A Singular Fold Model for Capacity-Constrained Dynamics

James Kovalenko james@progressfunction.org June 2026 *Preprint. Not peer reviewed.*

Abstract

We study a minimal singular fast-slow dynamical model for capacity-constrained compounding. The fast subsystem takes the smooth saddle-node normal form, while the slow vector field $g(x) = \alpha - \beta k/x$ is singular at the boundary $x = 0$.

Our results are local and explicit: a saddle-node bifurcation of the fast subsystem (Proposition 1); a conditional friction-scaling law $F \sim (Q - V)^{-1/2}$ along the attracting branch under the coupling $F(x) = k/x$ (Proposition 2); a barrier monotonicity of $L(V) = -F(x^*(V))$ on the congestion regime of the reduced slow flow (Proposition 3); a boundedness equivalence between bounded friction and bounded-away-from-fold trajectories on the reduced flow (Criterion 1); a fold-repulsion and exit-time theorem for the full singular system on strict congestion subsets at *every* fixed fold distance $\delta > 0$ (Theorem 1), with an explicit smallness threshold $\varepsilon \le \varepsilon_0(\delta, r)$; and an isolating-block certificate covering all initial conditions in the slow-manifold tube (Proposition 5). Composing the two (Corollary 1), every congestion-tube trajectory exits the cooperative face without fold approach in explicit time $\lesssim 1/(\gamma_0 \mu_0 \delta \varepsilon)$. The minimum-distance threshold $\delta_{\min}(r)$ is an artifact of the Lyapunov normalization and is removed by a weighted construction.

We do not prove a global sufficiency result on the full singular system, a Krupa–Szmolyan classification, or uniform-in-δ behavior as $\delta \to 0$ at fixed ε (the fold neighborhood in the joint limit). The blow-up regularisation of the $1/x$ singularity is not performed here; we record what a resolution would require.

We derive four fold-consistent CSD predictions for exogenous approach and one barrier prediction (P5) that tests the autonomous certificate directly, with explicit identifiability assumptions and competing mechanisms, and specify a falsification protocol.

Contents

---

1. Status of Claims

We separate proved local results from open singular-perturbation items.

Claim	Status
Saddle-node bifurcation of the fast subsystem (Prop. 1)	**Proved.** Standard normal-form verification.

Claim	Status
Friction scaling $F \sim (Q-V)^{-1/2}$ on attracting branch (Prop. 2)	**Proved**, conditional on coupling $F(x) = k/x$.
Barrier monotonicity on the congestion regime of the reduced flow (Prop. 3)	**Proved** on the reduced slow flow restricted to the attracting branch.
Boundedness equivalence (Crit. 1)	**Proved.** Equivalence between bounded friction and bounded-away-from-fold on the reduced flow.
Fold-repulsion funnel and exit-time bound on strict congestion subsets of the full singular system (Thm. 1)	**Proved** for every fixed $\delta > 0$, with $\varepsilon \le \varepsilon_0(\delta, r)$. Weighted Lyapunov certificate; the $\delta_{\min}(r)$ threshold is removed (Remark 5).
Isolating-block certificate: tube exit only through the cooperative-side face (Prop. 5)	**Proved.** Face-flux computation; covers all initial conditions in the tube.
Rated exit for all tube initial conditions (explicit exit time)	**Proved** (Theorem 1 + Proposition 5; Corollary 1).
Uniform-in-δ certification at fixed ε (joint limit $\delta \to 0$; fold neighborhood)	**Open.** $\varepsilon_0(\delta, r) \to 0$ as $\delta \to 0$; requires blow-up resolution.
Local asymptotic stability of the cooperative equilibrium (x_0, V_0)	**Proved** (Proposition 4). Hurwitz linearisation plus Lyapunov indirect method.
Explicit overlap between Proposition 4's basin and Theorem 1's certified region across $V = V_0 + r$	**Open.** Quantitative matching requires Lyapunov-function scaling not pursued here.
Krupa–Szmolyan jump-point classification	**Open.** Standard hypotheses fail at $x = 0$.
Full-divisor blow-up regularisation	**Open.** Not attempted here; requirements recorded in §6.

Table 1: Status of claims. The fold-repulsion / exit-time theorem and the isolating-block certificate promote the reduced-flow admissibility equivalence to the full singular system at every fixed fold distance. The remaining open items are localized: the joint limit $\delta \to 0$ at fixed ε, and the cooperative equilibrium neighborhood.

2. Introduction

Capacity-constrained systems appear across domains: scientific publishing pipelines, organisations operating under fixed verification bandwidth, AI deployment where generation outpaces evaluation, and ecological regimes near tipping points [7, 8]. Each features a load variable growing against a bounded capacity, with a coordination cost increasing nonlinearly as load approaches capacity.

This paper is a self-contained mathematical result with stated empirical implications. The theorems and empirical predictions below follow from the local model alone, without dependence on any broader framework. We study a minimal local model in which the fast dynamics realise a smooth saddle-node bifurcation and the slow vector field is singular at the bifurcation boundary. The combination is natural for capacity-constrained mechanisms in which per-unit coordination cost diverges as available slack vanishes, but it lies outside the scope of standard smooth fast-slow theorems [3, 10, 15]. Our contribution is to identify a strict congestion regime, bounded away from both the fold and the slow equilibrium, on which a local Lyapunov certificate for fold-avoidance holds for the full singular system. The certificate is parametrically explicit; its reach is an ε-dependent gap $\delta \gtrsim \varepsilon^{1/4}$ rather than any parameter-fixed minimum distance, and the joint limit $\delta \to 0$ at fixed ε is left open.

Section 3 verifies the fast bifurcation and the conditional friction-scaling law. Section 4 analyses the reduced slow flow on the attracting branch, formalises the friction barrier as Lyapunov monotonicity on the congestion regime, and states the boundedness equivalence as a reduced-flow result. Section 5 promotes the admissibility equivalence to a fold-repulsion and exit-time theorem for the full singular system on strict congestion subsets, and shows the limitation is an ε-threshold degenerating as $\delta \to 0$ rather than a parameter-fixed minimum

distance. Section 6 addresses the $x = 0$ singularity and presents the formal blow-up heuristic. Sections 7 and 8 state fold-consistent empirical predictions with explicit competing mechanisms.

---

3. Local Model: Fast Subsystem

3.1 System Definition

Let $x \in \mathbb{R}$ and $V \in \mathbb{R}$. We consider the fast-slow system

$$\varepsilon \dot{x} = f(x, V) := (Q - V) - x^2, \tag{1}$$

$$\dot{V} = g(x) := \alpha - \beta(k/x), \tag{2}$$

with friction coupling $F(x) = k/x$, where $\varepsilon > 0$, $0 < \varepsilon \ll 1$, $Q > 0$, $\alpha > 0$, $\beta > 0$, $k > 0$.

The variable x represents structural slack: the margin between capacity and load. The physical domain is $x > 0$. The fold of the fast subsystem occurs at the boundary $x = 0$.

The fast vector field f is smooth on $\mathbb{R}^2$. The slow vector field g is smooth on $\{x > 0\}$ but singular at $x = 0$: $g(x) \to -\infty$ as $x \to 0^+$. This singularity is the central technical feature of the model; we return to it in Section 6.

3.2 Fast-Subsystem Bifurcation

Proposition 1 *(Saddle-node bifurcation of the fast subsystem).* The fast subsystem $\varepsilon \dot{x} = f(x, V)$ defined by (1) satisfies the standard sufficient conditions for a saddle-node (fold) bifurcation at $(x, V) = (0, Q)$.

Proof. Equilibria of the fast subsystem satisfy $f(x, V) = 0$, giving $x = \pm\sqrt{Q - V}$. The attracting branch $x^*(V) = +\sqrt{Q - V}$ has $f_x(x^*, V) = -2x^* < 0$; the repelling branch $x^-(V) = -\sqrt{Q - V}$ has $f_x(x^-, V) > 0$. As $V \to Q^-$ the branches collide at $(0, Q)$. Verifying the standard saddle-node conditions [1, 2]: $f(0, Q) = 0$, $f_x(0, Q) = 0$, $f_{xx}(0, Q) = -2 \neq 0$, $f_V(0, Q) = -1 \neq 0$. By the normal-form theorem applied to the smooth fast subsystem, it is locally topologically equivalent to $\varepsilon \dot{u} = \lambda - u^2$ near $(0, Q)$. □

Remark 1. Proposition 1 is a normal-form verification for the smooth fast subsystem. It does not by itself provide a fast-slow fold theorem for the full system (1)–(2), because the slow vector field is singular at the fold point. The full-system result is established for compact subsets of the congestion regime in Section 5.

Proposition 2 *(Conditional friction divergence).* Under the coupling $F(x) = k/x$, the friction along the attracting fast-subsystem equilibrium branch satisfies $F(x^*(V)) \to +\infty$ as $V \to Q^-$, with rate $F \sim (Q - V)^{-1/2}$.

Proof. $F(V) = k/x^*(V) = k/\sqrt{Q - V}$, so $\lim_{V \to Q^-} F(V) = +\infty$ with rate $(Q - V)^{-1/2}$. □

Remark 2 *(Scope of universality and what the empirical protocol tests).* The branch scaling $x^* \sim (Q - V)^{1/2}$ is universal for any saddle-node bifurcation under smooth coordinate changes. The critical-slowing-down observables that follow from this geometry, recovery time τ and stationary variance Var, inherit the universal exponent $-1/2$ independent of the friction coupling form (Section 7). The friction exponent of Proposition 2 itself is model-dependent: $F = k/x^a$ would give $F \sim (Q - V)^{-a/2}$, and the inverse coupling $a = 1$ is a structural assumption when x represents available slack and F represents per-unit coordination cost. The $a = 1$ form has a queueing pedigree: for an M/M/1 server the mean sojourn time is $1/(\text{service} - \text{load})$, so per-unit delay scales as k/slack [20], and the hyperbolic divergence is less arbitrary than a bare postulate. The central hypothesis is qualitative rather than exponent-level: that coordination cost *diverges* as slack vanishes, rather than saturating or diverging only logarithmically. This hyperbolic divergence is what places the model outside the smooth Fenichel and Krupa–Szmolyan theory; the model presumes it rather than adjudicates it, and its direct empirical status is the friction-exponent measurement of Proposition 2, distinct from the CSD

protocol. The empirical protocol of Appendix A tests the universal τ and Var exponents, which falsify the saddle-node classification when violated; the friction exponent of Proposition 2 is a separate prediction that would require direct measurement of per-unit coordination cost across the observation window, and is not part of the universal CSD test.

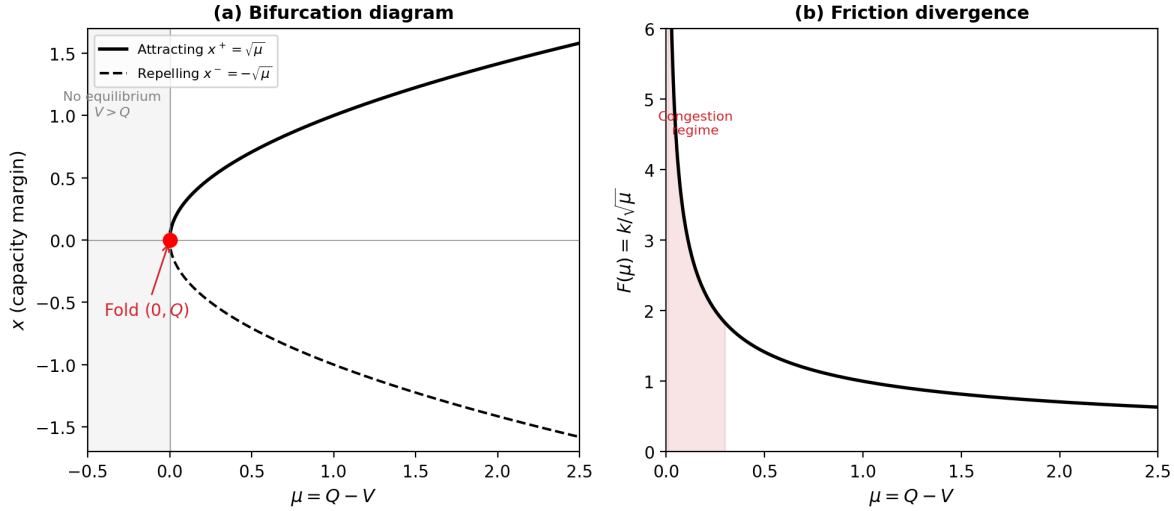

Figure 1. Both panels are plotted against $\mu = Q - V$. (a) Bifurcation diagram of the fast subsystem. The attracting branch $x^*(V) = +\sqrt{Q-V}$ (solid) and the repelling branch $x^-(V) = -\sqrt{Q-V}$ (dashed) collide at the fold $(0, Q)$, i.e., at $\mu = 0$. The left-side shading ($\mu < 0$, equivalently $V > Q$) marks the region where no real equilibrium exists. (b) Friction $F(V) = k/\sqrt{Q-V}$ along the attracting branch, diverging as $V \to Q^-$ with exponent $-1/2$ in $\mu = Q - V$. The left-side shading (small μ) marks the congestion regime $V_0 < V < Q$ where the attracting branch x^* approaches the fold $x = 0$; the divergence in μ corresponds to a hyperbolic divergence $F = k/x$ in the slack variable along the same branch.

---

4. Reduced Slow Flow on the Attracting Branch

We analyse the slow equation $\dot{V} = g(x) = \alpha - \beta(k/x)$ along the attracting branch $x = x^*(V) = \sqrt{Q-V}$ as a one-dimensional flow in V, well-defined for $V < Q$.

4.1 Regime Partition

Define $x_0 = \beta k/\alpha$ and $V_0 = Q - (\beta k/\alpha)^2$. The reduced flow $\dot{V} = \alpha - \beta k/x^*(V)$ partitions the attracting branch into:

- **Cooperative regime** ($x^* > x_0$, equivalently $V < V_0$): $\dot{V} > 0$. Load grows because friction is below baseline growth rate α.
- **Congestion regime** ($0 < x^* < x_0$, equivalently $V_0 < V < Q$): $\dot{V} < 0$. Load contracts because friction overwhelms growth.

The reduced flow has a stable fixed point at V_0.

4.2 Barrier Structure of the Friction

Definition 1 *(Barrier function for fold-avoidance).* On the attracting branch, define

$$L(V) := -F(x^*(V)) = -\frac{k}{\sqrt{Q-V}}.$$

L is smooth and strictly decreasing in V on $(-\infty, Q)$, with $L(V) \to -\infty$ as $V \to Q^-$. The certificate W assembled from L below is a barrier certificate in the sense of [17], the object the control-barrier-function literature [18] uses for safety-critical control.

Proposition 3 *(Barrier monotonicity on the congestion regime).* On the congestion regime $\{V_0 < V < Q\}$, the reduced slow flow satisfies $\dot{V} < 0$, and consequently

$$\frac{dL}{dt} > 0$$

along trajectories of the reduced flow. Equivalently, F is strictly decreasing along reduced trajectories in the congestion regime: the trajectory moves away from the fold.

Proof. Differentiating along the reduced flow:

$$\frac{dL}{dt} = -\frac{d}{dt}\left(\frac{k}{\sqrt{Q-V}}\right) = -\frac{k\dot{V}}{2(Q-V)^{3/2}}.$$

For $V < Q$, the denominator is positive, so dL/dt has the sign of $-\dot{V}$. On the congestion regime, $\sqrt{Q-V} < \beta k/\alpha$ gives $\dot{V} = \alpha - \beta k/\sqrt{Q-V} < 0$, hence $dL/dt > 0$. □

Remark 3. The barrier monotonicity is active in the congestion regime, where the trajectory is close to the fold and the barrier drives it back. In the cooperative regime, $\dot{V} > 0$ and the dynamics are growth-dominated; the trajectory may approach the boundary $V = V_0$ from below.

4.3 Structural Admissibility Equivalence

Criterion 1 *(Boundedness equivalence).* On the reduced slow flow restricted to the attracting branch, the condition

$$\exists\, M < \infty \text{ such that } F(x^*(V(t))) \le M \text{ for all } t \ge 0$$

is equivalent to the trajectory remaining bounded away from the fold:

$$\exists\, \delta > 0 \text{ such that } x^*(V(t)) \ge \delta \text{ for all } t \ge 0.$$

Proof. If $F(x^*(V)) = k/\sqrt{Q-V} \le M$, then $x^*(V) \ge k/M$; take $\delta = k/M$. Conversely if $x^*(V) \ge \delta$, then $F(x^*(V)) \le k/\delta$; take $M = k/\delta$. □

Remark 4. Criterion 1 is a structural equivalence on the reduced flow. The promotion to a fold-repulsion and exit-time statement for the full two-dimensional singular system is given in Section 5, restricted to strict congestion subsets bounded away from both the fold and the cooperative equilibrium.

---

5. Fold Repulsion in the Full System

We promote the reduced-flow admissibility equivalence to a fold-repulsion and exit-time statement for the full singular system on a compact subset of the congestion regime, bounded away from both the fold and the slow equilibrium. The result is a local Lyapunov-decay certificate for fold-avoidance in the full two-dimensional phase plane.

5.1 Setup

We work in fast time t, with the system

$$\dot{x} = (Q - V) - x^2, \qquad \dot{V} = \varepsilon\, g(x), \quad g(x) := \alpha - \frac{\beta k}{x}.$$

Let $x_0 = \beta k/\alpha$ and $V_0 = Q - x_0^2$. For $r > 0$ and $\delta \in (0, x_0)$, define the strict congestion interval

$$J(r, \delta) := [V_0 + r,\, Q - \delta^2],$$

assumed nonempty (i.e., $V_0 + r < Q - \delta^2$, equivalently $\delta^2 < x_0^2 - r$). The slow manifold over J is $x = x^*(V) = \sqrt{Q - V}$, with $x^*(V) \in [\delta,\, \sqrt{x_0^2 - r}]$.

For constants $C > 0$ and $\varepsilon > 0$ to be specified, define the **slow-manifold tube**

$$T_{C,\varepsilon}(r, \delta) := \{(x, V) : V \in J(r, \delta),\, |x - x^*(V)| \le C\sqrt{\varepsilon}\}.$$

The deviation variable is $y := x - x^*(V)$.

What the certificate must control. The nontriviality is not that the reduced slow flow avoids the fold: on the manifold it flows toward V_0, away from $x = 0$. It is that the full-system fast-fiber dynamics, coupled to the slow direction through the diverging $g(x)$, do not open an escape route back toward $x = 0$. Term I supplies the fast restoration, the fast-fiber decay $-y^2(x + x^*)/\varepsilon$ that dominates the residual $\mathrm{III}_{\mathrm{res}}$; the weight w then makes the weighted barrier (Term II) dominate the worst-case cross-coupling (Term IV, through the tube-wide constant $\Gamma_\flat = 2\beta k/\delta - \alpha$). Making the singularity's dynamical load explicit is what separates this from a smooth Fenichel reduction.

5.2 The Lyapunov Function

Define

$$W(x, V) := w \cdot \frac{k}{\sqrt{Q - V}} + \frac{y^2}{2\varepsilon} + b(V)\, y, \qquad b(V) := \frac{g(x^*(V))}{4(Q - V)},$$

with a constant weight $w > 0$ fixed in the proof. The first term is the reduced-flow barrier $-L(V)$ of Definition 1, scaled by w, penalizing approach to the fold; the weight is the construction's key degree of freedom (Remark 5). The second term is a fast-fiber penalty at the Fenichel scaling. The third term is a linear-in-y correction whose role is to eliminate the leading-order cross coupling between the fast and slow directions; it does not involve w. The function W is proper and bounded below on the tube: completing the square, $W = \frac{1}{2\varepsilon}(y + \varepsilon b)^2 + wk/\sqrt{Q - V} - \frac{\varepsilon}{2} b(V)^2$, a sum of a nonnegative fast-fiber term, the positive barrier $wk/\sqrt{Q - V}$, and a benign $O(\varepsilon)$ shift, so its sublevel sets within the tube are compact, and on them the decay estimate is a Lyapunov-decay certificate.

5.3 Structural Parameters

The proof uses the following ε-free constants depending on $(r, \delta, \alpha, \beta, k)$:

- **Congestion gap on the slow manifold.** For $V \in J$, $x^*(V) \le \sqrt{x_0^2 - r} < x_0$, so $g(x^*(V)) \le -\gamma_0$ with

$$\gamma_0(r) := \frac{\beta k}{\sqrt{x_0^2 - r}} - \alpha > 0.$$

- **Fast restoration.** For $(x, V) \in T_{C,\varepsilon}$ with $C\sqrt{\varepsilon} \le \delta/2$, $x + x^*(V) \ge 2\delta - C\sqrt{\varepsilon} \ge \delta$.

- **Cross-coupling bounds.** On the tube, $g(x) < 0$ with $|g(x)| = \beta k/x - \alpha$, provided $x < x_0$ throughout the tube (this is ensured by the constraint on ε stated below). For $C\sqrt{\varepsilon} \le \delta/2$, $x \ge \delta/2$, giving the tube-wide bound

$$|g(x)| \le \Gamma_\flat(\delta) := \frac{2\beta k}{\delta} - \alpha;$$

on the slow manifold over J, where $x^*(V) \ge \delta$, the sharper bound

$$|g(x^*)| \le \Gamma_\sharp(\delta) := \frac{\beta k}{\delta} - \alpha$$

holds. Both are used in Term IV; no $\varepsilon \to 0$ limit is invoked anywhere in the proof.

- **Slow repulsion lower bound.** $(Q - V)^{-3/2} \ge (x_0^2 - r)^{-3/2} =: \mu_0$ for $V \in J$.

5.4 Main Result

Theorem 1 *(Lyapunov decay and fold-repulsion funnel).* Fix $r > 0$ and $\delta \in (0, \sqrt{x_0^2 - r})$ (equivalently $V_0 + r < Q - \delta^2$, so $J(r, \delta)$ is nonempty). Set the slow-term weight

$$w := \frac{\Gamma_\sharp(\delta)\,\Gamma_\flat(\delta)}{k\,\gamma_0(r)}.$$

There exist constants $C > 0$, $\varepsilon_0 > 0$, and $\kappa > 0$ (depending on r, δ and the model parameters) such that for all $\varepsilon \in (0, \varepsilon_0]$,

$$\frac{dW}{dt} \le -\kappa\,\varepsilon \quad \text{on } T_{C,\varepsilon}(r, \delta).$$

In particular, write $L_- := wk/\sqrt{x_0^2 - r}$. For any M in the nonempty window

$$L_- + O(\sqrt{\varepsilon}) \;<\; M \;<\; \min\{\, L_- + C^2/2,\; wk/\delta \,\} - O(\sqrt{\varepsilon}),$$

the funnel $\{W \le M\} \cap T_{C,\varepsilon}$ is nonempty, and every trajectory initialized in it has W strictly decreasing at rate $\kappa\varepsilon$; it therefore never reaches the fast-fiber boundaries $\{|y| = C\sqrt{\varepsilon}\}$ or the fold-side boundary $\{V = Q - \delta^2\}$, stays in the tube with $x \ge \delta/2$, and exits through the cooperative-side face $\{V = V_0 + r\}$ in time at most $(M - L_- + O(\varepsilon))/(\kappa\varepsilon) \lesssim 1/(\kappa\varepsilon)$. The funnel is not forward-invariant: the fifth constraint gives $\dot V = \varepsilon g(x) < 0$ throughout $T_{C,\varepsilon}$, so every trajectory leaves in finite time (Proposition 5). This is a Lyapunov-decay and exit-time certificate; invariance can hold only in the cooperative-equilibrium region of Remarks 4a and 9, which this theorem does not reach.

Proof. Compute $dW/dt = (\partial_x W)\dot x + (\partial_V W)\dot V$ with

$$\partial_x W = \frac{y}{\varepsilon} + b(V), \qquad \partial_V W = \frac{w\,k}{2(Q - V)^{3/2}} + \frac{y}{2\varepsilon\sqrt{Q - V}} + b'(V)\,y + \frac{b(V)}{2\sqrt{Q - V}},$$

and the dynamics $\dot x = -y(x + x^*)$ (using $(Q - V) - x^2 = -(x - x^*)(x + x^*)$) and $\dot V = \varepsilon g(x)$.

Substituting and grouping, with $u := \sqrt{Q - V}$:

$$\frac{dW}{dt} = \underbrace{-\frac{y^2(x + x^*)}{\varepsilon}}_{\text{I}} + \underbrace{\frac{\varepsilon\,w\,k\,g(x)}{2(Q - V)^{3/2}}}_{\text{II}} + \underbrace{\left[\frac{g(x)}{2u} - b(V)(x + x^*)\right] y}_{\text{III}_{\text{lin}}} + \underbrace{\frac{\varepsilon\,b(V)\,g(x)}{2u}}_{\text{IV}} + \underbrace{\varepsilon\,b'(V)\,g(x)\,y}_{\text{V}}.$$

Only Term II carries the weight w; the cancellation below and Terms I, III, V are w-free.

Cancellation of the linear cross term. With $b(V) = g(x^*)/[4(Q-V)]$, write $x = x^*+y$ so $x+x^* = 2u+y$, and use $g(x^*+y) - g(x^*) = \beta k y/[u(u+y)]$:

$$\frac{g(x)}{2u} - \frac{g(x^*)(2u+y)}{4u^2} = \frac{2u[g(x^*+y) - g(x^*)] - y\, g(x^*)}{4u^2} = \frac{y}{4u^2}\left[\frac{2\beta k}{u+y} - g(x^*)\right].$$

Hence

$$\mathrm{III}_{\mathrm{lin}} = \frac{y^2}{4u^2}\left[\frac{2\beta k}{u+y} - g(x^*)\right] =: \mathrm{III}_{\mathrm{res}}.$$

The linear-in-y term has been eliminated; the residual is quadratic.

Bounding the residual. On the tube, $u \geq \delta$ and $u + y \geq \delta/2$ (for $C\sqrt{\varepsilon} \leq \delta/2$). On congestion, $-g(x^*) \leq \Gamma_\sharp(\delta) = \beta k/\delta - \alpha < \beta k/\delta$, combined with $2\beta k/(u+y) \leq 4\beta k/\delta$:

$$|\mathrm{III}_{\mathrm{res}}| \leq \frac{y^2}{4\delta^2}\left[\frac{4\beta k}{\delta} + \Gamma_\sharp\right] \leq \frac{5\beta k\, y^2}{4\delta^3}.$$

Domination by Term I. $|\mathrm{I}| \geq \delta y^2/\varepsilon$. For $\varepsilon \leq 2\delta^4/(5\beta k)$:

$$\frac{\mathrm{I}}{2} + \mathrm{III}_{\mathrm{res}} \leq -\frac{\delta y^2}{2\varepsilon} + \frac{5\beta k\, y^2}{4\delta^3} \leq 0,$$

leaving $\mathrm{I}/2 \leq -\delta y^2/(2\varepsilon)$ available.

Bounding Term IV. On the tube, $|g(x^*)| \leq \Gamma_\sharp$ (the manifold bound, since $x^*(V) \geq \delta$ on J) and $|g(x)| \leq \Gamma_\flat$ (the tube-wide bound, since $x \geq \delta/2$; both negative subject to $x < x_0$ throughout the tube, see the ε constraint below). Hence

$$|\mathrm{IV}| = \frac{\varepsilon\, |g(x^*)|\, |g(x)|}{8(Q-V)^{3/2}} \leq \frac{\varepsilon\, \Gamma_\sharp \Gamma_\flat}{8(Q-V)^{3/2}}.$$

This bound holds at finite ε; no limit is taken.

Negativity of Term II. g is Lipschitz on $[\delta/2, 2x_0]$ with $|g'(x)| = \beta k/x^2 \leq 4\beta k/\delta^2 =: L_g$. For ε small enough that $L_g C\sqrt{\varepsilon} \leq \gamma_0/2$, $g(x) \leq -\gamma_0/2$ on the tube. Hence

$$\mathrm{II} \leq -\frac{\varepsilon\, w\, k\gamma_0}{4(Q-V)^{3/2}}.$$

Combining II and IV.

$$\mathrm{II} + \mathrm{IV} \leq \frac{\varepsilon}{(Q-V)^{3/2}}\left[-\frac{w\, k\gamma_0}{4} + \frac{\Gamma_\sharp \Gamma_\flat}{8}\right].$$

By the choice $w = \Gamma_\sharp \Gamma_\flat/(k\gamma_0)$ in the theorem statement,

$$-\frac{wk\gamma_0}{4} + \frac{\Gamma_\sharp \Gamma_\flat}{8} \leq -\frac{\Gamma_\sharp \Gamma_\flat}{8} < 0$$

for *every* $\delta \in (0, \sqrt{x_0^2 - r})$: no further condition on δ arises. At the pinned weight this gives $\mathrm{II} + \mathrm{IV} \le -\varepsilon\mu_0\Gamma_\sharp\Gamma_\flat/8$ directly, so set $\kappa_0 := \mu_0\Gamma_\sharp\Gamma_\flat/8 > 0$, depending only on $(\alpha, \beta k, \delta, r)$ (split-invariant, Remark 5), so that $\mathrm{II} + \mathrm{IV} \le -\kappa_0\varepsilon$ on the tube.

Bounding Term V. Compute

$$b'(V) = -\frac{\beta k}{8(Q-V)^{5/2}} + \frac{g(x^*)}{4(Q-V)^2},$$

so, writing $u := \sqrt{Q-V} \ge \delta$ for the fiber's fold distance and using the manifold value $|g(x^*(V))| = \beta k/u - \alpha \le \beta k/u$, $|b'(V)| \le \beta k/(8u^5) + |g(x^*(V))|/(4u^4) \le \beta k/(8u^5) + \beta k/(4u^5) = 3\beta k/(8u^5)$ pointwise, and $|g(x)| \le 2\beta k/u$ for the tube state (constraint 1 gives $x \ge u - C\sqrt{\varepsilon} \ge u/2$). Hence, fiberwise,

$$|\mathrm{V}| \le \varepsilon\,|b'(V)|\,|g(x)|\,|y| \le \frac{3C\beta^2k^2}{4}\,\frac{\varepsilon^{3/2}}{u^6}.$$

The proof already establishes $\mathrm{II} + \mathrm{IV} \le -\varepsilon\,\Gamma_\sharp\Gamma_\flat/(8u^3)$ pointwise (before $\mu_0 \le u^{-3}$ is applied). Requiring $|\mathrm{V}| \le |\mathrm{II} + \mathrm{IV}|/2$ fiberwise gives $\varepsilon^{1/2} \le \Gamma_\sharp\Gamma_\flat\, u^3/(12C\beta^2k^2)$, worst at $u = \delta$; hence, once $\varepsilon \le \varepsilon_1 := \big(\Gamma_\sharp\Gamma_\flat\,\delta^3/(12C\beta^2k^2)\big)^2$, we have $|\mathrm{V}| \le |\mathrm{II} + \mathrm{IV}|/2$ fiberwise, so $\mathrm{II} + \mathrm{IV} + \mathrm{V} \le (\mathrm{II} + \mathrm{IV})/2 \le -\kappa_0\varepsilon/2$. Since $\varepsilon_1 \sim \delta^2$, Term V is dominated by the δ^4 constraints below and does not bind the certified reach.

Assembly. Choose $C = 1$. The bounds above require:

- $C\sqrt{\varepsilon} \le \delta/2$, i.e., $\varepsilon \le \delta^2/4$, so that the tube stays in $\{x \ge \delta/2\}$ and the fast-restoration / residual bounds apply.
- $\varepsilon \le 2\delta^4/(5\beta k)$, so that Term I dominates the residual.
- $L_g C\sqrt{\varepsilon} \le \gamma_0/2$, i.e., $\varepsilon \le (\gamma_0/(2L_g))^2$, so that Term II is negative.
- $\varepsilon \le \varepsilon_1 = (\Gamma_\sharp\Gamma_\flat\delta^3/(12C\beta^2k^2))^2 \sim \delta^2$, so that Term V is subdominant (dominated by the δ^4 constraints above).
- $C\sqrt{\varepsilon} \le x_0 - \sqrt{x_0^2 - r}$, equivalently $\varepsilon \le \big(r/(x_0 + \sqrt{x_0^2 - r})\big)^2$, so that $x < x_0$ throughout the tube and the bounds on $|g(x)|$ and $|g(x^*)|$ in terms of $\Gamma_\sharp$ remain valid. (Without this, the tube can cross into $\{x \ge x_0\}$ at the upper-V end of J, where $g(x) \ge 0$ and the closure condition no longer applies.)
- $\varepsilon \le \varepsilon_{\mathrm{win}}(\delta, r)$, so that the M-window is nonempty and the boundary minimum of W sits on the cooperative face (Lemma 1).

Set

$$\varepsilon_0 := \min\left\{\frac{\delta^2}{4},\ \frac{2\delta^4}{5\beta k},\ \left(\frac{\gamma_0}{2L_g}\right)^2,\ \varepsilon_1,\ \left(\frac{r}{x_0 + \sqrt{x_0^2 - r}}\right)^2,\ \varepsilon_{\mathrm{win}}\right\}, \qquad \varepsilon_{\mathrm{win}}(\delta, r) := \min\left\{\frac{16w\delta^5}{3\beta^2k},\ \left(\frac{4w\delta^2}{3\beta C}\right)^2,\ \frac{C^2(x_0^2 - r)^2}{\gamma_0^2}\right\}$$

For $\varepsilon \in (0, \varepsilon_0]$:

$$\frac{dW}{dt} \le \underbrace{\frac{\mathrm{I}}{2}}_{\le 0} + \underbrace{(\mathrm{II} + \mathrm{IV} + \mathrm{V})}_{\le -\kappa_0\varepsilon/2} + \underbrace{(\mathrm{I}/2 + \mathrm{III}_{\mathrm{res}})}_{\le 0} \le -\frac{\kappa_0\varepsilon}{2}.$$

Setting $\kappa := \kappa_0/2$ gives the claimed bound.

Proposition 4 *(Local asymptotic stability of the cooperative equilibrium).* The point $(x_0, V_0) = (\beta k/\alpha,\ Q - (\beta k/\alpha)^2)$ is an asymptotically stable equilibrium of (1)–(2) for every $\varepsilon > 0$.

Proof. Linearisation around (x_0, V_0) with $u = x - x_0$, $\eta = V - V_0$ gives, in fast time,

$$\dot{u} = -2x_0\,u - \eta + O(u^2,\, u\eta), \qquad \dot{\eta} = \frac{\varepsilon\alpha}{x_0}\,u + O(\varepsilon u^2).$$

The Jacobian is $J(\varepsilon) = \begin{pmatrix} -2x_0 & -1 \\ \varepsilon\alpha/x_0 & 0 \end{pmatrix}$, with $\operatorname{tr} J = -2x_0 < 0$ and $\det J = \varepsilon\alpha/x_0 > 0$ for every $\varepsilon > 0$. Both eigenvalues have strictly negative real part (Hurwitz). By Lyapunov's indirect method, (x_0, V_0) is locally asymptotically stable. The eigenvalues separate as $\varepsilon \to 0$: $\lambda_{\text{fast}} \approx -2x_0$ and $\lambda_{\text{slow}} \approx -\varepsilon\alpha/(2x_0^2)$, exhibiting the standard fast–slow time-scale separation. □

Remark 4a *(Matching to Theorem 1).* Proposition 4 supplies a certified neighborhood $N(\varepsilon)$ around (x_0, V_0); Theorem 1 supplies a certified region $T_{C,\varepsilon}(r,\delta)$ on $V \in [V_0 + r,\, Q - \delta^2]$. The neighborhood $N(\varepsilon)$ is forward-invariant under the autonomous flow via the quadratic Lyapunov function of the Hurwitz linearization at the cooperative equilibrium (x_0, V_0) (Proposition 4); on $T_{C,\varepsilon}(r,\delta)$, by contrast, $\dot{V} < 0$ forces finite-time exit, so there Theorem 1 gives a fold-repulsion funnel rather than invariance. Whether $N(\varepsilon)$ can be chosen large enough to overlap $T_{C,\varepsilon}(r,\delta)$ at the lower-V boundary $V = V_0 + r$ depends on the quantitative Lyapunov-function scaling at the cooperative equilibrium, which decreases in the slow direction as $\varepsilon \to 0$. Because the attracting branch is normally hyperbolic near the cooperative equilibrium, where g is smooth, Fenichel's theorem [3] supplies the qualitative matching; only the quantitative overlap range is left as the open item in the Status table.

Summary of structural constants. The proof above uses the ε-free constants depending on $(r, \delta, \alpha, \beta, k)$, the rate constants κ_0 and κ, and the sufficiency threshold ε_0. They are summarised in Table 2.

Constant	Definition	Role
$\gamma_0(r)$	$\beta k/\sqrt{x_0^2 - r} - \alpha$	Lower bound on $-g(x^*(V))$ over $V \in J$; congestion gap on the slow manifold.
$\Gamma_\sharp(\delta)$	$\beta k/\delta - \alpha$	Upper bound on $\lvert g(x^*)\rvert$ on the slow manifold over J.
$\Gamma_\flat(\delta)$	$2\beta k/\delta - \alpha$	Upper bound on $\lvert g(x)\rvert$ on the tube; worst-case cross-coupling magnitude at finite ε.
$\mu_0(r)$	$(x_0^2 - r)^{-3/2}$	Lower bound on $(Q-V)^{-3/2}$ for $V \in J$; slow-repulsion strength.
$L_g(\delta)$	$4\beta k/\delta^2$	Lipschitz constant for g on the tube.
$w(\delta, r)$	$\Gamma_\sharp\Gamma_\flat/(k\gamma_0)$	Slow-term weight closing II + IV for every δ (see Remark 5).
κ_0	$\mu_0\Gamma_\sharp\Gamma_\flat/8$	Tube-wide slow-decay constant (II + IV closure; Remark 5).
κ	$\kappa_0/2$	Net decay rate $dW/dt \le -\kappa\varepsilon$; sets the exit-time bound.

Table 2. Structural constants in the proof of Theorem 1.

Fold-repulsion funnel and exit time. Boundary values of W on $\partial T_{C,\varepsilon}$:

- Upper-V boundary $\{V = Q - \delta^2\}$: $W \ge wk/\delta$, large.
- Lower-V boundary $\{V = V_0 + r\}$: $W \le wk/\sqrt{x_0^2 - r} + C^2/2 + O(\sqrt{\varepsilon})$, bounded.
- Fast-fiber boundaries $\{|y| = C\sqrt{\varepsilon}\}$: $W \ge w(-L(V)) + C^2/2 - O(\sqrt{\varepsilon})$, where $-L(V) = k/\sqrt{Q - V} \ge k/\sqrt{x_0^2 - r}$.

The non-emptiness condition $V_0 + r < Q - \delta^2$ is equivalent to $\delta < \sqrt{x_0^2 - r}$, so the lower-$V$ value of W is strictly less than the upper-V value (the gap grows with w). With $L_- := wk/\sqrt{x_0^2 - r}$, the fast-fiber and fold-side faces sit at $W \ge L_- + C^2/2 - O(\sqrt{\varepsilon})$ and $W \ge wk/\delta - O(\varepsilon)$ respectively, while the

cooperative-side face sits at $W \le L_- + C^2/2 + O(\sqrt{\varepsilon})$. Fix any M in the nonempty window $L_- + O(\sqrt{\varepsilon}) < M < \min\{L_- + C^2/2,\, wk/\delta\} - O(\sqrt{\varepsilon})$. Since $dW/dt \le -\kappa\varepsilon < 0$ on $T_{C,\varepsilon}$, a trajectory initialized in $\{W \le M\} \cap T_{C,\varepsilon}$ has W strictly decreasing, so it never reaches the fast-fiber or fold-side faces and satisfies $x \ge \delta - C\sqrt{\varepsilon} \ge \delta/2$ until it exits. Because $\dot V = \varepsilon g(x) < 0$ throughout the tube, the sublevel set is not forward-invariant: the trajectory leaves through the cooperative-side face $\{V = V_0 + r\}$ in time at most $(M - L_- + O(\varepsilon))/(\kappa\varepsilon) \lesssim 1/(\kappa\varepsilon)$. This is a decay-and-exit certificate rather than an invariance statement. □

The floor location and window nonemptiness are secured by the following lemma, whose constants enter $\varepsilon_{\mathrm{win}}$.

Lemma 1 *(Floor location and window nonemptiness).* With w pinned and $u_+ := \sqrt{x_0^2 - r}$, if $\varepsilon \le \varepsilon_{\mathrm{win}}(\delta, r)$ then the boundary minimum of W over $\partial T_{C,\varepsilon}$ is attained on the cooperative face $u = u_+$, and the M-window is nonempty.

Proof. With $b = g(x^*)/(4u^2)$ and $|g(x^*)| = \beta k/u - \alpha \le \beta k/u$, $|(b^2)'| = |g(x^*)|\,(\beta k/u + 2|g(x^*)|)/(8u^5) \le 3\beta^2 k^2/(8u^7)$ and $|\mathrm{d}|b|/\mathrm{d}u| = (\beta k/u + 2|g(x^*)|)/(4u^3) \le 3\beta k/(4u^4)$. The interior floor is $h(u) := wk/u - (\varepsilon/2)b^2$, nonincreasing once $(\varepsilon/2)\cdot 3\beta^2k^2/(8u^7) \le wk/u^2$, i.e. $\varepsilon \le 16wu^5/(3\beta^2 k)$, worst at $u = \delta$ (first entry of $\varepsilon_{\mathrm{win}}$). The fiber-face floor is the envelope $\varphi(u) := wk/u + C^2/2 - |b(u)|\,C\sqrt{\varepsilon}$, and φ is nonincreasing once $C\sqrt{\varepsilon}\cdot 3\beta k/(4u^4) \le wk/u^2$, i.e. $\varepsilon \le (4wu^2/(3\beta C))^2$, worst at $u = \delta$ (second entry). Both floors are therefore minimized at u_+, where $|g(x^*)| = \gamma_0$, $|b| \le \gamma_0/(4u_+^2)$, and $(\varepsilon/2)b^2 \le \varepsilon\gamma_0^2/(32u_+^4)$; so the boundary minimum of W sits on the cooperative face. For nonemptiness, when $wk(1/\delta - 1/u_+) \ge C^2/2$ the cap is $L_- + C^2/2$ and the fiber-floor correction stays below the $C^2/2$ gap once $\gamma_0 C\sqrt{\varepsilon}/(4u_+^2) \le C^2/4$, i.e. $\varepsilon \le C^2u_+^4/\gamma_0^2$ (third entry); when $wk(1/\delta - 1/u_+) < C^2/2$ the cap is wk/δ and the window width is the fold gap $wk(1/\delta - 1/u_+)$ itself, positive by the same h-monotonicity ($h(\delta) - h(u_+)$ equals the gap minus $O(\varepsilon)$). At the pinned weight the three $\varepsilon_{\mathrm{win}}$ entries scale as δ^3, δ^0, δ^0 (the last two r-degenerate), each weaker than the δ^4 assembly constraints, so the reach $\delta \gtrsim \varepsilon^{1/4}$ is untouched. □

5.5 Isolating-Block Certificate for the Whole Tube

Theorem 1's funnel consists of the sublevel sets $\{W \le M\}$ with M below the fast-fiber and fold-side floors of W; initial conditions near the fold-side face or the fast-fiber boundary lie above those floors and belong to no such set. The following face-flux argument closes that gap: it certifies *every* initial condition in the tube, at the cost of giving no decay rate.

Proposition 5 *(Exit only through the cooperative-side face).* Fix $r > 0$, $\delta \in (0, \sqrt{x_0^2 - r})$, and $C > 0$. Suppose

$$\varepsilon < \varepsilon_{\mathrm{blk}} := \min\left\{ \frac{\delta^2}{4C^2},\ \left(\frac{2C\delta^2}{\Gamma_\flat(\delta)}\right)^2,\ \left(\frac{r}{C\,(x_0 + \sqrt{x_0^2 - r})}\right)^2 \right\}.$$

Then on $T_{C,\varepsilon}(r,\delta)$: the fold-side face $\{V = Q - \delta^2\}$ and both fast-fiber faces $\{y = \pm C\sqrt{\varepsilon}\}$ are strictly inflowing, and $\dot V < 0$ holds uniformly. Consequently every trajectory starting in $T_{C,\varepsilon}$ remains in it, with $x \ge \delta/2$, until it exits in finite time through the cooperative-side face $\{V = V_0 + r\}$. No trajectory starting in the tube approaches the fold before it exits; global non-return is not certified here (Remark 4b).

Proof. The deviation variable satisfies $\dot y = \dot x - (x^*)'(V)\,\dot V = -y(x + x^*) + \varepsilon g(x)/(2u)$. The third constraint keeps the tube inside $\{x < x_0\}$, so $g(x) < 0$ with $|g(x)| \le \Gamma_\flat$ (first constraint: $C\sqrt{\varepsilon} \le \delta/2$, so $x \ge \delta/2$). On the upper fiber face $y = +C\sqrt{\varepsilon}$: $\dot y = -C\sqrt{\varepsilon}(x + x^*) + \varepsilon g(x)/(2u) < 0$, both terms negative, so inflowing. On the lower fiber face $y = -C\sqrt{\varepsilon}$, using $x + x^* \ge \delta$ and $u \ge \delta$:

$$\dot y \ \ge\ C\sqrt{\varepsilon}\,\delta - \frac{\varepsilon\,\Gamma_\flat}{2\delta} \ >\ 0$$

by the second constraint ($\sqrt{\varepsilon} < 2C\delta^2/\Gamma_\flat$) so inflowing. On the fold-side face $\{V = Q - \delta^2\}$: $\dot V = \varepsilon g(x) < 0$, so inflowing. Finally, on the tube $x \le \sqrt{x_0^2 - r} + C\sqrt{\varepsilon} < x_0$, so $\dot V = \varepsilon g(x) \le \varepsilon\, g(\sqrt{x_0^2 - r} + C\sqrt{\varepsilon}) < 0$

uniformly; V therefore decreases at a uniform rate and reaches $V_0 + r$ in finite time. The only non-inflowing face is the cooperative-side face, and while in the tube $x \geq x^*(V) - C\sqrt{\varepsilon} \geq \delta/2$. □

Corollary 1 *(Rated exit for all tube initial conditions).* For $\varepsilon \leq \varepsilon_0(\delta, r)$ and $\varepsilon < \varepsilon_{\text{blk}}$, every trajectory starting in $T_{C,\varepsilon}(r, \delta)$ stays in the tube with $x \geq \delta/2$ and exits through the cooperative-side face $\{V = V_0 + r\}$, without fold approach, in time at most

$$\frac{\sup_T W - \inf_T W}{\kappa\varepsilon} \lesssim \frac{16}{\gamma_0\, \mu_0\, \delta\, \varepsilon}.$$

Proof. Theorem 1 gives $dW/dt \leq -\kappa\varepsilon$ at *every* point of $T_{C,\varepsilon}$, so W decreases at rate at least $\kappa\varepsilon$ along any trajectory while it remains in the tube; Proposition 5 confines every tube trajectory to $T_{C,\varepsilon}$ until it exits through $\{V = V_0 + r\}$. Hence the residence time is at most $(W_0 - \inf_T W)/(\kappa\varepsilon) \leq (\sup_T W - \inf_T W)/(\kappa\varepsilon)$. With $\sup_T W - \inf_T W \leq wk/\delta + O(1/\delta)$ (the $O(1/\delta)$ is the fiber b-correction $\leq \sqrt{2\beta k/5}\, C/(4\delta)$ under the δ^4 constraint), $\kappa = \kappa_0/2$, and $\kappa_0 = \mu_0\Gamma_\sharp\Gamma_\flat/8$, the pinned weight $w = \Gamma_\sharp\Gamma_\flat/(k\gamma_0)$ cancels $\Gamma_\sharp\Gamma_\flat$, giving the split-invariant bound $\lesssim 16/(\gamma_0\mu_0\delta\varepsilon)$. □

Remark 4b *(Division of labor).* Theorem 1's decay rate $dW/dt \leq -\kappa\varepsilon$ holds at every point of the tube, but its rated sublevel funnels form only an $O(\delta^2)$ collar hugging the cooperative face (their V-extent is $\approx C^2(x_0^2 - r)^{3/2}\gamma_0/(\Gamma_\sharp\Gamma_\flat)$, shrinking like $1/w$ at the pinned weight). Corollary 1 removes this restriction: composing the tube-wide rate with Proposition 5's confinement rates *every* tube initial condition, so the certificate covers the whole congestion tube rather than the collar alone. Every trajectory entering the strict congestion tube exits only through the cooperative-side face $\{V = V_0 + r\}$ without fold approach. Beyond that face the trajectory enters $[V_0, V_0 + r)$, a region the full system certifies through neither result: the cooperative restoring force $\dot{V} > 0$ there is a reduced-flow fact that Section 4 declines to promote to the full 2D system. Matching the exit to the basin of Proposition 4 across $[V_0, V_0 + r)$ remains open (Remark 4a, Remark 9).

5.6 Remarks

Remark 5 *(The weight w, and the removal of the $\delta_{\min}$ threshold).* An unweighted construction fixes the slow-term weight at $w = 1$ (slow term $k/\sqrt{Q - V}$); the closure of II + IV then required $\Gamma_\sharp^2 < 2k\gamma_0$, producing a minimum-distance threshold $\delta_{\min}(r) = \beta k/(\alpha + \sqrt{2k\gamma_0(r)})$ and an "uncertified fold neighborhood" $[Q - \delta_{\min}^2, Q]$. That threshold was an artifact of the normalization, with a dimensional tell: the closure comparison involved k *alone*, while the dynamics (1)–(2) depend on (α, β, k) only through α and the product βk: two parameterizations with identical dynamics ($\beta \to \beta/c$, $k \to ck$) received different $\delta_{\min}$. A certificate threshold that depends on how βk is split between β and k cannot be structural.

Scaling the slow term by w removes the threshold at no cost: the cross-term cancellation, Terms I, III, V, and every ε constraint in the assembly are w-free, while the slow-closure margin $\mu_0(wk\gamma_0/4 - \Gamma_\sharp\Gamma_\flat/8)$ for the family with w free grows in w; pinning $w = \Gamma_\sharp\Gamma_\flat/(k\gamma_0)$ globally (no $\max\{1, \cdot\}$, since the II + IV closure gives $-\Gamma_\sharp\Gamma_\flat/8$ at this value regardless and properness needs only $w > 0$) gives the w-independent $\kappa_0 = \mu_0\Gamma_\sharp\Gamma_\flat/8$ and makes the entire Lyapunov function – barrier $\Gamma_\sharp\Gamma_\flat/(\gamma_0\sqrt{Q - V})$, b, and $y^2/(2\varepsilon)$ – depend on $(\alpha, \beta k, \delta, r)$ alone, so W is itself split-invariant and the Corollary 1 bound has leading constant exactly 16 (i.e. $\lesssim 16/(\gamma_0\mu_0\delta\varepsilon)$) on the whole parameter range. The weight does double duty: the same $wk/\sqrt{Q - V}$ barrier whose δ^{-3} growth closes II + IV also pins the boundary minimum of W to the cooperative end (floor-location, §5.4), while the rated funnel's spatial extent shrinks like $1/w$ – immaterial once Corollary 1 rates the whole tube. What survives as the limitation is the ε threshold: $\varepsilon_0(\delta, r) \to 0$ as $\delta \to 0$ (the assembly constraints scale as δ^2 and δ^4), so at fixed ε the certificate still cannot reach the fold. That joint limit, rather than any parameter-fixed distance, is the blow-up frontier of Section 6.

Remark 6 *(No parameter-fixed threshold).* For every fixed $\delta \in (0, \sqrt{x_0^2 - r})$ the certificate closes, so no parameter-dependent minimum distance from the fold survives at the level of the reduced geometry. The binding quantity is $\varepsilon_0(\delta, r)$, whose δ-scaling ($\varepsilon_0 \lesssim \delta^4$ from the residual-domination constraint $\varepsilon \leq 2\delta^4/(5\beta k)$ and the Term-II Lipschitz constraint $\varepsilon \leq (\gamma_0\delta^2/(8\beta k))^2$, both δ^4, the latter with an r-dependent prefactor through γ_0) quantifies how the certificate degenerates as the fold is approached at fixed ε: the assembly

requires $\delta \gtrsim (5\beta k\,\varepsilon/2)^{1/4}$. (The trade-off curve $\delta_{\min}(r)/x_0$ describes the $w = 1$ subset of the construction rather than a property of the dynamics, and is not reported as a prediction.)

Remark 7 *(Relation to Criterion 1).* Criterion 1 of Section 4.3 is the reduced-flow structural equivalence: bounded friction on the reduced flow corresponds to bounded distance from the fold. Theorem 1 and Proposition 5 are its full-system promotion under the strict-congestion truncation: bounded friction with explicit gap δ from the fold is fold-repelling for the full singularly perturbed flow, with every congestion trajectory kept at $x \geq \delta/2$ until it exits, for every $\delta > 0$, provided $\varepsilon \leq \varepsilon_0(\delta, r)$.

Remark 8 *(Why no ε-free version).* The system has two boundary regimes, the fold singularity at $x = 0$ and the cooperative equilibrium at V_0, that each strain a different aspect of any Lyapunov certificate. A basic $W_0 = -L + y^2/(2\varepsilon)$ leaves an uncontrolled linear cross-coupling term; the linear-corrected $W_0 + b(V)y$ eliminates it; and the weighted family $w(-L)+y^2/(2\varepsilon)+b(V)y$ closes the slow-direction comparison for every δ. The resulting threshold might appear irremovable by Lyapunov engineering; it is not: the weight removes it (Remark 5). The obstruction that no choice within this family removes is the degeneration $\varepsilon_0(\delta) \to 0$ as $\delta \to 0$: every constant in the assembly is polynomial in δ, and the fast-restoration margin that powers Term I vanishes at the fold. A certificate uniform in δ at fixed ε would require resolving the singularity itself (Section 6).

Remark 9 *(Open: behavior at $V = V_0$).* Theorem 1 excludes a neighborhood $[V_0, V_0 + r)$ of the slow equilibrium. The reduced flow has a stable fixed point at (x_0, V_0) on the manifold; trajectories in the cooperative regime approach V_0 from below. A full statement combining cooperative-regime convergence to V_0 with congestion-regime repulsion from the fold would require either matching arguments at $V = V_0$ or a global Lyapunov function valid on both regimes. This is not provided here.

Remark 10 *(Open: uniform-in-δ behavior).* The certificates apply at every fixed $\delta > 0$ but with $\varepsilon_0(\delta, r) \to 0$ as $\delta \to 0$. At fixed ε, the assembly requires $\delta \gtrsim \varepsilon^{1/4}$, so dynamics within an $O(\varepsilon^{1/4})$-neighborhood of the fold remain uncertified. Resolving behavior there is the Krupa–Szmolyan blow-up question (Section 6), which remains open. Measured in the load variable $\mu = Q - V$, the weighted Lyapunov certificate reaches to $\mu \gtrsim \varepsilon^{1/2}$ (since $\delta \gtrsim \varepsilon^{1/4}$ and $\mu = \delta^2$ on the fold-side face), whereas for a smooth slow field the elementary reduction reaches $\mu \gtrsim \varepsilon^{2/3}$ before Krupa–Szmolyan blow-up is required. Since $1/2 < 2/3$, at $\varepsilon = 10^{-2}$ the singular field leaves $\mu \lesssim 0.10$ for the blow-up against 0.046 in the smooth case; and unlike the smooth case, where the blow-up resolves that region, here it is left open. The obstruction of Section 6 is therefore a worsening of the smooth case rather than a certificate artifact.

---

6. The Singularity and Its Open Resolution

The full fast-slow system (1)–(2) is singular at $x = 0$ because $g(x) = \alpha - \beta k/x$ blows up there. Standard smooth fast-slow theorems (Fenichel's invariant manifold theorem [3] and the Krupa–Szmolyan classification of fold points [4]) require the slow vector field to be smooth in a neighbourhood of the fold; this hypothesis fails for (2). Application of the Krupa–Szmolyan apparatus to the present system therefore requires either regularising the singular slow vector field, or extending the Krupa–Szmolyan analysis to singular vector fields of the form $\alpha - \beta k/x$. Neither is performed in this paper. Multiplying the slow equation by x recasts the system in impasse form $A(x)\dot{z} = F(z)$ with impasse set $\{\det A = 0\} = \{x = 0\}$, placing it in the constrained-system programme [21], [22] and its singular-perturbation extension [23], [24]; the non-smooth alternative is the blow-up regularization of [26]. The desingularization below is that programme's standard construction rather than a device introduced here.

The standard smooth fold blow-up is the $(1, 2, 3)$-weighting in (x, μ, ε) with fold-centred coordinate $\mu = Q - V$ [4], whose classical inner scales are $x \sim \varepsilon^{1/3}$, $\mu \sim \varepsilon^{2/3}$; it does not accommodate the $1/x$ singularity directly. The canonical impasse desingularization [21] rescales time by $ds = dt/x$ (multiplying through by $\det A = x$, an orbital equivalence on $x > 0$), giving the polynomial system $dx/ds = x(\mu - x^2)$, $d\mu/ds = \varepsilon(\beta k - \alpha x)$, in which the $1/x$ term has resolved before blow-up. Quasi-homogeneous balance with $\deg x = 1$, $\deg \mu = 2$ makes $x\mu$ and x^3 both degree 3, forcing weight 4 on ε; the resulting $(1, 2, 4)$ divisor field $\bar{x}' = \bar{x}\bar{\mu} - \bar{x}^3$, $\bar{\mu}' = \bar{\varepsilon}\,\beta k$ is regular on the whole divisor including $\bar{x} = 0$, which is what the $1/x$ term requires. Its inner

region sits at $x \sim \varepsilon^{1/4}$, $\mu \sim \varepsilon^{1/2}$, matching the certificate frontier $\delta \gtrsim \varepsilon^{1/4}$, $\mu \gtrsim \varepsilon^{1/2}$ (Remark 10's 1/2-vs-2/3 comparison is $2/a$ at $a = 4$ against the smooth $a = 3$). A Krupa–Szmolyan jump-point classification would still require verification of the blown-up vector field across all three charts (K_1, K_2, K_3), regularity on the entire exceptional divisor including chart transitions, the divisor-to-original-time bookkeeping (the rescaling degenerates on $\{x = 0\}$), and rigorous determination of the weights. None of this is performed here; the Section 1 Status of Claims table records the blow-up question as open. Closing this item is the natural follow-up to the present paper and would, if successful, resolve the uniform-in-δ behavior at fixed ε left open in Section 5 (Remark 10).

This programme is active. The equivalence between impasse dynamics and multiple-timescale singular perturbation is developed in [23], [24], and a recent planar parasite–host model [25] resolves a frequency-dependent extinction singularity by the blow-up method, with a regime split ($\beta - d$, equivalently $\mathcal{R}_0 - 1$, of order 1 versus order ε) parallel to the cooperative/congestion partition here. Those analyses are qualitative – orbit passage, homoclinic families, and the canard-like or *funneling* trapping of orbits near the repelling manifold [25] – in ecological and epidemiological settings, with order-of-magnitude ($O(1)$, $O(\varepsilon)$) scales and, at most, a LaSalle argument rather than a finite-ε estimate. The present paper contributes the quantitative finite-ε certificate in the capacity-constrained setting: the funnel here is a sublevel set of an explicit Lyapunov function with a written-down decay rate and exit-time bound, a distinct object from the qualitative canard funneling of [25], and the pole-field blow-up itself is left open.

---

7. Empirical Predictions and Identifiability

Scope and relation to the certificates. Theorem 1 and Proposition 5 certify fold repulsion and finite-time exit on the strict congestion subsets $J(r, \delta) = [V_0 + r,\, Q - \delta^2]$ for *every* fixed $\delta > 0$, with the smallness threshold $\varepsilon_0(\delta, r)$ degenerating as $\delta \to 0$ (Remark 10). The predictions below concern $V \to Q^-$, the pre-transition regime under *exogenous* forcing. Proposition 4 supplements the certificates with local stability of the cooperative equilibrium (x_0, V_0); the matching between Proposition 4's basin and the certified tube is open (Remark 4a). Predictions P1–P4 do not depend on Theorem 1, Proposition 4, or Proposition 5; prediction P5 below tests the barrier certificate directly.

The division of labor is structural. The certificates concern the *autonomous* slow dynamics: the friction barrier of Proposition 3 lifts to the full singular flow, so endogenous dynamics near the slow manifold cannot reach the fold from any fixed distance (at sufficiently small ε); off-manifold transients lie outside the certificate. The empirical predictions P1–P4 are about *exogenous* fold approach: a process outside the autonomous slow equation (external load growth, structural capacity reduction, quasi-static perturbation) drives V toward Q. The CSD signatures are derived from the saddle-node linearisation of the fast subsystem around the attracting branch $x^*(V) = \sqrt{Q - V}$, which exists for every $V < Q$; the linearisation requires none of the certificates and applies on the whole branch.

The two results combine to tell the same story from two sides: autonomous dynamics near the slow manifold cannot reach the fold within the certified tube (the friction barrier holds), and when something else forces the system toward Q, the pre-transition signature is fold-consistent CSD (P1–P4). The blow-up of Section 6 would extend the certified region uniformly in δ; pending that, the present split is the conservative accounting of what each result establishes.

We state four fold-consistent predictions for exogenous approach (P1–P4), with explicit identifiability and competing mechanisms, and one barrier prediction for the autonomous regime (P5).

P1 (thresholded transition). The regime transition is thresholded rather than gradual. *Falsified if:* observed transitions are smooth and proportional to load. *Competing mechanisms to rule out:* gradual capacity degradation, smooth performance decline.

P2 (fold signature). Pre-transition dynamics are dominated by real-eigenvalue (saddle-node) signatures rather than complex-eigenvalue (Hopf) signatures. *Falsified if:* observed pre-transition variability shows cyclical instability with complex eigenvalue dominance. *Competing mechanisms to rule out:* oscillatory bottlenecks, periodic forcing, delay-induced oscillation.

P3 (delayed recovery, conjectural). The transition is abrupt and recovery is delayed: returning to admissible operation requires V significantly below the collapse threshold. *Falsified if:* recovery occurs at the same load level as collapse with no delay or overshoot. *Competing mechanisms to rule out:* hysteresis from independent causes, memory effects in load measurement. *Promotion conditions:* P3 would move from conjectural to derived if a slow-equation augmentation modelling structural damage (state-dependent decrement of Q following a fold crossing) were specified and shown to produce a hysteresis loop with a quantitatively predicted minimum separation between collapse and recovery loads.

P4 (fold-consistent scaling). Near transition, recovery time and stationary fluctuation variance scale with the same fold exponent:

$$\tau \propto |\mu|^{-1/2}, \qquad \mathrm{Var} \propto |\mu|^{-1/2},$$

where $\mu = Q - V$. *Falsified if:* observed exponents are systematically different from $-1/2$ for both observables, or if τ and Var exhibit systematically different exponents. *Competing mechanisms to rule out:* nonnormality, delay-induced critical slowing, Hopf proximity (giving $\tau \propto |\mu|^{-1}$), noise-driven near-bifurcation behaviour without underlying saddle-node.

Derivation of P4. Linearising the fast subsystem around the attracting branch, $\delta\dot{x} = f_x(x^*, V)\,\delta x = -2x^*\,\delta x = -2\sqrt{\mu}\,\delta x$. The relaxation eigenvalue is $\lambda(\mu) = 2\sqrt{\mu}$. Standard Ornstein–Uhlenbeck analysis under additive noise of constant intensity D gives:

$$\tau(\mu) = \frac{1}{\lambda(\mu)} = \frac{1}{2\sqrt{\mu}}, \qquad \mathrm{Var}(\delta x) = \frac{D}{2\lambda(\mu)} = \frac{D}{4\sqrt{\mu}}.$$

Both observables share the exponent $-1/2$. This exponent is the universal saddle-node CSD signature; it does not depend on the friction-coupling form $F = k/x^a$. The friction-exponent prediction of Proposition 2 (which does depend on a) is a distinct claim requiring direct measurement of per-unit coordination cost and is not tested by the protocol of Appendix A. The ratio $\mathrm{Var}/\tau = D/2$ is independent of μ and provides an internal cross-check: under the fold mechanism with constant-intensity noise, the two CSD indicators must scale with the same exponent. Departures from this proportionality indicate either non-constant noise intensity or a different bifurcation class.

The empirical protocol for testing P4 is specified in Appendix A.

Remark 11. We use *fold-consistent* prediction rather than *fold* prediction to mark that the observed patterns are consistent with the fold model but do not uniquely identify it without ruling out the listed competitors.

P5 (barrier: autonomous relaxation). Predictions P1–P4 concern exogenous fold approach and do not test Theorem 1. The friction-barrier certificate makes a separate, testable prediction. In the strict congestion regime $V_0 < V < Q$ the autonomous slow flow points *away* from the fold, toward the standoff V_0; so a capacity-constrained system in that regime, with exogenous drivers held stationary, relaxes toward V_0 rather than drifting to Q. *Falsified if:* the slow variable is observed in autonomous monotone approach $V \to Q$ through the congestion regime while the exogenous controls (α, β, k, Q) are demonstrably stationary across the window; the falsifying observation must lie in the certified region $[V_0 + r,\, Q - \delta^2]$ with the state near its slow manifold, excluding the uncertified sliver $\mu \lesssim \varepsilon^{1/2}$ of Remark 10 and far-off-manifold transients.

Operationalisation (the identification bottleneck). The barrier is compatible with fold approach *under forcing*; only autonomous approach falsifies it, so a valid test must certify driver-stationarity. This is the same exogenous-drive-versus-autonomous-drift distinction that §A.6 already carries for P1–P4. Two designs are admissible: (a) a naturally driver-stationary window, in which the controls are measured or argued constant and the sign of $\dot{V}$ in the congestion regime is read directly; (b) a matched-control comparison of $\dot{V}$ under a known-forcing condition against a no-forcing condition, the barrier predicting relaxation only in the latter. Absent driver-stationarity the observation is uninformative about the barrier. P5 is therefore falsifiable in principle and demanding in practice: it states the observable that makes the certificate earn its empirical keep, and defers closing the operationalisation to the identification programme of §A.6.

8. Discussion

8.1 What the Model Contributes

1. **An explicit finite-ε certificate (principal contribution).** Theorem 1 and Proposition 5 give a quantitative Lyapunov/barrier funnel with every constant written down: on strict congestion subsets the flow is fold-repelling with reach $\delta \gtrsim \varepsilon^{1/4}$ (load $\mu \gtrsim \varepsilon^{1/2}$), net decay rate $\kappa = \kappa_0/2$, and a closed-form exit-time bound $\lesssim 1/(\kappa\varepsilon)$. The constrained/impasse and singular-perturbation literature treats this class of singularities by qualitative blow-up (phase portraits, orbit passage, canards); a finite-ε certificate with explicit reach and exit time is the methodological difference here.
2. **A fold on an impasse set.** Multiplying the slow equation by x puts the system in impasse form $A(x)\dot{z} = F(z)$ with impasse set $\{\det A = 0\} = \{x = 0\}$; the configuration situates within the constrained-system and impasse-plus-singular-perturbation programme [21]–[24], and the $ds = dt/x$ desingularization of §6 is that programme's canonical move rather than a new device. What is specific here is the pole field $g(x) = \alpha - \beta k/x$: its desingularized branch flow $\dot{x} = (\beta k - \alpha x)/(2x^2)$ carries a double pole with fixed repelling sign, a different local object from the impasse normal forms classified in [22].
3. **The reach exponent as a sharpness signal.** The load reach $\mu \gtrsim \varepsilon^{1/2}$, against $\varepsilon^{2/3}$ for a smooth slow field, coincides after §6 with the $(1, 2, 4)$-versus-$(1, 2, 3)$ quasi-homogeneous inner scales – evidence that the reach is set by the pole rather than by the Lyapunov family.
4. **A self-organized operating margin and dynamic fold-shielding.** The stable standoff from the critical point $x = 0$ is $x_0 = \beta k/\alpha$ (with $V_0 = Q - x_0^2$), the growth-to-friction ratio. Within the certified tube (fixed δ, $\varepsilon \le \varepsilon_0$ and $\varepsilon < \varepsilon_{\mathrm{blk}}$) the autonomous flow is fold-repelling, so the fold is reachable only under exogenous forcing (P5). The claim is conditional: off-tube transients, the $\mu \lesssim \varepsilon^{1/2}$ sliver, and $[V_0, V_0 + r)$ are uncertified, and with the drivers (Q, α, β, k) frozen the Ashwin [16] B-tipping notion, which needs a drifting parameter, does not literally apply.
5. **The certified-versus-open boundary.** The certificate of item 1 rests on the admissibility equivalence (barrier monotonicity, Proposition 3; Criterion 1) and an isolating-block certificate covering all tube initial conditions; its threshold $\varepsilon_0(\delta, r)$, with δ^4 scaling, locates the certified-versus-open boundary at the joint limit $\delta \to 0$, ε fixed.
6. **A conditional scaling law.** The friction-divergence rate $F \sim (Q - V)^{-1/2}$ on the attracting branch is a derived, conditional, falsifiable prediction (Proposition 2); the underlying $a = 1$ coupling has a queueing pedigree [20].

Convergent rediscovery. The same fold-with-CSD structure has been identified independently in ecological tipping-point research (slow-recovery indicators as generic precursors of nearby catastrophic shifts [7], rising variance as a leading indicator of ecological transition [9]) and in climate-bifurcation detection (degenerate fingerprinting applied to the Atlantic Meridional Overturning Circulation [8]), within the critical-transitions framework [11] whose critical-slowing-down lineage traces to [19]. The friction-barrier mechanism and the Lyapunov certificate developed here are framework contributions specific to the capacity-constrained setting; the CSD scaling laws they imply belong to a longer convergent tradition.

8.2 Open Issues and Scope

1. The full fast-slow system is singular and the Krupa–Szmolyan classification is not applied as a theorem; the certificates of Section 5 hold at every fixed $\delta > 0$ but with $\varepsilon_0(\delta, r) \to 0$ as $\delta \to 0$, so uniform-in-δ behavior at fixed ε remains open.
2. The Krupa–Szmolyan classification is not performed; Section 6 records only what a resolution would require.
3. Theorem 1 excludes a neighborhood of the cooperative equilibrium V_0 via the parameter r; matching analysis between cooperative-regime convergence and congestion-regime repulsion is not provided.
4. Domain-specific application requires operationalising slack x, load V, capacity Q, the friction coupling, and ruling out the competing mechanisms of Section 7.

The rolling-window variance estimator used in empirical applications *under*-estimates the stationary variance near the fold (by up to $\approx 2\times$ at $\mu = 0.02$ in the simulation of Appendix B.5), precisely the pre-transition regime in which the variance early-warning signal is read; empirical users must either calibrate the window length against system relaxation times or replace the rolling-window estimator with an AR(1)-based variance reconstruction, which corrects the relaxation-time under-estimate.

A candidate embedding of this local model into a wider verification-geometric framework is developed in a separate document; the present paper does not depend on that embedding for any of its theorems or empirical predictions.

8.3 Distinguishing Predictions and Scope of Novelty

What this framework predicts that standard saddle-node CSD does not (with reference points [5, 6] in the standard CSD literature)**.**

1. *Friction-divergence exponent $-1/2$ under inverse coupling.* Standard CSD predicts critical slowing down and variance inflation near a saddle-node but is silent on per-unit coordination cost. Proposition 2 derives $F \sim (Q - V)^{-1/2}$ as a conditional prediction tied to the structural assumption $F = k/x$. A different coupling $F = k/x^a$ would give exponent $-a/2$, so the specific value is empirically informative about coupling form.
2. *Friction-barrier inadmissibility.* Proposition 3 and Criterion 1 establish that, on the reduced slow flow, endogenous dynamics are *repelled* from the fold in the congestion regime. Standard CSD treats the fold as a destination of slow drift; this framework treats it as accessible only under exogenous forcing past a parameter-dependent barrier.
3. *Cross-observable consistency (σ^2/τ flatness).* The two CSD observables co-scale with exponent $-1/2$ in the present framework, so their ratio is μ-independent. A systematic μ-dependence of σ^2/τ falsifies the constant-intensity-noise assumption *before* it falsifies the fold classification.
4. *Three-model falsification structure.* The M1/M2/M3 design distinguishes between "no divergence," "fold-exponent divergence," and "generic power-law divergence." Standard CSD diagnostics typically report only the presence of slowing down rather than the discrimination between exponents.

What this framework does not claim to do better than standard CSD. Detection rates for approaching transitions; sensitivity to noise-driven near-bifurcation behaviour; classification at the fold itself (open, see §6); behaviour beyond the inverse-coupling structural assumption $F = k/x$. (The $\delta_{\min}(r)$ trade-off curve is not a distinguishing prediction: the threshold is an artifact of the Lyapunov normalization rather than a property of the dynamics; see Remark 5.)

Pre-registered failure conditions. Apply the protocol of Appendix A to an observational dataset for which: $\mu = Q - V$ is operationalisable with proxies meeting the standards of §A.6; $n \geq 30$ paired (μ_i, τ_i) observations are available; $\mathrm{CV}_\mu \leq 0.20$ and $\mathrm{CV}_\tau \leq 0.20$; the system is in a pre-transition regime with monotone-decreasing μ.

Saddle-node classification falsified if any of the following hold:

(i) M2 (linear, no divergence) is preferred over M1 by both AIC and BIC, with corroboration per A.3.

(ii) A Wald test of H_0: slope $= -1/2$ rejects at $\alpha = 0.05$ on the log-log fit of τ vs μ in two independent datasets meeting the protocol preconditions (or in two operationalisations of the same series; §A.6).

Conditions (i) and (ii) falsify the saddle-node classification of the local mechanism.

Test-voiding condition (variance branch). If the σ^2/τ log-log slope is significantly non-zero at $p < 0.01$ in both of two independent perturbation subsets of the same series, the constant-noise-intensity assumption underlying the OU *variance* has failed, which voids the variance branch of the test rather than falsifying the classification. Because $\tau = 1/\lambda$ carries no diffusion constant, the recovery-time branch remains admissible, and classification inference is then restricted to the τ branch.

The two-dataset requirement in (ii) applies symmetrically to support: a claim of fold-consistent support likewise requires the $-1/2$ exponent to replicate across two independent datasets (or two operationalisations

of one series; §A.6), rather than a single near$-1/2$ fit, since operationalisation error in the μ proxy can manufacture spurious support as readily as spurious rejection. The symmetry is deliberate: a protocol resting on falsifiability discipline should not admit support on a weaker evidential gate than it demands for rejection. Single-dataset results are reported as suggestive rather than confirmatory, and the competing-mechanism checks of §A.3 apply on top of the two-dataset gate rather than in place of it. The classification falsifiers (i) and (ii) do not separately falsify the friction-coupling assumption $F = k/x$ (Proposition 2), which the CSD protocol does not directly test. A test of the coupling form requires direct measurement of per-unit coordination cost F across the observation window and a log-log fit of F vs μ; the predicted slope is $-a/2$ under $F = k/x^a$, with $a = 1$ being the framework's structural assumption. Falsification of $a = 1$ does not falsify the saddle-node classification; the two failure modes are tested by separate empirical operations, and either may obtain without the other.

Appendix A: Empirical Protocol for Testing P4

A.1 Purpose

A protocol for testing $\tau \propto |\mu|^{-1/2}$ and Var $\propto |\mu|^{-1/2}$ near a capacity-constrained transition, where $\mu = Q - V$.

A.2 Variable Mapping

Control parameter μ: the margin between capacity and load, $\mu = Q - V$. Recovery time τ: time to return to pre-perturbation state after a small shock. Variance σ^2: variance of the system's key output metric over a rolling window.

A.3 Primary Inference: Log-Log Regression

Under multiplicative observation noise, $y_{\text{obs}} = y_{\text{true}}(1 + \epsilon)$ with ϵ small, $\log y_{\text{obs}} \approx \log y_{\text{true}} + \epsilon$. The log-log regression of an observable against μ is therefore correctly specified: the slope estimates the exponent directly, with approximately Gaussian residuals.

For τ and σ^2, the fold predictions in log-log space are

$$\log \tau = a_\tau - \tfrac{1}{2} \log \mu, \qquad \log \sigma^2 = a_\sigma - \tfrac{1}{2} \log \mu.$$

Fit by OLS in log-log coordinates:

$$\hat{s} = -\widehat{\text{slope}}(\log y_{\text{obs}} \text{ vs } \log \mu_{\text{obs}}),$$

with the 95% CI obtained from the regression standard error on the slope.

Model comparison. Fit three models:

- M1 (fold): slope constrained to $-1/2$; intercept by OLS in log-log space.
- M2 (linear, no divergence): $y = c\,\mu + d$ in linear-linear space.
- M3 (free exponent): slope unconstrained in log-log space.

Inference structure. M1 vs. M3 is a nested comparison (M1 imposes slope $= -1/2$). The appropriate test is a Wald test on the log-log slope: reject H_0: slope $= -1/2$ if $|\widehat{\text{slope}} + 1/2|/\text{SE}(\widehat{\text{slope}}) > z_{1-\alpha/2}$ at the chosen significance level α, equivalently if the $(1 - \alpha)$ CI for the slope excludes $-1/2$. M1 vs. M2 is non-nested (different functional forms); compare via AIC/BIC evaluated on the observable in its original space (residuals $y_{\text{obs}} - \hat{y}$, not $\log y_{\text{obs}} - \log \hat{y}$), so both models are scored on the same loss. For confirmatory M2-preference the AIC/BIC ranking is corroborated by out-of-sample predictive loss or a Vuong non-nested test; the original-space AIC/BIC is otherwise descriptive (B.2).

Decision rule:

- M2 dominates M1 by both AIC and BIC, corroborated per the inference-structure paragraph ⇒ P4 falsified (no fold-consistent divergence).
- Wald test fails to reject H_0 at $\alpha = 0.05$ AND M1 preferred over or comparable to M2 by AIC/BIC ⇒ fold-consistent scaling supported in this dataset (suggestive; confirmatory support requires two-dataset replication, §8.3).
- Wald test rejects H_0 at $\alpha = 0.05$ ⇒ saddle-node classification falsified at this significance level; the framework's CSD prediction is rejected for this dataset. The result does not by itself identify the alternative bifurcation class (Hopf gives slope -1, cusp gives a different exponent), nor does it speak to the friction-coupling form a in $F = k/x^a$, which the protocol does not directly test.

Cross-check. Under P4, σ^2/τ is constant in μ. Plot $\hat{\sigma}_i^2/\hat{\tau}_i$ against μ_i in log-log space and test for slope = 0. Systematic dependence on μ falsifies the constant-noise-intensity assumption underlying the OU derivation, even if the individual scaling fits look consistent.

Competing-mechanism checks. Before accepting M1: confirm the 95% CIs of $\hat{s}_\tau$ and $\hat{s}_\sigma$ both contain 1/2 and are well separated from 1 (Hopf gives slope magnitude 1); verify monotone rather than transient-amplified recovery (nonnormality check); estimate dominant timescales and check they do not match system delays (delay check).

Alternative linearisation forms. The algebraically equivalent forms $\tau^{-2} = a\mu + b$ and $\sigma^{-4} = a'\mu + b'$ acquire β-dependent heteroscedasticity under multiplicative noise on the observable, leading to anti-conservative profile-likelihood CIs. The log-log form above is preferred for inference; the inverse-square forms remain useful for visual inspection because they map fold-consistent data to straight lines in linear-linear axes.

Caveat on the collapse forecast. The M1 fit yields a forecast collapse location as the margin at which the fitted curve reaches a chosen recovery-time threshold τ_c: $\hat{\mu}^*(\tau_c) = \exp(2(a_\tau - \log \tau_c))$ under the fold slope $-1/2$. The threshold must be stated: the form $\exp(2a_\tau)$ corresponds to $\tau_c = 1$ in whatever units τ is measured, a hidden unit convention. This is a point estimate with intrinsic identifiability limits: the observed Q is a proxy for the true dynamical capacity, subject to systematic measurement bias, and a μ^* derived from the fit absorbs that bias along with the model. The decomposition $Q_{\text{true}} = Q_{\text{observed}} + \text{bias}$ is *internal to the protocol* and cannot, by itself, distinguish structural slack from regression artifact. Independent evidence on the proxy error (e.g., an audit measurement of true capacity, or a controlled perturbation experiment) is required to license a structural reading of $\hat{\mu}^*$. Without that evidence, treat $\hat{\mu}^*$ as a calibrated warning rather than a forecast.

A.4 Protocol Steps

1. Operationalise $\mu = Q - V$; verify monotonicity over the observation window.
2. Measure τ from perturbation events. Minimum $n \geq 10$ (CI coverage is approximately nominal from $n = 10$ at the noise levels of Appendix B; Table 5); $n \geq 30$ recommended for discriminating power against M2.
3. Plot $\log \tau_i$ vs. $\log \mu_i$. If σ^2 data are also available, plot $\log \sigma_i^2$ vs. $\log \mu_i$ in parallel.
4. Fit M1, M2, M3 per A.3.
5. Apply the three-way decision rule, the σ^2/τ flatness cross-check, and the competing-mechanism checks.
6. If M1 is supported in this dataset (suggestive pending two-dataset replication, §8.3), report $\hat{s}_\tau$ and $\hat{s}_\sigma$ with regression 95% CIs and the cross-check slope. Report $\hat{\mu}^*$ with the structural-reading caveat from A.3.

Independence of recovery-time observations. The Wald inference of A.3 assumes the residuals of the log-log fit are uncorrelated. For τ, this requires the perturbation events generating successive (μ_i, τ_i) pairs to be independent. When perturbations are separated by at least one relaxation time and the system returns to baseline between events, the assumption holds. When the system exhibits slow trends, autocorrelated state, or perturbations spaced more closely than the relaxation time, recovery-time observations may be serially correlated and the OLS standard errors are biased downward (anti-conservative on the Wald test). Two corrections: (a) cluster-robust standard errors with clustering by perturbation epoch; or (b) block bootstrap

with block length matched to the estimated autocorrelation timescale. Report the larger of the OLS and corrected CIs.

A.5 Variance-Only Variant and Autocorrelation Warning

If recovery-time data are unavailable, the variance test alone can be used. Rolling-window variance estimates are serially correlated by construction: adjacent windows share $w-1$ observations. OLS standard errors on $\sigma^{-4} = a'\mu + b'$ are therefore biased downward, and significance tests are anti-conservative. Two corrections: (a) Newey–West standard errors with bandwidth equal to w; or (b) subsample at interval w so windows do not overlap, reducing effective sample size by a factor of w. If the scaling result survives both corrections, it is robust. Recommended window length $w \in [5, 10]$, sensitivity-checked against $w = 3$ and $w = 15$.

A.6 Identification and Operationalisation Standards

The regression estimates a statistical association between μ and the CSD observables rather than a causal structural law. The structural interpretation (fold scaling) requires that variation in μ is not driven by an unobserved common cause that independently affects τ or σ^2. In domains where controlled perturbation experiments are feasible, exogenous variation in μ can be introduced directly. In observational settings, the assumption that μ is the dominant channel should be stated and its plausibility assessed domain by domain.

Because Q (capacity) and V (load) are rarely directly observable, the operationalisation of these proxies is itself a methodological commitment that must be discharged before the protocol is run. Without minimum standards the test risks measurement-construct circularity: the operationalisation choice may itself be informed by the prediction being tested, and the slope estimate $\hat{s}$ inherits that dependence.

Minimum operationalisation standards.

1. *Pre-registration.* The proxies for Q and V must be specified, with the source data identified, before the data are examined. Adjusting the proxies after seeing $\hat{s}$ collapses the falsification structure.
2. *Independence from the framework's predictions.* The proxy for Q must not be defined by reference to past episodes of the system's collapse (e.g., $Q \neq V_{\text{observed-at-collapse}}$). The proxy for V must not invoke quantities defined by the friction coupling under test. Both proxies must be expressible without invoking the model whose predictions the protocol is testing.
3. *Alternative-operationalisation robustness.* At least one alternative proxy choice must be specified in advance and the protocol run on both. If the fold mechanism holds, $\hat{s}$ from the two operationalisations should overlap in 95% CI; both contain $-1/2$ or both exclude it. Sensitivity of $\hat{s}$ to the proxy choice is evidence against the structural interpretation, independent of the Wald-test outcome.

Standard (2) catches direct dependence between the proxy specification and the framework's predictions. In domains where the framework's vocabulary has already entered conventional measurement practice, indirect dependence may persist even when no explicit reference to the framework appears in the proxy definition; in such cases, operationalisations should be defensible without invoking the framework, and standard (3) becomes the primary safeguard against indirect circularity.

The collapse forecast $\hat{\mu}^*$ is a point estimate with uncertainty. It should not be treated as a precise prediction of collapse timing but as a structural warning: the system's dynamics are consistent with a fold at this distance, and the margin is shrinking.

Appendix B: Synthetic Validation

B.1 Setup

Synthetic observations are generated from the OU linearisation of the fast subsystem with $Q = 1$, $D = 1$. For $i = 1, \ldots, n = 40$, μ_i is swept linearly from 0.70 down to 0.02; the true CSD observables are $\tau_{\text{true}}(\mu) = 1/(2\sqrt{\mu})$ and $\text{Var}_{\text{true}}(\mu) = D/(4\sqrt{\mu})$. Multiplicative observation noise is added with $\text{CV}_\mu = 0.08$, $\text{CV}_\tau = 0.12$, $\text{CV}_\sigma = 0.15$.

B.2 Log-Log Inference (Primary)

For a representative seed (single realisation, $n = 40$):

Observable	$\hat{s}$	95% CI	ΔBIC_{M2-M1}	ΔBIC_{M3-M1}	R^2 (M1)
τ	0.47	$[0.42, 0.52]$	$+37.5$	-8.4	0.870
σ^2	0.47	$[0.42, 0.53]$	$+81.5$	$+11.7$	0.928

Table 3. Single-realisation log-log inference. Both 95% CIs contain the fold prediction $s = 1/2$. Slope estimates and CIs are read from the single pinned seed of Figure 2; the τ row matches panel (a). The ΔBIC and R^2 columns await regeneration from that seed. M2 (linear, no divergence) was rejected for both observables in that run. The ΔBIC_{M3-M1} column is realization-noisy by construction: across the 200 seeds of B.3 its distribution for τ is centred near zero (mean $\approx +1$, sd ≈ 7, with $\approx 10\%$ of seeds below -8), because original-space BIC scoring after log-space fitting is a noisy instrument for the nested M1-vs-M3 comparison. The displayed seed's -8.4 for τ is a lower-decile draw rather than evidence against M1; the Wald test on the log-log slope (A.3) is the primary inference for that comparison. The same pseudo-likelihood caveat applies to the M1-vs-M2 comparison: BIC differences after log-space fitting are not a proper likelihood ratio, so for empirical use M2 rejection should be corroborated by out-of-sample cross-validation or a Vuong non-nested test rather than by BIC alone; those instruments are recommended for the protocol and are not computed here.

Cross-check. The ratio σ^2/τ regressed on $\log \mu$ has log-log slope -0.006 with $p = 0.85$, consistent with the fold prediction that the ratio is μ-independent.

B.3 Robustness Across Seeds

To convert the single-seed result into a statement about protocol operating characteristics, the validation was repeated for 200 independent seeds at $n = 40$:

Observable	$\bar{\hat{s}}$	$\text{sd}(\hat{s})$	CI coverage of 1/2
τ	0.498	0.024	95.5%
σ^2	0.494	0.028	95.5%

Table 4. Distribution of log-log slope estimates across 200 seeds. CI coverage hits the nominal 95% rate, confirming that the log-log inference is well-calibrated.

B.4 Power

Coverage of the fold slope $-1/2$ inside the 95% CI as a function of n (300 replications per n):

n	$\Pr(\text{CI}_{95\%} \ni -1/2)$
10	92%
20	93%
30	91%
40	95%
60	94%
80	94%

Table 5. Log-log CI coverage as a function of sample size for the variance observable. Coverage is approximately nominal (91–95%) from $n = 10$ onward at these noise levels; the discriminating power against M2 (linear) saturates near $n \approx 30$.

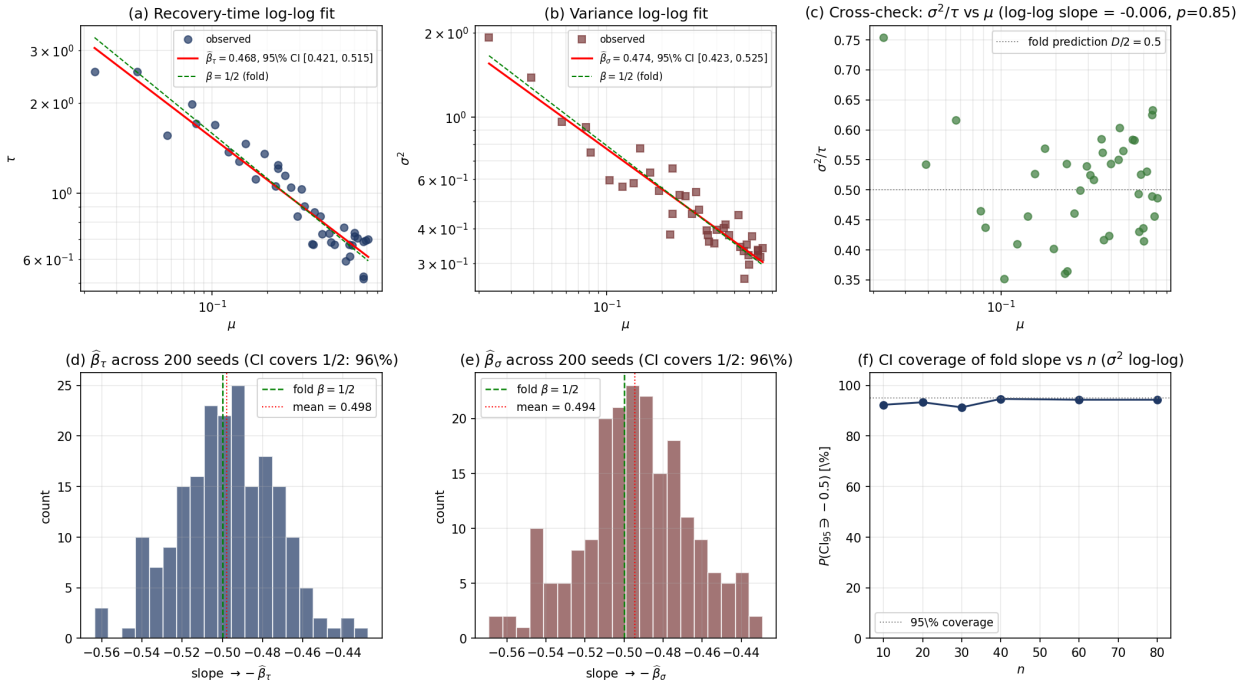

Figure 2. Log-log validation results. (a) τ vs. μ on log-log axes with the M3 free-slope fit (red) and the fold-constrained M1 fit (green dashed); both overlap visibly. (b) Same for σ^2. (c) The σ^2/τ flatness cross-check on log-log axes; the log-log slope is consistent with zero. (d, e) Distribution of $\hat{s}_\tau$ and $\hat{s}_\sigma$ across 200 seeds; both centre on the fold prediction (green dashed), each covering 1/2 at 95.5% (Table 4; rounded to 96% in-panel). (f) CI coverage of the fold slope as a function of n, approximately nominal (91–95%) from $n \geq 10$.

B.5 Rolling-Window Estimator Bias (Caveat)

The protocol as written requires a stationary-variance estimate per μ_i; in real data this is approximated by a rolling-window variance estimator. A direct OU simulation of the linearised fast subsystem ($T = 400$ time units per μ, sampling step 1 time unit, window $w = 8$, 200 seeds), compared against the analytic stationary variance $\mathrm{Var}_{\mathrm{true}} = D/(4\sqrt{\mu})$ of §7, shows that the rolling-window estimator carries a finite-sample bias that *under*-estimates the stationary variance, monotonically in μ: the ratio (estimator / $\mathrm{Var}_{\mathrm{true}}$) rises from ≈ 0.52 near the fold ($\mu = 0.02$, relaxation time $1/\lambda \approx 3.5$) to ≈ 0.94 far from it ($\mu = 0.70$, $1/\lambda \approx 0.6$), never exceeding 1. The under-estimate is largest near the fold because the relaxation time there exceeds the window span, so an 8-sample window cannot accumulate the full stationary variance; far from the fold the samples decorrelate and the estimate approaches the unbiased value. The effect is exact in closed form: for a discrete OU process with one-step autocorrelation $a = e^{-\lambda}$, the Bessel-corrected windowed sample variance satisfies $\mathbb{E}[\widehat{V}]/\mathrm{Var}_{\mathrm{true}} = 1 - \frac{1}{w(w-1)} \sum_{i \neq j} a^{|i-j|} < 1$, increasing monotonically toward 1 as λ grows (Figure 3). The identity is the standard serial-correlation bias of the sample variance: writing $s^2 = \frac{1}{w-1} \sum_i (x_i - \bar{x})^2$ for a zero-mean stationary process with variance σ^2 and autocorrelation $\rho(k) = a^{|k|}$, one has $\mathbb{E}\big[\sum_i (x_i - \bar{x})^2\big] = (w-1)\sigma^2 - \frac{\sigma^2}{w} \sum_{i \neq j} \rho(|i-j|)$, whence $\mathbb{E}[s^2] = \sigma^2 \big(1 - \frac{1}{w(w-1)} \sum_{i \neq j} a^{|i-j|}\big)$. Evaluated at $w = 8$ and $\lambda(\mu) = 2\sqrt{\mu}$, the closed form gives:

μ	0.02	0.10	0.30	0.50	0.70
relaxation $1/\lambda$	3.54	1.58	0.91	0.71	0.60
ratio $\mathbb{E}[\widehat{V}]/\mathrm{Var}_{\mathrm{true}}$	0.52	0.76	0.88	0.92	0.94

The simulation points of Figure 3 match these values to within seed noise. The validation of B.2-B.4 used the analytic stationary variance to isolate the inference method from this estimator bias; empirical applications combining the two should calibrate the window length w against the longest plausible relaxation time and report sensitivity across $w \in \{3, 5, 8, 15\}$, or replace the rolling-window estimator with an AR(1) coefficient

ρ and back out $\mathrm{Var} = \mathrm{Var}(\varepsilon)/(1-\rho^2)$, which removes the relaxation-time under-estimate near the fold.

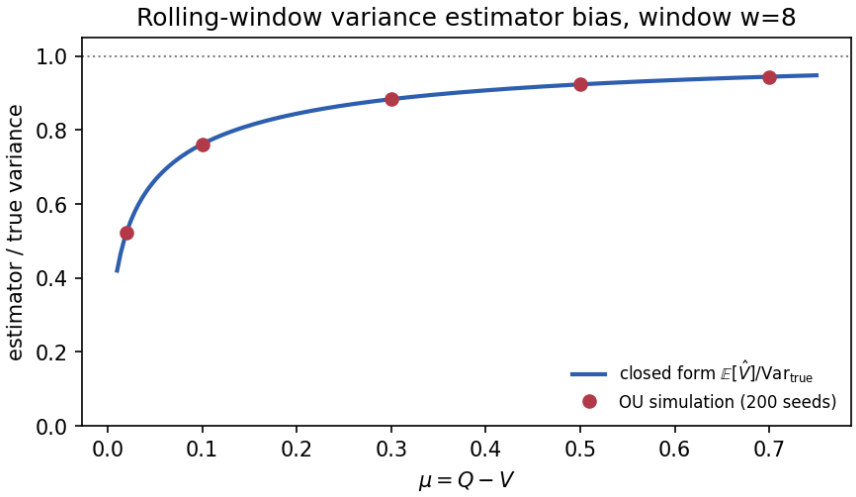

Figure 3. Finite-sample bias of the rolling-window variance estimator (window $w = 8$, sample step one time unit) relative to the OU stationary variance $D/(4\sqrt{\mu})$: the estimator under-estimates monotonically, most severely near the fold. Solid curve, closed-form $\mathbb{E}[\hat{V}]/\mathrm{Var}_{\mathrm{true}}$; points, OU simulation (200 seeds).

References

[1] S. Strogatz, *Nonlinear Dynamics and Chaos*, Westview Press, 2015.

[2] J. Guckenheimer and P. Holmes, *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*, Springer, 1983.

[3] N. Fenichel, "Geometric singular perturbation theory for ordinary differential equations," *J. Differential Equations* 31 (1979), 53–98.

[4] M. Krupa and P. Szmolyan, "Extending geometric singular perturbation theory to non-hyperbolic points: fold and canard points in two dimensions," *SIAM J. Math. Anal.* 33 (2001), 286–314.

[5] M. Scheffer, J. Bascompte, W. A. Brock, et al., "Early-warning signals for critical transitions," *Nature* 461 (2009), 53–59.

[6] C. Boettiger and A. Hastings, "Quantifying limits to detection of early warning for critical transitions," *J. R. Soc. Interface* 9 (2012), 2527–2539.

[7] E. H. van Nes and M. Scheffer, "Slow recovery from perturbations as a generic indicator of a nearby catastrophic shift," *American Naturalist* 169 (2007), 738–747.

[8] H. Held and T. Kleinen, "Detection of climate system bifurcations by degenerate fingerprinting," *Geophysical Research Letters* 31 (2004), L23207.

[9] S. R. Carpenter and W. A. Brock, "Rising variance: a leading indicator of ecological transition," *Ecology Letters* 9 (2006), 311–318.

[10] C. K. R. T. Jones, "Geometric singular perturbation theory," in *Dynamical Systems (Montecatini Terme, 1994)*, Lecture Notes in Math. 1609, Springer, 1995, 44–118.

[11] C. Kuehn, "A mathematical framework for critical transitions: bifurcations, fast–slow systems and stochastic dynamics," *Physica D* 240 (2011), 1020–1035.

[12] C. Kuehn, *Multiple Time Scale Dynamics*, Applied Mathematical Sciences 191, Springer, 2015.

[13] C. Kuehn, "Normal hyperbolicity and unbounded critical manifolds," *Nonlinearity* 27 (2014), 1351–1366.

[14] N. Berglund and B. Gentz, *Noise-Induced Phenomena in Slow–Fast Dynamical Systems*, Springer, 2006.

[15] M. Wechselberger, *Geometric Singular Perturbation Theory Beyond the Standard Form*, Frontiers in Applied Dynamical Systems 6, Springer, 2020.

[16] P. Ashwin, S. Wieczorek, R. Vitolo, and P. Cox, "Tipping points in open systems: bifurcation, noise-induced and rate-dependent examples in the climate system," *Phil. Trans. R. Soc. A* 370 (2012), 1166–1184.

[17] S. Prajna and A. Jadbabaie, "Safety verification of hybrid systems using barrier certificates," in *Hybrid Systems: Computation and Control*, Lecture Notes in Comput. Sci. 2993, Springer, 2004, 477–492.

[18] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: theory and applications," in *Proc. 18th European Control Conference (ECC)*, 2019, 3420–3431.

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

[20] L. Kleinrock, *Queueing Systems, Volume 1: Theory*, Wiley, 1975.

[21] J. Sotomayor and M. Zhitomirskii, "Impasse singularities of differential systems of the form $A(x)x' = F(x)$," *J. Differential Equations* 169 (2001), 567–587.

[22] M. Zhitomirskii, "Local normal forms for constrained systems on 2-manifolds," *Bol. Soc. Bras. Mat.* 24 (1993), 211–232.

[23] P. T. Cardin, P. R. da Silva, and M. A. Teixeira, "Implicit differential equations with impasse singularities and singular perturbation problems," *Israel J. Math.* 189 (2012), 307–322.

[24] P. T. Cardin and M. A. Teixeira, "A geometric singular perturbation theory approach to constrained differential equations," *Math. Nachr.* 292 (2019), 892–904.

[25] J. Borsotti, H. Jardón-Kojakhmetov, and M. Sensi, "Slow–fast dynamics in a planar parasite–host model with an extinction singularity," arXiv:2602.19257 (2026).

[26] K. U. Kristiansen and S. J. Hogan, "Regularizations of two-fold bifurcations in planar piecewise smooth systems using blowup," *SIAM J. Appl. Dyn. Syst.* 14 (2015), 1731–1786.