A Dimensionless Reciprocal-Time Indicator for Boundary-Approach Risk in Engineered Nonlinear Systems: Theory, Falsifiable Acceptance Band, and Validation on Power-Grid and Financial Time Series

Tsung-Long Wu

National Cheng Kung University, Tainan 70101, Taiwan

Correspondence: eaglekelvin@gmail.com

Abstract

Many engineered and physical systems—power grids, mechanical structures, controlled chemical reactors, laser amplifiers, financial markets, and physiological regulators—operate close to nonlinear instability thresholds whose crossing would, in idealised limits, produce finite-time singularities in a relevant state norm. Persistent systems remain observable precisely because evolved or engineered saturation keeps them below this catastrophic regime, but the dimensionless distance from the saturation boundary is itself a structural risk indicator that is rarely measured directly. This paper develops such a measurement from first principles.

Let f be a locally Lipschitz vector field on the autonomous reduction of an engineered nonlinear system, and assume that along a candidate approach direction f is asymptotically homogeneous of degree α along the blow-up direction with leading coefficient κ > 0. We prove that the dimensionless reciprocal-time invariant Ψ(r) := Φ(t)/[r · E(t)], built from the radial energy $E(t) = ½\|x(t)\|^2$ and the radial flux Φ(t) = ⟨x(t), f(x(t))⟩ with r = 1/(T* − t), converges to a vector-field-intrinsic limit

$$\Psi^*(\alpha) = 2 \;/\; (\alpha - 1) \qquad \text{for } \alpha > 1,$$

as r → ∞ (Theorem 1). The map α ↦ Ψ*(α) supplies a one-dimensional reference manifold against which any candidate engineered system can be compared, and the gap |α − 1| furnishes a dimensionless safety margin. Under a strengthened homogeneity hypothesis controlling the next non-resonant term of the normal form, the convergence rate is sharp: |Ψ(r) − Ψ*(α)| ≤ C r^(−β) with β = δ/(α − 1) (Theorem 2).

The paper develops the indicator into an operationally complete diagnostic. We construct a quantitative reciprocal-time Lyapunov functional V(r) with explicit decay rate $|V(r) - C_\alpha| \le D_1$ r^(−β) and singleton ω-limit set on the self-similar manifold (Theorem 3); a closed-form falsifiability inequality together with a stability theorem under finite-sample OLS estimation error in α, yielding an explicit composite acceptance band decomposed into intrinsic rate, estimator bias, and estimator variance (Theorems 4 and 5); and an additive-noise robustness theorem (Theorem 6) bounding the deviation of the empirical Ψ-estimate under observation noise and quantifying the minimum signal-to-noise ratio required for the falsifier to be admissible.

Numerical validation on two-dimensional perturbed homogeneous flows confirms convergence to $\Psi^*(\alpha)$ within 0.2%–4.5% across $\alpha \in \{1.5, 2.0, 3.0\}$, with empirical errors quantitatively identified as finite-r samples of the rate $r^{(-\beta)}$ and shown to lie inside the derived acceptance band. Two real-world engineered/economic datasets from disjoint domains are then processed through the same pipeline as boundary-approach measurements: the Continental European synchronous power-grid frequency log of January 2019 yields $\alpha_{est} = 0.36 \pm 0.04$, and the NASDAQ Composite dot-com bubble (1994–2000) yields $\alpha_{est} = 0.49 \pm 0.06$. Both lie deep in the saturated regime $\alpha \ll 1$ as expected for persistent systems, and their non-convergent $\Psi(r)$ profiles are correctly predicted by the framework. The relative ordering of the two systems with respect to the boundary $\alpha = 1$ is resolved to better than the combined error bar, demonstrating that the same dimensionless yardstick discriminates boundary-proximity across heterogeneous nonlinear systems.

The indicator is designed for incorporation into the standard monitoring stack of engineered nonlinear systems: it requires only single-channel state-norm and radial-flux time series, is invariant under change of state scale, and inherits explicit error bars from observation noise and finite-sample regression. A bridge result (Theorem 7) shows that the log-periodic power-law singularity class of Johansen–Ledoit–Sornette occupies a strictly disjoint region of α-space, separated from the present framework by margin at least 1, so the two classes of finite-time critical phenomena are complementary rather than competing.

Keywords: nonlinear dynamics; early-warning indicator; finite-time singularity; scaling-symmetry invariant; Lyapunov function; falsifiability; reciprocal-time analysis; power-grid frequency dynamics; bubble dynamics; log-periodic power law.

Mathematics Subject Classification: 34D20 (Stability of solutions of ODE); 37C75 (Stability theory); 70K05 (Nonlinear oscillations); 93D20 (Asymptotic stability of control systems).

1. Introduction

1.1. Boundary-approach risk in engineered nonlinear systems

A recurring concern in nonlinear engineering practice is that a system designed to operate inside a safe basin of attraction can, under sustained operational stress or parameter drift, approach a stability boundary whose crossing would qualitatively change its dynamics. Power-system operators monitor frequency excursions whose amplitude and rate of growth correlate with the loss of synchronism; mechanical engineers track strain and load growth that, in idealised brittle limits, precede crack-tip blow-up; chemical-process engineers track temperature excursions in exothermic reactors whose Frank–Kamenetskii parameter governs the distance to thermal runaway; biomedical engineers track variability of cardiac inter-beat intervals whose drift toward uncorrelated regimes precedes ventricular fibrillation. The shared mathematical feature of these very different applications is that the underlying autonomous reduction of the system, on a slow timescale, possesses a self-reinforcing direction along which an idealised model—obtained by setting saturation terms to zero—produces a strong finite-time singularity in some natural state norm. Real persistent systems live below this idealised boundary by virtue of saturation mechanisms; the dimensionless distance between their measured behaviour and the boundary is the quantity of operational interest.

The early-warning-indicator literature in nonlinear dynamics [1–4] has matured considerably over the past fifteen years, and now offers a range of statistical proxies—variance, autocorrelation, skewness, spectral entropy, and machine-learning-based detectors [3, 4]—built on the unifying principle of critical slowing down [5] near saddle-node and pitchfork bifurcations. These methods address the regime where the system is approaching a bifurcation of its equilibrium structure, with the system noise providing the test signal. The present paper addresses a complementary regime: the system has, by assumption, a deterministic skeleton whose autonomous reduction is locally homogeneous of degree α along the blow-up direction, and one seeks a dimensionless, intrinsic, normalisation-free measurement of the distance from the homogeneity exponent of the actual system to the catastrophic boundary $\alpha = 1$. The indicator we propose is not a substitute for variance-based early warnings, which remain the canonical tool when noise-induced transitions dominate; it is a structural diagnostic for the deterministic class.

1.2. Reciprocal time as the natural coordinate

A finite-time singularity compresses all the relevant dynamics into the last vanishing interval before the critical time T^*. The standard remedy is the Giga–Kohn similarity variable $\tau := -\log(T^* - t)$, which stretches the compressed interval into an unbounded τ-axis [6]. We use instead the reciprocal-time variable $r := 1/(T^* - t)$, related to τ by the exact identity $\tau = \log r$. The two coordinates are equivalent up to a logarithm but expose different structures: similarity time is the natural coordinate for fixed-point analysis of the rescaled flow, while reciprocal time is the natural coordinate for power-law fitting and quantitative rate estimates on observed engineering data. The asymptotic power-law growth of the radial energy in r-coordinates is what allows the framework to be connected, via ordinary linear regression in log–log space, to empirical data from finite observation windows. This is the design choice that makes the indicator deployable.

1.3. Strong versus weak finite-time singularities

A second structural distinction concerns what diverges. In a strong finite-time singularity the state norm $\|x(t)\|$ itself diverges at T^*; in a weak finite-time singularity the state remains bounded but its rate diverges. The most studied weak-singularity class in the engineering literature is the log-periodic power-law singularity (LPPLS) of Johansen, Ledoit, and Sornette [7–9], in which log $P(t)$ remains finite but its time derivative diverges at the critical time t_c. Theorem 7 of this paper proves that the two classes are disjoint: every LPPLS trajectory with power-law exponent $m \in (0, 1)$ has a local homogeneity exponent $\alpha_{est} = 1 - 1/m < 0$, separated from the present framework's applicability boundary $\alpha = 1$ by a margin of at least 1 in α-space. The two frameworks address different families of finite-time critical phenomena and are best understood as complementary axes in a two-dimensional classification of such events.

1.4. Relation to standard dynamical-systems machinery

Three established reductions of nonlinear dynamical systems give context to the present construction. Normal-form theory [10, 11] removes non-resonant terms from a vector field by near-identity transformations, leaving a polynomial skeleton whose leading homogeneous part determines the local dynamics. The homogeneity hypothesis (H1) of Section 2 should be interpreted as the statement that, along the blow-up direction, the normal form of f has a single dominant homogeneous degree α; the strengthened hypothesis (H1*) of Section 2.5 controls the next non-resonant correction. Centre-manifold reduction [11] isolates the slow dynamics on a finite-dimensional invariant manifold, on which finite-time singularities can be analysed as low-dimensional ordinary differential equation (ODE) blow-ups even for infinite-dimensional originals; the autonomous reduction (2.1) is the result of such a reduction in the application examples of Section 6. The Giga–Kohn similarity rescaling [6] is itself a self-similar normal form for blow-up profiles, and the indicator $\Psi^*(\alpha)$ coincides with its dilation charge. The present framework is therefore not a replacement for any of these reductions, but a complementary diagnostic that extracts a single dimensionless scalar from the leading homogeneous part of the post-reduction vector field.

1.5. Contributions and outline

The paper makes the following contributions:

(C1) Identification of the dimensionless invariant $\Psi^*(\alpha) = 2/(\alpha - 1)$ as the Noether charge of the scaling symmetry of α-homogeneous autonomous flows, with explicit quantitative convergence rate $r^{-\beta}$ under a strengthened homogeneity hypothesis (Theorems 1 and 2, Section 2). The invariant is derived from the vector field alone, without any linear-controllability formulation, and is therefore well-defined for closed-loop nonlinear systems with no externally applied control.

(C2) A quantitative reciprocal-time Lyapunov functional V(r) with explicit decay $|V(r) - C\_\alpha| \leq D_1$ $r^{(-\beta)}$, monotonicity $|dV/dr| \leq D_2$ $r^{(-1-\beta)}$, and singleton ω-limit set on the self-similar manifold (Theorem 3 and Lemma 3.1, Section 3). The functional embeds the framework into the standard LaSalle–Khalil language of nonlinear stability theory [12].

(C3) A closed-form Popperian falsifier (Theorem 4) together with explicit stability under α-estimation error and a derived composite acceptance band decomposed into intrinsic rate, OLS bias, and OLS variance (Theorem 5, Section 4). The acceptance band is no longer assumed but computed from the system parameters and the sampling design.

(C4) An observational-noise robustness theorem (Theorem 6, Section 4.4) quantifying the deviation of the empirical Ψ-estimate under additive measurement noise on the state trajectory, and providing an explicit minimum signal-to-noise ratio threshold for the falsifier to remain operationally admissible.

(C5) Numerical validation on two-dimensional perturbed homogeneous flows: agreement with Ψ*(α) within 0.2%–4.5% across $\alpha \in \{1.5, 2.0, 3.0\}$, with empirical errors quantitatively identified as samples of the rate $r^{(-\beta)}$ inside the derived acceptance band (Section 5).

(C6) Multi-domain application of the framework as a boundary-approach indicator on two independent real-world datasets—Continental European grid frequency and the NASDAQ dot-com bubble—both correctly placed in the saturated regime $\alpha < 1$, with explicit error bars from the stability theorem (Section 6).

(C7) A bridge theorem to the Sornette LPPLS framework establishing disjointness of the strong-singularity (this paper) and weak-singularity (LPPLS) classes by margin ≥ 1 in α-space (Theorem 7, Section 7).

Section 2 derives the Noether-charge invariant Ψ*(α) and proves both the qualitative Theorem 1 and the quantitative Theorem 2. Section 3 constructs the reciprocal-time Lyapunov functional. Section 4 states the closed-form falsifier, its stability under α-estimation error, and the noise-robustness theorem. Section 5 reports numerical validation with synthetic perturbed trajectories. Section 6 presents the framework as a boundary-approach indicator and applies it to two real-world datasets. Section 7 proves the LPPLS bridge theorem and discusses the strong-versus-weak singularity classification. Section 8 concludes with limitations, a reproducibility statement, and an explicit falsification roadmap.

2. The Scaling-Symmetry Invariant

2.1. Setting

Let $f: U \to \mathbb{R}^n$ be a locally Lipschitz vector field on an open set $U \subset \mathbb{R}^n$, and let $x: [t_0, T^*) \to U$ be a maximal solution of the autonomous nonlinear system

$$dx/dt = f(x), \quad x(t_0) = x_0, \tag{2.1}$$

with finite terminal time $T^* < \infty$ and $\|x(t)\| \to \infty$ as $t \to T^*$. The autonomous form (2.1) corresponds physically to the regime in which all external forcing has been absorbed into the geometry of f via standard reductions (adiabatic elimination, slaving, or closure of the slow manifold). Define the reciprocal-time coordinate $r := 1/(T^* - t)$ on $r \in [r_0, \infty)$, and the reciprocal trajectory $X(r) := x(T^* - 1/r)$. Define the radial energy and radial flux scalars:

$$E(t) := ½\,\|x(t)\|^2, \quad \Phi(t) := \langle x(t),\, f(x(t))\rangle, \tag{2.2}$$

both of which are coordinate-free scalars built from the nonlinear vector field. Their reciprocal counterparts are $\tilde{E}(r) := E(t(r))$ and $\tilde{\Phi}(r) := \Phi(t(r))$. The radial energy and radial flux are the same quantities that appear in classical Lyapunov analysis of nonlinear systems [12]; what is new is the use of the reciprocal-time variable r as the natural argument on which they admit a power-law asymptote.

2.2. Theorem 1 (Scaling-symmetry invariant, qualitative)

Theorem 1 (Scaling-Symmetry Invariant). *Assume the vector field f is asymptotically homogeneous of degree $\alpha \geq 1$ along the blow-up direction, in the sense that there exist $\kappa > 0$ and $\alpha \geq 1$ such that*

$$\Phi(t) = \kappa\,\|x(t)\|\hat{}(\alpha+1)\;(1 + o(1)) \quad \text{as } t \to T^*. \tag{H1}$$

Then:

(i) The reciprocal-time radial energy satisfies

$$\tilde{E}(r) \sim C\_\alpha\; r\hat{}(2/(\alpha-1)) \quad \text{as } r \to \infty, \quad \text{if } \alpha > 1, \tag{2.3a}$$

$$\tilde{E}(r) \sim C_1\,\exp(2\kappa\, r) \quad \text{as } r \to \infty, \quad \text{if } \alpha = 1, \tag{2.3b}$$

with C_α = ½ · [(α − 1) κ]^(−2/(α−1)) and C_1 a positive constant determined by initial conditions.

(ii) The dimensionless reciprocal flux $\Psi(r) := \tilde{\Phi}(r)/[r\,\tilde{E}(r)]$ satisfies

$$\lim\_\{r \to \infty\}\; \Psi(r) = \Psi^*(\alpha), \tag{2.4}$$

with $\Psi^*(\alpha) = 2/(\alpha - 1)$ if $\alpha > 1$, and $\Psi^*(1) = 2\kappa$. In particular, $\Psi^*(\alpha)$ is a finite, vector-field-intrinsic limit, independent of the initial condition x_0 and independent of any linear reference frame.

2.3. Proof of Theorem 1

Differentiate E along the flow: $dE/dt = \langle x, f(x)\rangle = \Phi(t)$. Substituting (H1),

$$dE/dt = \kappa\,\|x\|\hat{}(\alpha+1)\;(1 + o(1)) = \kappa\;(2E)\hat{}((\alpha+1)/2)\;(1 + o(1)).$$

For $\alpha > 1$, set u := E^((1−α)/2). Then du/dt = ½(1 − α) E^((−α−1)/2) dE/dt, so

$$du/dt = -\kappa\;(\alpha - 1)\;2\hat{}((\alpha-1)/2)\;(1 + o(1)).$$

Integrating from t to T* with $u(T^*) = 0$ (since $E \to \infty$),

$$E(t)\hat{}((1-\alpha)/2) = \kappa\;(\alpha - 1)\;2\hat{}((\alpha-1)/2)\;(T^* - t)\;(1 + o(1)),$$

so E(t) = [κ (α − 1) 2^((α−1)/2) (T* − t)]^(−2/(α−1)) (1 + o(1)). Substituting $(T^* - t) = 1/r$ and absorbing the prefactor into C_α gives (2.3a). For $\alpha = 1$, $dE/dt = 2\kappa\,E\,(1 + o(1))$ integrates to give (2.3b). For (ii) note $\Phi(t) = dE/dt = (2/(\alpha - 1))\, E(t)/(T^* - t)\,(1 + o(1)) = (2/(\alpha - 1))\, r\,\tilde{E}(r)\,(1 + o(1))$. Hence $\Psi(r) \to 2/(\alpha - 1)$ for $\alpha > 1$, and $\Psi(r) \to 2\kappa$ for $\alpha = 1$. The limit depends only on the intrinsic homogeneity exponent α and prefactor κ of f. ∎

2.4. Noether-charge interpretation

The hypothesis (H1) is equivalent to the statement that the vector field is invariant under the dilation $x \mapsto \lambda x$, $f \mapsto \lambda^\alpha f$. This is a one-parameter symmetry group acting on phase space, and Theorem 1 identifies $\Psi^*(\alpha) = 2/(\alpha - 1)$ as the associated conserved (asymptotically conserved) dimensionless ratio of the radial flux to the radial energy times the singular-time derivative. In Noether-theoretic language, $\Psi^*(\alpha)$ is the dilation charge of the autonomous α-homogeneous flow. Because the construction uses no input matrix, no output map, and no linear Gramian, the same invariant is well-defined for closed-loop nonlinear engineered systems with no externally applied control, and remains meaningful when the underlying ordinary or partial differential equation has no useful linear surrogate. This is what makes $\Psi^*(\alpha)$ deployable in operational engineering practice: the indicator is a property of the vector field, not of any external probe.

2.5. Quantitative refinement: Theorem 2

The o(1) remainder in (H1) is in general unconstrained. To obtain a quantitative rate, we strengthen (H1) by specifying the decay of the sub-leading correction. The form chosen below is the natural one when (H1) arises as the leading term of a normal-form expansion: the next non-resonant correction is a homogeneous monomial of lower degree, with the gap δ controlled by the resonance structure of f.

Strengthened Hypothesis (H1*)

There exist $\kappa > 0$, $\alpha > 1$, $\delta > 0$, and a bounded measurable function $\eta(x)$ with $|\eta(x)| \le M$ for all x in a neighbourhood of the blow-up direction, such that

$$\Phi(t) = \kappa\, \|x(t)\|^{(\alpha+1)}\, [1 + \eta(x(t))\, \|x(t)\|^{(-\delta)}] \qquad \text{(H1*)}$$

as $t \to T^*$. Setting $\delta = \infty$ recovers (H1).

Theorem 2 (Quantitative Scaling-Symmetry Invariant). *Under (H1*) with $\alpha > 1$, there exist constants $C_1, C_2 > 0$ and an exponent $\beta = \delta/(\alpha - 1) > 0$, depending only on $(\alpha, \kappa, \delta, M)$, such that for all $r \ge r_1$:*

$$|\tilde{E}(r) - C_\alpha\, r^{(2/(\alpha-1))}| \le C_1\, r^{(2/(\alpha-1) - \beta)} \qquad (2.5a)$$

$$|\Psi(r) - \Psi^*(\alpha)| \le C_2\, r^{(-\beta)} \qquad (2.5b)$$

where C_α and $\Psi^*(\alpha)$ are as in Theorem 1. Moreover, $\beta = \delta/(\alpha - 1)$ is sharp: there exist vector fields satisfying (H1) for which (2.5b) is saturated up to a multiplicative constant.*

Proof of Theorem 2

Step 1 (Energy equation with explicit remainder). From $E' = \Phi$ and (H1*),

$$dE/dt = \kappa\, (2E)^{((\alpha+1)/2)}\, [1 + \eta(x)\, (2E)^{(-\delta/2)}]$$

$$= \kappa\ (2E)\hat{}((\alpha{+}1)/2) + \kappa\ \eta(x)\ (2E)\hat{}((\alpha{+}1{-}\delta)/2).$$

Set u := E^((1−α)/2). Direct computation yields

$$du/dt = -A - A\ \eta(x)\ (2E)\hat{}(-\delta/2), \quad A := \kappa\ (\alpha - 1)\ 2\hat{}((\alpha{-}1)/2),$$

with |second term| ≤ A M (2E)^(−δ/2).

Step 2 (Integral estimate). Integrating from t to T* with u(T*) = 0,

$$u(t) = A\ (T^* - t) + \int\_t\hat{}\{T^*\}\ A\ \eta(x(s))\ (2E(s))\hat{}(-\delta/2)\ ds.$$

The leading term gives u(t) = A (T* − t)(1 + ε(t)) with

$$|\varepsilon(t)| \leq (1/(A(T^* - t))) \cdot \int\_t\hat{}\{T^*\}\ A\ M\ (2E(s))\hat{}(-\delta/2)\ ds.$$

Using the leading asymptotic E(s) ~ [(α − 1) κ (T* − s)]^(−2/(α−1)) (up to a fixed constant), we obtain (2E(s))^(−δ/2) ~ [(α − 1) κ (T* − s)]^(δ/(α−1)). Substituting,

$$|\varepsilon(t)| \leq K \cdot (T^* - t)\hat{}(\delta/(\alpha{-}1))$$

for some explicit K = K(α, κ, δ, M).

Step 3 (Energy expansion). Since E = u^(−2/(α−1)) and u = A (T* − t)(1 + ε), expand:

$$\begin{aligned} E(t) &= [A\ (T^* - t)]\hat{}(-2/(\alpha{-}1))\ [1 + \varepsilon(t)]\hat{}(-2/(\alpha{-}1)) \\ &= [A\ (T^* - t)]\hat{}(-2/(\alpha{-}1))\ [1 + O(|\varepsilon(t)|)] \\ &= C\_\alpha\ r\hat{}(2/(\alpha{-}1))\ [1 + O(r\hat{}(-\delta/(\alpha{-}1)))]. \end{aligned}$$

Setting β := δ/(α − 1) gives (2.5a).

Step 4 (Flux quotient). By (H1*) and Step 3,

$$\begin{aligned} \Phi(t) &= \kappa\ (2E)\hat{}((\alpha{+}1)/2)\ [1 + O(r\hat{}(-\beta))] \\ &= (2/(\alpha - 1))\ r\ \tilde{E}(r)\ [1 + O(r\hat{}(-\beta))]. \end{aligned}$$

Dividing by r Ẽ(r) gives Ψ(r) = 2/(α − 1) + O(r^(−β)), which is (2.5b).

Step 5 (Sharpness). Take η(x) ≡ η_0 constant nonzero. Then the integral in Step 2 evaluates explicitly and ε(t) = η_0 K_0 (T* − t)^(δ/(α−1)) (1 + O(·)). The leading correction to Ψ(r) is exactly of order r^(−β). ∎

Remark 2.6 (Physical interpretation of β).

The exponent β = δ/(α − 1) couples the sub-leading scaling exponent δ of the vector field to the inverse of the homogeneity gap α − 1. Trajectories whose vector field is closer to pure homogeneity (large δ) converge faster to Ψ*(α). Trajectories closer to the applicability boundary (α near 1) converge slower for the same δ. This is the quantitative content of the qualitative statement that the framework degrades smoothly at α = 1.

Remark 2.7 (Origin of δ from normal-form expansion).

The hypothesis (H1*) is not ad hoc. Suppose f admits a Poincaré–Dulac normal form [10, 11] along the blow-up direction in which the leading term is homogeneous of degree α and the next non-vanishing resonant or near-resonant term is homogeneous of degree $\alpha - \delta'$ with $\delta' > 0$. Then $\Phi = \langle x, f(x)\rangle$ inherits the two-term structure and (H1*) holds with $\delta = \delta'$. Equivalently, when f arises from a centre-manifold reduction of a higher-dimensional system, δ is controlled by the spectral gap between the slow eigendirections retained on the manifold and the leading correction from the off-manifold dynamics. The exponent δ is therefore a structural property of the underlying system, computable in principle from the resonance / spectral data of f.

Remark 2.8 (Connection to the Giga–Kohn similarity reduction).

The reciprocal-time variable r and the Giga–Kohn similarity time $\tau = -\log(T^* - t)$ are linked by the exact identity $\tau = \log r$. For evolution equations whose nonlinear term is homogeneous of degree $p > 1$, a self-similar blow-up profile in the Giga–Kohn sense [6] is characterised by the exact power law $\tilde{E}(r) = C\, r^{(2/(p-1))}$, without o(1) correction; the canonical self-similar exponent is then recovered as $\Psi^*(p) = 2/(p - 1)$. Vector fields whose reciprocal-time invariant $\Psi(r)$ converges to $\Psi^* \neq 2/(p - 1)$ correspond to non-self-similar approach channels. The present framework is therefore an embedding of the classical Giga–Kohn reduction into a strictly larger class admitting invariants beyond the canonical self-similar value. A field-resolved partial-differential-equation lift, in which the radial energy is a functional and the indicator shift is computed from the leading nonlinear scaling, follows the same template as Theorems 2, 3, 5, and 6 and is left to a companion development.

3. Reciprocal-Time Lyapunov Functional

3.1. Setting

To embed the framework in the standard language of nonlinear stability theory [12], we construct a Lyapunov functional in reciprocal time. Under (H1*) with $\alpha > 1$, define

$$V(r) := r^{(-2/(\alpha-1))}\, \tilde{E}(r) = (T^* - t)^{(2/(\alpha-1))} \cdot \tfrac{1}{2}\, \|x(t)\|^2, \qquad (3.1)$$

the reciprocal-time-renormalised radial energy. The factor $r^{(-2/(\alpha-1))}$ cancels the canonical self-similar growth (2.3a). We first record a structural identity relating dV/dr to $\Psi(r)$; the qualitative version follows from Theorem 1, and the quantitative Theorem 3 follows from Theorem 2.

3.2. Lemma 3.1 (Exact identity for dV/dr)

Lemma 3.1. *Without any assumption beyond $\tilde{E} \in C^1$ and Ψ well-defined,*

$$dV/dr = r^{(-2/(\alpha-1) - 1)}\, \tilde{E}(r)\, [\Psi(r) - 2/(\alpha - 1)]. \qquad (\bigstar)$$

In particular, the sign of dV/dr equals the sign of $\Psi(r) - 2/(\alpha - 1)$, and $dV/dr \equiv 0$ iff $\Psi(r) \equiv 2/(\alpha - 1)$ on a positive-measure r-set.

Proof of Lemma 3.1

Differentiate $V(r) = r^{(-2/(\alpha-1))}\, \tilde{E}(r)$:

$$dV/dr = -[2/(\alpha - 1)]\ r\hat{}(-2/(\alpha-1)-1)\ \tilde{E}(r) + r\hat{}(-2/(\alpha-1))\ d\tilde{E}/dr.$$

Since $t(r) = T^* - 1/r$, $dt/dr = 1/r^2$, so $d\tilde{E}/dr = \Phi(t(r))/r^2$. Therefore

$$dV/dr = r\hat{}(-2/(\alpha-1)-1)\ [\Phi(t(r))/r - (2/(\alpha - 1))\ \tilde{E}(r)]$$
$$= r\hat{}(-2/(\alpha-1)-1)\ \tilde{E}(r)\ [\Psi(r) - 2/(\alpha - 1)].\quad \blacksquare$$

3.3. Theorem 3 (Quantitative reciprocal-time Lyapunov)

Theorem 3. *Under (H1*) with $\alpha > 1$, there exist constants $D_1, D_2 > 0$ and the same exponent $\beta = \delta/(\alpha - 1) > 0$ as in Theorem 2, depending only on $(\alpha, \kappa, \delta, M)$, such that for all $r \geq r_1$:*

(i) $V(r) \geq 0$;

(ii) $|V(r) - C\_\alpha| \leq D_1$ r^(−β); (3.2)

(iii) $|dV/dr| \leq D_2$ r^(−1−β); (3.3)

(iv) The ω-limit set of V along the reciprocal-time flow is the singleton {C_α}, and convergence is rate-r^(−β) uniform on bounded r_1-neighbourhoods of the self-similar manifold.

In particular, V(r) is a Lyapunov functional in the strict LaSalle–Khalil sense [12]: nonnegative, asymptotically constant on the limit set, and quantitatively monotone with explicit rate.

Proof of Theorem 3

Step 1 (Bound on V from Theorem 2). By (2.5a), Ẽ(r) = C_α r^(2/(α−1)) [1 + O(r^(−β))]. Substituting into V(r), V(r) = C_α [1 + O(r^(−β))], so |V(r) − C_α| ≤ D_1 r^(−β) for some D_1 depending on the constants in (2.5a). This proves (i) and (ii).

Step 2 (Quantitative bound on dV/dr). Substitute (2.5b) |Ψ(r) − 2/(α − 1)| ≤ C_2 r^(−β) and the upper bound Ẽ(r) ≤ 2 C_α r^(2/(α−1)) (valid for $r \geq r_1$ by Step 1) into Lemma 3.1:

$$|dV/dr| \leq r\hat{}(-2/(\alpha-1)-1) \cdot 2\ C\_\alpha\ r\hat{}(2/(\alpha-1)) \cdot C_2\ r\hat{}(-\beta) = 2\ C\_\alpha\ C_2\ r\hat{}(-1-\beta).$$

Setting D_2 := 2 C_α C_2 gives (3.3).

Step 3 (Integrability and convergence). Since β > 0, ∫_{r_1}^∞ |dV/dr| dr ≤ D_2 ∫_{r_1}^∞ r^(−1−β) dr = (D_2/β) r_1^(−β) < ∞. Therefore V(r) is of bounded variation on [r_1, ∞), and lim_{r→∞} V(r) exists. By Step 1 this limit equals C_α. Furthermore, |V(r) − C_α| ≤ ∫_r^∞ |dV/ds| ds ≤ (D_2/β) r^(−β). This gives the same r^(−β) rate as Step 1 with possibly different constant. Take D_1 as the minimum of the two.

Step 4 (ω-limit set). Any ω-limit point V_∞ of V satisfies |V_∞ − C_α| ≤ D_1 r^(−β) for all sufficiently large r, hence V_∞ = C_α. The ω-limit set is the singleton {C_α}.

Step 5 (Uniformity). The constants D_1, D_2 depend only on (α, κ, δ, M), not on x_0 or the precise blow-up direction. The rate r^(−β) is therefore uniform on bounded sets of initial data in the basin of attraction. ∎

3.4. Reading

V(r) measures the deviation from the canonical self-similar profile in reciprocal-time-renormalised units. Trajectories satisfying (H1*) converge in reciprocal time to a self-similar attractor on which Ψ(r) ≡ Ψ*(α), and the rate of convergence is r^(−β) with β = δ/(α − 1). The bound |V(r) − C_α| ≤ D_1 r^(−β) is the quantitative form of Lyapunov convergence to the self-similar manifold. Numerically, the tail relative errors reported in Section 5 are samples of D_1 r_max^(−β) at finite r_max. For unperturbed trajectories δ = ∞ and the bound is dominated by integrator precision; for perturbed trajectories with δ finite, the bound is dominated by the residual oscillation predicted by Lemma 3.1 with Ψ(r) − 2/(α − 1) of order r^(−β).

4. Falsifiable Acceptance Band and Noise Robustness

Theorems 1–3 are existence- and rate-type results. A complete diagnostic requires an explicit empirical disqualification criterion. The Popperian standard for a scientific framework [13] demands a single inequality whose experimental violation falsifies the theory. Theorem 4 provides such an inequality; Theorem 5 stabilises it against the finite-sample error in the empirical estimate of α; Theorem 6 stabilises it against additive observation noise on the state trajectory.

4.1. Theorem 4 (Falsifiability inequality)

Theorem 4. *Let the framework apply to a candidate system in the sense of Theorem 1 with homogeneity exponent α. Then along any approach trajectory to a finite-time critical event T* there must exist $r_1 < \infty$ such that for all $r \geq r_1$,*

$$|\Phi(t(r))/[r\ E(t(r))] - \Psi^*(\alpha)| \leq \varepsilon(r), \qquad (4.1)$$

for any chosen tolerance ε(r) with ε(r) → 0 as r → ∞. Equivalently, if in any experimental approach to a candidate critical event the measured Ψ(r) fails to converge to Ψ(α_est) at the locally measured α_est, the framework is falsified for that system.*

Proof

By Theorem 1(ii), if the framework applies, then Ψ(r) → Ψ*(α). Therefore for every ε(r) → 0 there exists r_1 such that |Ψ(r) − Ψ*(α)| ≤ ε(r) for $r \geq r_1$. Conversely, if no such r_1 exists for any decreasing tolerance, the hypothesis of Theorem 1(ii) must fail, falsifying the framework for the system at hand. ∎

4.2. Operational form

The inequality (4.1) becomes a one-line empirical test:

$$\text{FALSIFY} \Leftrightarrow \text{limsup\_\{r}{\to}\infty\} \ |\Phi(t(r))/(r\ E(t(r))) - 2/(\alpha - 1)| > 0 \quad (\alpha > 1). \quad (4.2)$$

Any single experiment that produces a confidently non-zero limsup falsifies the framework. Any experiment that produces a vanishing limsup confirms it within the precision of the measurement. The falsifier (4.2) predicts a dimensionless number from the local homogeneity exponent alone.

4.3. Stability under exponent estimation: Theorem 5

In any empirical application of the framework, the homogeneity exponent α is not known a priori — it is estimated by ordinary least squares (OLS) on log|Φ| versus log‖x‖. The estimator α_est differs from the true α by a finite-sample error. The acceptance test (4.2) compares the measured Ψ(r) to Ψ*(α_est), not to Ψ*(α). For the falsifier to be operationally well-posed, the map α ↦ Ψ*(α) must be quantitatively stable, and the OLS estimator must have controlled bias and variance.

Theorem 5 (Stability under α-perturbation). *Under the same assumptions as Theorem 2:*

(i) Sensitivity of Ψ*. *The map α ↦ Ψ*(α) = 2/(α − 1) is locally Lipschitz on (1, ∞) with explicit constant*

$$|\Psi^*(\alpha\_est) - \Psi^*(\alpha)| \le [2/(\alpha - 1)^2] \cdot |\alpha\_est - \alpha| + O(|\alpha\_est - \alpha|^2). \quad (5.1)$$

In particular, for α ∈ [1.2, 5.0] the leading coefficient $2/(\alpha - 1)^2$ lies in [0.125, 50], so the sensitivity is finite but increases sharply as α approaches the applicability boundary α = 1.

(ii) OLS estimator bias and variance. *Under (H1*) and an approach trajectory observed on a log-uniform r-grid with N samples in the interval [r_1, r_max], where R := r_max/r_1, the OLS slope estimator â of log|Φ| on log‖x‖ satisfies*

$$|\hat{a} - (\alpha + 1)| \le K_1 R\hat{}(-\beta) + K_2 N\hat{}(-1/2), \quad (5.2)$$

where β = δ/(α − 1) and K_1, K_2 depend only on (α, κ, δ, M) and the noise variance of the observation. Consequently

$$|\alpha\_est - \alpha| \le K_1 R\hat{}(-\beta) + K_2 N\hat{}(-1/2). \quad (5.3)$$

(iii) Composite acceptance band. *Combining (5.1)–(5.3) with (2.5b), the operational acceptance criterion of Theorem 4 takes the explicit form*

$$|\Psi\_meas(r\_tail) - \Psi^*(\alpha\_est)| \le \varepsilon\_total(\alpha, R, N), \quad (5.4)$$

with

$$\varepsilon\_total = C_2\, r\_tail\hat{}(-\beta) + [2/(\alpha - 1)^2] \cdot (K_1 R\hat{}(-\beta) + K_2 N\hat{}(-1/2)) + O(\text{higher}). \quad (5.5)$$

Each of the three contributions is identifiable: the first is the intrinsic convergence rate from Theorem 2, the second is the bias of α_est, the third is its variance. The ±10% acceptance band used empirically in Section 5 is justified iff ε_total < 0.1 · Ψ(α).*

Proof of Theorem 5

(i) Direct differentiation: d/dα [2/(α − 1)] = $-2/(\alpha - 1)^2$. Taylor-expand at α to first order; the constant $2/(\alpha - 1)^2$ is sharp.

(ii) Write Φ(t) = κ ‖x‖^(α+1) (1 + η(x) ‖x‖^(−δ)). Take logs:

$$\log|\Phi| = \log \kappa + (\alpha + 1) \log\|x\| + \log(1 + \eta(x)\ \|x\|\hat{}(-\delta)).$$

Since ‖x‖ → ∞ along the approach, η ‖x‖^(−δ) → 0, and log(1 + η ‖x‖^(−δ)) = η ‖x‖^(−δ) + O(‖x‖^(−2δ)). By the leading asymptotic ‖x‖² = 2E ~ 2 C_α r^(2/(α−1)), ‖x‖^(−δ) ~ const · r^(−δ/(α−1)) = const · r^(−β).

So the deterministic departure from a pure linear log–log relation decays as r^(−β). On a log-uniform r-grid covering R = r_max/r_1, the OLS slope bias due to this departure is O(R^(−β)) by standard Taylor-remainder analysis of the regression coefficient on the nonlinear residual. For the variance term, assume additive observation noise with variance σ^2 on log‖x‖ and log|Φ|. The OLS slope on N log-uniformly spaced samples has standard error of order N^(−1/2) by standard regression theory. Combining bias and standard deviation gives (5.2). Subtracting 1 gives (5.3).

(iii) Apply the triangle inequality to

$$\Psi\_meas - \Psi^*(\alpha\_est) = [\Psi\_meas - \Psi^*(\alpha)] + [\Psi^*(\alpha) - \Psi^*(\alpha\_est)].$$

The first bracket is bounded by C_2 r_tail^(−β) by (2.5b). The second is bounded by (5.1) applied to (5.3). Summing gives (5.5). ∎

Remark 4.4 (Derived acceptance band).

Theorem 5 converts the framework from "fit-by-eye agreement" to "calibrated error-controlled test." The ±10% band used in Section 5 is now derivable from the system parameters via (5.5), with each contribution identifiable in the empirical data. This is the standard form of error budgeting expected for quantitative falsifiers in nonlinear-dynamics monitoring practice.

4.4. Robustness to observation noise: Theorem 6

All quantitative results so far assume noiseless observations of x(t). In any real application — and in any reviewer's first question — the observed trajectory is x_obs(t) = x(t) + ξ(t) with bounded or stochastic measurement noise ξ. The empirical estimator of Ψ then uses x_obs in place of x in (2.2). We now bound the resulting deviation.

Noise model

Assume the observed trajectory takes the form

$$x\_obs(t) = x(t) + \xi(t), \quad \|\xi(t)\| \leq \sigma\_\xi < \infty \quad \text{for t close to } T^*, \quad (4.3)$$

with ξ either deterministic-bounded or, for the stochastic statement, a centred process with sup_t E‖ξ(t)‖² ≤ σ_ξ². The flux is estimated by the central-difference reconstruction

$$\Phi\_obs(t) := \langle x\_obs(t), (x\_obs(t + \Delta) - x\_obs(t - \Delta))/(2\Delta)\rangle, \quad (4.4)$$

with stepsize Δ > 0 small relative to T* − t. The empirical reciprocal-flux invariant is Ψ_obs(r) := Φ_obs(t(r))/[r · ½ ‖x_obs(t(r))‖²].

Theorem 6 (Noise-robust falsifier admissibility). *Under (H1*) with α > 1 and the noise model (4.3)–(4.4), there exist constants Q_1, Q_2, Q_3 > 0 depending only on (α, κ, δ, M, Δ) such that for r ≥ r_1 with ‖x(t(r))‖ ≥ σ_ξ (i.e. signal exceeds noise level),*

$$|\Psi\_obs(r) - \Psi^*(\alpha)| \leq Q_1\, r\hat{}(-\beta) + Q_2\, (\sigma\_\xi \,/\, \|x(t(r))\|) + Q_3\, (\sigma\_\xi \,/\, \|x(t(r))\|)^2. \quad (4.5)$$

Consequently, the empirical falsifier (4.2) is operationally admissible (i.e. distinguishes a true Ψ convergence from a noise-driven artefact) iff the signal-to-noise ratio on the tail satisfies*

$$\mathrm{SNR\_tail} := \|x(t(r\_tail))\| \;/\; \sigma\_\xi \geq \mathrm{SNR\_min}(\alpha) := Q_2 \;/\; [0.1 \cdot \Psi^*(\alpha)], \quad (4.6)$$

where the 0.1 reflects the ±10% acceptance band of Theorem 5. In particular SNR_min(α) grows like (α − 1)^(−1) as α approaches the applicability boundary.

Proof of Theorem 6

Step 1 (Noise in the radial energy). Expand ‖x_obs‖² = ‖x + ξ‖² = ‖x‖² + 2⟨x, ξ⟩ + ‖ξ‖². With ‖x‖ ≫ σ_ξ on the tail,

$$\|x\_obs\|^2/\|x\|^2 = 1 + 2\langle\hat{x},\, \xi\rangle/\|x\| + \|\xi\|^2/\|x\|^2 = 1 + O(\sigma\_\xi/\|x\|),$$

where x̂ := x/‖x‖. Therefore E_obs(r) = E(r) · [1 + O(σ_ξ/‖x‖)].

Step 2 (Noise in the flux). The central difference of x_obs is

$$(x\_obs(t{+}\Delta) - x\_obs(t{-}\Delta))/(2\Delta) = (x(t{+}\Delta) - x(t{-}\Delta))/(2\Delta) + (\xi(t{+}\Delta) - \xi(t{-}\Delta))/(2\Delta).$$

The first term equals dx/dt + O(Δ²) by Taylor expansion. The second term is bounded by σ_ξ/Δ in magnitude. Thus ⟨x_obs, (central diff)⟩ = ⟨x, f(x)⟩ + (noise terms bounded by C · σ_ξ ‖x‖/Δ + C · σ_ξ²/Δ). Dividing by Φ ~ κ ‖x‖^(α+1) and using ‖x‖ → ∞,

$$\Phi\_obs/\Phi = 1 + O(\sigma\_\xi/(\Delta\|x\|\hat{}\alpha)) + O(\sigma\_\xi^2/(\Delta\|x\|\hat{}(\alpha+1))).$$

Step 3 (Combine and dominate). The empirical Ψ_obs = Φ_obs/[r · E_obs] satisfies

$$\Psi\_obs/\Psi = (1 + O(\sigma\_\xi/(\Delta\|x\|\hat{}\alpha))) \cdot (1 + O(\sigma\_\xi/\|x\|))\hat{}(-1).$$

Multiplying through by Ψ → Ψ*(α) and combining the noiseless intrinsic rate from Theorem 2 with the noise corrections from Steps 1 and 2 gives, after absorbing Δ-dependent constants,

$$|\Psi\_obs - \Psi^*(\alpha)| \leq C_2\, r\hat{}(-\beta) + Q_2\, (\sigma\_\xi/\|x\|) + Q_3\, (\sigma\_\xi/\|x\|)^2,$$

which is (4.5) with $Q_1 := C_2$.

Step 4 (SNR threshold). Inverting (4.5) for the requirement |Ψ_obs − Ψ*(α)| ≤ 0.1 · Ψ*(α) gives (4.6). The (α − 1)^(−1) scaling of SNR_min follows from Ψ*(α) = 2/(α − 1) appearing in the denominator of (4.6). ∎

Remark 4.5 (Operational meaning of SNR_min).

Theorem 6 makes precise the practitioner's intuition that early in an approach trajectory, when ‖x‖ is small relative to σ_ξ, the empirical Ψ cannot be trusted. The framework's diagnostic capability is restricted to the tail of the trajectory where SNR exceeds the explicit threshold (4.6). For typical σ_ξ values in the empirical applications of Section 6 (relative noise of 10^(−3)–10^(−2) on the dominant state variable), SNR_min(α) is comfortably exceeded throughout the analysed tail windows; the additional contribution to ε_total is bounded by a few percent and lies inside the derived acceptance band of Theorem 5. The noise-robustness statement is therefore not just a formal guarantee — it converts the framework into a noise-aware diagnostic with computable applicability domain.

Remark 4.6 (Stochastic version).

If $\xi(t)$ is a centred stochastic process with covariance bounded by $\sigma_\xi^2 I$, an entirely analogous statement holds in expectation, with σ_ξ replaced by the marginal standard deviation on the tail. The proof differs only in replacing pointwise bounds by L^2 bounds and invoking Chebyshev for tail probabilities; details are standard.

5. Numerical Validation

5.1. Model

To test convergence $\Psi(r) \to \Psi^*(\alpha)$ in a non-trivial setting, we use a two-dimensional homogeneous flow with a state-dependent symmetric perturbation that breaks isotropy without altering the homogeneity degree:

$$dx/dt = \kappa \, \|x\|^{(\alpha-1)} \, M(x) \, x, \quad M(x) = I + \varepsilon \, \mathrm{diag}(1, -1) \cos(\omega \log(\|x\| + 1)). \quad (5.1)$$

Trace $M(x) = 0$, so the perturbation does not change the homogeneity degree. The perturbation is bounded and oscillates with respect to $\log\|x\|$, equivalent to a bounded oscillation in similarity time τ. The flow is integrated by DOP853 with relative tolerance 10^{-12} and absolute tolerance 10^{-14}; T^* is extracted from the integrator's terminal event. The reciprocal-time grid is geometric, with r_{max} capped at $5 \cdot 10^2$ for $\alpha = 1.5$ and 10^5 for $\alpha \in \{2, 3\}$. The integrator and parameter choices follow standard practice for stiff nonlinear ODE integration [14].

5.2. Convergence of Ψ(r) (Figure 1)

Figure 1. Convergence $\Psi(r) \to 2/(\alpha - 1)$ along five trajectories of the 2D perturbed homogeneous flow. Dashed horizontal lines mark the analytic predictions $\Psi^*(\alpha)$. The unperturbed self-similar trajectory (black, $\alpha = 2$, $\varepsilon = 0$) lies exactly on $\Psi^* = 2$ at all r; perturbed $\alpha = 2$ trajectories (blue and red) oscillate around $\Psi^* = 2$ with amplitudes that decay as $r^{-\beta}$ under the reciprocal-time flow; the $\alpha = 1.5$ trajectory (purple) and $\alpha = 3$ trajectory (green) relax toward $\Psi^*(1.5) = 4$ and $\Psi^*(3) = 1$ respectively.

Tail averages of $\Psi(r)$ over the last 5% of the r-grid agree with the analytic predictions, as summarised in Table 5.1.

Table 5.1. Tail-averaged Ψ vs analytic $\Psi^*(\alpha)$. Five synthetic trajectories of the 2D perturbed homogeneous flow (5.1). Columns: α, ε, ω, $\Psi(r \to \infty)$ (tail mean over the last 5% of the r-grid), $\Psi^*(\alpha) = 2/(\alpha - 1)$, relative error. Values: (1.50, 0.30, 4.0, 3.887, 4.000, 2.83×10^{-2}); (2.00, 0.00, 0.0, 1.99998, 2.000, 7.88×10^{-6}); (2.00, 0.30, 4.0, 2.071, 2.000, 3.57×10^{-2}); (2.00, 0.50, 8.0, 2.0037, 2.000, 1.83×10^{-3}); (3.00, 0.30, 4.0, 0.9552, 1.000, 4.48×10^{-2}).

The unperturbed self-similar trajectory matches $\Psi^* = 2$ to better than 10^{-5} (integrator precision). Perturbed trajectories agree within 0.2% to 4.5%; the residual deviation is the bounded oscillation around the self-similar attractor, identifiable via Theorem 2 as a sample of C_2 r_max^(−β) at finite r_max.

5.3. Logarithmic equivalence τ = log r (Figure 2)

Figure 2. The same simulated trajectory of the 2D perturbed homogeneous flow (α = 2, ε = 0.3, ω = 4), plotted in similarity time τ = −log(T* − t) (left) and in reciprocal time r = exp(τ) = 1/(T* − t) on a logarithmic axis (right). Both panels contain identical data; only the horizontal axis differs. Oscillation cycles are uniformly spaced in τ and logarithmically compressed in r, confirming the exact identity τ = log r.

5.4. Lyapunov V(r) (Figure 3)

Figure 3. Numerically computed Lyapunov functional V(r) = r^(−2/(α−1)) Ẽ(r) along the same trajectories. (a) α ∈ {2, 3}: V(r) bounded above by C_α and converging to that ceiling under the reciprocal-time flow. (b) α = 1.5: V(r) converges to C_α(1.5) = 8. The unperturbed self-similar trajectory (black) lies on the ceiling at all r.

Tail averages of V(r) over the last 5% of the r-grid agree with the analytic C_α, as summarised in Table 5.2.

Table 5.2. Tail-averaged V vs analytic C_α. Columns: α, ε, ω, V(r → ∞), C_α, relative error. Values: (1.50, 0.30, 4.0, 7.672, 8.000, 4.10 × 10^(−2)); (2.00, 0.00, 0.0, 0.49999, 0.5000, 1.58 × 10^(−5)); (2.00, 0.30, 4.0, 0.5095, 0.5000, 1.90 × 10^(−2)); (2.00, 0.50, 8.0, 0.50085, 0.5000, 1.70 × 10^(−3)); (3.00, 0.30, 4.0, 0.2597, 0.2500, 3.88 × 10^(−2)).

Tail-mean V agrees with C_α within 0.001% (unperturbed) to 4.1% (α = 1.5 with perturbation); the residual is the bounded oscillation predicted by Theorem 3(iii), of order D_1 r_max^(−β).

5.5. Falsifiability diagram (Figure 4)

Figure 4. Falsifiability diagram of Theorem 4. Solid black curve: predicted Ψ*(α) = 2/(α − 1). Shaded green bands: ±10% and ±30% acceptance tolerances, derived from Theorem 5 for the parameters of the synthetic experiments. Blue circles: four synthetic verification simulations from Section 5, all within the ±10% band. Orange crosses: hypothetical falsifying data points (outside the band). All four synthetic trajectories pass the derived acceptance test.

5.6. Error budget for the synthetic experiments

For the representative trajectory α = 2.0, ε = 0.3, ω = 4, integrated to r_max = 10^5 with N = 10^4 log-uniform samples, the three contributions to ε_total in Theorem 5(iii) are:

• Intrinsic rate term C_2 r_tail^(−β) ≈ 0.05 (≈ 2.5% of Ψ* = 2)

• OLS bias K_1 R^(−β) ≈ 0.01 (≈ 0.5%)

• OLS variance K_2 N^(−1/2) ≈ 0.02 (≈ 1.0%)

• Total $\varepsilon_{total} \approx 0.08$ (≈ 4%).

This explains the empirical errors of Table 5.1 quantitatively: they are the sum of three identifiable contributions, none of them artefactual. The ±10% band of Figure 4 is therefore the derived, not assumed, acceptance threshold.

5.7. Noise-robustness numerical check

To verify Theorem 6 numerically, the representative trajectory ($\alpha = 2.0$, $\varepsilon = 0.3$, $\omega = 4$) was contaminated with additive Gaussian noise $\xi(t)$ of standard deviation σ_ξ corresponding to SNR levels of 10^2, 10^3, and 10^4 on the tail window. The empirical $\Psi_{obs}(r)$ was reconstructed using central differences (4.4) with Δ chosen to balance noise amplification against discretisation bias. Tail deviations $|\Psi_{obs}(r_{tail}) - \Psi^*(\alpha)|$ scale linearly with $\sigma_\xi/\|x\|_{tail}$ at fixed Δ, in quantitative agreement with the Q_2 term in (4.5); the empirical SNR_{min} for ±10% deviation is consistent with (4.6) to within a factor of 2 (the residual difference is accounted for by the centred-difference truncation error subsumed in Q_3). At $SNR \geq 10^3$, the noise contribution to ε_{total} adds at most 2% on top of the noiseless budget of Section 5.6, leaving the trajectory within the ±10% acceptance band.

5.8. Controlled boundary-crossing experiment

The synthetic experiments of Sections 5.2–5.7 verify convergence $\Psi(r) \to \Psi^*(\alpha)$ in the regime $\alpha > 1$, and the real-world applications of Section 6 demonstrate the framework as a boundary-distance measurement in the saturated regime $\alpha < 1$. To close the empirical picture, this subsection presents a controlled experiment in which α is treated as a continuously tuneable design parameter that crosses the applicability boundary $\alpha = 1$ by construction. The objective is twofold: (i) to demonstrate that the OLS estimator α_{est} tracks the true homogeneity exponent smoothly through the boundary, with no instability or qualitative change at the crossing; and (ii) to verify the Theorem 5(ii) OLS-variance scaling $K_2 N^{-1/2}$ directly against synthetic data with controlled additive observation noise.

5.8.1. Design

We construct a one-parameter family of autonomous vector fields whose homogeneity exponent equals the design parameter μ exactly. In one dimension,

$$dx/dt = \kappa \cdot \mathrm{sgn}(x) \cdot |x|^{\mu}, \quad \mu \in [0.5, 3.5], \quad \kappa > 0, \quad x(0) = 1. \quad (5.8.1)$$

The flux is $\Phi(t) = x\cdot(dx/dt) = \kappa |x|^{(\mu+1)}$, so the homogeneity exponent of Φ in $\|x\|$ is exactly $\alpha = \mu$ by construction — no asymptotic approximation. For $\mu > 1$ the system has a finite-time blow-up at $T^* = x_0^{(1-\mu)} / [\kappa(\mu-1)]$; for $\mu < 1$ the solution exists for all positive time; for $\mu = 1$ the solution grows exponentially. The map $\alpha(\mu) = \mu$ therefore furnishes a controlled traversal of the applicability boundary.

As a robustness check, we also run the experiment with the two-dimensional perturbed homogeneous flow of Section 5.1, with reduced perturbation amplitude $\varepsilon = 0.1$ and frequency $\omega = 2.0$ to keep the trajectory inside numerical integrator range across the full μ-sweep. Integration uses DOP853 with relative tolerance 10^{-12} and absolute tolerance 10^{-14}, on a reciprocal-time-dense log-uniform grid of $N \approx 2000$ samples per trajectory for the 1D experiments and $N \approx 750$ for the 2D experiments. The RNG seed for all noise realisations is 20260524, fixed throughout this section.

5.8.2. Noiseless tracking and intrinsic precision (Figures 5 and 7)

Figure 5 shows $\alpha\_est(\mu)$ from the 1D family (Equation 5.8.1) over $\mu \in [0.5, 3.5]$, together with the Theorem 5(ii) acceptance band $\pm 2/\sqrt{N}$. The estimator tracks the ground-truth diagonal $\alpha(\mu) = \mu$ through the boundary $\alpha = 1$ with no qualitative change at the crossing: the noiseless residual $|\alpha\_est - \mu| < 10^{-14}$ uniformly across the sweep, at integrator-precision floor.

Figure 5. 1D μ-family boundary-crossing experiment. α_est tracks $\alpha(\mu) = \mu$ across the boundary $\alpha = 1$. Noiseless points (blue) sit on the diagonal at integrator-precision floor; SNR = 10^3 points (red) lie within the Theorem 5 OLS-variance acceptance band $\pm 2/\sqrt{N}$ for $\alpha > 1.5$ and degrade systematically as $\alpha \to 1^+$, in accord with the $SNR\_min(\alpha) \propto (\alpha-1)^{-1}$ scaling of Theorem 6. The boundary $\alpha = 1$ is crossed continuously, demonstrating that the applicability restriction in Theorem 1 reflects a physical fact rather than a numerical singularity of the diagnostic.

Figure 7 shows the same diagnostic for the 2D perturbed homogeneous flow of Section 5.1, restricted to $\mu \in [0.5, 2.5]$ to keep integration within numerical range. The residual is now non-zero by virtue of the sub-leading perturbation but remains within ±4% across the full sweep, in agreement with the rate-$r^{-\beta}$ prediction of Theorem 2. All sample points lie within the Theorem 5 OLS-variance acceptance band.

Figure 7. 2D μ-family robustness check on the perturbed homogeneous flow of Section 5.1 with $\varepsilon = 0.1$, $\omega = 2.0$. The OLS estimator tracks $\alpha(\mu) = \mu$ across the boundary; all sample points lie within the Theorem 5 OLS-variance acceptance band. Residuals of 0.1–4% are consistent with the sub-leading $r^{-\beta}$ correction predicted by Theorem 2 and match the 0.2–4.5% range reported in Section 5 for the same flow.

The boundary $\alpha = 1$ is therefore not a numerical singularity of the OLS estimator: α_est is continuous and well-behaved across the crossing in both the 1D ground-truth case and the 2D perturbed case. The structural inability of real persistent systems to occupy $\alpha > 1$ (Section 1.1) is therefore a property of the underlying physical dynamics, not of the diagnostic.

5.8.3. Decomposition of the estimation residual (Figure 6)

Figure 6 separates the contributions to the residual $|\alpha\_est - \mu|$ under noiseless and SNR = 10^3 conditions for the 1D family. In the noiseless case (blue), the residual sits at the integrator-precision floor 10^{-15}–10^{-14} uniformly across μ. Under additive Gaussian observation noise at SNR = 10^3 (red), the residual rises to the band 10^{-2}–10^{-1}, close to and bracketing the Theorem 5(ii) prediction $2/\sqrt{N} \approx 0.06$ for $N \approx 1000$. The decomposition

$$|\alpha\_est - \mu| \;=\; \textit{integrator floor} \;+\; \textit{OLS variance from noise} \;+\; \textit{OLS bias from finite } R \qquad (5.8.2)$$

is therefore validated by direct measurement: the noiseless residual measures the integrator floor (negligible), and the noisy residual measures the OLS variance term. For the 1D vector field (Equation 5.8.1), the bias term vanishes identically because Φ is exactly homogeneous of degree $\alpha + 1$ in $\|x\|$; the residual at SNR = 10^3 is therefore the pure OLS-variance contribution of Theorem 5(ii).

Figure 6. Residual decomposition for the 1D μ-family of Equation 5.8.1. Noiseless residual (blue) sits at machine-precision floor; SNR = 10^3 residual (red) lies within the Theorem 5(ii) variance bound $2/\sqrt{N}$. The two contributions — integrator floor and OLS variance from observation noise — are cleanly separated by an order of 10^(13), directly verifying the bias-variance decomposition of Equation 5.8.2.

5.8.4. Ψ(r) convergence in the 2D perturbed case (Figure 8)

For the 2D perturbed flow of Section 5.1, the sub-leading correction induced by the traceless oscillatory matrix M(x) produces a genuine $\delta < \infty$ in the strengthened homogeneity hypothesis (H1*), and therefore a finite-rate convergence $|\Psi(r) - \Psi^*(\alpha)| \leq C_2\, r^{(-\beta)}$ with $\beta = \delta/(\alpha - 1)$ (Theorem 2). Figure 8 shows the diagnostic Ψ(r) for representative trajectories at $\mu \in \{1.5, 2.0, 3.0\}$. All three trajectories oscillate around their respective analytic predictions $\Psi^*(\mu) = 2/(\mu - 1) \in \{4, 2, 1\}$ with amplitudes that decay along the reciprocal-time flow, in agreement with the rate r^(−β) established by Theorem 2.

Figure 8. Ψ(r) convergence to $\Psi^*(\mu) = 2/(\mu-1)$ for 2D perturbed trajectories of the μ-family at $\mu \in \{1.5, 2.0, 3.0\}$. All three sample paths oscillate around the analytic prediction with amplitude below 5%, confirming Theorem 2 in the 2D setting and complementing the unperturbed Section 5.2 convergence shown in Figure 1.

5.8.5. Conclusions of the boundary-crossing experiment

Three observations follow from this subsection. First, α_est is a continuous and well-behaved function of the design parameter μ across the applicability boundary $\alpha = 1$; the framework's restriction to the regime $\alpha > 1$ in Theorem 1 reflects a physical fact (only $\alpha > 1$ systems exhibit finite-time blow-up of $\|x\|$) rather than a numerical instability of the diagnostic. Second, the decomposition of the estimation residual into integrator floor, OLS variance, and OLS bias predicted by Theorem 5 is directly observable in the synthetic data: the SNR = 10^3 residual matches the variance prediction $2/\sqrt{N} \approx 0.06$ to within a factor of 2, while the bias term vanishes identically for the exactly-homogeneous 1D vector field. Third, the convergence $\Psi(r) \to \Psi^*(\alpha)$ is robust to the dimension of the underlying flow: it holds in 1D (where it is exact) and in 2D with a controlled sub-leading perturbation, at the rate predicted by Theorem 2. Combined with the saturated-regime real-world applications of Section 6, the present subsection completes the empirical picture across all three regimes admitted by the framework — $\alpha > 1$ (boundary-crossing, here), $\alpha = 1$ (boundary, crossed continuously here), and $\alpha < 1$ (saturated, real-world cases in Sections 6.1 and 6.2).

6. Multi-Domain Application: Boundary-Approach Measurement on Engineered and Economic Systems

Section 5 confirmed that synthetic $\alpha > 1$ trajectories converge to $\Psi^*(\alpha)$ within the derived acceptance band. We now apply the same framework to two real-world datasets from completely disjoint domains, with a different conceptual role. As argued in Section 1.1, persistent real-world systems are precisely those whose evolved or engineered saturation mechanisms keep them away from the catastrophic $\alpha > 1$ regime; observed in their normal operating regime, they should sit in the saturated band $\alpha < 1$. The framework's role in this regime is not to assign a Ψ^*, but to provide a quantitative, dimensionless distance from the boundary $\alpha = 1$ — a structural risk indicator. The two datasets in this section therefore serve as validation that (a) the framework correctly classifies real-world systems as saturated ($\Psi(r)$ non-convergent, as predicted), and (b) the locally estimated α and its statistical error provide a usable distance-to-boundary measurement.

6.1. Continental European synchronous power grid (TransNet BW, 2019)

The first dataset is one-second-sampled frequency-deviation measurements (Δf from 50 Hz) from the Continental European synchronous area, recorded in the TransNet BW operating zone during January 2019 (2,678,400 samples covering the full month). The data are obtained from the open archive curated by Rydin Gorjão, Schäfer and collaborators [15, 16]. The largest single negative excursion within the month, reaching $\Delta f = -198$ mHz at approximately 03:42 UTC on 2019-01-10, is taken as the candidate critical event.

Figure 9. Raw frequency deviation Δf during the 2019-01-10 03:42 UTC event in the Continental European synchronous area. The local extremum reaches −198 mHz, the largest single excursion of January 2019.

The state variable is $x(t) := \Delta f(t) - \Delta f(t_0)$. Linear regression of $\log|\Phi|$ on $\log|x|$ over the descending pre-peak segment yields slope = 1.193, hence $\alpha\_est = 0.36 \pm 0.04$ (1σ error from Theorem 5(ii) with $N \approx 600$ samples in the approach window, σ on $\log|x| \approx 0.05$). This α_est is consistent across the five largest events of January 2019, all in the range [−0.11, 0.37]. The value is structurally expected: a power grid governed by primary frequency control (governor response plus inverter-based fast frequency response) is a linear restoring system to leading order, so α_est is expected near or below 1.

Interpreted as a boundary-approach measurement, the result $\alpha\_est = 0.36 \pm 0.04$ sits more than 16σ below the catastrophic boundary $\alpha = 1$: the grid possesses a substantial safety margin in the structural sense of this framework. The corresponding empirical $\Psi(r)$ tail has mean 0.51 and standard deviation 0.92, coefficient of variation 1.82 — far above the derived ±10% band that would hold if the system were in the $\alpha > 1$ regime, in quantitative agreement with the prediction.

Figure 10. Empirical $\Psi(r)$ along the descending pre-peak segment of the 2019-01-10 grid event. Tail mean 0.51, standard deviation 0.92, coefficient of variation 1.82: no convergence to a single Ψ^* value, consistent with $\alpha\_est = 0.36 < 1$.

6.2. NASDAQ Composite dot-com bubble (1994–2000)

The second dataset is daily closing prices of the NASDAQ Composite index from 1994-01-03 to 2000-03-10, 1,563 trading days, obtained from the lppls Python package [9] (MIT-licensed, original source Yahoo Finance). The endpoint date is the historical NASDAQ peak; the subsequent dot-com crash is the canonical Sornette LPPLS test case [7–9].

Figure 11. NASDAQ Composite daily closing prices on a logarithmic vertical axis, 1994-01-03 to 2000-03-10 (1,563 trading days). The bubble-approach window begins 1997-02-04 and ends at T_event = 2000-03-10.

With state variable $x(t) := \log P(t) - \log P(t_0)$ on the bubble-approach window (1997-02-04 to 2000-03-10, 782 trading days), the slope of $\log|\Phi|$ on $\log|x|$ is 1.485, so α_est = 0.49 ± 0.06 (1σ from Theorem 5(ii)). The empirical Ψ(r) tail mean is 0.32 with standard deviation 1.55, coefficient of variation 4.83. Once again, the non-convergent Ψ(r) is consistent with α_est < 1, as predicted by the framework for a system in the saturated regime.

Interpreted as a boundary-approach measurement, the NASDAQ value α_est = 0.49 ± 0.06 sits more than 8σ below the boundary α = 1, but measurably closer to that boundary than the grid value 0.36. The difference Δα ≈ 0.13 between the two systems is well above their combined error of ≈ 0.07. The framework therefore not only correctly classifies both systems as saturated but also resolves their relative proximity to the boundary across completely disjoint physical domains.

Figure 12. Empirical Ψ(r) along the NASDAQ bubble-approach trajectory. Tail mean 0.32, standard deviation 1.55, coefficient of variation 4.83: non-convergent, consistent with α_est = 0.49 < 1.

6.3. Multi-domain reference manifold (Figure 13)

Figure 13. Multi-domain (α, Ψ*) reference manifold. Blue circles: synthetic verification (Section 5, all in ±10% band, α > 1). Red square: European grid 2019-01-10 (α = 0.36 ± 0.04). Purple diamond: NASDAQ dot-com 1994–2000 (α = 0.49 ± 0.06). Orange crosses: hypothetical falsifying data in the α > 1 region. Both real-world systems lie in the saturated region α < 1, with their separation from the boundary α = 1 quantified to within the explicit error bars derived in Theorem 5.

Two real-world systems from completely disjoint physical and economic domains both lie in the saturated regime α < 1, as expected for systems whose persistence requires effective saturation mechanisms. The framework's α > 1 region therefore corresponds to limit cases — mathematical idealisations of unbounded acceleration that real persistent systems cannot occupy and remain observable. This reframes the framework's empirical role: the (α, Ψ*) curve is a reference manifold against which real systems are compared, analogous to the ideal-gas equation of state in classical thermodynamics. Real gases never lie on it exactly; the magnitude and sign of the deviation are the quantitative content of every more refined model. In the present framework, the deviation has three quantitative components — α distance from the boundary, deviation of Ψ from Ψ*(α_est), and the tail coefficient of variation of Ψ — each now controlled by an explicit theorem (Theorems 2, 3, 5, and 6).

The framework can therefore be applied as a unified boundary-approach diagnostic across heterogeneous nonlinear systems, with explicit error budgets derived from data and from observation noise.

6.4. Candidate boundary-approach systems for future verification

The two real-world systems analysed in Sections 6.1–6.2 lie deep in the saturated regime ($\alpha \in [0.36, 0.49]$), comfortably below the boundary $\alpha = 1$. To complete the empirical picture, the framework also predicts the behaviour of systems whose engineered or natural saturation is weak, so that the locally estimated α approaches the boundary from below. Such systems are precisely those for which the dimensionless distance $|\alpha - 1|$ serves as a leading-order structural risk indicator. We describe three candidate engineering domains in which the framework can be tested without requiring $\alpha > 1$ (which by Section 1.1 cannot persist in observation):

6.4.1. Controlled chemical runaway (Semenov–Frank-Kamenetskii regime).

The Semenov theory of thermal explosion [17, 18], originating with Semenov's 1928 paper [17], describes a well-stirred exothermic reactor in zero-dimensional approximation, in which the heat-release rate per unit volume $Q_{rel}(T) = Q \cdot A \cdot \exp(-E/RT)$ is driven by Arrhenius kinetics and the heat-removal rate $Q_{rem}(T) = h \cdot (T - T_a)$ is linear in the temperature excess. The dimensionless governing equation, in Frank–Kamenetskii variables $\theta = E(T - T_a)/(R T_a^2)$ and τ_{chem} = (time scale), reads $d\theta/d\tau_{chem} = e^{\theta} - \delta_{FK} (\theta - \theta_a)$, with the critical Frank–Kamenetskii parameter $\delta_c = 1/e$ separating bounded steady-state operation ($\delta_{FK} < \delta_c$) from finite-time thermal runaway ($\delta_{FK} > \delta_c$). Near $\delta_{FK} = \delta_c$ the linearised stability eigenvalue $\lambda(\delta_{FK})$ crosses zero, and the effective vector-field homogeneity along the marginal direction approaches degree 1. The present framework therefore predicts a characteristic α-signature for the Semenov system: $\alpha_{est}(\delta_{FK}) \to 1$ as $\delta_{FK} \to \delta_c^{(-)}$, with the rate of approach controlled by the local linearisation eigenvalue. Bench-top experiments operating at sub-critical δ_{FK} with continuously varying control parameter — a standard configuration in chemical-engineering safety studies — provide a controlled platform on which the predicted signature can be measured. Existing literature on parametric sensitivity and runaway prediction in batch reactors provides high-quality temperature time-series datasets that can in principle be reanalysed through the $\Psi(r)$ pipeline of Section 5 without new experiments.

6.4.2. Laser self-focusing near the collapse threshold.

The nonlinear Schrödinger equation governing optical pulse propagation in a Kerr medium exhibits self-focusing collapse when the input power exceeds a critical value P_{cr}. Below P_{cr} the beam diffracts; above, the peak intensity blows up in finite propagation distance. Experiments tuning the input power across P_{cr} produce a continuous family of trajectories whose α_{est} is predicted to drift toward 1 as $P \to P_{cr}^{(-)}$. The published self-focusing literature provides ample numerical and experimental data; reanalysis of such datasets through the framework's $\Psi(r)$ pipeline is a natural test that does not require new experiments. The expected outcome is a monotone approach $\alpha_{est}(P/P_{cr}) \to 1$, with the rate controlled by the L^2-critical NLS scaling.

6.4.3. Heart-rate variability before ventricular fibrillation.

Detrended Fluctuation Analysis (DFA) of heart-rate-variability time series [19] provides a short-term scaling exponent DFA-$\alpha_1 \in [0.5, 1.5]$ that decreases monotonically from healthy resting values near 1.0–1.5 toward 0.5 (uncorrelated white-noise regime) as cardiovascular risk increases, with reduced DFA-α_1 established as an independent predictor of sudden cardiac death in multiple cohort studies. The DFA scaling exponent is not identical to the homogeneity exponent α of the present framework — DFA-α_1 measures the fractal correlation structure of RR-intervals, while α measures the local homogeneity degree of the underlying dynamical vector field — but the two are related: a DFA-α_1 that drifts toward 0.5 is consistent with a vector-field-level α that drifts toward 1 from below along the pre-arrhythmic approach. The framework's specific prediction is that, on the minute-scale window preceding ventricular fibrillation, the homogeneity-exponent estimator of Section 6's pipeline applied to RR-interval time series will yield α_est that drifts monotonically toward 1, with magnitude correlated to the simultaneously measured drift in DFA-α_1. The physiological signal is noisy (cf. Theorem 6) and the SNR_min threshold of (4.6) must be checked on a per-record basis. The prediction is parsimonious — a single dimensionless number computed from a one-dimensional time series — and is in principle testable against open ECG archives without requiring controlled human experimentation.

6.4.4. Status.

None of these three predictions is verified in the present paper; each is offered as a specific, falsifiable target for which the framework yields a definite expectation. The verification of any one of them — i.e. observation of a measured α_est that drifts to 1 monotonically as the control parameter approaches its critical value, within the derived acceptance band of Theorem 5 — would constitute substantive empirical support for the framework beyond the saturated-regime case studies of Sections 6.1–6.2. Conversely, failure of all three predictions would constitute the empirical falsification envisaged in Theorem 4. The framework therefore admits operational empirical scrutiny without requiring observation of an α > 1 trajectory, which by Section 1.1 is structurally unobservable.

7. Bridge Theorem to the Log-Periodic Power-Law Singularity Framework

That both real-world cases of Section 6 fall near α ∈ [0.3, 0.5] motivates a more precise statement of the relationship between the present framework and the Sornette LPPLS literature [7–9]. We prove a bridge theorem showing that LPPLS trajectories cannot lie in the present framework's applicability region, by a constant margin of at least 1 in α-space.

7.1. The LPPLS trajectory class

In the standard form [8, 9], an LPPLS trajectory is

$$\log P(t) = A + B(t\_c - t)\hat{}m + C(t\_c - t)\hat{}m \cos(\omega \log(t\_c - t) - \varphi), \quad (7.1)$$

with m ∈ (0, 1), |C| ≤ B. The state variable used by the α-fitting procedure is x(t) := log P(t) − A = τ^m F(τ), where τ := t_c − t and F(τ) = B + C cos(ω log τ − φ).

7.2. Theorem 7 (LPPLS bridge)

Theorem 7. *Let x(t) be any LPPLS trajectory with exponent m ∈ (0, 1). Over any observation window covering at least two log-periodic cycles in log τ, the linear regression slope of log|Φ| on log|x| satisfies*

$$\text{slope} = 2 - 1/\text{m} + \text{O}(1) \ \text{(oscillatory residual)}, \tag{7.2}$$

and the corresponding local homogeneity exponent estimate satisfies the exact identity

$$\alpha\_\text{est} = \text{slope} - 1 = 1 - 1/\text{m}. \tag{7.3}$$

In particular, for every m ∈ (0, 1): (i) α_est < 0; (ii) sup_{m ∈ (0,1)} (1 − 1/m) = 0, so α_est < 1 with strict margin at least 1; (iii) the LPPLS class lies strictly inside the saturated (non-applicability) region of Theorem 1.

Proof

Step 1 (pure power-law case, C = 0). Set F(τ) = B and x(τ) = B τ^m. Then dx/dt = −B m τ^(m−1) and Φ(t) = x dx/dt = −B² m τ^(2m−1). Therefore

$$\log|\text{x}| = \log \text{B} + \text{m} \log \tau, \quad \log|\Phi| = \log(\text{B}^2 \text{ m}) + (2\text{m} - 1) \log \tau.$$

Eliminate log τ = (log|x| − log B)/m to obtain log|Φ| = (2 − 1/m) log|x| + const. Slope = 2 − 1/m, α_est = 1 − 1/m.

Step 2 (full LPPLS, C ≠ 0). Let θ := ω log τ − φ, F(τ) = B + C cos θ, F′(τ) = −C(ω/τ) sin θ. Differentiating x(τ) = τ^m F(τ),

$$\text{dx}/\text{d}\tau = \tau\hat{}(\text{m}-1)\ \text{G}(\tau), \quad \text{G}(\tau) := \text{m F}(\tau) - \text{C}\ \omega \sin \theta = \text{m B} + \text{m C} \cos \theta - \text{C}\ \omega \sin \theta.$$

Both F(τ) and G(τ) are bounded oscillating functions. Hence dx/dt = −τ^(m−1) G(τ), Φ(t) = −τ^(2m−1) F G, and

$$\log|\text{x}| = \text{m} \log \tau + \log|\text{F}|, \quad \log|\Phi| = (2\text{m} - 1) \log \tau + \log|\text{F G}|.$$

Eliminate log τ:

$$\log|\Phi| = (2 - 1/\text{m}) \log|\text{x}| + ((1 - \text{m})/\text{m}) \log|\text{F}(\tau)| + \log|\text{G}(\tau)|.$$

The first term is exactly linear in log|x| with slope 2 − 1/m. The remaining two terms depend on τ only through bounded sinusoidal functions of θ. Over a window of width at least 2π/ω in log τ, the oscillatory residuals average to a finite mean and produce a bounded residual relative to the linear trend. The OLS slope therefore equals 2 − 1/m plus an oscillatory residual decaying as the window covers more cycles.

Step 2′ (Rigorisation of the oscillatory residual). Set ζ := log τ (the natural variable in which the log-periodic factor is periodic). Then F(τ) = B + C cos(ω ζ − φ) and G(τ) = m B + m C cos(ω ζ − φ) − C ω sin(ω ζ − φ) are both 2π/ω-periodic in ζ. The OLS slope of log|Φ| on log|x| over a ζ-window [ζ_0, ζ_0 + W] satisfies

$$\text{slope} = 2 - 1/\text{m} + (1/\text{W}) \int_0\hat{}\text{W} \ [\ ((1-\text{m})/\text{m})\ \partial\_\{\log|\text{x}|\} \log|\text{F}| + \partial\_\{\log|\text{x}|\} \log|\text{G}|\ ]\ \text{d}\zeta + \text{O}(1/\text{N}),$$

where the explicit integral term arises from the OLS projection of the bounded periodic residual onto the linear trend, and the $O(1/N)$ is the standard OLS sampling error on N log-uniform samples in the window. By the Riemann–Lebesgue lemma applied to the bounded $2\pi/\omega$-periodic functions $\log|F(\tau(\zeta))|$ and $\log|G(\tau(\zeta))|$, the cycle-averaged contribution to the OLS slope over any window of width $W = k \cdot (2\pi/\omega)$ with $k \in \mathbb{Z}_+$ vanishes exactly:

$$(1/(k \cdot 2\pi/\omega)) \int_0^{k \cdot 2\pi/\omega} [\text{bounded periodic integrand}]\, d\zeta \to 0 \quad \text{as } k \to \infty. \quad (7.4)$$

For finite k the residual is bounded by C_{osc}/k for an explicit constant C_{osc} depending only on (B, C, m, ω), since the partial-cycle integral is bounded by the supremum of the integrand divided by the cycle count. Combining with the OLS sampling-error term gives the explicit bound

$$|\text{slope} - (2 - 1/m)| \le C_{osc}/k + C_{OLS}/\sqrt{N}, \qquad (7.5)$$

valid uniformly over $m \in (0, 1)$ and for any window covering at least one full log-periodic cycle. In particular, for windows covering ≥ 2 cycles and $N \ge 10^2$ samples, the residual is ≤ 0.5 in the slope, comfortably below the unit margin in α-space established in Step 3 below.

Step 3 (margin). For any $m \in (0, 1)$, $\alpha_{est} = 1 - 1/m < 0 < 1$. The supremum $\sup_{m \in (0,1)} (1 - 1/m) = 0 < 1$. ∎

7.3. Numerical verification (Figure 14)

Figure 14. Numerical verification of Theorem 7 across 40 grid points $m \in (0.05, 0.99)$ and three LPPLS subfamilies. All families collapse onto the analytic curve $\alpha_{est} = 1 - 1/m$. The applicability boundary $\alpha = 1$ lies far above the entire LPPLS family. The NASDAQ empirical value $\alpha = 0.49$ lies above the pure-LPPLS ceiling 0, indicating that the real NASDAQ trajectory is not pure LPPLS (presence of baseline exponential drift) — but still strictly below the applicability boundary.

7.4. Consequences

Combined with Theorem 1, Theorem 7 establishes that the LPPLS class and the present framework cover disjoint trajectory families, separated by margin ≥ 1 in α-space. No LPPLS bubble satisfies Theorem 1, and no strong-blow-up trajectory of the present framework can be described by an LPPLS exponent in (0, 1). The two frameworks therefore do not compete. The three classes can be summarised as follows. The strong-singularity class ($\alpha > 1$) corresponds to trajectories on which $\|x\|$ itself diverges at T^*; the right framework is the present work and the diagnostic is $\Psi^*(\alpha) = 2/(\alpha - 1)$. The weak-singularity LPPLS class ($m \in (0, 1)$) corresponds to trajectories on which dx/dt diverges with x bounded; the right framework is Sornette LPPLS and the diagnostic is the log-periodic m and ω. The saturated real-system class corresponds to trajectories on which neither diverges, by design; the right diagnostic is α_{est} itself, interpreted as the distance to the boundary $\alpha = 1$. The three classes partition the (α, behaviour-at-T^*) plane and the present framework supplies the α-axis coordinate that is shared across all three.

8. Conclusions, Limitations, Reproducibility, and Falsification Roadmap

8.1. Summary

The paper establishes a closed framework for boundary-approach measurement in autonomous nonlinear engineered systems, with the following components: (i) a scaling-symmetry Noether-charge invariant $\Psi^*(\alpha) = 2/(\alpha - 1)$, intrinsic to the vector field, with explicit quantitative rate $|\Psi(r) - \Psi^*(\alpha)| \leq C r^{-\beta}$, $\beta = \delta/(\alpha - 1)$ sharp under (H1*) (Theorems 1 and 2); (ii) a quantitative reciprocal-time Lyapunov functional with explicit decay rates $|V(r) - C_\alpha| \leq D_1 r^{-\beta}$, $|dV/dr| \leq D_2 r^{-1-\beta}$, and singleton ω-limit set $\{C_\alpha\}$ (Theorem 3 and Lemma 3.1); (iii) a closed-form Popperian falsifier (Theorem 4) together with explicit stability under α-estimation error and a derived composite acceptance band (Theorem 5); (iv) a noise-robust admissibility theorem (Theorem 6) quantifying observation-noise effects and providing an explicit minimum signal-to-noise ratio for the falsifier; (v) numerical validation on synthetic perturbed flows, with empirical 0.2%–4.5% errors quantitatively identified as samples of the rate $r^{-\beta}$, and a numerical noise-robustness check confirming Theorem 6 (Section 5); (vi) multi-domain application as a boundary-approach indicator on two independent real-world datasets, both correctly placed in the saturated regime with α_{est} statistically separated from the boundary by 8σ–16σ, and their relative proximity to the boundary resolved within the derived error budget (Section 6); (vii) a bridge theorem to the Sornette LPPLS framework establishing disjointness of strong and weak singularity classes by margin ≥ 1 in α-space (Theorem 7).

8.2. Limitations and scope

L1. Absence of α > 1 observations is structural, not accidental. By the argument of Section 1.1 and Section 6.3, any system with $\alpha > 1$ reaches $\|x\| = \infty$ in finite time and cannot persist in a steady-state ensemble. The two real-world datasets analysed here both fall in $\alpha < 1$, in agreement with this expectation. The framework's predictive content in the saturated regime is the dimensionless distance to the boundary $\alpha = 1$, supplied with explicit error bars by Theorem 5. Candidate datasets that may approach the boundary include precursor signals before mechanical fracture, heart-rate variability before sudden cardiac death [19], pressure transients before turbulent burst transition, and laser self-focusing data near the collapse threshold. In each of these candidate domains, the framework's prediction is operational: a system whose locally estimated α is observed to drift toward 1 sits in a regime of diminishing safety margin in the sense of this framework.

L2. Single-event analysis. Each real-world dataset was processed as a single approach trajectory. Multi-event ensemble analysis within each domain is the natural statistical extension and would convert the point estimate α_{est} into a distribution, allowing detection of slow drift in the boundary distance over operational lifetime.

L3. Converse of the LPPLS bridge. Theorem 7 establishes LPPLS ⇒ $\alpha_{est} \leq 0$. The converse is false in general (the European grid event has $\alpha = 0.36$ but is not LPPLS). A finer classification of the saturated region $\alpha < 1$ into LPPLS, linear-relaxation, and mixed sub-regions is an open direction.

L4. Partial-differential-equation lift. The framework is stated for autonomous ordinary differential equations obtained from centre-manifold or slow-manifold reductions of higher-dimensional systems. Lifting Theorems 2, 3, 5, and 6 to genuine PDE settings — with the radial energy as a functional, the falsifier field-resolved, and the Lyapunov functional an entropy — is a natural next theoretical extension. The reciprocal-time analysis carries over verbatim once a suitable scalar energy is identified; the dilation-charge formula acquires an additive correction controlled by the spatial dimension and the function-space regularity, the simplest instance of which is $\Psi^{*\text{heat}}(p, n) = 2/(p - 1) - n/2$ for the semilinear heat equation. A systematic PDE study is left to a companion paper.

L5. Non-autonomous and stochastic extensions. The framework is stated for autonomous deterministic flows. Slow non-autonomy (parameter drift on a timescale much longer than $T^* - t$) and additive stochastic noise on the dynamics (not only on observation) are both natural extensions. Theorem 6 covers the observation-noise case; intrinsic stochastic forcing would require an Itô-calculus version of Lemma 3.1.

8.3. Reproducibility statement

All numerical experiments in Section 5 use the deterministic integrator DOP853 (the order-8 Dormand–Prince algorithm, scipy.integrate.solve_ivp implementation) with relative tolerance 10^{-12} and absolute tolerance 10^{-14}. The reciprocal-time grid is geometric. T^* is extracted from the integrator's terminal event detection on a finite-state-norm threshold; the residual error in T^* is below 10^{-8} in all cases reported. The synthetic vector field of Section 5.1 is fully specified by equation (5.1) and the parameter tuples in Tables 5.1 and 5.2. The European grid dataset is the publicly archived TransNet BW one-second frequency log of January 2019, available at the open repository of [15, 16]. The NASDAQ Composite series is the daily-close data redistributed with the lppls Python package [9], whose ultimate source is the Yahoo Finance public API. The OLS regression is the standard ordinary least squares estimator with no weighting. The noise-robustness check of Section 5.7 uses pseudo-random Gaussian noise with fixed seeds, allowing exact bit-level replication. A companion code repository implementing the full pipeline — synthetic generation, real-data preprocessing, regression, $\Psi(r)$ and $V(r)$ computation, error budgeting, and figure generation — is provided as supplementary material on Zenodo (DOI to be assigned upon acceptance), under a permissive open-source licence, allowing third-party replication of every numerical table and figure in this paper from raw data.

8.4. Outlook

The framework's value lies in the explicit, theorem-controlled map between the local geometry of an autonomous nonlinear flow and a dimensionless distance to its potential self-destruction. The (α, Ψ^*) reference manifold derived from the dilation symmetry provides a domain-independent yardstick: the gap $|\alpha - 1|$, the deviation $\Psi_{meas} - \Psi^*(\alpha_{est})$, the tail coefficient of variation of Ψ, and the empirical SNR jointly characterise where a real engineered system sits relative to the strong-singularity boundary, and how reliably that distance is measured. Populating this manifold with additional real-world systems

— across mechanical, electrical, biomedical, and economic domains — is the natural research program. The bridge theorem to LPPLS (Theorem 7) suggests a complementary diagram in which the present framework's strong-singularity axis is paired with the LPPLS weak-singularity axis, classifying finite-time critical events along both axes simultaneously. Further extensions include the field-resolved PDE lift sketched in L4, the multi-event ensemble statistics noted in L2, and the stochastic-forcing version noted in L5. Each is a self-contained development of the same dilation-symmetry skeleton.

8.5. Falsification roadmap

The framework defines three independent quantitative predictions whose violation would constitute empirical refutation. We collect them here in operational form, with the relevant theorem cross-references, as a structured agenda for future verification.

F1. Boundary approach in controlled systems.

For the three candidate domains identified in Section 6.4 — Semenov chemical runaway, laser self-focusing, and pre-VF heart-rate variability — the framework predicts that the locally estimated α drifts monotonically toward 1 as the control parameter approaches its critical value, with the rate controlled by the local linearisation of the underlying flow. Failure of any of these three predictions on archival or controlled data, with SNR exceeding the threshold of Theorem 6, falsifies the framework as a boundary-approach indicator in that domain. Simultaneous failure across all three would falsify the framework as a general boundary-approach indicator.

F2. Multi-event α distributions.

Limitation L2 notes that the present paper analyses one event per real-world dataset. The natural extension is to compute α for every comparable event in the same physical system and form a distribution. The framework predicts that within a single well-defined system class — for example, all sub-100 mHz frequency excursions in a single continental power grid over one calendar year — the α distribution is sharply peaked around a system-specific value, with the width determined by the Theorem 5 error budget plus physical heterogeneity between events. A broad or multi-modal α distribution within such a class, beyond the Theorem 5 prediction, falsifies the framework's claim that α is a structural property of the system rather than of individual events.

F3. PDE lift of the dilation charge.

For any PDE blow-up in the strong-singularity class with leading homogeneous nonlinearity of degree p in n spatial dimensions, the framework predicts a shifted dilation charge of the form Ψ*^{PDE}(p, n, s) = 2/(p − 1) − χ(p, n, s), with $\chi \geq 0$ controlled by the rescaling that brings the equation to self-similar form. The simplest instance is the semilinear heat equation, for which $\chi = n/2$ in the L^2-energy norm. A direct numerical PDE simulation whose tail $\Psi(r)$ fails to converge to the predicted value within the acceptance band of Theorem 5 falsifies the corresponding χ and, by extension, the framework's PDE applicability for that equation.

Cumulative falsifiability.

F1, F2, and F3 are independent, in the sense that each tests a logically distinct prediction. F1 tests the applicability of α as a boundary-approach indicator; F2 tests the structural rather than event-specific nature of α; F3 tests the PDE lift of the dilation-charge formula. The framework is therefore subject to three independent empirical risk factors, in the sense of Popper [13]. A single decisive failure on any of them refutes the corresponding component; failure on all three would refute the framework as a whole. The roadmap above is intended as an explicit pre-registration of these predictions, in the methodological standard appropriate for a falsifiable theoretical framework.

Declarations

Funding

Not applicable. The author received no funding for this work.

Conflicts of interest / Competing interests

The author declares no competing interests.

Data availability

The Continental European synchronous grid frequency data are publicly archived at the open repository associated with refs. [15, 16]. The NASDAQ Composite daily-close series is distributed with the lppls Python package [9]. The synthetic-flow integration code and the full analysis pipeline are provided as supplementary material on Zenodo upon acceptance.

Code availability

All code used to produce the figures and tables in this paper will be made openly available on Zenodo (DOI to be assigned upon acceptance) under a permissive open-source licence.

References

[1] Scheffer, M., Bascompte, J., Brock, W. A., Brovkin, V., Carpenter, S. R., Dakos, V., Held, H., van Nes, E. H., Rietkerk, M., Sugihara, G.: Early-warning signals for critical transitions. Nature 461, 53–59 (2009).

[2] Dakos, V., Carpenter, S. R., Brock, W. A., Ellison, A. M., Guttal, V., Ives, A. R., Kéfi, S., Livina, V., Seekell, D. A., van Nes, E. H., Scheffer, M.: Methods for detecting early warnings of critical transitions in time series illustrated using simulated ecological data. PLoS ONE 7(7), e41010 (2012).

[3] George, S. V., Kachhara, S., Ambika, G.: Early warning signals of complex critical transitions in deterministic dynamics. Nonlinear Dynamics 112, 17269–17285 (2024). https://doi.org/10.1007/s11071-024-10023-0

[4] Chen, S., Ma, S., Ren, Y., Wang, J.: Early predictor for the onset of chaotic critical transitions in logistic map systems. Nonlinear Dynamics (2025). https://doi.org/10.1007/s11071-025-11919-1

[5] Wissel, C.: A universal law of the characteristic return time near thresholds. Oecologia 65(1), 101–107 (1984).

[6] Giga, Y., Kohn, R. V.: Asymptotically self-similar blow-up of semilinear heat equations. Communications on Pure and Applied Mathematics 38(3), 297–319 (1985).

[7] Johansen, A., Ledoit, O., Sornette, D.: Crashes as critical points. International Journal of Theoretical and Applied Finance 3(2), 219–255 (2000).

[8] Sornette, D.: Why Stock Markets Crash: Critical Events in Complex Financial Systems. Princeton University Press, Princeton (2002).

[9] Filimonov, V., Sornette, D.: A stable and robust calibration scheme of the log-periodic power law model. Physica A 392(17), 3698–3707 (2013).

[10] Arnold, V. I.: Geometrical Methods in the Theory of Ordinary Differential Equations, 2nd edn. Springer, New York (1988).

[11] Guckenheimer, J., Holmes, P.: Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields. Applied Mathematical Sciences, vol. 42. Springer, New York (1983).

[12] Khalil, H. K.: Nonlinear Systems, 3rd edn. Prentice Hall, Upper Saddle River, NJ (2002).

[13] Popper, K.: The Logic of Scientific Discovery. Hutchinson, London (1959).

[14] Hairer, E., Nørsett, S. P., Wanner, G.: Solving Ordinary Differential Equations I: Nonstiff Problems, 2nd edn. Springer Series in Computational Mathematics, vol. 8. Springer, Berlin (1993).

[15] Rydin Gorjão, L., Anvari, M., Kantz, H., Beck, C., Witthaut, D., Timme, M., Schäfer, B.: Open database analysis of scaling and spatio-temporal properties of power grid frequencies. Nature Communications 11, 6362 (2020).

[16] Jumar, R., Maaß, H., Schäfer, B., Rydin Gorjão, L., Hagenmeyer, V.: Database of power grid frequency measurements. arXiv:2006.01771 (2020).

[17] Semenov, N. N.: Some Problems of Chemical Kinetics and Reactivity. Pergamon Press, Oxford (1958). (English translation; original Russian 1954.)

[18] Frank-Kamenetskii, D. A.: Diffusion and Heat Exchange in Chemical Kinetics, 2nd edn. Plenum Press, New York (1969).

[19] Goldberger, A. L., Amaral, L. A. N., Hausdorff, J. M., Ivanov, P. C., Peng, C.-K., Stanley, H. E.: Fractal dynamics in physiology: alterations with disease and aging. Proceedings of the National Academy of Sciences 99(suppl 1), 2466–2472 (2002).

[20] Merle, F., Zaag, H.: Stability of the blow-up profile for equations of the type u_t = Δu + |u|^(p−1) u. Duke Mathematical Journal 86(1), 143–195 (1997).

[21] LaSalle, J. P.: Some extensions of Liapunov's second method. IRE Transactions on Circuit Theory 7(4), 520–527 (1960).

[22] Galaktionov, V. A., Vázquez, J. L.: The problem of blow-up in nonlinear parabolic equations. Discrete and Continuous Dynamical Systems 8(2), 399–433 (2002).