T_m , with censored cases marked separately. (f) Mixing time parity for jointly uncensored realizations; the deadline-status disagreements are reported separately.

D. Sensitivity to numerical and modeling choices

Sensitivity is evaluated on the independent 30-trajectory cavity test set for $r \in \{4, 6, 8, 10, 12\}$, $P \in \{1, 2, 3\}$, and $N \in \{25, 50, 75\}$, corresponding to retention horizons of 5, 10, and 15. For $r = 4$ –10, the Spearman correlation between G_N and T_m remains between 0.983 and 0.990, with AUC 1.000. At $r = 12$, the correlation and AUC decrease to 0.697 and 0.879. For $P = 1$ –3, the correlation remains between 0.981 and 0.984, with AUC 1.000. Across the three retention horizons, it ranges from 0.970 to 0.984, again with AUC 1.000. The selected configuration therefore lies in the stable part of the tested range. Figure S4 and Table S3 of the [supplementary material](#) contain the full results.

Once trained, the reduced model replaces a full-grid advection–diffusion solve by phase-conditioned operator evaluation and propagation in the POD space. This makes dense evaluation of the mixing time tail and retention score practical. More importantly, the operator model provides the least-damped scalar directions, which cannot be obtained from direct interpolation of mixing times.

VII. DISCUSSION AND CONCLUSIONS

The results show that mixing reliability is controlled by more than the nominal or mean decay rate of scalar variance. In the

uncertain flows considered here, most parameter realizations mix rapidly, but a smaller subset produces delayed homogenization. These slow realizations are not random outliers in the scalar variance curve. They are associated with coherent scalar structures that remain organized over finite times and therefore delay the crossing of the prescribed mixing threshold. This interpretation is consistent with the classical lamellar picture of mixing. Fast realizations generate thin scalar structures and strong gradients early enough for diffusion to act efficiently. Slow realizations delay this deformation and accumulate less scalar dissipation over the finite observation window. The connection between deformation history and scalar decay is central to lamellar descriptions of mixing.¹⁰ The retention diagnostic complements this viewpoint. Lamellar models describe how a prescribed deformation history controls diffusion, whereas the present framework identifies which members of an uncertain flow family retain scalar structures long enough to populate the upper tail of the mixing time distribution. This interpretation is also consistent with reduced descriptions based on diffusive scalar elements.¹¹

The main physical message is that rare slow mixing is governed by parameter-conditioned scalar retention. A coherent structure may be visible in a deterministic flow, but it becomes relevant for reliability only when it is sustained by parameter values that occur within the

uncertain family. Thus, the important object is not simply a transport barrier of the nominal flow. It is a scalar structure selected by the parameter-dependent transport operator. This view is consistent with finite-time transport analysis, but shifts the emphasis from a single deterministic flow to an ensemble of possible transport operators.^{14,16}

The finite-time retention score $G_N(\xi)$ provides a compact way to detect these slow mixing realizations. In the full passive-scalar equation, the L^2 scalar variance is non-increasing because of diffusion. The relevant mechanism is therefore not growth of total scalar variance above its initial value. It is weak finite-time contraction, or retention of scalar-energy directions, relative to the typical parameter realization. In the reduced POD coordinates, this retention appears through large singular directions of the finite-time propagator. This is analogous to non-normal finite-time behavior, where eigenvalues or asymptotic decay rates alone do not determine the finite-time response.^{22,23}

The propagated retention modes connect the operator diagnostic back to the fluid-mechanics picture. A large value of $G_N(\xi)$ is not only a numerical indicator. The leading right singular vector identifies the least-damped initial direction, and its propagated final state is reconstructed to reveal the scalar structure retained over the selected horizon. For slow mixing realizations, these decoded structures align with coherent unmixed regions observed in the high-fidelity simulations. This provides a physical explanation for why particular parameter values populate the slow tail of the mixing time distribution.

The Koopman polynomial-chaos model is useful because it gives a finite-dimensional representation of an uncertain family of transport operators. This is consistent with a broader trend in fluid mechanics toward data-informed reduced models for uncertainty quantification and statistical prediction.³⁶ The Koopman approximation advances the reduced POD coordinates, while the polynomial-chaos expansion represents how this operator varies with uncertain parameters. Once trained, the model can be sampled cheaply over many parameter realizations. This makes it possible to estimate the tail probability $P_f(\tau)$ and to compute $G_N(\xi)$ without repeating a full advection-diffusion simulation for every parameter sample. The use of polynomial chaos is appropriate when the dependence of the reduced operator on the parameters is sufficiently regular.^{26,27}

Several limitations should be kept in mind. First, the method assumes that the reduced operator varies smoothly enough with ξ to be represented by a moderate-order polynomial-chaos expansion. If the flow undergoes abrupt topological changes, such as the creation or destruction of barriers over small parameter intervals, a global polynomial representation may converge slowly. Piecewise polynomial expansions or adaptive sampling would then be more appropriate.

Second, the reduced model depends on the quality of the POD representation. If the truncated POD basis removes fine scalar filaments that are important for the mixing threshold, the predicted mixing time and retention score may be biased. This issue becomes more serious at high Peclet number, where scalar gradients can become very sharp. Reconstruction error alone is therefore not sufficient. The POD space must also preserve scalar variance decay and the structures that control late-time mixing.

Third, the present study uses two-dimensional benchmark flows. These flows are useful because they allow the slow mixing mechanism to be visualized clearly. Three-dimensional or turbulent scalar mixing

would introduce additional difficulties, including broader spatial scales, stronger intermittency, more complex passive-scalar structures, and more expensive training data.³⁷ The same operator viewpoint remains applicable, but the POD model and the retention calculation would need to be tested carefully in that setting.

The analysis also has a design implication, although no feedback controller is designed in this work. A conventional mixing design may aim to maximize a mean stretching rate or minimize an average mixing time. The present results show that this is not sufficient when reliability matters. A design can have good average performance while still allowing parameter values that retain coherent scalar structures for too long. Reliable operation should therefore also avoid regions of parameter space with large finite-time retention scores, since these regions are more likely to populate the slow tail of the mixing time distribution. A full closed-loop control formulation, including online measurements, input constraints, and controller comparison, is left for future work.

SUPPLEMENTARY MATERIAL

See the [supplementary material](#) for additional numerical validation, sensitivity analyses, reduced-model diagnostics, and supplementary figures and tables for the double-gyre and periodically stirred-cavity benchmarks.

ACKNOWLEDGMENTS

This project is supported by the U.S. Food and Drug Administration under the FDA BAA-22-00123 program, Award No. 75F40122C00200, and the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation), SFB 1615, Project No. 503850735. The author sincerely thanks Prof. Richard D. Braatz (MIT) and Dr. Saikat Mukherjee (IIT Kharagpur) for valuable scientific discussions and insightful suggestions that helped shape this work.

AUTHOR DECLARATIONS

Conflict of Interest

The author has no conflicts to disclose.

Author Contributions

Srimanta Santra: Data curation (lead); Formal analysis (lead); Investigation (lead); Resources (lead); Supervision (lead); Validation (lead); Visualization (lead); Writing – original draft (lead); Writing – review & editing (lead).

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

REFERENCES

- ¹H. Aref, "Stirring by chaotic advection," *J. Fluid Mech.* **143**, 1–21 (1984).
- ²J. M. Ottino, *The Kinematics of Mixing: Stretching, Chaos, and Transport* (Cambridge University Press, 1989), Vol. 3.
- ³V. Toussaint, P. Carriere, J. Scott, and J.-N. Gence, "Spectral decay of a passive scalar in chaotic mixing," *Phys. Fluids* **12**, 2834–2844 (2000).

- ⁴R. T. Pierrehumbert, "Tracer microstructure in the large-eddy dominated regime," *Chaos, Solitons Fractals* **4**, 1091–1110 (1994).
- ⁵J.-L. Thiffeault, "Using multiscale norms to quantify mixing and transport," *Nonlinearity* **25**, R1–R44 (2012).
- ⁶D. R. Fereday and P. H. Haynes, "Scalar decay in two-dimensional chaotic advection and Batchelor-regime turbulence," *Phys. Fluids* **16**, 4359–4370 (2004).
- ⁷P. H. Haynes and J. Vanneste, "What controls the decay of passive scalars in smooth flows?," *Phys. Fluids* **17**, 097103 (2005).
- ⁸P. Meunier and E. Villiermaux, "The diffusive strip method for scalar mixing in two dimensions," *J. Fluid Mech.* **662**, 134–172 (2010).
- ⁹D. Martínez-Ruiz, P. Meunier, B. Favier, L. Duchemin, and E. Villiermaux, "The diffusive sheet method for scalar mixing," *J. Fluid Mech.* **837**, 230–257 (2018).
- ¹⁰E. Villiermaux, "Mixing versus stirring," *Annu. Rev. Fluid Mech.* **51**, 245–273 (2019).
- ¹¹P. Meunier and E. Villiermaux, "The diffuselet concept for scalar mixing," *J. Fluid Mech.* **951**, A33 (2022).
- ¹²D. Domínguez-Vázquez and T. Aquino, "A hyperbolic description of mixing," *J. Fluid Mech.* **1027**, A40 (2026).
- ¹³S. C. Shadden, F. Lekien, and J. E. Marsden, "Definition and properties of Lagrangian coherent structures from finite-time Lyapunov exponents in two-dimensional aperiodic flows," *Physica D* **212**, 271–304 (2005).
- ¹⁴G. Haller, "Lagrangian coherent structures," *Annu. Rev. Fluid Mech.* **47**, 137–162 (2015).
- ¹⁵G. Froyland, S. Lloyd, and N. Santitissadeekorn, "Coherent sets for nonautonomous dynamical systems," *Physica D* **239**, 1527–1541 (2010).
- ¹⁶G. Froyland, "An analytic framework for identifying finite-time coherent sets in time-dependent dynamical systems," *Physica D* **250**, 1–19 (2013).
- ¹⁷W. Tang and C. Luna, "Dependence of advection–diffusion–reaction on flow coherent structures," *Phys. Fluids* **25**, 106602 (2013).
- ¹⁸E. Gouillart, O. Dauchot, J.-L. Thiffeault, and S. Roux, "Open-flow mixing: Experimental evidence for strange eigenmodes," *Phys. Fluids* **21**, 023603 (2009).
- ¹⁹E. Gouillart, O. Dauchot, and J.-L. Thiffeault, "Measures of mixing quality in open flows with chaotic advection," *Phys. Fluids* **23**, 013604 (2011).
- ²⁰I. Mezić, "Analysis of fluid flows via spectral properties of the Koopman operator," *Annu. Rev. Fluid Mech.* **45**, 357–378 (2013).
- ²¹M. O. Williams, I. G. Kevrekidis, and C. W. Rowley, "A data-driven approximation of the Koopman operator: Extending dynamic mode decomposition," *J. Nonlinear Sci.* **25**, 1307–1346 (2015).
- ²²L. N. Trefethen, A. E. Trefethen, S. C. Reddy, and T. A. Driscoll, "Hydrodynamic stability without eigenvalues," *Science* **261**, 578–584 (1993).
- ²³P. J. Schmid, "Nonmodal stability theory," *Annu. Rev. Fluid Mech.* **39**, 129–162 (2007).
- ²⁴B. R. Noack, I. Mezić, G. Tadmor, and A. Banaszuk, "Optimal mixing in recirculation zones," *Phys. Fluids* **16**, 867–888 (2004).
- ²⁵O. Gorodetskyi, M. F. M. Speetjens, and P. D. Anderson, "An efficient approach for eigenmode analysis of transient distributive mixing by the mapping method," *Phys. Fluids* **24**, 053602 (2012).
- ²⁶D. Xiu and G. E. Karniadakis, "The Wiener–Askey polynomial chaos for stochastic differential equations," *SIAM J. Sci. Comput.* **24**, 619–644 (2002).
- ²⁷O. Le Maître and O. M. Knio, *Spectral Methods for Uncertainty Quantification: With Applications to Computational Fluid Dynamics* (Springer, 2010).
- ²⁸G. Berkooz, P. Holmes, and J. L. Lumley, "The proper orthogonal decomposition in the analysis of turbulent flows," *Annu. Rev. Fluid Mech.* **25**, 539–575 (1993).
- ²⁹L. Sirovich, "Turbulence and the dynamics of coherent structures. Part I: Coherent structures," *Q. Appl. Math.* **45**, 561–571 (1987).
- ³⁰C. W. Rowley, T. Colonius, and R. M. Murray, "Model reduction for compressible flows using POD and Galerkin projection," *Physica D* **189**, 115–129 (2004).
- ³¹S. Santra, "Polynomial chaos expansion and feature-space regression surrogates for stabilization of uncertain coupled PDE–ODE reactor processes," *J. Process Control* **164**, 103769 (2026).
- ³²B. van Leer, "Towards the ultimate conservative difference scheme. V. A second-order sequel to Godunov's method," *J. Comput. Phys.* **32**, 101–136 (1979).
- ³³S. Gottlieb, C.-W. Shu, and E. Tadmor, "Strong stability-preserving high-order time discretization methods," *SIAM Rev.* **43**, 89–112 (2001).
- ³⁴E. B. Wilson, "Probable inference, the law of succession, and statistical inference," *J. Am. Stat. Assoc.* **22**, 209–212 (1927).
- ³⁵L. D. Brown, T. T. Cai, and A. DasGupta, "Interval estimation for a binomial proportion," *Stat. Sci.* **16**(2), 101–117 (2001).
- ³⁶A. Charalampopoulos and T. Sapsis, "Uncertainty quantification of turbulent systems via physically consistent and data-informed reduced-order models," *Phys. Fluids* **34**(7), 075120 (2022).
- ³⁷J. Buchmeier, A. Bußmann, X. Gao, and I. Bermejo-Moreno, "Geometry and dynamics of passive scalar structures in compressible turbulent mixing," *Phys. Fluids* **33**, 105126 (2021).