

Deterministic and Stochastic Dynamics of the Two-State-Variable Spring-Slider with Bayesian Data Assimilation and Forecasting

Thomas Holland

3311 words with 5 figures (not including supplemental appendices)

Received April 2026

SUMMARY

We apply a grid-based exact Bayesian filter and Rauch–Tung–Striebel (RTS) smoother to the two-state-variable quasi-static spring-slider of Becker (2000), a low-dimensional chaotic model of frictional fault slip. The system is weakly chaotic, with Lyapunov time $\tau_L \approx 216$ time units and Kaplan–Yorke dimension $D_{KY} = 2.057$. The near-2D attractor supports exact propagation of the posterior on a 50^3 grid.

State estimation and forecasting are evaluated under three truth scenarios using a deterministic forward model: a matched ODE, an SDE with additive noise on slip velocity (x_1), and an SDE with additive noise on the slow state variables (x_2, x_3). Scalar observations of x_1 are assimilated at $\Delta t_{\text{obs}} = 4$ with $\sigma_{\text{obs}} = 0.05$.

Tracking skills remain within 0.62–0.65 across all cases, but calibration diverges: SDE(x_1) is overdispersed (z-score std 1.435), while SDE(x_2, x_3) is near-calibrated (1.090). The RTS smoother improves posterior accuracy throughout, especially for unobserved variables.

Propagating the final posterior forward without updates yields skill horizons of 0.3–0.7 τ_L and coverage horizons of 0.4–0.6 τ_L , with SDE(x_2, x_3) shortest on both. This is consistent with the random-attractor picture of Gualandi et al. (2023) and this study finds noise on slow fault-state variables damages forecast skill more than equivalent noise on slip velocity. The grid approach scales exponentially with dimension, but within its regime sustains skill horizons approaching the Lyapunov limit.

PLAIN LANGUAGE SUMMARY

Earthquakes are hard to predict because fault systems are chaotic: tiny uncertainties grow rapidly until forecasts become meaningless. This sets a hard ceiling on predictability, called the Lyapunov time. We ask how close we can get to it using a simple model of a fault, the spring-slider, and whether the small random stress fluctuations acting on real faults push the useful forecast window even shorter.

Rather than tracking a single best guess of the fault's state, we track the full range of plausible states and update that range exactly as the system evolves. Even when our model ignores the random forcing on the true fault, it tracks behaviour reliably and produces useful forecasts across roughly two-thirds of the theoretical predictability window. For real seismic monitoring, the implication is that slow, hidden changes in fault friction are the main bottleneck on forecasting, more so than noise in our measurements of how fast the fault is slipping. When a forecasting system fails to account for these unseen disturbances, its forecasts degrade much faster, suggesting that unobserved creep is the real limit on earthquake prediction.

Key words: Friction; Non-linear differential equations; Bayesian inference; Instability analysis; Probabilistic forecasting; Earthquake interaction, forecasting, and prediction.

1 INTRODUCTION

This paper applies a grid-based exact Bayesian data assimilation framework to the rate-and-state spring-slider model with stochastic forcing. The spring-slider is a low-dimensional but dynamically rich representation of fault slip that captures the essential features of frictional instability and stick-slip recurrence, and has been used extensively as a testbed for the dynamics of earthquakes and slow slip (Becker 2000; Gualandi et al. 2023; Gauriau et al. 2023). We evaluate the robustness of filtering and smoothing, characterise the performance of grid-based methods for forecasting, and quantify how stochastic forcing degrades both.

The two-state-variable quasi-static spring-slider used throughout this paper, formulated by Becker (2000), consists of a block of unit mass connected via a spring of finite stiffness to a load point moving at constant velocity, with the frictional interface described by a rate- and state-dependent friction law (Dieterich 1979; Ruina 1983). For sufficiently compliant springs and an appropriate range of loading velocities the system exhibits deterministic chaos, and recent work has mapped the structure of these chaotic regimes in parameter space (Gauriau et al. 2023; Wang 2024). Gualandi et al. (2020) showed that a decade of Cascadia slow-slip events is consistent with a low-dimensional chaotic attractor, inferring positive Lyapunov exponents via extreme value statistics on the geodetic record. Gualandi et al. (2023) extended this picture using laboratory stick-slip experiments, finding that the observed cycles are best explained by a random attractor: rate-and-state friction forced by small stochastic stress perturbations ($\sim 0.05\%$) producing coefficient-of-variation changes of a few percent. This motivates the central question: how does a Bayesian data assimilation scheme behave when the truth is a stochastic differential equation (SDE) but the forward model remains deterministic?

Data assimilation for rate-and-state fault systems has been dominated by ensemble methods, beginning with the Ensemble Kalman Filter (EnKF) applied to a visco-elasto-plastic continuum (van Dinther et al. 2019) and extended to frictional parameter recovery (Hirahara & Nishikiori 2019) and 1-D sequences spanning both slow-slip and earthquake regimes (Diab-Montero et al. 2023). The non-Gaussian transitions at coseismic onset violate the Gaussian closure underlying these methods, producing underdispersed posteriors and degraded skill in chaotic regimes; recent responses adopt non-Gaussian filters (Diab-Montero et al. 2025) or learned reduced-order models (Kaveh et al. 2025).

Our approach is methodologically distinct. Following Al-Attar (2026), a grid-based exact Bayesian filter and Rauch-Tung-Striebel (RTS) smoother (Rauch et al. 1965) are implemented directly on the three-dimensional state space. Representing the posterior on a regular grid propagates non-Gaussian density exactly within discretisation error, sidestepping the closure assumptions that limit ensemble methods at coseismic transitions, at the cost of feasibility restricted to low-dimensional problems. The spring-slider's three state variables sit comfortably within this regime.

A subsidiary question concerns smoothing, which has been argued to be problematic for chaotic systems because the unstable subspace resists backward integration (Al-Attar 2026; Carrassi et al. 2022). For the weakly chaotic spring-slider with modest noise, the RTS smoother nonetheless recov-

ers useful information about past states up to a horizon set by grid resolution rather than Lyapunov instability. Forecasting beyond the assimilation window is the practical motivation. For a deterministically chaotic system the maximum Lyapunov exponent sets the e -folding rate of initial-condition errors, so that $\tau_L = 1/\lambda_{\max}$ provides a natural reference timescale, with the standard heuristic ceiling at a small multiple of τ_L (Gualandi et al. 2020). Prior studies have extended useful forecast horizons by sequentially constraining the state with observations (van Dinther et al. 2019; Diab-Montero et al. 2023; Kaveh et al. 2025), typically reporting skill as alarm duration relative to recurrence interval rather than in Lyapunov-time units. A coarse grid-based posterior, propagated forward without further observations across multiple stick-slip cycles, sustains forecast skill on timescales of order $0.4\text{--}0.8\tau_L$. Skill degrades to varying levels when the truth includes stochastic forcing.

2 METHODS

2.1 Deterministic spring-slider system

The two-state-variable quasi-static spring-slider of Becker (2000) provides the model system, governed by three coupled ordinary differential equations (ODEs) in the logarithmic slip-velocity and state variables $[x_1, x_2, x_3]$:

$$\dot{x}_1 = -\kappa(e^{x_1} - v_0) + e^{x_1} [x_1(\beta_1 - 1) + x_2 - x_3] + \rho e^{x_1} (\beta_2 x_1 + x_3), \quad (1)$$

$$\dot{x}_2 = -\kappa(e^{x_1} - v_0), \quad (2)$$

$$\dot{x}_3 = -\rho e^{x_1} (\beta_2 x_1 + x_3). \quad (3)$$

Parameter values are taken from Becker (2000) and A. Gualandi (personal comms.): $\beta_1 = 1$, $\beta_2 = 0.84$, $\rho = 0.048$, $v_0 = 1$, and $\kappa' = 0.8544$. To characterise the fault regime, we define the dimensionless stiffness ratio κ' relative to the second critical stiffness $\kappa_{\text{cr}}^{(2)}$:

$$\kappa' = \frac{\kappa}{\kappa_{\text{cr}}^{(2)}}, \quad (4)$$

$$\kappa_{\text{cr}}^{(2)} = \frac{1}{2(1 + \rho)} \left(\kappa_{\text{cr}}^{(1)} + \rho[2\beta_1 + (\beta_2 - 1)(2 + \rho)] + \sqrt{4\rho^2(\kappa_{\text{cr}}^{(1)} + \beta_2) + (\kappa_{\text{cr}}^{(1)} + \rho^2(\beta_2 - 1))^2} \right), \quad (5)$$

where $\kappa_{\text{cr}}^{(1)} = \beta_1 - 1$. Physically, κ' governs the fault regime. For $\kappa' > 1$, the spring is stiff: the loading rate cannot sustain runaway slip, and the system settles to steady sliding. For $\kappa' < 1$, the spring is compliant: it cannot stop accelerating slip, giving rise to stick-slip instability. The value $\kappa' = 0.8544$ places the system well inside the chaotic regime.

The ODE is integrated using the Vern9 algorithm (Verner 2010) within `DifferentialEquations.jl` (Rackauckas & Nie 2017), with absolute and relative tolerances 10^{-10} , from $t = 0$ to $t = 50\,000$ with initial condition $\mathbf{x}(0) = [0.05, 0, 0]^T$. The first 400 time units are discarded as transient startup, yielding a post-transient record of $\approx 49\,600$ time units ($\approx 230\tau_L$, where $\tau_L = 1/\lambda_1$ is the Lyapunov time). For stochastic differential equation (SDE) runs, a small radiation-damping correction $\nu = 10^{-3}$ is included in the denominator of \dot{x}_1 , replacing it with

$(1 + \nu e^{x_1})$; this regularises near-singular dynamics at high slip velocities.

Lyapunov exponents are computed using Benettin's algorithm (Benettin et al. 1980) via periodic QR re-orthonormalisation. We augment the system state with a 3×3 tangent frame Q , whose columns are perturbation vectors evolving under the linearised dynamics. At intervals of $\Delta t_{\text{reortho}}$, we decompose the tangent frame as $Q \rightarrow QR$ and re-orthonormalise its columns. This prevents all perturbation directions from collapsing onto the leading Lyapunov vector. We then estimate the Lyapunov exponents by accumulating the logarithms of the diagonal elements of R over the integration:

$$\lambda_i \approx \frac{1}{T} \sum_k \log |R_{ii}^{(k)}|, \quad (6)$$

where T is the total integration time and the sum runs over all re-orthonormalisation steps k .

To characterise the attractor geometry, the Kaplan–Yorke dimension D_{KY} (Kaplan & Yorke 1979) is estimated directly from the Lyapunov spectrum. Given the sorted exponents $\lambda_1 \geq \lambda_2 \geq \lambda_3$, we define it as

$$D_{\text{KY}} = j + \frac{\sum_{i=1}^j \lambda_i}{|\lambda_{j+1}|}, \quad (7)$$

where j is the largest index for which $\sum_{i=1}^j \lambda_i \geq 0$. Geometrically, D_{KY} interpolates between the integer dimension j and the next. Within dimension j , the partial sum of exponents is non-negative; at dimension $j + 1$, accumulated contraction overcomes expansion. The result is a fractal estimate of the attractor dimension.

The correlation dimension d_2 is also estimated using the Grassberger–Procaccia algorithm (Grassberger & Procaccia 1983). We compute the correlation sum $C(\varepsilon)$ over a subsampled trajectory by counting the fraction of point pairs whose Euclidean separation falls within radius ε :

$$C(\varepsilon) = \frac{2}{N(N-1)} \sum_{i < j} \Theta(\varepsilon - \|x_i - x_j\|), \quad (8)$$

where Θ is the Heaviside function and N is the number of trajectory points. To suppress spurious correlations between temporally adjacent points, a Theiler window (Theiler 1986) is applied, excluding pairs separated by fewer than w time steps. Two choices of w are considered: the lag at which the normalised autocorrelation of x_1 first crosses zero (the minimal defensible value), and the lag at which it first drops below $1/e$ (a more conservative standard commonly used in embedding analyses). d_2 under the zero-crossing window is reported in the main text; the sensitivity of d_2 to the choice of w is examined as a diagnostic in the supplemental material (Fig. A2). We then extract the correlation dimension as the scaling exponent

$$d_2 = \frac{d \log C(\varepsilon)}{d \log \varepsilon}, \quad (9)$$

estimated by linear regression over an automatically detected scaling region. The local slope $d(\log C)/d(\log \varepsilon)$ is computed at each ε via central differences; points where $C(\varepsilon) <$

0.005 (noise floor) or $C(\varepsilon) > 0.20$ (saturation and transition to the plateau) are excluded. The scaling region is detected automatically: candidate windows of initial width $k \approx n/4$ are discarded if the OLS slope falls below 0.5, R^2 below 0.99, or local-slope standard deviation σ_s exceeds 0.10; surviving candidates are grown rightward while all three constraints hold and ranked by

$$Q = \frac{\ell}{1 + \alpha \sigma_s}, \quad \alpha = 10, \quad (10)$$

where ℓ is window length, favouring long, flat plateaux over curved transition regions. Finite-sample convergence is assessed at six subsampling fractions (1–100%) and d_2 is recorded as w is swept up to one Lyapunov time; all evaluations share a single $O(N^2)$ pass over pair distances (Figs. A1 and A2).

2.2 Stochastic implementation

Two cases of stochastic forcing are considered. In both, additive Gaussian white noise enters through a scalar Wiener increment dW , with amplitude σ and diffusion vector g .

In the first case, noise is applied directly to the log slip-velocity x_1 :

$$dx = f(x) dt + \sigma g dW, \quad g = (1, 0, 0)^T. \quad (11)$$

In the second case, following the motivations outlined in Gualandi et al. (2023), noise is applied to the state variables x_2 and x_3 . This represents unresolved stress perturbations acting on the normal-stress and friction state rather than directly on slip velocity:

$$dx = f(x) dt + \sigma g dW, \quad g = (0, 1, 1)^T. \quad (12)$$

Both cases use $\sigma = 0.01$ for data-assimilation experiments. This motivates a labelling convention: x_1 forcing is termed *observed* noise and x_2/x_3 forcing *unobserved* noise, reflecting whether the noise acts directly on the observed variable.

Both SDE cases are integrated with the Euler–Heun scheme (Kloeden & Platen 1992), using a fixed step $\Delta t = 10^{-3}$ and saving every 0.1 time units.

200-member ensembles are generated for each SDE case at three noise levels $\sigma \in \{5 \times 10^{-3}, 10^{-2}, 2 \times 10^{-2}\}$. All members share the same on-attractor initial condition, the first post-transient point of the deterministic trajectory, simulating a settled system.

2.3 Bayesian filtering and smoothing: framework

Bayesian state estimation for a dynamical system alternates two operations. Given the posterior density $p(x_{t-1} | y_{1:t-1})$ at the previous observation time, the *prediction* step propagates it forward in time using the model dynamics,

$$p(x_t | y_{1:t-1}) = \int p(x_t | x_{t-1}) p(x_{t-1} | y_{1:t-1}) dx_{t-1}, \quad (13)$$

and the *analysis* step conditions on the new observation y_t via Bayes' rule,

$$p(x_t | y_{1:t}) \propto p(y_t | x_t) p(x_t | y_{1:t-1}). \quad (14)$$

Equations (13)–(14) define the optimal filter and are exact: every choice that follows is a representational decision about how to evaluate them. Ensemble methods replace p with a finite sample; Kalman methods assume Gaussianity and propagate only mean and covariance. Following Al-Attar (2026), p is instead represented as gridded values on a fixed lattice covering the attractor, and (13)–(14) are evaluated directly. This is feasible only because the spring-slider state space is three-dimensional; the cost is $O(N^3)$ per step, but no closure is imposed on the shape of p .

2.4 Bayesian data assimilation: deterministic truth

The posterior is discretised on a regular 50^3 grid covering the attractor’s bounding box with a 20% margin per axis. The prior is a broad isotropic Gaussian, $\mathcal{N}(0, \sigma_{\text{prior}}^2 I)$ with $\sigma_{\text{prior}} = 1.5$, centred away from the true initial state. This ensures convergence reflects the filter’s assimilation skill rather than a tuned initial guess.

Prediction step under deterministic flow. Because the forward model is a deterministic ODE $\dot{x} = f(x)$ with flow $\Phi_{\Delta t}$, the transition kernel collapses to a delta distribution, $p(x_t | x_{t-1}) = \delta(x_t - \Phi_{\Delta t}(x_{t-1}))$. The convolution in (13) then reduces to a change of variables:

$$p(x_t | y_{1:t-1}) = p(x_0 | y_{1:t-1}) \left| \det(\partial \Phi_{\Delta t}^{-1} / \partial x_t) \right|, \quad (15)$$

$$x_0 \equiv \Phi_{\Delta t}^{-1}(x_t). \quad (16)$$

For each grid point x_t , we therefore (i) integrate the negated ODE backward by $\Delta t_{\text{obs}} = 4$ time units (Vern9, tolerances 10^{-8} , ten sub-steps of 0.4) to obtain the pre-image x_0 , and (ii) accumulate the inverse-Jacobian determinant simultaneously. Letting $J(t) \equiv \left| \det(\partial \Phi_t^{-1} / \partial x_t) \right|$, Liouville’s theorem applied to the inverse flow gives

$$\frac{d \log J}{dt} = -\text{tr} \left(\frac{\partial f}{\partial x} \right), \quad (17)$$

which we integrate along the backward trajectory as a fourth augmented state. The pushed-forward density is then

$$p_t(x_t) = p_{t-1}(x_0) \exp(\log J), \quad (18)$$

with trilinear interpolation for off-node pre-images (Jazwinski 2007); grid points whose pre-images leave the domain are assigned zero density. In the quasi-static limit ($\nu = 0$, used for the forward model), $\dot{x}_2 = -\kappa(e^{x_1} - \nu_0)$ depends only on x_1 , so $J_{22} = \partial \dot{x}_2 / \partial x_2 = 0$ and the trace in (17) reduces to $J_{11} + J_{33}$. The truth integration uses $\nu = 10^{-3}$; this small mismatch makes even the “deterministic” case a mild model-truth mismatch experiment rather than a perfect-model one.

Analysis step. Synthetic scalar observations of x_1 are generated by adding $\mathcal{N}(0, \sigma_{\text{obs}}^2)$ noise with $\sigma_{\text{obs}} = 0.05$ to the true trajectory at $N_{\text{obs}} = 40$ equally spaced times. The likelihood factorises across grid points, so we implement (14) as pointwise multiplication of the predicted density by $\exp[-(y_t - x_1)^2 / (2\sigma_{\text{obs}}^2)]$, followed by renormalisation against the grid quadrature.

Skill diagnostics. At each analysis, a probability-mass score $q \in [0, 1]$ is recorded, defined as the fraction of total posterior mass at grid points with density *below* that at the nearest grid

point to the truth. For a Gaussian, $q > 0.5$ corresponds approximately to the truth lying within the 1σ contour; values near 1 indicate the truth sits at a density mode. The standardised innovation is also computed,

$$z_n = \frac{x_1^{\text{true}}(t_n) - \mu_n}{\sigma_n}, \quad (19)$$

where μ_n, σ_n are the posterior mean and standard deviation. A well-calibrated filter produces $z_n \sim \mathcal{N}(0, 1)$; deviations diagnose under- or over-confidence. The empirical z distribution is summarised with a Gaussian-kernel KDE using Silverman’s rule (Silverman 1986),

$$h = 0.9 \hat{\sigma} n^{-1/5}, \quad \hat{\sigma} = \min(\text{std}(z), \text{IQR}(z)/1.34), \quad (20)$$

which is robust to outliers in the IQR-based scale estimator.

RTS smoother. The fixed-interval smoother $p(x_t | y_{1:T})$ ($T \gg t$) follows the standard Rauch et al. (1965) recursion, which under the Markov assumption on the dynamics takes the form

$$p(x_t | y_{1:T}) = p(x_t | y_{1:t}) \int \frac{p(x_{t+1} | y_{1:T}) p(x_{t+1} | x_t)}{p(x_{t+1} | y_{1:t})} dx_{t+1}. \quad (21)$$

Under deterministic flow, the transition kernel collapses again and the integral reduces to a pointwise ratio: the smoothed posterior at x_t is the filter posterior multiplied by $p_{\text{smooth}}(\Phi_{\Delta t}(x_t)) / p_{\text{pred}}(\Phi_{\Delta t}(x_t))$. This is implemented as a backward pass over stored filter and prediction marginals. Where p_{pred} is small the ratio can blow up; two variants are run, one capped at $r_{\text{max}} = 100$ and one uncapped, and both are reported as a robustness check on the smoother’s behaviour under model-truth mismatch.

Filter and smoother are applied to all three truth cases (deterministic ODE; x_1 -noise SDE; x_2/x_3 -noise SDE) using the same deterministic forward model throughout, so that the SDE cases probe robustness to unresolved stochastic forcing.

2.5 Bayesian forecasting

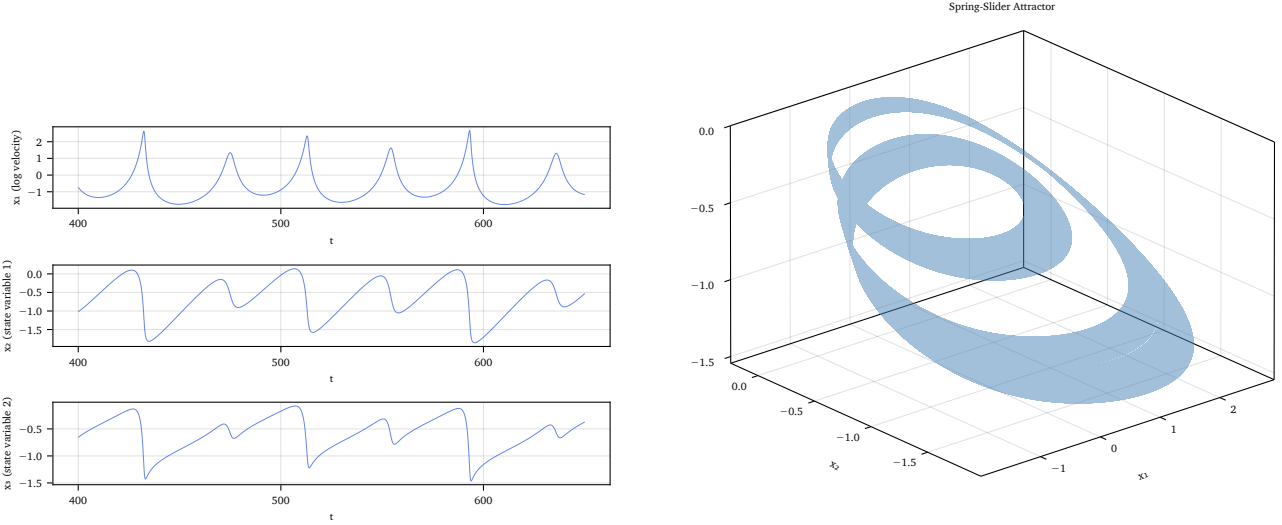
The assimilation window spans 15 stick-slip cycles ($t \in [430, 1550]$, $\approx 5.2 \tau_L$), identified as local maxima of $x_1 \geq 2.0$ in the post-transient trajectory. Within this window, observations are taken at fixed interval $\Delta t_{\text{obs}} = 10$, yielding $N_{\text{obs}} = 115$ observations with the same noise model as before. The following $N_{\text{forecast}} = 6$ cycles provide ground truth for forecast verification.

After assimilating all N_{obs} observations on the 50^3 grid, the final posterior is propagated forward by repeated application of (18) without further updates, with sub-steps of $\Delta t_{\text{sub}} = 1.0$. At each sub-step, the posterior mean $\hat{x}_1(t)$ and standard deviation $\sigma_{\text{post}}(t)$ are recorded.

Skill horizon. Chaotic stick-slip dynamics cause large but transient phase errors at event peaks. A small timing mismatch produces a brief spike in normalised error even when the forecast is otherwise skilful. To filter these spikes, we define the absolute, posterior-normalised error

$$z(t) = \frac{|\hat{x}_1(t) - x_1^{\text{true}}(t)|}{\sigma_{\text{post}}(t)}, \quad (22)$$

smooth it with a moving average over a quarter of the median inter-event cycle (≈ 20 time units), and define the *skill horizon*



(a) Post-transient time series of the three state variables over $t \in [400, 650]$, spanning approximately three stick-slip cycles. x_1 (log slip velocity) exhibits sharp seismic spikes separated by slow interseismic loading; x_2 and x_3 evolve more gradually. System parameters: $\kappa' = 0.8544$, $\beta_1 = 1.0$, $\beta_2 = 0.84$, $\rho = 0.048$.

(b) Strange attractor of the spring-slider ODE in (x_1, x_2, x_3) phase space, computed from the full post-transient trajectory ($t \in [400, 50000]$). The trajectory never closes, consistent with the fractal Kaplan–Yorke dimension $D_{KY} = 2.057$ and correlation dimension $d_2 = 1.507$ (Grassberger–Procaccia with Theiler window).

Figure 1. Characterisation of the deterministic spring-slider system ($\kappa' = 0.8544$, leading Lyapunov exponent $\lambda_1 = +4.62 \times 10^{-3}$, Lyapunov time $\tau_L \approx 216$ time units). *Left:* representative time series showing the stick-slip rhythm. *Right:* phase-space attractor confirming weakly fractal chaotic geometry ($D_{KY} = 2.057$, $d_2 = 1.507$).

as the lead time at which the smoothed z first exceeds 1.5. For a calibrated forecast, $|z| > 1.5$ corresponds to the truth leaving the central $\approx 87\%$ of posterior mass, a natural threshold for “the forecast has drifted outside its own claimed uncertainty”. Normalising by σ_{post} rather than a fixed climatological scale automatically accommodates the growth of forecast spread with lead time.

This is complemented by the cumulative mean $\bar{z}(T) = T^{-1} \int_0^T z dt$, a running calibration diagnostic; the lead time at which \bar{z} first reaches 1 marks the point where forecast error and posterior spread agree on average.

Coverage fraction. A complementary calibration diagnostic tracks the running $\pm 1\sigma$ coverage fraction: at each lead time T , the fraction of all forecast steps up to T for which $|\hat{x}_1(t) - x_1^{\text{true}}(t)| \leq \sigma_{\text{post}}(t)$. For a well-calibrated Gaussian posterior, this should remain at the nominal 68% level. The *coverage horizon* T_{cov} is defined as the last time at which the running fraction remains $\geq 68\%$; it measures how long the posterior spread provides honest coverage of the true trajectory, and is less sensitive to isolated phase-error spikes than the instantaneous z-score threshold. The posterior mean and spread, z-score, and coverage fraction make the three diagnostics used and are shown in Figs. 4 and 5.

3 RESULTS

3.1 Deterministic system characterisation

The deterministic spring-slider system, shown in Figure 1, exhibits a positive leading Lyapunov exponent $\lambda_1 = +4.62 \times 10^{-3}$,

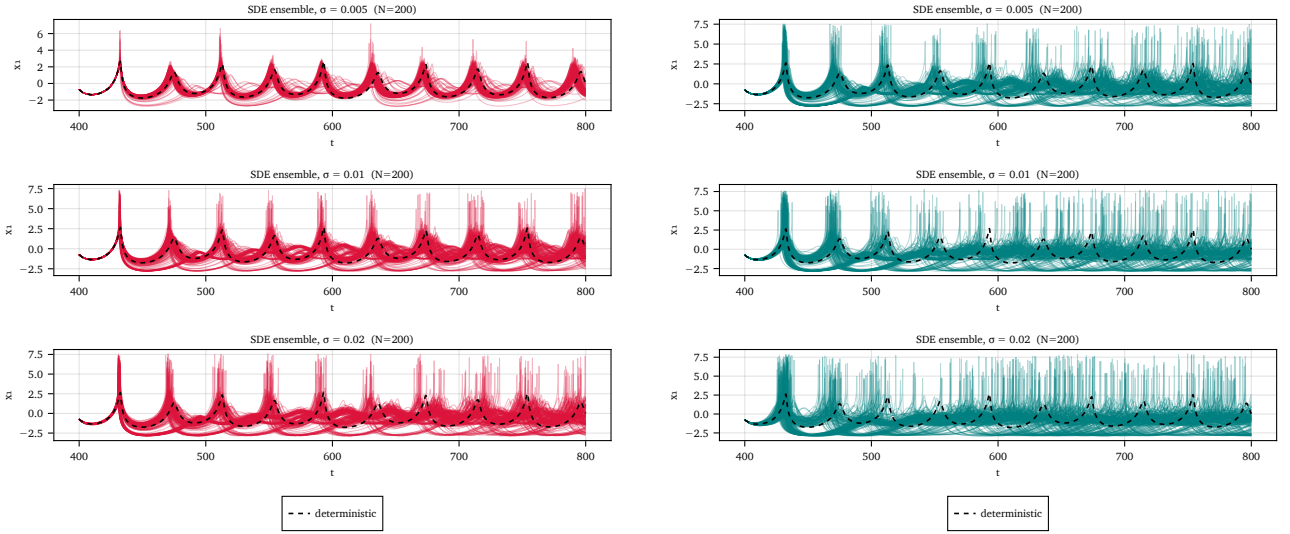
a near-zero second exponent $\lambda_2 \approx 0$ (-1.0×10^{-5}), and a negative third exponent $\lambda_3 = -0.081$, yielding a Lyapunov time of $\tau_L \approx 216$ time units. This confirms that the system is chaotic, albeit with relatively mild instability compared to more strongly chaotic systems.

The attractor geometry is characterised by a Kaplan–Yorke dimension of $D_{KY} = 2.057$, indicating that the attractor is confined to a near-2D subspace within the 3D state space. The correlation dimension, estimated via the Grassberger–Procaccia method with a Theiler window of 92 steps (the first zero-crossing of the x_1 autocorrelation), is $d_2 = 1.507$, which is consistent with the theoretical requirement that $d_2 \leq D_{KY}$ for strange attractors.

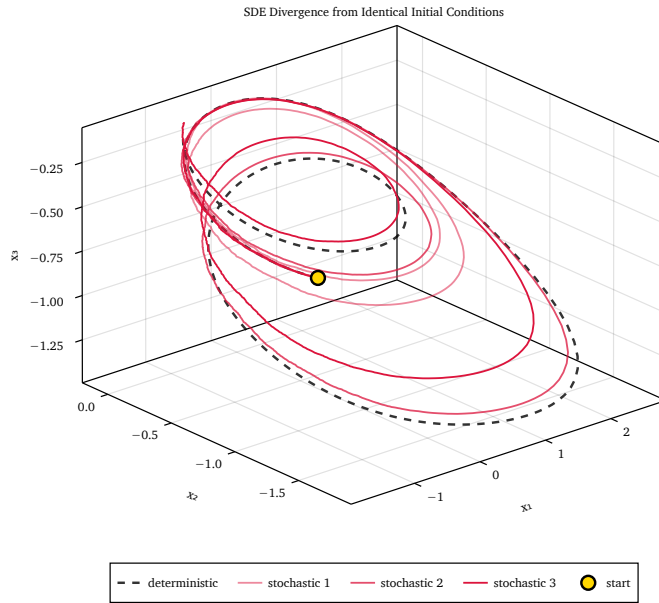
The negligible difference in d_2 with and without the Theiler window (1.506 vs 1.507) confirms that temporal correlations do not significantly bias the dimension estimate, and convergence analysis shows that d_2 stabilises from around 10% of the trajectory length onward. These results confirm that the attractor has a weakly fractal, near-2D geometry, which supports the feasibility of applying grid-based data assimilation on a 50^3 grid to resolve the relevant state space effectively.

3.2 SDE ensemble behaviour

Ensemble realisations of the spring-slider SDE, shown in Figure 2, reveal that the location of the stochastic forcing has a significant impact on the spread of the slip velocity x_1 . When additive noise is applied directly to x_1 , the mean ensemble spread $\bar{\sigma}(x_1)$ over $t \in [400, 800]$ ranges from 0.39 to 0.83 as the noise amplitude σ increases from 5×10^{-3} to 2×10^{-2} . In contrast, when the same amplitude of noise is applied to the



(a) 200-member SDE ensemble (x_1 component, additive noise on x_1) at three noise amplitudes $\sigma \in \{5 \times 10^{-3}, 10^{-2}, 2 \times 10^{-2}\}$ (light to dark), overlaid on the deterministic reference (black dashed). Mean ensemble spread $\bar{\sigma}(x_1) = 0.39, 0.75, 0.83$ respectively over $t \in [400, 800]$. (b) Same as (a) but with additive noise on the slow variables (x_2, x_3) at identical amplitudes. Mean spread $\bar{\sigma}(x_1) = 0.85, 0.90, 0.94$ — substantially larger than the x_1 -noise case, showing that slow-variable perturbations project efficiently onto slip velocity.



(c) Three individual SDE realisations ($\sigma = 10^{-2}$ on x_1 , coloured) launched from identical initial conditions in (x_1, x_2, x_3) phase space, alongside the deterministic reference (black). Trajectories diverge within one inter-event cycle (≈ 80 time units), illustrating how noise amplifies the underlying chaotic instability ($\tau_L \approx 216$ time units).

Figure 2. Stochastic ensemble behaviour of the spring-slider SDE ($N = 200$ members). *Top left/right:* x_1 spread under noise on x_1 vs. noise on (x_2, x_3) at $\sigma \in \{5 \times 10^{-3}, 10^{-2}, 2 \times 10^{-2}\}$; slow-variable noise produces greater spread at all amplitudes. *Bottom:* phase-space divergence of three individual realisations ($\sigma = 10^{-2}$, x_1 noise) from a shared initial condition, demonstrating chaotic amplification of stochastic perturbations.

slow variables (x_2, x_3), the mean spread in x_1 is substantially larger, ranging from 0.85 to 0.94 across the same noise amplitudes.

3.3 Bayesian data assimilation

Figure 3 shows the grid-based Bayesian filter and RTS smoother posterior mean $\pm 1\sigma$ for all three state variables under three truth scenarios with observations on x_1 alone.

The ODE case with matched deterministic truth (Fig. 3a)

achieves a mean tracking skill of 0.629, with a z-score standard deviation of 1.018, indicating that the filter is well-calibrated and does not suffer from ensemble collapse or underdispersion.

Under model-truth mismatch, the $SDE(x_1)$ (Fig. 3b) and $SDE(x_2x_3)$ (Fig. 3c) cases with a deterministic forward model yield mean tracking skills of 0.646 and 0.622 respectively; all three skills cluster within 0.62–0.65.

However, the z-score standard deviations, shown in Figure 5a, reveal differences in calibration: the $SDE(x_1)$ case is overdispersed (z-score std 1.435), while $SDE(x_2x_3)$ is near well-calibrated (z-score std 1.090). In all cases, the RTS smoother’s backward pass improves the posterior estimates.

In constraining the x_2 and x_3 posteriors, the smoother has a large impact on the accuracy of the unobserved variables. For the ODE and $SDE(x_1)$ cases, the filter and to a greater extent the smoother constrains the unobserved variables well, with a good estimation of uncertainty in the posterior. However, in the $SDE(x_2x_3)$ case, when there is a large jump in the slow variables to outside the typical attractor bounds, the filter PDF for these loses contact with the true trajectory, leading to larger errors. The smoother helps to partially recover the true state during these episodes.

3.4 Forecasting

Figure 4 shows the Bayesian forecast of x_1 following assimilation of $N_{\text{obs}} = 115$ observations at $\Delta t_{\text{obs}} = 10$ time units spanning 15 stick-slip events ($t \in [430, 1550]$, $\approx 5.2 \tau_L$; $\sigma_{\text{obs}} = 0.05$; 50^3 grid). All three cases track the unseen forward trajectory well initially. In the ODE case, the posterior mean dampens relative to the true trajectory after 3–4 events; the SDE cases lose phase alignment sooner, with $SDE(x_2x_3)$ drifting earliest.

The smoothed z-score $z(t)$, shown in the right column of Figure 4, crosses the skill threshold of 1.5 at skill horizons of 83 time units ($0.38 \tau_L$) for the ODE truth, 145 time units ($0.67 \tau_L$) for the $SDE(x_1)$ truth, and 69 time units ($0.32 \tau_L$) for the $SDE(x_2x_3)$ truth. The coverage horizons (Fig. 5), defined as the last time the running $\pm 1\sigma$ coverage fraction remains $\geq 68\%$, are 122 time units ($0.56 \tau_L$) for the ODE case, 120 time units ($0.55 \tau_L$) for the $SDE(x_1)$ case, and 80 time units ($0.37 \tau_L$) for the $SDE(x_2x_3)$ case.

4 DISCUSSION

The weakly chaotic, low-dimensional attractor ($D_{KY} = 2.057$, $d_2 = 1.507$) supports the feasibility of grid-based data assimilation: the near-2D geometry keeps the relevant state space tractable on a 50^3 grid. The matched ODE case (z-score std ≈ 1) demonstrates that the grid-based filter then tracks the true trajectory reliably, even at coseismic transitions where the posterior is strongly non-Gaussian. The RTS smoother’s backward pass adds value by incorporating future observations, though the theoretical validity of smoothing in chaotic systems remains an open question (e.g. Al-Attar 2026); the mild Lyapunov instability ($\tau_L \approx 216$) likely mitigates the typical difficulties.

Both stochastic cases remain within the 0.62–0.65 skill range. The $SDE(x_1)$ overdispersion (z-std=1.435) suggests the filter over-hedges against unmodelled stochastic increments, while $SDE(x_2x_3)$ is better calibrated (z-std=1.090) at

slightly lower skill (0.622). The narrow spread across all three cases may reflect sensitivity to the specific stochastic realisation used.

Skill horizons of $0.32\text{--}0.67 \tau_L$ indicate meaningful predictability under mismatch. The $SDE(x_1)$ outlier ($0.67 \tau_L$) likely reflects a favourable threshold-crossing rather than a structural increase in predictability.

The cumulative mean z-score places the ODE forecast horizon at 203 t.u. ($\approx 0.94 \tau_L$), ahead of both SDE cases ($\approx 0.75 \tau_L$), while the coverage fraction metric places ODE and $SDE(x_1)$ at near-identical horizons ($\approx 0.55\text{--}0.56 \tau_L$), both substantially longer than $SDE(x_2x_3)$ at $0.37 \tau_L$. The discrepancy reflects the ODE forecast’s sharper instantaneous skill collapse versus $SDE(x_1)$ ’s more gradual decline: the cumulative z-score captures timing, while coverage fraction captures holistic skill, and both metrics are needed to fully characterise the forecasts.

For τ_L 216 time units scaling to seismic cycle lengths (years to decades), a skill horizon of $0.3\text{--}0.7 \tau_L$ represents meaningful intermediate-term predictability.

5 CONCLUSIONS

This paper demonstrates grid-based exact Bayesian filtering and RTS smoothing for state estimation and forecasting in a low-dimensional chaotic system, achieving meaningful skill ($0.32\text{--}0.67 \tau_L$) under model-truth mismatch for both seen and unseen stochastic forcing. The robustness of the filter to unresolved stochastic forcing supports its use for historic state estimation on weakly chaotic fault systems.

The results are consistent with Gualandi et al. (2023) in that stochastic stress perturbations on slow fault-state variables reduce predictability more than equivalent noise on slip velocity. For real fault systems, unobserved slow creep fluctuations (x_2 , x_3 analogues) are therefore the primary limit on forecast skill, not stochastic fluctuations in seismic velocity alone.

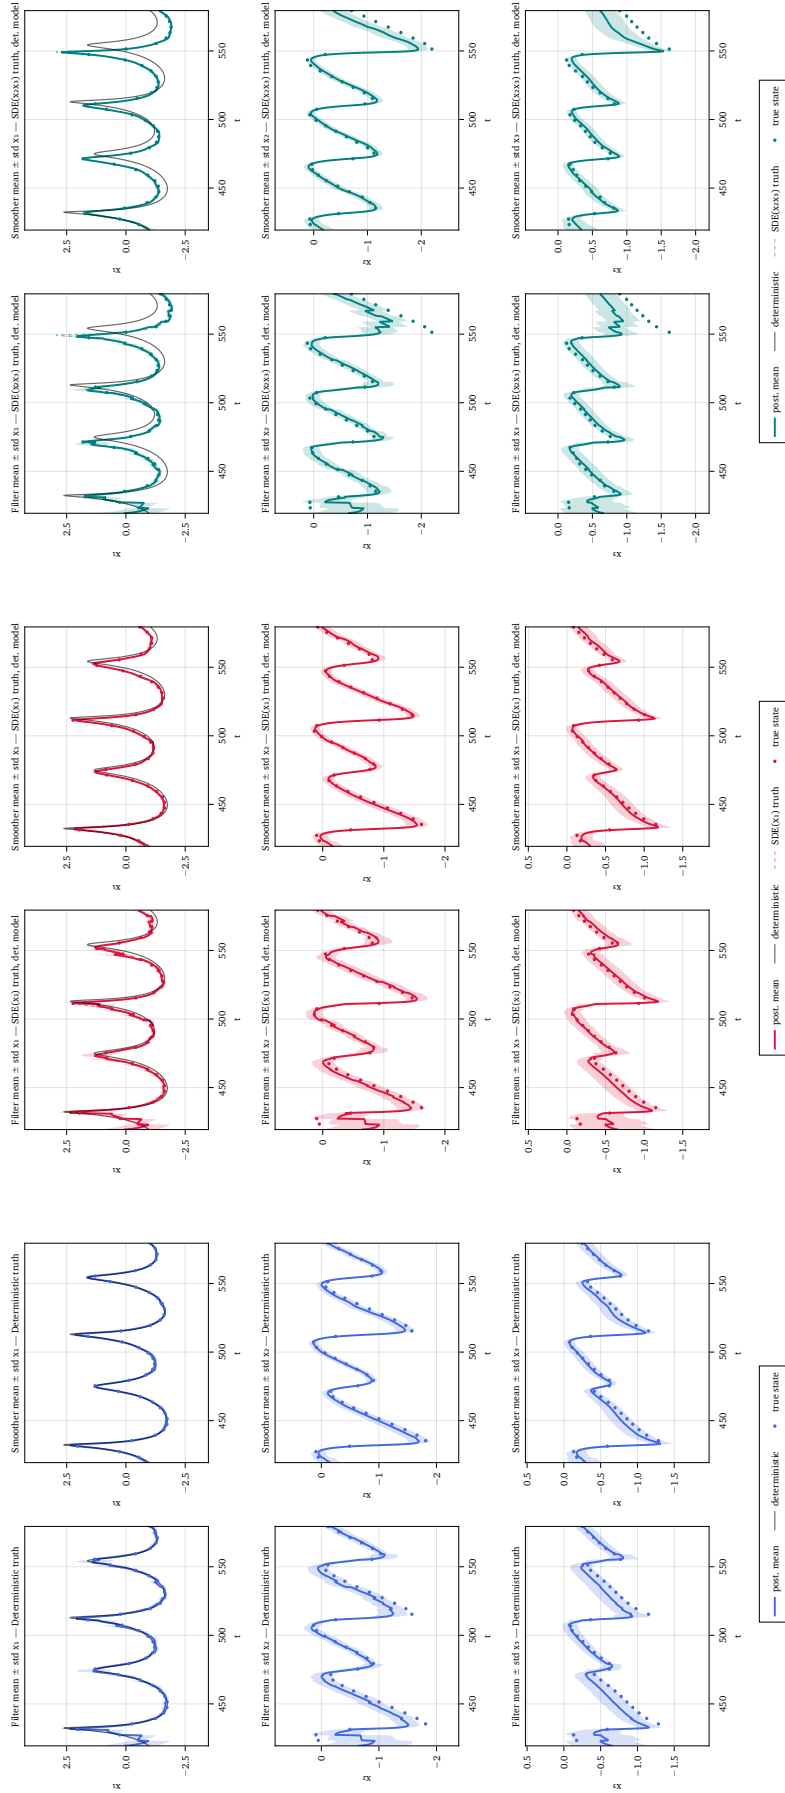
Forecast horizons of up to $0.67 \tau_L$ ($SDE(x_1)$ case) with a deliberately simple grid method suggest that exact Bayesian propagation of non-Gaussian posteriors show that these methods can be used on these systems, at least in the low-dimensional limit.

The primary limitation is the exponential scaling of grid-based assimilation with state-space dimension; dimension-reduction techniques would be required to extend this approach beyond low-dimensional systems.

Robustness assessments across initial conditions, stochastic realisations, grid resolutions, and mismatched model parameters would strengthen these conclusions and better bound the skill horizon under realistic conditions.

Future implementations could incorporate process noise in the forward model, compare with Gaussian-mixture or particle-flow filters (Diab-Montero et al. 2025), or couple the approach to machine-learned reduced-order models to extend applicability to higher-dimensional fault systems.

Grid-based exact Bayesian assimilation is robust to model error and long observation intervals in low-dimensional chaotic systems; practical seismic application will require careful design of grid resolution, observation strategy, and stochastic representation, with event-triggered observations offering a promising path forward for operational networks.



(a) Matched deterministic ODE truth (mean tracking skill 0.629). Rows: x_1, x_2, x_3 ; columns: filter posterior mean $\pm 1\sigma$ (left), RTS smoother (right). Dots: observations ($\sigma_{\text{obs}} = 0.05$). Full marginal PDFs in Fig. A3.

(b) SDE(x_1) truth ($\sigma_{\text{SDE}} = 0.01$) with deterministic forward model (mean tracking skill 0.646). Same layout as (a).

(c) SDE(x_2, x_3) truth ($\sigma_{\text{SDE}} = 0.01$ on x_2, x_3) with deterministic forward model (mean tracking skill 0.622). Same layout as (a). Noise entering the slow variables broadens x_2, x_3 posteriors.

Figure 3. Grid-based Bayesian filter and smoother posterior mean $\pm 1\sigma$ (50^3 grid, $N_{\text{obs}} = 40$, $\Delta t_{\text{obs}} = 0.05$) for all three state variables under three truth scenarios with observations on x_1 alone. Each panel shows six sub-panels (rows: x_1, x_2, x_3 ; columns: filter, smoother). (a) Matched ODE model (skill 0.629). (b) SDE(x_1) truth with deterministic model (skill 0.646). (c) SDE(x_2, x_3) truth with deterministic model (skill 0.622). Full marginal PDFs in Fig. A3.

Event-based Bayesian forecast: ODE vs SDE truth

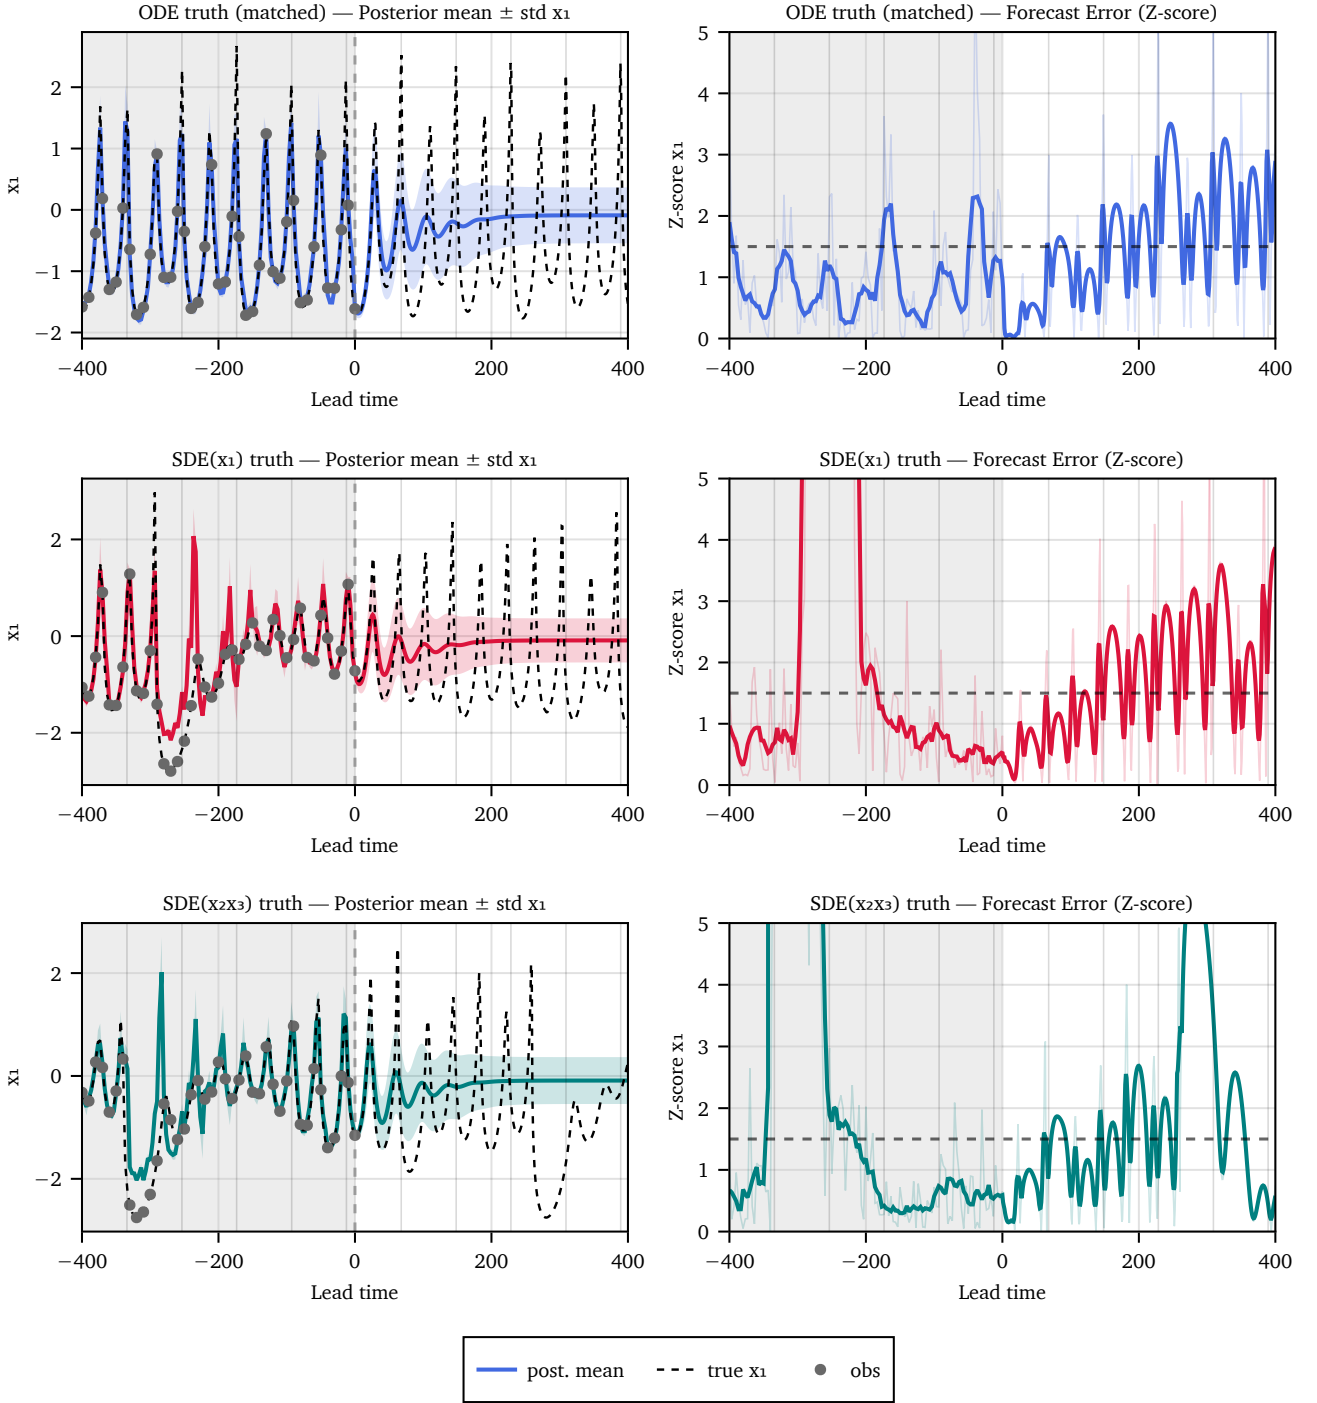
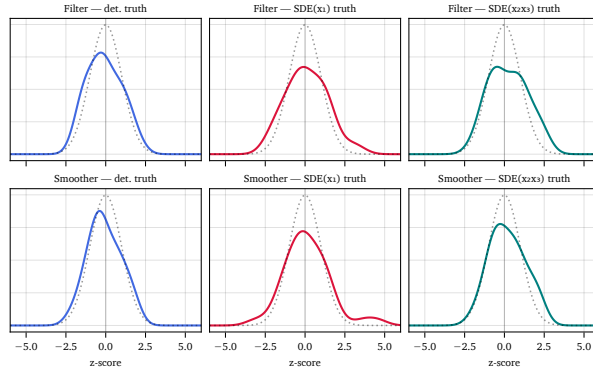
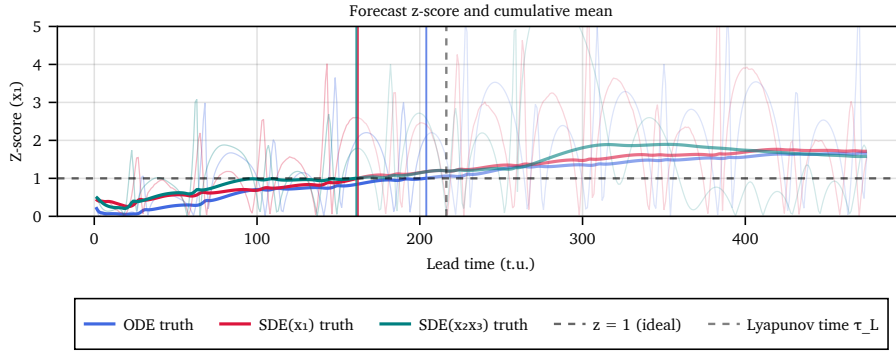


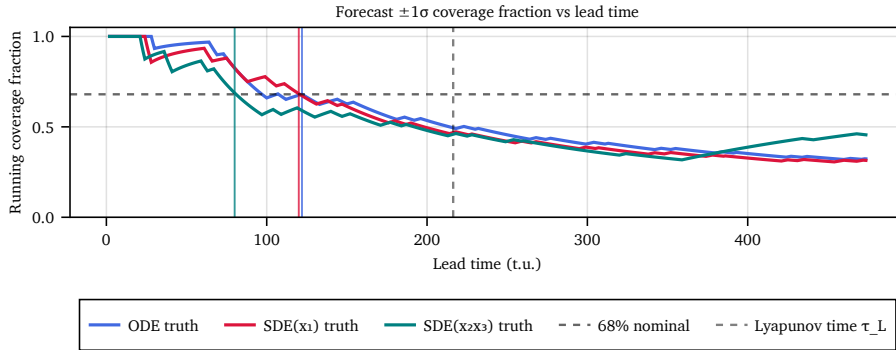
Figure 4. Free Bayesian forecast of x_1 following assimilation of $N_{\text{obs}} = 115$ observations at $\Delta t_{\text{obs}} = 10$ time units spanning 15 stick-slip events ($t \in [430, 1550]$, $\approx 5.2 \tau_L$; $\sigma_{\text{obs}} = 0.05$; 50^3 grid). Each row corresponds to one truth case. *Left:* posterior mean $\pm 1\sigma$ of x_1 vs lead time; grey shading marks the assimilation window (lead time < 0); observations (dots) and true x_1 (solid) are overlaid. *Right:* posterior-normalised z-score $z(t) = |\hat{x}_1 - x_1^{\text{true}}|/\sigma_{\text{post}}$ (faint: instantaneous; solid: rolling-mean smoothed over ≈ 20 time units); dashed horizontal line at $z = 1.5$ marks the skill horizon threshold. Skill horizons (smoothed z first exceeding 1.5): ODE truth 83 time units ($0.38 \tau_L$); SDE(x_1) truth 145 time units ($0.67 \tau_L$); SDE(x_2x_3) truth 69 time units ($0.32 \tau_L$). The running calibration diagnostics are shown in Fig. 5; full marginal probability distributions are in Fig. A4.



(a) Kernel density estimates of the posterior z-score $z = (x_1^{\text{true}} - \mu) / \sigma$ evaluated at each observation time, for the filter (solid) and smoother (dashed) under three truth scenarios: ODE (blue), $SDE(x_1)$ (red), $SDE(x_2x_3)$ (green). The reference $\mathcal{N}(0, 1)$ curve (grey) indicates ideal calibration. Broadened tails under SDE truth reflect episodes where the stochastic truth departs from the deterministic prior.



(b) Instantaneous z-score $z(t) = |\hat{x}_1 - x_1^{\text{true}}| / \sigma_{\text{post}}$ (faint) and cumulative mean z-score $\bar{z}(T) = T^{-1} \int_0^T z dt$ (solid) during the free forecast phase for three truth configurations. Vertical coloured lines mark where \bar{z} first reaches 1.0; horizontal dashed line is $z = 1$; black dashed vertical line is τ_L .



(c) Running $\pm 1\sigma$ coverage fraction: at each lead time T , the fraction of all forecast steps up to T for which $|\hat{x}_1 - x_1^{\text{true}}| \leq \sigma_{\text{post}}$. Dashed horizontal line at 68% (nominal $\pm 1\sigma$ level); coloured vertical lines mark the coverage horizon T_{cov} (last time the running fraction remains $\geq 68\%$): ODE 122 t.u. ($0.56 \tau_L$); $SDE(x_1)$ 120 t.u. ($0.55 \tau_L$); $SDE(x_2x_3)$ 80 t.u. ($0.37 \tau_L$).

Figure 5. Calibration diagnostics for the assimilation and forecasting results. *Top:* z-score distributions at observation times quantify filter and smoother fidelity during the assimilation phase ($N_{\text{obs}} = 40$, $\Delta t_{\text{obs}} = 4$ time units); model–truth mismatch broadens the tails relative to the ideal $\mathcal{N}(0, 1)$. *Middle:* cumulative mean z-score vs. lead time during the free forecast ($N_{\text{obs}} = 115$, $\Delta t_{\text{obs}} = 10$ time units); calibration crossing ($\bar{z} = 1$) arrives sooner under SDE truth. *Bottom:* running $\pm 1\sigma$ coverage fraction vs. lead time; ODE and $SDE(x_1)$ maintain $\geq 68\%$ coverage to ≈ 120 t.u., while $SDE(x_2x_3)$ drops below at 80 t.u. Together, the panels show that slow-variable model error (x_2, x_3 noise) degrades both analysis quality and forecast reliability most severely.

DATA AVAILABILITY STATEMENT

All code and data is viewable on [Github](#) with the output log files with the results in [compute_log.txt](#). Details on how to download and run the code contained within the [README.md](#).

REFERENCES

- Al-Attar, D., 2026. QCES Climate Data: Data Assimilation, https://github.com/da380/QCES_data_assimilation.
- Becker, T. W., 2000. Deterministic chaos in two state-variable friction sliders and the effect of elastic interactions, in *Geophysical Monograph Series*, vol. 120, pp. 5–26, eds Rundle, J. B., Turcotte, D. L., & Klein, W., American Geophysical Union, Washington, D. C., doi: 10.1029/GM120p0005.
- Benettin, G., Galgani, L., Giorgilli, A., & Strelcyn, J.-M., 1980. Lyapunov Characteristic Exponents for smooth dynamical systems and for hamiltonian systems; a method for computing all of them. Part 1: Theory, *Meccanica*, **15**(1), 9–20, doi: 10.1007/BF02128236.
- Carrassi, A., Bocquet, M., Demaeyer, J., Grudzien, C., Raanes, P., & Vannitsem, S., 2022. Data Assimilation for Chaotic Dynamics, in *Data Assimilation for Atmospheric, Oceanic and Hydrologic Applications (Vol. IV)*, pp. 1–42, eds Park, S. K. & Xu, L., Springer International Publishing, Cham, doi: 10.1007/978-3-030-77722-7_1.
- Diab-Montero, H. A., Li, M., van Dinther, Y., & Vossepoel, F. C., 2023. Estimating the occurrence of slow slip events and earthquakes with an ensemble Kalman filter, *Geophysical Journal International*, **234**(3), 1701–1721, doi: 10.1093/gji/ggad154.
- Diab-Montero, H. A., Stordal, A. S., Van Leeuwen, P. J., & Vossepoel, F. C., 2025. Ensemble Kalman, adaptive Gaussian mixture, and particle flow filters for optimized earthquake occurrence estimation, *Computers & Geosciences*, **196**, 105836, doi: 10.1016/j.cageo.2024.105836.
- Dieterich, J. H., 1979. Modeling of rock friction: 1. Experimental results and constitutive equations, *Journal of Geophysical Research: Solid Earth*, **84**(B5), 2161–2168, doi: 10.1029/JB084iB05p02161.
- Gauriau, J., Barbot, S., & Dolan, J. F., 2023. Islands of chaos in a sea of periodic earthquakes, *Earth and Planetary Science Letters*, **618**, 118274, doi: 10.1016/j.epsl.2023.118274.
- Grassberger, P. & Procaccia, I., 1983. Measuring the strangeness of strange attractors, *Physica D: Nonlinear Phenomena*, **9**(1-2), 189–208, doi: 10.1016/0167-2789(83)90298-1.
- Gualandi, A., Avouac, J.-P., Michel, S., & Faranda, D., 2020. The predictable chaos of slow earthquakes, *Science Advances*, **6**(27), eaaz5548, doi: 10.1126/sciadv.aaz5548.
- Gualandi, A., Faranda, D., Marone, C., Cocco, M., & Mengaldo, G., 2023. Deterministic and stochastic chaos characterize laboratory earthquakes, *Earth and Planetary Science Letters*, **604**, 117995, doi: 10.1016/j.epsl.2023.117995.
- Hirahara, K. & Nishikiori, K., 2019. Estimation of frictional properties and slip evolution on a long-term slow slip event fault with the ensemble Kalman filter: Numerical experiments, *Geophysical Journal International*, **219**(3), 2074–2096, doi: 10.1093/gji/ggz415.
- Jazwinski, A. H., 2007. *Stochastic Processes and Filtering Theory*, Dover Publications, Mineola, N.Y., unabridged republ edn., This Dover edition, first published in 2007, is an unabridged republication of the work originally published by Academic Press, New York, in 1970.
- Kaplan, J. L. & Yorke, J. A., 1979. Chaotic behavior of multidimensional difference equations, in *Functional Differential Equations and Approximation of Fixed Points (Proc. Summer School and Conf., Univ. Bonn, Bonn, 1978)*, vol. 730 of *Lecture Notes in Math.*, pp. 204–227, Springer, Berlin.
- Kaveh, H., Avouac, J. P., & Stuart, A. M., 2025. Data assimilation in machine-learned reduced-order model of chaotic earthquake sequences, *Geophysical Journal International*, **244**(2), ggaf518, doi: 10.1093/gji/ggaf518.
- Kloeden, P. E. & Platen, E., 1992. *Numerical Solution of Stochastic Differential Equations*, Springer Berlin Heidelberg, Berlin, Heidelberg, doi: 10.1007/978-3-662-12616-5.
- Rackauckas, C. & Nie, Q., 2017. DifferentialEquations.jl—a performant and feature-rich ecosystem for solving differential equations in Julia, *Journal of Open Research Software*, **5**(1).
- Rauch, H. E., Tung, F., & Striebel, C. T., 1965. Maximum likelihood estimates of linear dynamic systems, *AIAA Journal*, **3**(8), 1445–1450, doi: 10.2514/3.3166.
- Ruina, A., 1983. Slip instability and state variable friction laws, *Journal of Geophysical Research: Solid Earth*, **88**(B12), 10359–10370, doi: 10.1029/JB088iB12p10359.
- Silverman, B. W., 1986. *Density Estimation for Statistics and Data Analysis*, Monographs on Statistics and Applied Probability, Chapman and Hall, London ; New York, Spine title: Density estimation Includes indexes.
- Theiler, J., 1986. Spurious dimension from correlation algorithms applied to limited time-series data, *Physical Review A*, **34**(3), 2427–2432, doi: 10.1103/PhysRevA.34.2427.
- van Dinther, Y., Künsch, H. R., & Fichtner, A., 2019. Ensemble data assimilation for earthquake sequences: Probabilistic estimation and forecasting of fault stresses, *Geophysical Journal International*, **217**(3), 1453–1478, doi: 10.1093/gji/ggz063.
- Verner, J. H., 2010. Numerically optimal Runge–Kutta pairs with interpolants, *Numerical Algorithms*, **53**(2-3), 383–396, doi: 10.1007/s11075-009-9290-3.
- Wang, S., 2024. Toward quantitative characterization of simulated earthquake-cycle complexities, *Scientific Reports*, **14**(1), 16811, doi: 10.1038/s41598-024-67685-4.

APPENDIX A: SUPPLEMENTARY FIGURES

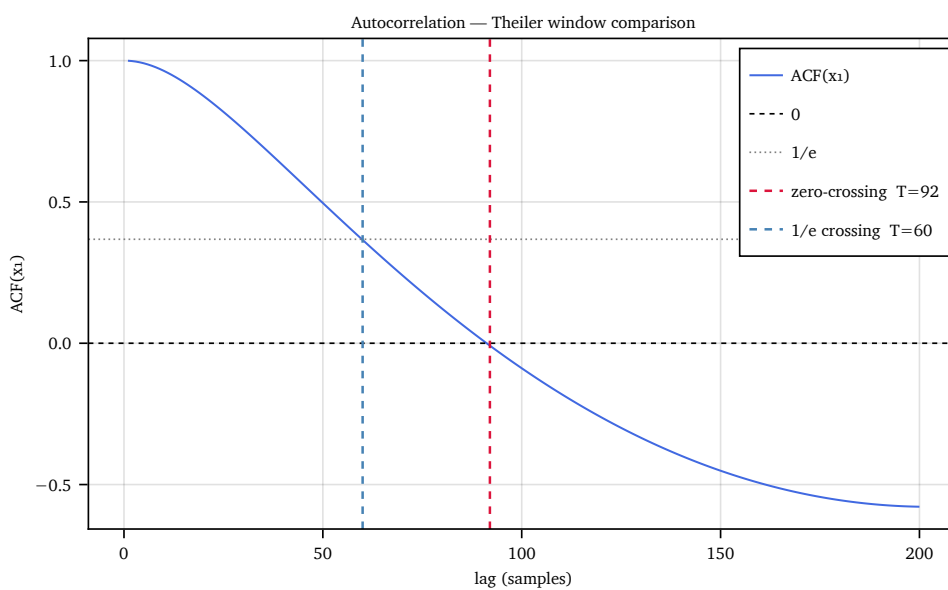


Figure A1. Autocorrelation function (ACF) of x_1 , used to determine the Theiler window. Vertical dashed lines mark the zero-crossing lag $w = T_{zero}$ (used in the main text) and the $1/e$ -crossing lag $w = T_{1/e}$ (more conservative). The zero-crossing window is the minimal defensible Theiler correction; the $1/e$ window provides the upper reference for the sensitivity sweep in Fig. A2.

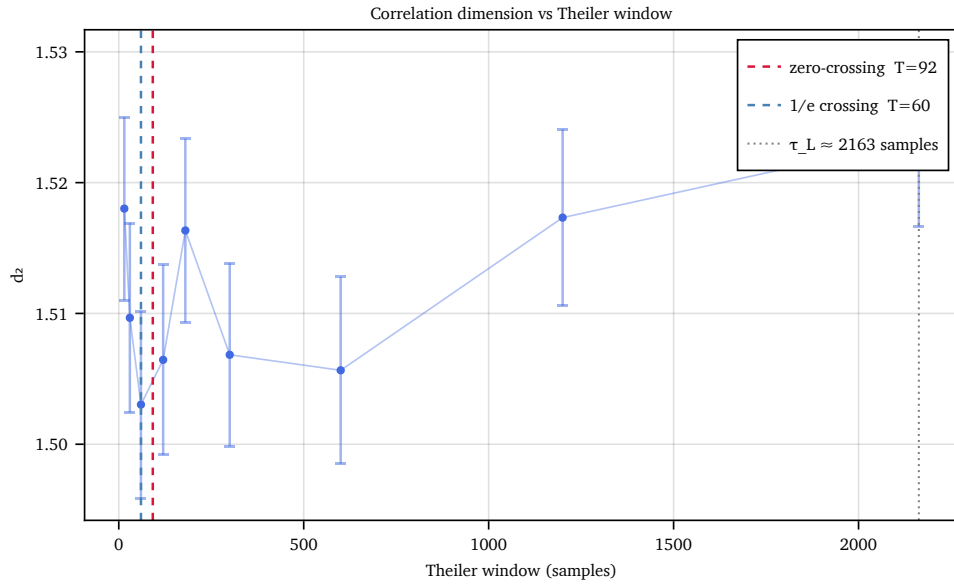
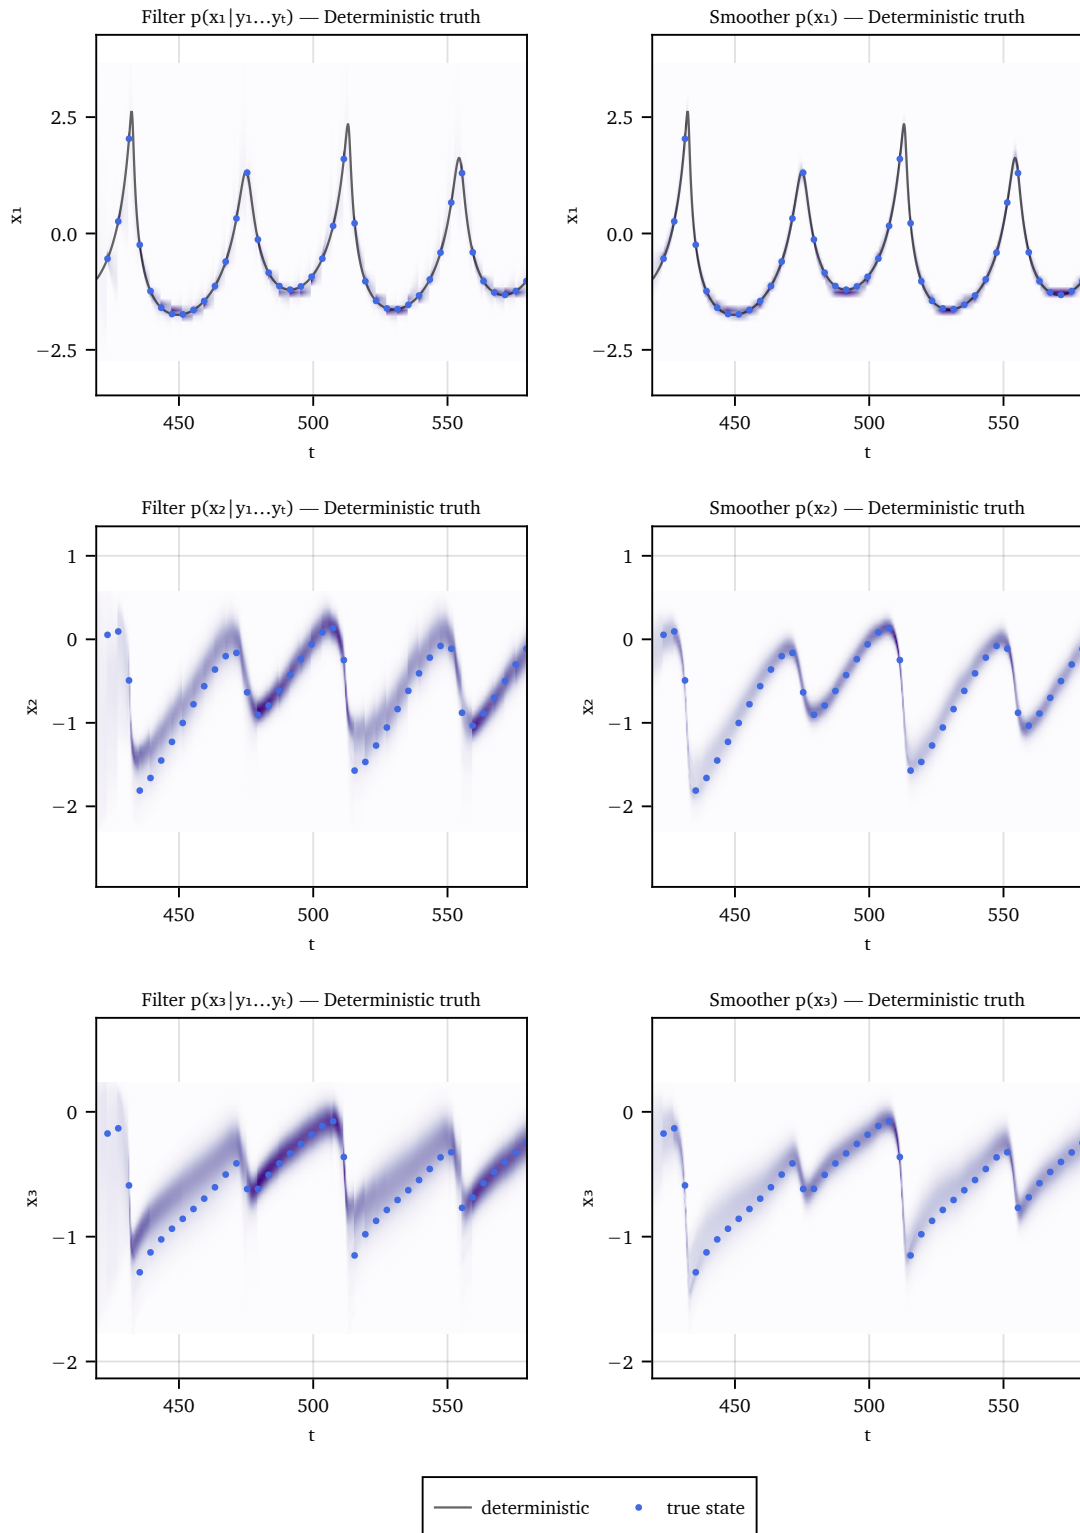
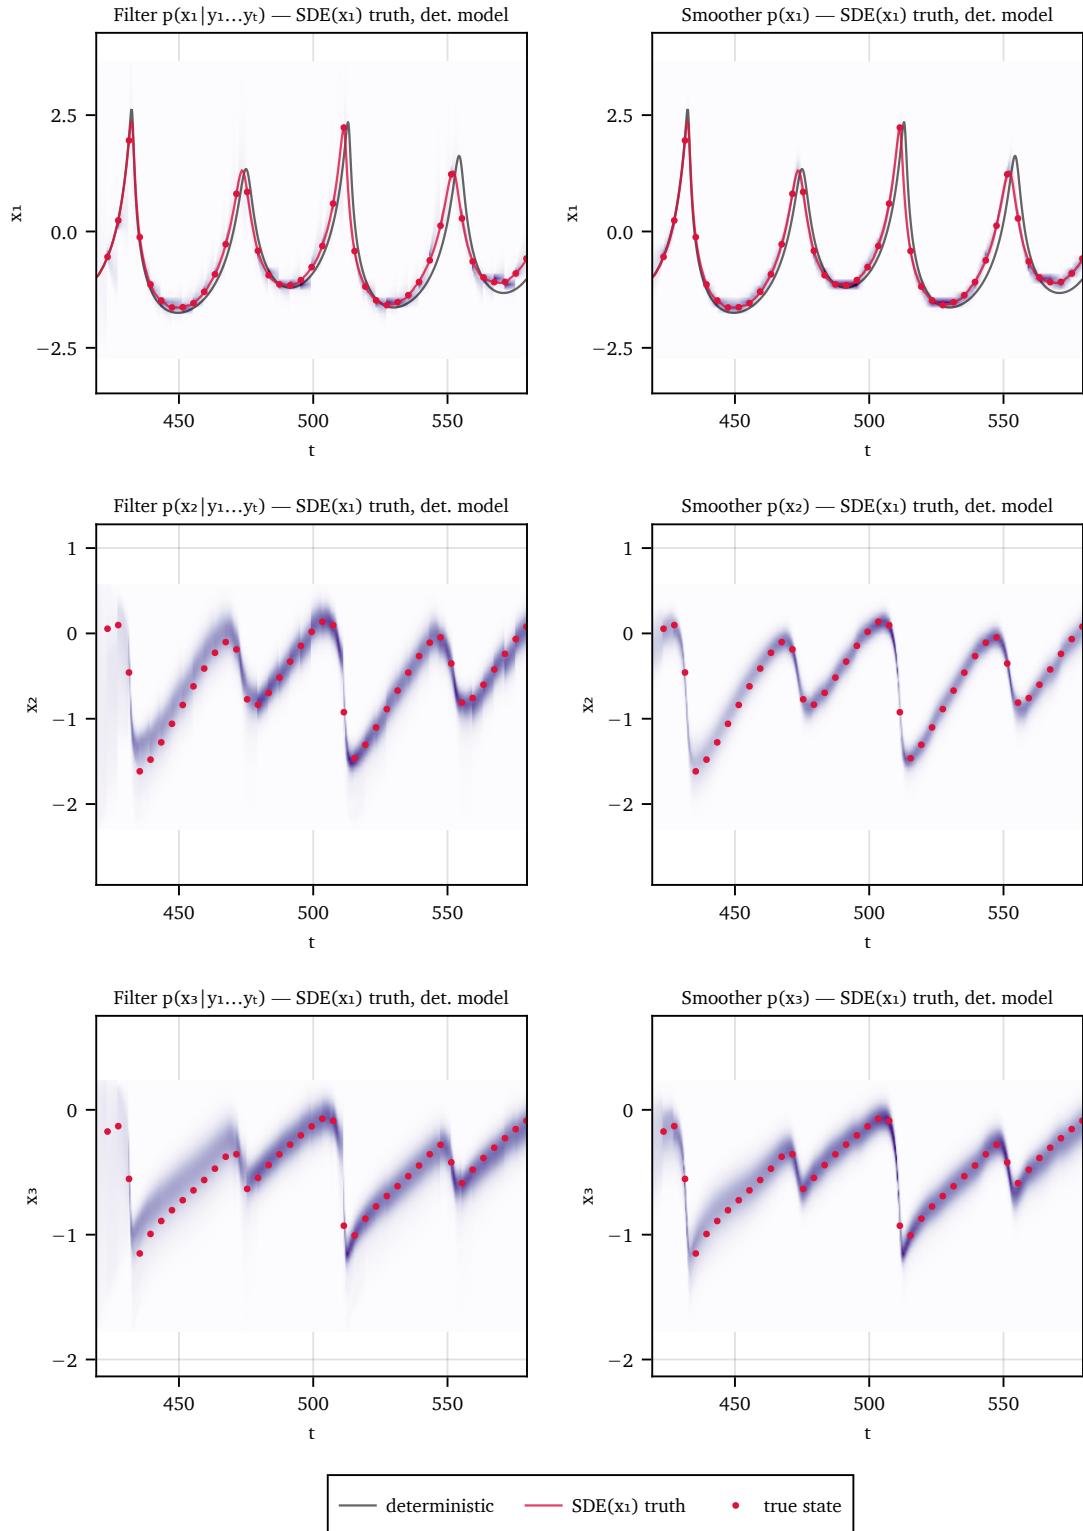


Figure A2. Sensitivity of the correlation dimension estimate d_2 to the Theiler window w , with $\pm 1\sigma$ OLS slope uncertainty. Vertical lines mark the zero-crossing window (used in the main text), the $1/e$ -crossing window, and the Lyapunov time τ_L . The plateau in d_2 across $[T_{\text{zero}}, \tau_L]$ confirms robustness to the precise window choice.



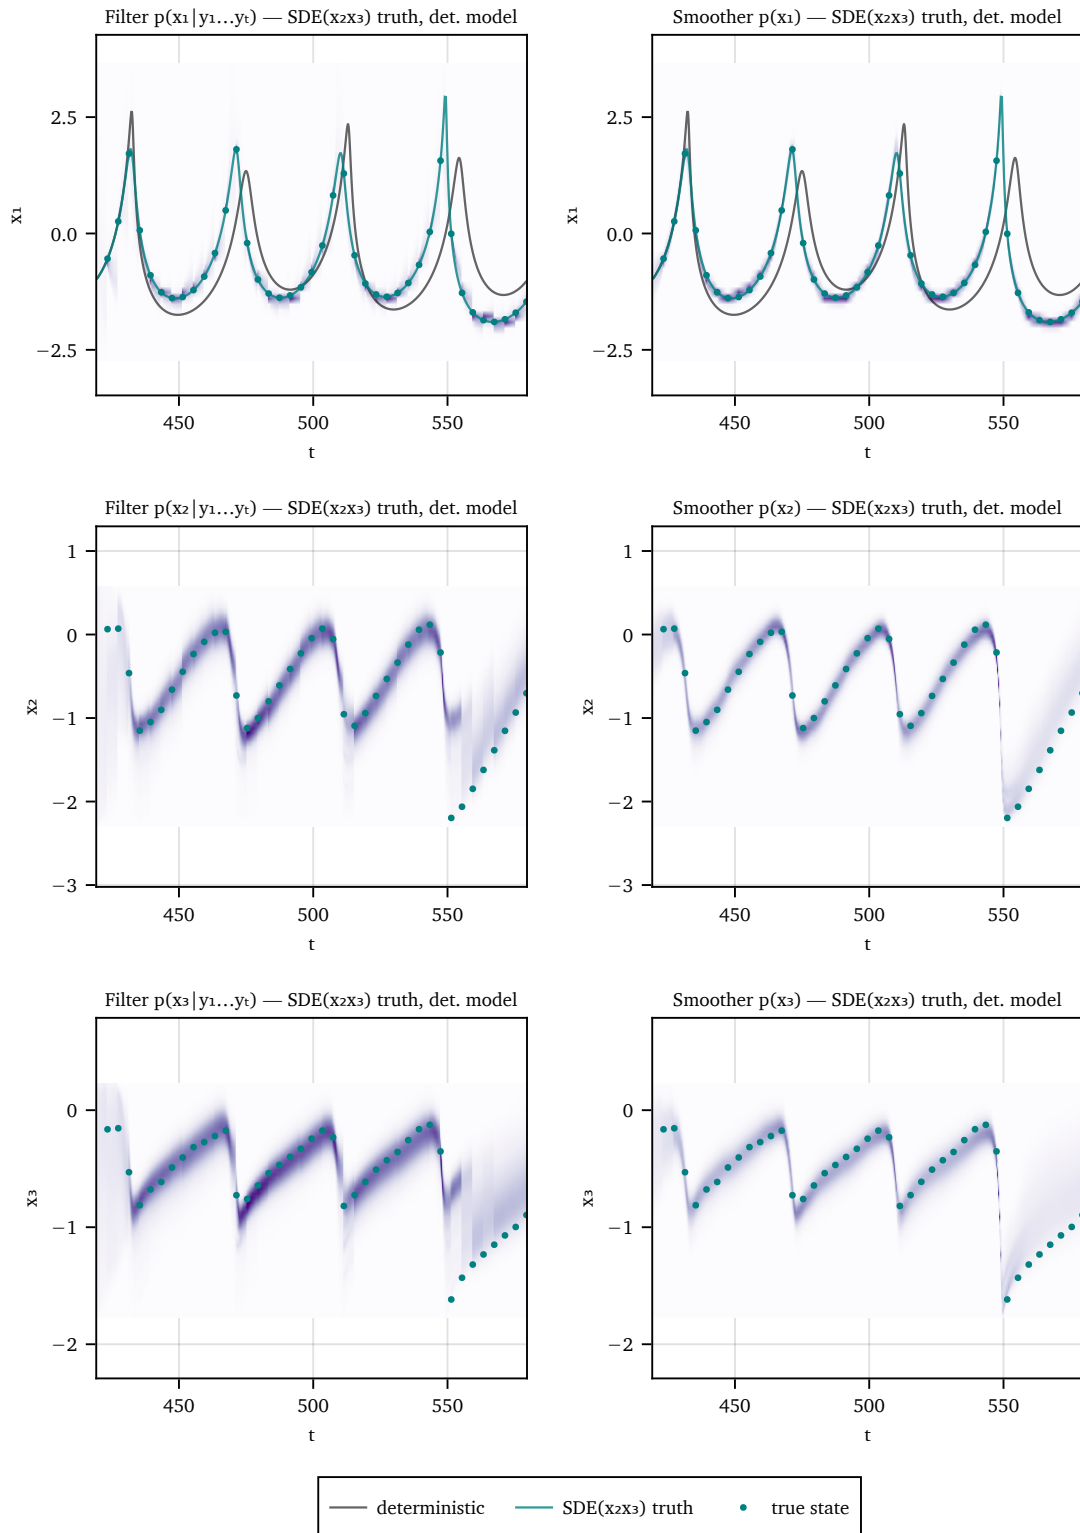
(a) Matched deterministic ODE truth (tracking skill 0.629). Rows: x_1 , x_2 , x_3 marginals; columns: filter posterior $p(x_i | y_1, \dots, y_t)$ (left) and smoother posterior (right). The posterior sharpens at each observation and spreads during free integration.

Figure A3. Full marginal posterior PDFs for all three state variables, corresponding to Fig. 3. Colour encodes probability density; observations (dots, $\sigma_{\text{obs}} = 0.05$) and true state (solid) are overlaid. Grid: 50^3 ; $N_{\text{obs}} = 40$; $\Delta t_{\text{obs}} = 4$ time units ($0.02 \tau_L$). Non-Gaussian and multi-modal structure during large slip events is visible in x_1 but not recoverable from the mean $\pm 1\sigma$ summary in the main text.



(b) $SDE(x_1)$ truth ($\sigma_{SDE} = 0.01$) with deterministic forward model (skill 0.646). Same layout as (a). Broader x_1 posterior reflects unmodelled stochastic increments; x_2, x_3 posteriors remain relatively tight as noise does not enter those channels.

Figure A3. Full marginal posterior PDFs for all three state variables, corresponding to Fig. 3. Colour encodes probability density; observations (dots, $\sigma_{obs} = 0.05$) and true state (solid) are overlaid. Grid: 50^3 ; $N_{obs} = 40$; $\Delta t_{obs} = 4$ time units ($0.02 \tau_L$). Non-Gaussian and multi-modal structure during large slip events is visible in x_1 but not recoverable from the mean $\pm 1\sigma$ summary in the main text.



(c) SDE($x_2 x_3$) truth ($\sigma_{\text{SDE}} = 0.01$ on x_2, x_3) with deterministic forward model (skill 0.622). Same layout as (a). Broader x_2, x_3 posteriors reflect noise entering through those channels; the filter PDF for x_1 occasionally loses contact with the true trajectory.

Figure A3 (cont.). Full marginal posterior PDFs for all three state variables, corresponding to Fig. 3. Colour encodes probability density; observations (dots, $\sigma_{\text{obs}} = 0.05$) and true state (solid) are overlaid. Grid: 50^3 ; $N_{\text{obs}} = 40$; $\Delta t_{\text{obs}} = 4$ time units ($0.02 \tau_L$). Non-Gaussian and multi-modal structure during large slip events is visible in x_1 but not recoverable from the mean $\pm 1\sigma$ summary in the main text.

Event-based Bayesian forecast: ODE vs SDE truth (PDF)

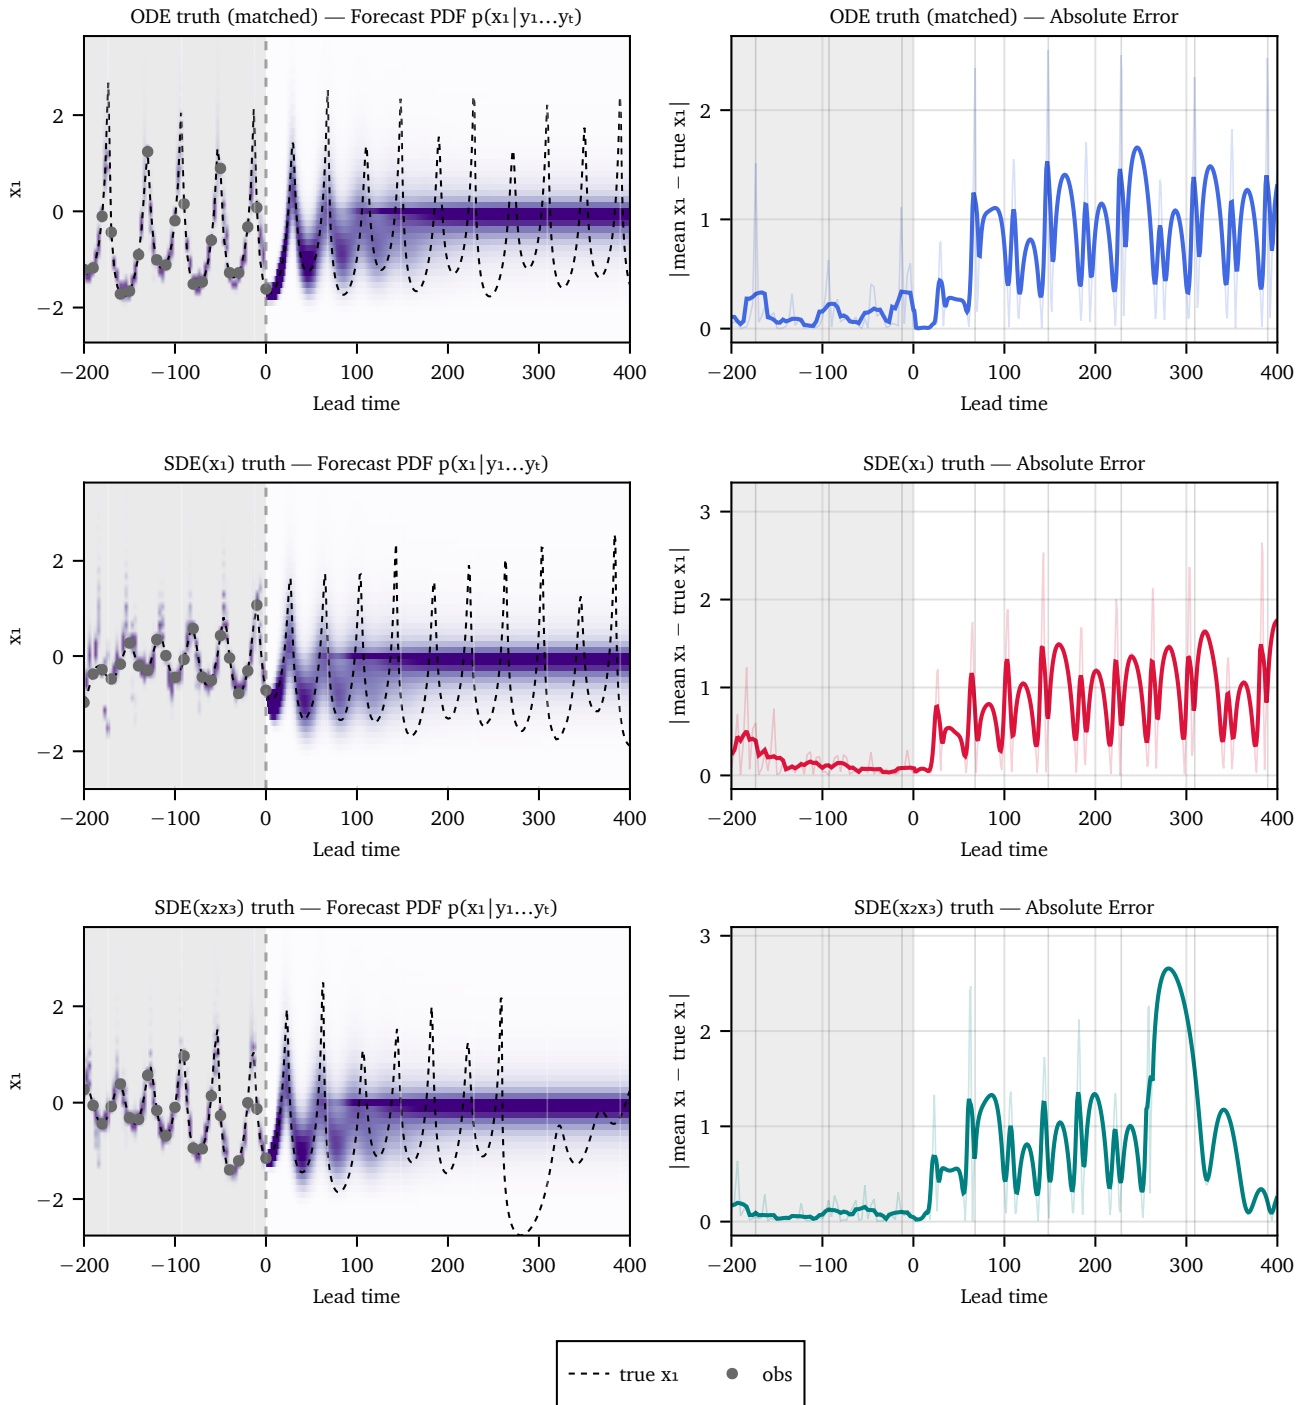


Figure A4. Full marginal forecast distributions corresponding to Fig. 4. Colour encodes probability density $p(x_1 | y_1, \dots, y_T)$ during assimilation (left of dashed line) and the forecast density thereafter, revealing non-Gaussian spreading and multi-modal structure that the mean \pm std summary in the main text cannot capture.

APPENDIX B: COMPUTE LOG

The following is the complete terminal output from the seven-stage Julia compute pipeline. It records the key numerical results at each stage and serves as a reproducibility record for the figures and statistics reported in the main text.

The same (within numerical precision) results should be reproducible by running the code contained in the Github repository. More details of how to run the code and how the codebase is structured are provided in the [README.md](#) file on Github.

```
=== Spring-Slider Compute Pipeline ===
```

```
Started: 2026-04-26 16:44:06
```

```
Threads: 32
```

[1/7] Deterministic Analysis

```
Parameters: =1.0, =0.84, =0.048, =0.068587, v=1.0
```

```
Post-transient trajectory: 496001 steps
```

```
Computing Lyapunov spectrum...
```

```
= [0.00462, -1.0e-5, -0.08088], sum = -0.076266
```

```
D_KY = 2.057, _L = 216.304
```

```
Computing correlation dimension...
```

```
Theiler window: zero-crossing=92 samples, 1/e crossing=60 samples
```

```
Theiler sweep: [15, 30, 60, 120, 180, 300, 600, 1200, 2163] samples (_L=2163 samples)
```

```
Running correlation_dimension_multi (n=496001, 11 Theiler values)...
```

```
correlation_dimension_multi: 5196.9s
```

```
d (full-data) = 1.5068 (Theiler=92), 1.5059 (no Theiler)
```

```
Theiler sweep d: [1.518, 1.5097, 1.503, 1.5065, 1.5163, 1.5068, 1.5057, 1.5173, 1.5233]
```

```
Fraction-based d convergence (n_full=496001)...
```

```
N=4960 (1.0%): d=1.5174 (0.8s)
```

```
N=49600 (10.0%): d=1.5061 (53.1s)
```

```
N=99200 (20.0%): d=1.5064 (215.3s)
```

```
N=198400 (40.0%): d=1.5068 (823.9s)
```

```
N=248000 (50.0%): d=1.5068 (1284.1s)
```

```
N=496001 (100.0%): d=1.5068 (from multi pass)
```

```
Saved deterministic.jld2
```

[2/7] SDE Ensemble Analysis

```
Running SDE(x) ensembles (t [400.0, 800.0], N=200)...
```

```
Running SDE(x x) ensembles (t [400.0, 800.0], N=200)...
```

```
Deterministic: 616 events, mean interval = 80.39
```

```
SDE(x) spread - =0.005:  $\bar{\tau}(x)$ =0.39, =0.01:  $\bar{\tau}(x)$ =0.747, =0.02:  $\bar{\tau}(x)$ =0.825
```

```
SDE(x x) spread - =0.005:  $\bar{\tau}(x)$ =0.853, =0.01:  $\bar{\tau}(x)$ =0.904, =0.02:  $\bar{\tau}(x)$ =0.94
```

```
Generating 3D divergence trajectories...
```

```
Phase-plane push (x)...
```

```
Phase-plane push (x x)...
```

```
Saved sde_ensembles.jld2
```

[3/7] Bayesian Data Assimilation

```
Observation window: t [419.4, 579.4]
```

```
Grid:  $50^3 = 125000$  points
```

```
Deterministic truth
```

```
Reanalysis (T_back=160)...
```

```
Computing RTS smoother (capped, max_ratio=100)...
```

```
Computing RTS smoother (uncapped, max_ratio=Inf)...
```

```
Saved bayes_det.jld2 (mean skill = 0.629)
```

```
SDE truth (noise on x)
```

```
Reanalysis (T_back=160)...
```

```
Computing RTS smoother (capped, max_ratio=100)...
```

```
Computing RTS smoother (uncapped, max_ratio=Inf)...
```

```
Saved bayes_sde.jld2 (mean skill = 0.646)
```

```
SDE truth (noise on x, x)
```

```
Reanalysis (T_back=160)...
```

```
Computing RTS smoother (capped, max_ratio=100)...
```

```
Computing RTS smoother (uncapped, max_ratio=Inf)...
```

```
Saved bayes_sde_x23.jld2 (mean skill = 0.622)
```

```
Mean skill - det: 0.629, SDE(x): 0.646, SDE(x x): 0.622
```

```
Smoothing z-score std (capped/uncapped):
```

det: capped=1.018 uncapped=1.019
 SDE(x): capped=1.435 uncapped=1.435
 SDE(x x): capped=1.09 uncapped=1.091

[4/7] Event Forecasting

Detecting stick-slip events...
 Assimilation events: 15, forecast events: 6
 Total observations: 115 ($\Delta t = 10.0$)
 _climate = 0.9776, _L = 216.304

Case: ODE truth
 Assimilation phase (115 obs)...
 Assimilation complete.
 Interpolating to forecast grid (50^3)...
 Free forecast phase...
 Forecast complete: 475 steps.
 Skill horizon: 83.0 t.u. (0.38 _L)
 Coverage horizon (68%): 122.0 t.u. (0.56 _L)

Case: SDE truth
 SDE solved: 16218 steps
 Assimilation phase (115 obs)...
 Assimilation complete.
 Interpolating to forecast grid (50^3)...
 Free forecast phase...
 Forecast complete: 475 steps.
 Skill horizon: 145.0 t.u. (0.67 _L)
 Coverage horizon (68%): 120.0 t.u. (0.55 _L)

Case: SDE_x23 truth
 SDE(x x) solved: 16218 steps
 Assimilation phase (115 obs)...
 Assimilation complete.
 Interpolating to forecast grid (50^3)...
 Free forecast phase...
 Forecast complete: 475 steps.
 Skill horizon: 69.0 t.u. (0.32 _L)
 Coverage horizon (68%): 80.0 t.u. (0.37 _L)
 Saved event_forecast.jld2

[5/7] Bifurcation Diagram

sweep: 0.2 \rightarrow 1.1, N = 1000
 Sweep complete. values with 1 maximum: 277
 Saved bifurcation.jld2

[6/7] Predictability Horizon

= 0.004623, _L = 216.3
 Grid: $50^3 = 125000$ points
 $H_{\infty} = -1.6822$ nats
 $H = -4.5343$, _init = [-0.9342, -0.0784, -0.2551]
 Push-forward: DT=5.0, T_MAX=865.2, N=173
 T = 35.0, _L = 216.3
 Measured dH/dt (early) = 0.018502 vs = 0.004623
 Saved predictability.jld2

[7/7] Interval Sensitivity

_L = 216.3
 Δt_{obs} values: [1, 2, 4, 8, 16, 32, 64]
 T_START_IDX = 1 (t = 400.0)
 Grid: $50^3 = 125000$ points

dt_obs = 1 (0.00462 _L)
 RMSE=0.0356 spread=0.0438 z-std=1.0478

dt_obs = 2 (0.00925 _L)
 RMSE=0.0658 spread=0.0449 z-std=1.7547

dt_obs = 4 (0.0185 _L)

20 *Thomas Holland*

RMSE=0.044 spread=0.0457 z-std=1.0394

dt_obs = 8 (0.037 _L)
RMSE=0.0369 spread=0.0441 z-std=0.8354

dt_obs = 16 (0.074 _L)
RMSE=0.0467 spread=0.0512 z-std=0.9872

dt_obs = 32 (0.148 _L)
RMSE=0.0586 spread=0.0489 z-std=1.3007

dt_obs = 64 (0.296 _L)
RMSE=0.0461 spread=0.0498 z-std=0.9757

Saved interval_sensitivity.jld2

Interval Sensitivity Summary

| dt_obs | dt/L | Grid | RMSE |
|--------|---------|--------|------|
| 1 | 0.00462 | 0.0356 | |
| 2 | 0.00925 | 0.0658 | |
| 4 | 0.0185 | 0.044 | |
| 8 | 0.037 | 0.0369 | |
| 16 | 0.074 | 0.0467 | |
| 32 | 0.148 | 0.0586 | |
| 64 | 0.296 | 0.0461 | |

=== All stages complete ===

Finished: 2026-04-26 19:04:32