

Selecting the Right Early-Warning Signal: Structural Failure of Critical-Slowing-Down Indicators at Scale-Mixture Transitions in Nonlinear Stochastic Dynamics

Tsung-Long Wu

National Cheng Kung University, Tainan, Taiwan

Abstract

The selection of the monitoring signal is the decisive design choice in any early-warning system, yet it has received little theoretical attention compared with the indicators developed for fold bifurcations by Scheffer and collaborators. We study this choice for *scale-mixture transitions* in nonlinear stochastic dynamics, in which the driving noise is conditionally Gaussian with a slowly ramping amplitude $\sigma(t)$, so that the increment law is locally Gaussian within any observation window but *aggregately heavy-tailed* across the ramp as a single amplitude parameter crosses a critical value σ_c . We compare five candidate indicators on a three-dimensional Itô reserve diffusion: (C2) an amplitude estimate from the quadratic variation of observation increments; (B2) sliding-window kurtosis; (A3) coupling-consistency residuals; and the two classical Scheffer indicators (Variance, AR(1)). Our central result is analytical, not merely empirical: we prove that the two critical-slowing-down (CSD) indicators *cannot* carry an early signal in this regime. Sliding-window kurtosis has a population value that is independent of the control parameter (Proposition 1); lag-1 autocorrelation approaches a bounded limit fixed by the relaxation rate rather than diverging, and its only dependence on the control parameter is an *inverted* one mediated by observation noise (Proposition 2). A Monte Carlo study over nine $(T_{\text{ramp}}, \sigma_{\text{obs}})$ configurations confirms both propositions: B2 never crosses threshold and AR(1) fires only at the highest observation noise, and then late. The C2 indicator, by contrast, is an unbiased amplitude estimator and delivers the earliest warning in the low-observation-noise regime, with a closed-form lead-time scaling law $L = \kappa(\sigma_{\text{obs}}) T_{\text{ramp}} - W dt/2$ that matches the simulations to within a few percent. The dominance is regime-dependent: sliding-window variance, although a biased and lagging amplitude proxy that C2 strictly improves upon at low observation noise, remains competitive and overtakes C2 once σ_{obs} becomes comparable to the dynamic range of the ramp signal. The contribution is a sharp, proof-backed delimitation of when classical EWS indicators fail by design and of the regime in which a noise-amplitude estimator is the correct alternative.

Keywords: early-warning signals; scale-mixture transition; quadratic variation; critical slowing down; nonlinear stochastic dynamics; observation noise; lead-time scaling.

1 Introduction

1.1 The signal-selection problem

Early-warning systems for dynamical regime shifts have become a central topic in nonlinear stochastic dynamics, following the programme of Scheffer and collaborators [15, 16]. The dominant paradigm rests on a small set of indicators—rising variance, rising lag-1 autocorrelation, increasing skewness—which derive from *critical slowing down* at a fold bifurcation, where the local restoring force vanishes and the relaxation timescale diverges [17, 18, 5]. These indicators have been validated on a range of empirical and synthetic systems.

There is, however, an unexamined premise: that the relevant transition is a continuous bifurcation of the deterministic skeleton. In many systems the transition is instead *structural*, arising not from the loss of stability of an attractor but from a change in the statistical character of the driving noise. In such transitions the attractor keeps its restoring property; the noise field shifts from a light-tailed to a heavy-tailed regime as an amplitude parameter is ramped. The critical slowing down on which classical indicators rely is then absent. This paper asks: which monitoring signal should an early-warning system use when the impending transition is of this kind?

1.2 Locally Gaussian, aggregately heavy-tailed

We make the question precise through a model class we call *scale-mixture transitions*. The driving noise is conditionally Gaussian given a slowly ramping amplitude $\sigma(t)$. Two facts coexist and are central to everything below:

1. *Aggregately heavy-tailed*. Pooled over the whole ramp—i.e. marginalised over the range of σ traversed—the increment law is a continuous scale mixture of normals and is therefore leptokurtic. The transition is a genuine light-to-heavy-tail change in the unconditional distribution.
2. *Locally Gaussian*. Within any single detection window, over which σ is quasi-constant, the increments are conditionally Gaussian. The heavy tail is an *aggregation* phenomenon and is invisible to a local, online statistic that sees only the current window and cannot anticipate the future regime.

This distinction is what keeps the negative result substantive rather than circular. A within-window statistic such as kurtosis is asked to detect a tail that, by construction, does not live inside any window. We confirm directly (Section 6) that even when genuine heavy-tailed *jumps* are injected *within* the window, the kurtosis indicator still fails to give early warning, because the tail co-emerges with the crossing rather than preceding it.

1.3 Contributions

- (i) We formulate signal selection for scale-mixture transitions as a calibrated benchmarking task on a canonical three-dimensional SDE (Sections 2–4).
- (ii) We *prove* that the two CSD indicators cannot carry an early signal in this regime: kurtosis has a control-parameter-independent population value (Proposition 1), and lag-1 autocorrelation has a bounded, inverted dependence on the control parameter (Proposition 2). We also show analytically that the quadratic-variation indicator C2 *does* carry a signal (Proposition 3).

- (iii) We confirm both propositions by Monte Carlo over nine $(T_{\text{ramp}}, \sigma_{\text{obs}})$ configurations, and we establish a closed-form lead-time scaling law for C2 that matches the data to within a few percent (Sections 5–6).
- (iv) We bound the result precisely. C2 gives the earliest warning only in the low-observation-noise regime; sliding-window variance, a biased lagging proxy that C2 improves upon there, remains competitive and overtakes C2 at high observation noise (Section 5.3).

1.4 Position in the literature

The EWS literature initiated by [15] and extended by Dakos, van Nes and others [6, 3, 11, 10] produces a taxonomy of indicators built almost entirely from autocorrelation, variance and higher moments, under the implicit assumption of a fold-bifurcation structure [4, 8]. Our model class is in the tradition of noise-induced transitions [9] and heavy-tail phenomenology [14]; the scale-mixture-of-normals construction is classical in statistics [2]. The quadratic-variation amplitude estimator is the realised-variance estimator of financial econometrics [1] and of SDE inference [7]. What is new is the application of such an estimator to the EWS problem together with a *proof* that the classical indicators fail by design for this transition class.

2 Model

2.1 State dynamics

On a complete filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ the state is a three-dimensional reserve vector $x(t) = (x_R, x_C, x_S)$ evolving as

$$dx_i = [-a_i x_i + b_i \tanh(x_i) - c_i x_j x_k] dt + \sigma(t) s_i dW_t^i, \quad (1)$$

with (i, j, k) any cyclic permutation of (R, C, S) . The drift contains a linear dissipation $-a_i x_i$, a bounded saturating term $b_i \tanh(x_i)$, and a cyclic coupling $-c_i x_j x_k$. The deterministic skeleton has a stable fixed point near the origin with restoring rate bounded below by $a_{\min} > 0$ uniformly along the ramp: crucially, there is *no* fold bifurcation, so no critical slowing down. We use $a_i = 1$, $b_i = 0.5$, $c_i = 0.1$, $s_i = 1$ (so $\sum_i s_i^2 = 3\sigma_{\text{diff}}^2$ with $\sigma_{\text{diff}} = 1$).

2.2 Scale-mixture driving noise

The amplitude $\sigma(t)$ is the control parameter whose crossing of a critical value σ_c the early-warning system must anticipate. Conditional on the amplitude path, the driving increments are Gaussian. The unconditional law obtained by pooling over the amplitudes traversed during the ramp,

$$p(dx) = \int p_{\mathcal{N}}(dx | \sigma) \mu(d\sigma), \quad (2)$$

is a scale mixture of normals and is leptokurtic whenever μ is non-degenerate. The “light-to-heavy-tail” transition is the growth of the upper portion of μ as the ramp carries σ through σ_c ; the kurtosis of (2) is an increasing functional of the spread of μ . The regime change is in the *statistical character* of the noise, not in the stability of the attractor.

2.3 Observation model

The monitor does not see $x(t)$. At discrete times $t_k = k \, dt$ with $dt = 0.04$ it receives

$$y(t_k) = x(t_k) + \xi_{\text{obs}}(t_k), \quad \xi_{\text{obs}} \sim \mathcal{N}(0, \sigma_{\text{obs}}^2 I_3), \quad (3)$$

with $\sigma_{\text{obs}} > 0$ known and ξ_{obs} independent across times and components. All indicators are computed from y alone.

2.4 Ramp protocol

To benchmark the indicators we drive the amplitude through a slow linear ramp,

$$\sigma(t) = \sigma_{\text{low}} + (\sigma_{\text{high}} - \sigma_{\text{low}}) t / T_{\text{ramp}}, \quad t \in [0, T_{\text{ramp}}], \quad (4)$$

with $\sigma_{\text{low}} = 0.20$, $\sigma_{\text{high}} = 1.00$, $\sigma_c = 0.60$ and $T_{\text{ramp}} \in \{20, 40, 80\}$. The critical crossing time is $t_c = T_{\text{ramp}}(\sigma_c - \sigma_{\text{low}}) / (\sigma_{\text{high}} - \sigma_{\text{low}}) = T_{\text{ramp}}/2$.

3 Early-warning indicators

All five indicators are computed on sliding windows of $W = 50$ steps ($W \, dt = 2$) of the observation trajectory.

C2: quadratic-variation amplitude estimate. With increments $dy(t_k) = y(t_k) - y(t_{k-1})$, the realised quadratic variation is $QV(t_k) = \frac{1}{W} \sum_{j=k-W+1}^k \|dy(t_j)\|^2$. Under (1)–(3), to leading order in dt ,

$$\mathbb{E}[QV] = 3 \sigma^2 \sigma_{\text{diff}}^2 dt + 6 \sigma_{\text{obs}}^2, \quad (5)$$

the second term arising because differencing white observation noise doubles its variance. Solving for σ gives the C2 indicator

$$S_{C2}(t_k) := \sqrt{\frac{\max(QV(t_k) - 6\sigma_{\text{obs}}^2, 0)}{3\sigma_{\text{diff}}^2 dt}}, \quad (6)$$

an unbiased estimate of the driving amplitude σ at the window centre.

B2: residual kurtosis. With window mean $\bar{y}(t_k)$ and residual $r(t_j) = y(t_j) - \bar{y}(t_k)$, pooled over the three components,

$$S_{B2}(t_k) := \frac{\frac{1}{3W} \sum_{j,i} r_i(t_j)^4}{\left(\frac{1}{3W} \sum_{j,i} r_i(t_j)^2\right)^2}. \quad (7)$$

For Gaussian residuals $S_{B2} = 3$; sustained departure is the classical signal.

A3: coupling-consistency residual. $S_{A3}(t_k) := \frac{1}{3} \sum_{(i,j)} \text{Var}(y_i y_j : l \in [k - W + 1, k])$, the sliding-window variance of the three observed cross-products.

Classical Variance and AR(1). S_{Var} is the pooled sliding-window sample variance of y ; S_{AR1} is the pooled sliding-window lag-1 autocorrelation of centred y .




Figure 1: Null distributions of the five indicators at constant $\sigma = 0.30$, $\sigma_{\text{obs}} = 0.05$. Alarm thresholds (dashed) are the 95th percentiles. C2 is centred at the true amplitude 0.30; B2 is centred near the Gaussian value 3 regardless of amplitude.

4 Calibration and lead-time metric

Null calibration. Each threshold is calibrated against a null at constant sub-critical $\sigma_{\text{null}} = 0.30$. The 95th percentile of each indicator over the steady-state half of the null run is the alarm threshold, fixing a uniform 5% false-alarm rate. Figure 1 shows the null distributions and thresholds at $\sigma_{\text{obs}} = 0.05$. The C2 null is centred near 0.30, confirming that C2 is an unbiased amplitude estimator; the B2 null sits near 3 (the finite-sample Gaussian value, slightly below 3 for windowed dependent samples).

Lead-time metric. For each $(T_{\text{ramp}}, \sigma_{\text{obs}})$ we run $N = 400$ independent paths, form the ensemble-mean indicator trajectory $\mathbb{E}[S(t)]$, and define the ensemble-mean lead time

$$L := t_c - \inf\{t : \mathbb{E}[S(t)] \geq \text{threshold}\}. \quad (8)$$

Positive L means the alarm fired before the critical crossing; “NEVER” means the ensemble-mean indicator did not cross.

5 Structural failure of critical-slowness indicators

We first establish the two negative results analytically; the numerics in Section 5.1–5.3 then confirm them. Throughout we use four standing assumptions: (H1) the driving noise is conditionally Gaussian given the amplitude path; (H2) the window is quasi-static, $|\Delta\sigma| = (\sigma_{\text{high}} - \sigma_{\text{low}})Wdt/T_{\text{ramp}} \ll 1$ (here $Wdt < T_{\text{ramp}}/10$); (H3) the restoring rate satisfies $a_i \geq a_{\text{min}} > 0$ uniformly (no fold bifurcation); (H4) observation noise is white and Gaussian, (3).

Proposition 1 (Kurtosis carries no control-parameter signal). *Under (H1),(H2),(H4), the population value of the windowed residual kurtosis is independent of the control parameter σ when the component diffusion scales are equal, and otherwise depends on σ only through a term of order $\mathcal{O}(\sigma_{\text{obs}}^2/\sigma^2)$ that decreases as σ grows. Consequently the excess-kurtosis signal is identically zero (or vanishing) along the ramp, and the ensemble-mean indicator cannot systematically exceed any threshold calibrated on a Gaussian null.*

Proof. Condition on the amplitude σ_k over the window; by (H2) it is constant to $\mathcal{O}(\Delta\sigma)$. Each $y_i(t_j)$ is the sum of a Gaussian state component and Gaussian observation noise (H1, H4), hence Gaussian; subtracting the window mean \bar{y}_i , a Gaussian linear combination, leaves residuals $r_i(t_j)$ that are marginally Gaussian. For any univariate Gaussian, $\mathbb{E}[r^4] = 3(\mathbb{E}[r^2])^2$, independent of its

variance and of the serial correlation (kurtosis is a marginal property). Writing v_i for the residual variance of component i , ergodicity of the stationary Gaussian window gives, as $W \rightarrow \infty$,

$$S_{B2} \rightarrow \frac{9 \sum_i v_i^2}{(\sum_i v_i)^2},$$

which equals 3 when $v_1 = v_2 = v_3$ and is ≥ 3 otherwise by Cauchy–Schwarz. This ratio is homogeneous of degree zero in (v_1, v_2, v_3) . By (H1),(H4), $v_i = \sigma^2 w_i + \sigma_{\text{obs}}^2$ with w_i independent of σ . At $\sigma_{\text{obs}} = 0$ the factor σ^2 cancels identically, so the limit is *strictly independent* of σ ; for $\sigma_{\text{obs}} > 0$ the residual dependence is $\mathcal{O}(\sigma_{\text{obs}}^2/\sigma^2)$ and decreasing in σ . Because the null ($\sigma = 0.30$) is itself Gaussian, the threshold is the finite-sample 95th percentile of this same constant; the ensemble-mean signal is a horizontal line (or one with the wrong, decreasing slope) and never crosses. \square

Proposition 2 (AR(1) has no critical-slowness signal; its control-parameter dependence is inverted). *Under (H3),(H4), the windowed lag-1 autocorrelation of centred observations does not diverge toward 1 as $\sigma \rightarrow \sigma_c$; it approaches a bounded limit $\rho_x = e^{-a dt} < 1$ fixed by the relaxation rate. Its only dependence on σ is through observation-noise dilution, and that dependence strengthens as σ_{obs} increases—the opposite of a genuine early-warning signal.*

Proof. Linearising (1) about the fixed point, the latent state is an Ornstein–Uhlenbeck process whose discrete lag-1 autocorrelation is $\rho_x = e^{-a dt}$, independent of σ (scaling the noise rescales the amplitude of x but not its correlation time). By (H3), $1/a$ is bounded, so ρ_x is constant along the ramp and does not diverge. Adding white observation noise (H4),

$$\rho_y = \rho_x \frac{\text{Var}(x)}{\text{Var}(x) + \sigma_{\text{obs}}^2} = \rho_x \frac{1}{1 + \sigma_{\text{obs}}^2/\text{Var}(x)}.$$

With $\text{Var}(x) \propto \sigma^2$, increasing σ raises ρ_y monotonically toward the fixed ceiling ρ_x (not toward 1). This rise is pure signal-to-noise dilution, not slowing down: the relaxation time $1/a$ never grows. Its magnitude is $\mathcal{O}(\sigma_{\text{obs}}^2/\sigma^2)$ and therefore *larger* for larger σ_{obs} . A true EWS should become more sensitive as observation noise decreases; AR(1) does the reverse, which is the operational signature of a spurious signal. \square

Proposition 3 (C2 carries a signal). *Under (H1),(H2),(H4), $\mathbb{E}[S_{C2}] = \sigma$ to $\mathcal{O}(\Delta\sigma)$, so the indicator crosses any threshold $\sigma_{\text{th}} \in (\sigma_{\text{low}}, \sigma_c)$ strictly before t_c , with*

$$L = T_{\text{ramp}} \frac{\sigma_c - \sigma_{\text{th}}}{\sigma_{\text{high}} - \sigma_{\text{low}}} - \frac{W dt}{2}, \quad (9)$$

valid in the regime $\sigma_{\text{obs}}^2 \ll \frac{1}{2} \sigma_c^2 \sigma_{\text{diff}}^2 dt$.

Proof. By (5), $\mathbb{E}[QV] - 6\sigma_{\text{obs}}^2 = 3\sigma^2 \sigma_{\text{diff}}^2 dt$ exactly, so (6) is unbiased for σ up to the $\mathcal{O}(\Delta\sigma) = \mathcal{O}((d\sigma/dt)Wdt)$ centring bias. The estimator output tracks $\sigma(t_k - Wdt/2)$; setting this equal to σ_{th} and using the linear ramp (4) gives the alarm time $t_{\text{alarm}} = Wdt/2 + T_{\text{ramp}}(\sigma_{\text{th}} - \sigma_{\text{low}})/(\sigma_{\text{high}} - \sigma_{\text{low}})$, and (9) follows from $L = t_c - t_{\text{alarm}}$. The bias term $6\sigma_{\text{obs}}^2$ must be subdominant to the signal term $3\sigma_c^2 \sigma_{\text{diff}}^2 dt$ at the crossing, giving the stated regime. \square




Figure 2: Ensemble-mean indicator trajectories at $(T_{\text{ramp}} = 40, \sigma_{\text{obs}} = 0.05)$, normalised by their thresholds. C2 crosses first; A3 and Variance follow; B2 and AR(1) never reach their thresholds. The dotted vertical line is t_c .

T_{ramp}	σ_{obs}	t_c	C2	A3	Variance	B2	AR(1)
20	0.02	10	+5.64	+3.76	+3.60	NEVER	NEVER
20	0.05	10	+4.72	+3.48	+3.44	NEVER	NEVER
20	0.10	10	+1.80	+3.72	+3.52	NEVER	-0.32
40	0.02	20	+12.24	+8.80	+8.72	NEVER	NEVER
40	0.05	20	+10.76	+8.04	+7.84	NEVER	NEVER
40	0.10	20	+4.52	+8.68	+8.28	NEVER	+1.60
80	0.02	40	+25.60	+19.40	+18.08	NEVER	NEVER
80	0.05	40	+22.32	+19.20	+17.68	NEVER	NEVER
80	0.10	40	+10.96	+18.64	+17.44	NEVER	+4.80

Table 1: Ensemble-mean lead times (time units). Positive entries: alarm before the critical crossing. “NEVER”: the ensemble-mean indicator did not cross. B2 never fires; AR(1) fires only at the highest observation noise, confirming Propositions 1–2.

6 Numerical results

6.1 Five-indicator comparison

Figure 2 shows ensemble-mean trajectories at the central configuration $(T_{\text{ramp}} = 40, \sigma_{\text{obs}} = 0.05)$, normalised by each indicator’s threshold. C2 crosses first (near $t \approx 9$, a lead of ≈ 11); A3 and Variance follow a few units later; B2 sits at ≈ 2.85 throughout and AR(1) plateaus just below its threshold—neither crosses, exactly as Propositions 1 and 2 require.

6.2 Lead-time grid and scaling law

Table 1 collects ensemble-mean lead times for all nine configurations. B2 never fires; AR(1) fires only at $\sigma_{\text{obs}} = 0.10$, and then with a small or negative lead, in agreement with Proposition 2. The working indicators C2, A3 and Variance all scale linearly in T_{ramp} (Figure 3A). A least-squares fit of the C2 lead time gives $\kappa(0.02) = 0.32$, $\kappa(0.05) = 0.28$, $\kappa(0.10) = 0.13$. These match the closed-form prediction (9): reading σ_{th} from the calibrated C2 null (0.337, 0.371, 0.477 at $\sigma_{\text{obs}} = 0.02, 0.05, 0.10$) gives $\kappa = (\sigma_c - \sigma_{\text{th}})/0.80 = 0.329, 0.286, 0.154$, and the predicted lead times $L = \kappa T_{\text{ramp}} - 1.0$ reproduce the simulated values to within a few percent (e.g. at $\sigma_{\text{obs}} = 0.05$: predicted 4.7/10.4/21.9 vs. measured 4.7/10.8/22.3).




Figure 3: (A) Lead time vs. ramp duration at $\sigma_{obs} = 0.05$; the dashed line is the C2 fit $L = 0.28 T_{ramp}$. (B) Regime of validity at $T_{ramp} = 80$: the C2 lead falls with observation noise and is overtaken by Variance and A3 near $\sigma_{obs} \approx 0.065$.

6.3 Regime of validity: the C2–Variance crossover

The advantage of C2 is regime-dependent. Figure 3B fixes $T_{ramp} = 80$ and varies σ_{obs} . The C2 lead time falls steeply with observation noise—because the bias term $6\sigma_{obs}^2$ in (5) inflates the calibrated threshold σ_{th} and shrinks κ via (9)—and crosses below Variance and A3 near $\sigma_{obs} \approx 0.065$. At $\sigma_{obs} = 0.10$, Variance (+17.4) and A3 (+18.6) *outperform* C2 (+11.0). Variance is not a CSD indicator in this regime: it carries a signal simply because the windowed variance scales with σ^2 . C2 is best understood as a debiased, trend-robust amplitude estimator that strictly improves on raw variance *when observation noise is small*, and that degrades to parity (and below) once σ_{obs} is comparable to the dynamic range. The result is therefore stated narrowly: the *critical-slowness* indicators (kurtosis, AR(1)) fail structurally; among the indicators that do work, C2 gives the earliest warning in the low-noise regime, with variance as a robust fallback.

6.4 Genuine within-window heavy tails do not rescue kurtosis

A natural objection is that B2 fails only because the simulated increments are locally Gaussian. We test this directly by injecting genuine heavy-tailed jumps *inside* the window: with probability $w(\sigma) = [1 + e^{-\gamma(\sigma - \sigma_c)}]^{-1}$ each innovation is drawn from a unit-variance Student- $t_{3.5}$ (per-draw sample kurtosis ≈ 34 , population fourth moment infinite) rather than a Gaussian. Figure 4A shows the result: the kurtosis trajectory now rises, but only to ≈ 3.0 and only *after* t_c , never approaching the 3.87 threshold. The reason is structural: $w(\sigma)$ is tied to the same crossing the system is trying to anticipate, so the tail co-emerges with the transition rather than preceding it; and a fourth-moment statistic of infinite-population-kurtosis data has no stable target to converge to within a window. The lead-time table for the jump experiment is unchanged for B2 (NEVER) while C2 still leads (+10.56). Heavy tails that are real but not *anticipatory* do not make kurtosis an early-warning indicator.

7 Discussion

7.1 Why C2 succeeds and CSD indicators fail

The contrast is structural. C2 estimates the latent control parameter σ directly; its alarm is a level crossing of an amplitude estimate. The CSD indicators are *response* statistics: they measure how




Figure 4: Evidence for the two propositions. (A) Proposition 1: B2 sits at ≈ 2.85 under conditionally Gaussian driving (solid) and, even with within-window heavy jumps (dashed), rises only to ≈ 3.0 after t_c , never reaching threshold. (B) Proposition 2: AR(1), normalised by its own per- σ_{obs} threshold, crosses only in the highest-noise case and then near t_c ; the low-noise cases plateau below threshold—the signal is inverted.

the state responds to the noise. For fold-bifurcation systems the response is amplified at criticality (critical slowing down), which is exactly what makes those statistics effective. For scale-mixture transitions the response does not diverge—the restoring rate stays bounded, so the autocorrelation does not grow (Proposition 2), and the local residuals stay Gaussian, so the kurtosis does not rise (Proposition 1). The choice between “monitor the state” (CSD) and “monitor the noise” (C2) is the decisive signal-selection decision, and it is dictated by the nature of the impending transition.

7.2 Scope of the result

The result is deliberately bounded. The structural-failure statement applies to the critical-slowing-down indicators—kurtosis and lag-1 autocorrelation—for which Propositions 1 and 2 supply the mechanism. Sliding-window variance lies outside that statement: in this model it is not a pure response statistic but tracks the amplitude through its σ^2 scaling, and so retains skill. The advantage of C2 is likewise scoped to the low-observation-noise regime $\sigma_{\text{obs}}^2 \ll \frac{1}{2} \sigma_c^2 \sigma_{\text{diff}}^2 dt$, within which it gives the earliest warning and strictly improves on variance; beyond it, C2 degrades to parity with, and below, variance (Section 5.3). Finally, the mechanism concerns transitions whose heavy tail is an aggregation phenomenon rather than an anticipatory within-window one; Section 6 shows that genuinely anticipatory local tails fall outside it.

7.3 Application domains and limitations

The natural domains are those in which the impending transition is plausibly a noise-amplitude (volatility) regime change rather than a bifurcation: engineering systems under mixed nominal/anomalous loads; volatility regime switches; geophysical forcing with structural intermittency. Three limitations remain. First, C2 requires a known σ_{obs} to subtract the bias $6\sigma_{\text{obs}}^2$; a joint estimator of $(\sigma, \sigma_{\text{obs}})$ is the natural next step. Second, (9) assumes an exactly linear ramp; non-monotonic or stochastic drivers call for a change-point layer (CUSUM [13, 12]) on the C2 output. Third, a single global amplitude is assumed; multi-amplitude systems require per-component estimates and a joint detection rule.

Declarations

Funding. The author received no funding for this work.

Conflict of interest. The author declares no competing interests.

Data and code availability. All experiments are reproducible from the parameters in Sections 2–4; the Monte Carlo implementation is available from the author on reasonable request.

References

- [1] Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P. (2003). Modeling and forecasting realized volatility. *Econometrica* 71, 579–625.
- [2] Andrews, D.F., Mallows, C.L. (1974). Scale mixtures of normal distributions. *J. R. Stat. Soc. B* 36, 99–102.
- [3] Boettiger, C., Hastings, A. (2012). Early warning signals and the prosecutor’s fallacy. *Proc. R. Soc. B* 279, 4734–4739.
- [4] Carpenter, S.R., Brock, W.A. (2006). Rising variance: a leading indicator of ecological transition. *Ecol. Lett.* 9, 311–318.
- [5] Dakos, V., Scheffer, M., van Nes, E.H., Brovkin, V., Petoukhov, V., Held, H. (2008). Slowing down as an early warning signal for abrupt climate change. *PNAS* 105, 14308–14312.
- [6] Dakos, V., Carpenter, S.R., Brock, W.A., et al. (2012). Methods for detecting early warnings of critical transitions in time series. *PLOS One* 7, e41010.
- [7] Genon-Catalot, V., Jacod, J. (1993). On the estimation of the diffusion coefficient for multidimensional diffusion processes. *Ann. Inst. H. Poincaré* 29, 119–151.
- [8] Held, H., Kleinen, T. (2004). Detection of climate system bifurcations by degenerate fingerprinting. *Geophys. Res. Lett.* 31, L23207.
- [9] Horsthemke, W., Lefever, R. (1984). *Noise-Induced Transitions*. Springer.
- [10] Kuehn, C. (2011). A mathematical framework for critical transitions. *Physica D* 240, 1020–1035.
- [11] Lenton, T.M. (2011). Early warning of climate tipping points. *Nat. Clim. Change* 1, 201–209.
- [12] Lorden, G. (1971). Procedures for reacting to a change in distribution. *Ann. Math. Stat.* 42, 1897–1908.
- [13] Page, E.S. (1954). Continuous inspection schemes. *Biometrika* 41, 100–115.
- [14] Resnick, S.I. (2007). *Heavy-Tail Phenomena*. Springer.
- [15] Scheffer, M., Bascompte, J., Brock, W.A., et al. (2009). Early-warning signals for critical transitions. *Nature* 461, 53–59.
- [16] Scheffer, M., Carpenter, S.R., Lenton, T.M., et al. (2012). Anticipating critical transitions. *Science* 338, 344–348.
- [17] van Nes, E.H., Scheffer, M. (2007). Slow recovery from perturbations as a generic indicator of a nearby catastrophic shift. *Am. Nat.* 169, 738–747.
- [18] Wissel, C. (1984). A universal law of the characteristic return time near thresholds. *Oecologia* 65, 101–107.