Sharp Uniqueness and Bistability in Driven Adaptive Spectral Carriers

Jeremy Rodgers

Reserved DOI: 10.5281/zenodo.23932350
everythingequation.com

3 November 2026

Abstract

The same constant input can sustain different stable amplitudes, memories and responses in a finite reciprocal adaptive carrier. We determine when this occurs without changing the constitutive law or its external power contract. Stationary memory elimination gives a sharp curvature constant $\chi = (2-\sqrt{3})\mathrm{e}^{-\sqrt{3}}$. We prove global asymptotic attraction to one equilibrium for every constant complex input under $\mu \geq \chi \|L\|$, including equality, where uniform strong convexity is not guaranteed. On the three-node path the exact all-input threshold is $\mu \geq 3\chi$, making the universal coefficient optimal. For every $0 < \mu < 3\chi$ suitable highest-mode inputs sustain two fully stable states. Within the fold interval and for $\chi < \mu < 3\chi$, a transverse convexity argument proves that the complete stationary set consists of those two minima and an index-one saddle. A mobility–inertia lemma classifies every amplitude and memory direction, including zero eigenvalues, and both folds are nondegenerate saddle-nodes in the completely counted regime. Independent modal cubics retain all eight eigenvalues. A full-state tilted-energy barrier gives sufficient basin regions, and prescribed numerical pulses switch between the stable states before returning to identical input. Physical work is accounted for separately. Equal mean, stationary power and clamped static norm can conceal different stable full-feedback responses. Exact scalar certificates, separate numerical formulations and complete replay checks support the worked example; physical implementation and quasistatic hysteresis limits remain open.

1 Introduction

An adaptive carrier changes the operator through which its own amplitudes evolve. Constant forcing can maintain such a carrier against dissipation, but maintenance alone does not determine whether the retained state is unique. Two preparations subject to the same input may return to one common state, or preserve different memories and different responses. The distinction requires the dynamics of the complete source, rather than stability inside a chosen symmetry sector or positivity of a clamped carrier.

The canonical spectral-curvature field (CSCF) programme organizes coherent and dissipative dynamics around a common positive operator [Rodgers, 2026b]. Reciprocal adaptive carriers make that operator state dependent [Rodgers, 2026a]. The preceding sustained-state analysis introduced a fixed external rate input, a physical power balance and a strict sufficient condition for global attraction [Rodgers, 2026e]. The present paper resolves the boundary of that condition and develops a complete bistable example in the same finite model. It adds no controller, reservoir, saturation or change of nonlinearity.

There are four mathematical bridges. First, eliminating memory at stationarity produces an effective potential with a bounded negative edge curvature. At the sharp curvature endpoint, a line argument supplies strict convexity although the Hessian need not be uniformly positive. Second, a scalar path branch law becomes a complete equilibrium theorem only after restoring phase, collective amplitude, reflection-odd directions and memory. Third, stationary curvature classifies the full dynamical Jacobian through a positive-symmetric-part mobility lemma. Fourth, a tilted-potential barrier gives basin membership for complete source perturbations, while physical supplied work remains governed by the original energy.

The information retained by an observation is equally consequential. The source-to-readout architecture of Shadow Theory asks which future operations an observable can predict [Rodgers, 2026c,d]. Here two fully stable states have identical mean, stationary power and clamped static operator norm, yet different full-feedback responses to subsequent probes. The obstruction is stated for those probes and that coarse summary. Complete calibrated stationary amplitude data can instead reconstruct memory. Neither statement permits eliminating evolving memory from a general trajectory.

The central results are analytic. Computations evaluate a fixed example, check the complete matrices against independent cubic formulas, and execute finite switching schedules with energy ledgers. The example is not unseen empirical evidence. No atomic, image or Kerr result is reinterpreted as a new test of this adaptive model. Its contribution is a precise constitutive mechanism for maintained state dependence and the mathematical and measurement contracts needed to assess it.

2 Complete source and external energy contract

2.1 Constitutive equations and units

Let $z \in \mathbb{C}^n$ be an amplitude and $m \in \mathbb{R}^p$ a dimensionless edge memory. Real incidence vectors $b_e \in \mathbb{R}^n$ define

$$B_e = b_e b_e^T, \qquad L = \sum_e B_e, \qquad K(m) = \mu I + \sum_e \mathrm{e}^{m_e} B_e, \qquad \mu > 0. \tag{1}$$

The amplitude metric is Euclidean; all vector norms and induced operator norms below refer to this fixed metric and its realification. Baseline edge stiffness is normalized to one. Positive weights k_e can be included by replacing b_e with $\sqrt{k_e} b_e$ everywhere, including in L and the curvature bound. In particular $K(m) \geq \mu I$, including on collective modes.

The physical stored energy and memory effort are

$$F(z, m) = \tfrac{1}{2} z^* K(m) z + \tfrac{\kappa}{2} |m|^2, \qquad \kappa > 0, \tag{2}$$

$$g_e(z, m) = (F_m)_e = \kappa m_e + \tfrac{1}{2} \mathrm{e}^{m_e} |b_e^T z|^2. \tag{3}$$

With $c = \gamma + i\alpha$, $\gamma, \alpha, \eta > 0$, the complete source is

$$\dot{z} = -cK(m) z + f, \qquad \dot{m} = -\eta g(z, m). \tag{4}$$

Unless a pulse is explicitly specified, $f \in \mathbb{C}^n$ is constant. Memory responds to the same F that generates the amplitude effort. This reciprocal constitutive assumption is part of the model, rather than a conclusion about an unspecified microscopic medium.

If amplitude has units A, energy has units E and time has units T, then K has units E/A^2, κ has units E, f has units A/T, γ and α have units $A^2/(ET)$, and η has units $1/(ET)$. Reference scales

A_0, E_0, T_0 give the dimensionless coordinates used in the joint Hessian and numerical calculations. Comparisons involving amplitude and memory norms are made only after that normalization. The dimensionless graph convention in (1) is used throughout.

2.2 Power and the constant-input tilt

The input is an amplitude-rate flow with conjugate effort Kz. Its signed instantaneous power is

$$P = \mathrm{Re}\,[(Kz)^* f], \qquad \dot{F} = -\gamma \|Kz\|^2 - \eta|g|^2 + P. \tag{5}$$

This is an effort–flow convention of dissipative systems theory [van der Schaft and Jeltsema, 2014]. The ideal actuator lies outside the source state and can supply or absorb power. No finite energy store for the actuator has been included. The reactive term does no work because $\mathrm{Re}[i\alpha \|Kz\|^2] = 0$.

For constant input define $\ell = f/c$ and

$$V(z, m) = F(z, m) - \mathrm{Re}(\ell^* z). \tag{6}$$

Let $q = (\mathrm{Re}\, z, \mathrm{Im}\, z) \in \mathbb{R}^{2n}$, $K_R = \mathrm{diag}(K, K)$ and

$$J = \begin{pmatrix} 0 & -I \\ I & 0 \end{pmatrix}, \quad G = \gamma I + \alpha J, \quad D = \mathrm{diag}(G, \eta I_p), \quad x = (q, m). \tag{7}$$

Then $\dot{x} = -D\nabla V$, with $\mathrm{Sym}\, D = \mathrm{diag}(\gamma I, \eta I) > 0$, and

$$\dot{V} = -\gamma \|Kz - \ell\|^2 - \eta|g|^2. \tag{8}$$

The tilt is a mathematical Lyapunov function, not the stored energy of the driven source or its actuator. Equation (5), rather than a difference of tilted barriers, accounts for supplied work.

Proposition 2.1 (Boundedness and stationary limit sets)**.** *For every initial state and constant f, (4) has a bounded global forward solution. Every point of its omega-limit set is stationary. If the stationary set is finite, each solution converges to a single stationary point. There is no nonstationary recurrent trajectory at fixed input.*

Proof. Young's inequality and $K \geq \mu I$ give

$$V \geq \tfrac{\mu}{4} \|z\|^2 + \tfrac{\kappa}{2}|m|^2 - \frac{\|\ell\|^2}{\mu}. \tag{9}$$

Thus every V sublevel is compact. The smooth vector field is locally Lipschitz, and (8) confines its solution to the initial compact sublevel. Local continuation excludes finite-time escape. This argument uses no curvature condition.

The nonincreasing V has a finite limit. If an omega-limit point were nonstationary, the continuous loss in (8) would be bounded below by a positive number on a small ball about it. On the containing compact set the velocity is bounded. Infinitely many visits to a smaller ball therefore give infinitely many disjoint time intervals of a fixed positive duration in the larger ball. Their cumulative loss contradicts the finite limit of V. Hence all limit points are stationary. This is the finite-dimensional strict-dissipation argument underlying the usual invariance principle [LaSalle, 1976].

The omega-limit set is nonempty, compact and connected: it is the intersection of nested compact connected closures of continuous forward tails. A connected subset of a finite stationary set is one point, which proves convergence. Finally, recurrence to a nonstationary initial point would require returning to its initial V value after a strict decrease. Monotonicity makes this impossible. In particular nonconstant periodic or recurrent motion cannot occur at fixed input. □

The finiteness hypothesis is needed for the stated single-equilibrium convergence argument; it will be proved in the path regimes below. Changing-input experiments do not have one fixed V. A finite sequence of constant pulses is a concatenation of globally defined segments, with continuous state and physical F at each jump. At an input jump the tilt changes by $-\operatorname{Re}[(\ell_{\text{new}} - \ell_{\text{old}})^* z]$.

3 Stationary reduction and endpoint uniqueness

3.1 Memory elimination and its exact curvature

For fixed z, V is strictly convex in m, since

$$(F_{mm})_{ee} = \kappa + \tfrac{1}{2}\mathrm{e}^{m_e}|b_e^T z|^2 \geq \kappa. \tag{10}$$

The unique minimizing memory is $m_e = -s_e$, where

$$s_e \mathrm{e}^{s_e} = \frac{|b_e^T z|^2}{2\kappa}, \qquad s_e = W_0\left(\frac{|b_e^T z|^2}{2\kappa}\right) \geq 0. \tag{11}$$

Only the nonnegative real principal branch is needed [Corless et al., 1996]. The effective stationary potential is

$$V_{\text{eff}}(z) = \tfrac{\mu}{2}\|z\|^2 + \sum_e \kappa\left(s_e + \tfrac{1}{2}s_e^2\right) - \operatorname{Re}(\ell^* z). \tag{12}$$

Critical points of V_{eff} lift bijectively to complete stationary points by (11). This is an exact stationary reduction. It does not replace the dynamical equation for memory away from stationarity.

To compute curvature, view $d = b_e^T z$ as a two-dimensional real vector. With $w = \mathrm{e}^{-s}$ the edge gradient is wd. Its Hessian has tangential eigenvalue w and radial eigenvalue

$$w\frac{1-s}{1+s} = -Q(s), \qquad Q(s) = \mathrm{e}^{-s}\frac{s-1}{s+1}. \tag{13}$$

Indeed differentiating $s\mathrm{e}^s = |d|^2/(2\kappa)$ gives $\mathrm{d}s/\mathrm{d}|d| = 2s/[(1+s)|d|]$, which yields the radial expression, with the continuous limit at $d = 0$. The sharp negative curvature follows from

$$Q'(s) = \mathrm{e}^{-s}\frac{3-s^2}{(1+s)^2}, \qquad \chi = \max_{s\geq 0} Q(s) = (2-\sqrt{3})\mathrm{e}^{-\sqrt{3}}. \tag{14}$$

The maximum is unique at $s = \sqrt{3}$. In real amplitude coordinates,

$$\nabla^2 V_{\text{eff}} \geq \mu I - \chi L_R, \qquad L_R = \operatorname{diag}(L, L). \tag{15}$$

Theorem 3.1 (Closed sufficient uniqueness condition)**.** *For a finite real incidence network, positive $\kappa, \gamma, \alpha, \eta$ and*

$$\mu \geq \chi\,\|L\|, \qquad \mu > 0, \tag{16}$$

every constant complex input has one stationary state, globally asymptotically stable in the complete source. At equality the effective potential is strictly convex and coercive, but uniform strong convexity and a positive general exponential-rate bound do not follow.

Proof. When $\mu > \chi\,\|L\|$, (15) gives strong convexity. At equality suppose first $L \neq 0$, and restrict V_{eff} to any nonconstant real line $z(t) = z_0 + t\xi$. Its second derivative can be written

$$\mu\,\|\xi\|^2 - \chi\sum_e |b_e^T\xi|^2 + \sum_e \left(\mathrm{d}^2\phi_e(d_e(t))[\delta_e, \delta_e] + \chi|\delta_e|^2\right), \tag{17}$$

where $\delta_e = b_e^T \xi$ and ϕ_e is the edge term in (12). Every displayed term is nonnegative. If the first term is positive, the line is strictly convex. Otherwise, since $\mu > 0$ and $\xi \neq 0$, some $\delta_e \neq 0$. For that edge its remainder can vanish only when $s_e = \sqrt{3}$ and δ_e lies in the radial direction. The first requirement alone is

$$|d_e(0) + t\delta_e|^2 = 2\kappa\sqrt{3}\mathrm{e}^{\sqrt{3}}. \tag{18}$$

The left side is a nonconstant quadratic with positive leading coefficient. Equality holds at at most two real t. Thus (17) is positive except possibly at a finite set, and the first derivative of the line restriction is strictly increasing on every nontrivial interval. Every nonconstant line restriction is strictly convex. If $L = 0$, the positive quadratic $\mu \|z\|^2 /2$ already proves the claim.

The bound (9) survives minimizing over memory, so V_{eff} is coercive. It has exactly one critical point, its strict global minimum. The stationary correspondence gives exactly one complete stationary point. Proposition 2.1 gives convergence of every trajectory. Moreover $V(x) \geq V_{\mathrm{eff}}(z)$, with equality only at the minimizing memory, so the lifted point is the unique strict global minimum of V. Compact sublevels and positive definiteness of $V - V_*$ give Lyapunov stability. Attraction and stability together prove global asymptotic stability. □

At equality the Hessian can have a zero direction at the minimum, as the path endpoint below demonstrates. Strict convexity is sufficient for uniqueness, but does not give a uniform quadratic lower curvature or exponential rate. Conversely, failure of (16) alone says nothing about uniqueness at a particular input. At $f = 0$, for example, the only stationary point is $(0, 0)$ for every $\mu > 0$.

4 Full stability from stationary curvature

4.1 Complete Hessian and Schur return

At a stationary point, let $H = \nabla^2 V$. In the coordinate order (q, m),

$$H = \begin{pmatrix} K_R & B \\ B^T & H_m \end{pmatrix}, \qquad B_{:e} = \mathrm{e}^{m_e} \operatorname{diag}(B_e, B_e) q, \quad H_m = \operatorname{diag}\Big(\kappa + \tfrac{1}{2}\mathrm{e}^{m_e} |b_e^T z|^2\Big). \tag{19}$$

The Schur complement is

$$S = K_R - B H_m^{-1} B^T = \nabla^2 V_{\mathrm{eff}}. \tag{20}$$

Completing the square in the memory perturbation gives the exact congruence

$$(\xi, \nu)^T H (\xi, \nu) = \xi^T S \xi + (\nu + H_m^{-1} B^T \xi)^T H_m (\nu + H_m^{-1} B^T \xi). \tag{21}$$

Thus the full Hessian has the negative and zero counts of S, plus p additional positive directions. This returns all eliminated memory coordinates before making a stability claim. Positivity of K alone is inadequate.

Lemma 4.1 (Mobility–inertia correspondence)**.** *Let H be real symmetric and D real with* $\operatorname{Sym} D > 0$. *For $A = -DH$, the number of eigenvalues in the open right half-plane equals the number of negative eigenvalues of H; the open left-half-plane count equals its positive count. Counts include algebraic multiplicities over $\mathbb{C}$. Zero has multiplicity* $\dim \ker H$ *and is semisimple. There is no nonzero imaginary eigenvalue.*

Proof. The matrix D is invertible, and

$$\operatorname{Sym}(D^{-1}) = D^{-T} (\operatorname{Sym} D) D^{-1} > 0. \tag{22}$$

If $Av = i\omega v$ for $v \in \mathbb{C}^N \setminus \{0\}$ and real $\omega \neq 0$, then $Hv \neq 0$ and

$$(Hv)^* D(Hv) = -i\omega v^* Hv. \tag{23}$$

The real part on the left is strictly positive, whereas the right has zero real part because H is Hermitian. This contradiction excludes every nonzero imaginary eigenvalue, including those with nonreal eigenvectors.

Invertibility of D gives $\ker A = \ker H$. If a Jordan chain at zero existed, there would be nonzero $v_0 \in \ker H$ and w with $Aw = v_0$. Then $Hw = -D^{-1}v_0$, so

$$0 = v_0^* Hw = -v_0^* D^{-1} v_0, \tag{24}$$

contradicting (22). Hence zero is semisimple with the asserted algebraic multiplicity.

Finally set $D_t = (1-t)I + tD$ for $0 \leq t \leq 1$. Its symmetric part stays positive. The preceding arguments apply throughout this homotopy: no eigenvalue crosses the nonzero imaginary axis, and the zero multiplicity stays fixed at $\dim \ker H$. A nonzero eigenvalue cannot reach zero without increasing that algebraic multiplicity. Continuity of roots of the characteristic polynomial preserves the open-half-plane counts. At $t = 0$ the matrix is $-H$, where the counts are immediate. This proves all statements, including multiplicities. □

For (4) the complete linearization is $A = -DH$. The lemma and (21) classify the entire amplitude–memory system from S. A positive H gives a Hurwitz Jacobian and nonlinear local exponential stability. To see the rate without misusing a global curvature bound, fix $H_* > 0$, put $\sigma = \min(\gamma, \eta)$ and $h_* = \lambda_{\min}(H_*)$. For a perturbation v, the fixed quadratic $E = v^T H_* v/2$ satisfies

$$\dot{E} \leq -2\sigma h_* E + O(E^{3/2}). \tag{25}$$

For every $0 < r < \sigma h_*$ a sufficiently small E sublevel is invariant and $\|v(t)\|_{H_*} \leq \mathrm{e}^{-rt} \|v(0)\|_{H_*}$. The linearized estimate has rate σh_*. This local argument remains valid when the global bound $\mu - \chi \|L\|$ is negative.

4.2 What varying mobility holds fixed

At fixed H, changing α or any other mobility while maintaining $\operatorname{Sym} D > 0$ cannot change the index. It can change eigenvalues, transient motion and decay rates, but cannot produce a Hopf crossing or stabilize a negative Hessian direction. This is a fixed-state, fixed-load statement. To keep a stationary state fixed while changing c, one must adjust $f = c\ell$ with ℓ fixed. Holding the physical input f fixed instead changes $\ell = f/c$ and generally changes the equilibrium and its Hessian. Section 5 gives the resulting fold condition explicitly.

5 Three-node path: exact count and sharp threshold

5.1 Modes and scalar stationary law

Take $b_1 = (1, -1, 0)^T$, $b_2 = (0, 1, -1)^T$. An orthonormal eigenbasis of L is

$$v_0 = \frac{(1,1,1)^T}{\sqrt{3}}, \quad v_1 = \frac{(1,0,-1)^T}{\sqrt{2}}, \quad v_3 = \frac{(1,-2,1)^T}{\sqrt{6}}, \quad Lv_\lambda = \lambda v_\lambda. \tag{26}$$

For $f = av_\lambda$, $a > 0$, $\lambda \in \{1, 3\}$, the equal-memory stationary sector has $z = uv_\lambda$, $m_1 = m_2 = -s$. Each edge satisfies $|b_e^T z|^2 = \lambda |u|^2/2$, so

$$R^2 := |u|^2 = \frac{4\kappa}{\lambda} s\mathrm{e}^s, \qquad u = \frac{a/c}{\mu + \lambda \mathrm{e}^{-s}}, \tag{27}$$

$$\frac{a^2}{|c|^2} = H_\lambda(s) := \frac{4\kappa}{\lambda} s\mathrm{e}^s(\mu + \lambda \mathrm{e}^{-s})^2. \tag{28}$$

The phase is fixed by a/c; R alone does not introduce a free phase direction. These sectors arise from reflection symmetry: the low mode uses signed reflection, the highest mode ordinary reflection, with zero collective amplitude at stationarity.

Differentiating the complete scalar law gives

$$H'_\lambda(s) = \frac{4\kappa}{\lambda}(\mu + \lambda \mathrm{e}^{-s})\mathrm{e}^s(1 + s)\,[\mu - \lambda Q(s)]. \tag{29}$$

When $\mu \geq \lambda\chi$, H_λ is strictly increasing from zero to infinity. At equality it has one isolated flat point, which does not destroy strict monotonicity. When $0 < \mu < \lambda\chi$, the equation $Q(s) = \mu/\lambda$ has precisely two roots

$$1 < s_- < \sqrt{3} < s_+. \tag{30}$$

The first is a strict local maximum of H_λ and the second a strict local minimum. Define

$$a_{\mathrm{down}} = |c|\sqrt{H_\lambda(s_+)}, \qquad a_{\mathrm{up}} = |c|\sqrt{H_\lambda(s_-)}. \tag{31}$$

There are three distinct positive scalar roots when $a_{\mathrm{down}} < a < a_{\mathrm{up}}$, two distinct roots with one double root at either endpoint, and one root outside the closed interval. This follows from the three monotonicity intervals and $H_\lambda(0) = 0$, $H_\lambda(s) \to \infty$.

At an equal-edge equilibrium a common phase rotation makes the amplitude real. The radial and tangential parts of S are respectively $\mu I - Q(s)L$ and $\mu I + \mathrm{e}^{-s}L$. Its six eigenvalues are

$$\mu, \quad \mu - Q, \quad \mu - 3Q; \qquad \mu, \quad \mu + w, \quad \mu + 3w, \quad w = \mathrm{e}^{-s}. \tag{32}$$

The collective coordinates are retained, not neutral: with $\bar{z} = (z_1 + z_2 + z_3)/3$,

$$\dot{\bar{z}} = -c\mu\bar{z} + \bar{f}. \tag{33}$$

Zero-mean forcing has $\bar{z}_* = 0$ and two decaying collective quadratures.

5.2 A sector-stable candidate that is fully unstable

For low-mode forcing the scalar fold condition is $Q = \mu$. At that point the transverse highest-mode radial curvature is $\mu - 3Q = -2\mu < 0$. A fold seen in the low-mode sector therefore occurs while the complete state already has an unstable direction.

The retained exact stationary candidate $\mu = 0.01$, $\kappa = 1$, $s = 5$, $a = |c|\sqrt{H_1(5)}$ is more explicit. Here $Q(5) = \frac{2}{3}\mathrm{e}^{-5}$ and rational exponential bounds give $1/300 < Q(5) < 1/100$. Consequently the sector radial curvature is positive while the transverse curvature is negative. Numerical evaluation gives approximately 0.00550804 and -0.00347589, respectively; the complete Jacobian has one unstable eigenvalue. Positivity of its clamped carrier and stability within the prepared sector do not rescue complete stability. This counterexample rejects that candidate, not every possible low-mode bistable regime. In particular the calculation does not rule out two sufficiently far outer low-mode states with positive full curvature.

5.3 Completeness for the highest-mode preparation

Theorem 5.1 (Complete path bistability)**.** *Suppose $\chi < \mu < 3\chi$ and $f = av_3$, $a > 0$. All complete equilibria are given by (27)–(28) with $\lambda = 3$, $z = uv_3$, $m_1 = m_2 = -s$. In the open interval (31) there are exactly three: two locally exponentially stable minima and one index-one saddle. Outside the closed interval there is one globally asymptotically stable equilibrium. At each endpoint the colliding pair has one zero and seven decaying linear modes. Every constant-input trajectory converges to one of the finitely many equilibria in these regimes.*

Proof. Write $z = tv_0 + uv_3 + hv_1$ with complex t, u, h. Edges do not see t; the input has no collective component. The collective part of V_{eff} is $\mu|t|^2/2$, so every critical point has $t = 0$.

Let reflection interchange nodes 1 and 3. It fixes v_3, reverses v_1, and interchanges the edges up to their irrelevant signs. For fixed u, the effective potential is even under $h \mapsto -h$. Its real two-dimensional Hessian in h obeys

$$\nabla_h^2 V_{\text{eff}} \geq (\mu - \chi) I_2 > 0, \tag{34}$$

because $v_1^T L v_1 = 1$ and the edge lower bound is $-\chi$. Thus $h = 0$ is its unique critical point for every u. This rules out all reflection-broken stationary points globally, without a root search.

At $t = h = 0$, the two edge magnitudes coincide. The unique stationary memory therefore has $m_1 = m_2 = -s$. The amplitude equation requires $u = (a/c)/(\mu + 3\mathrm{e}^{-s})$, fixing phase and reducing exactly to $H_3(s) = a^2/|c|^2$. The scalar count is consequently the complete count.

On either outer branch, (29) gives $Q < \mu/3$, so every eigenvalue in (32) is positive. Equation (21) gives $H > 0$, and Lemma 4.1 gives complete local exponential stability. On the middle branch $Q > \mu/3$, whereas $Q \leq \chi < \mu$. Exactly one eigenvalue in (32) is negative. The full Jacobian has exactly one right-half-plane eigenvalue. At either fold $Q = \mu/3$, all other Schur directions are positive, hence the zero is simple and all other seven modes decay.

All listed stationary sets are finite, so Proposition 2.1 proves convergence. In a one-root regime coercivity of V guarantees a global minimum, which must be that sole stationary point. It is strict; the same sublevel argument as in Theorem 3.1 gives Lyapunov stability and therefore global asymptotic stability. At a fold this conclusion is not asserted for the nonhyperbolic colliding point. □

Corollary 5.2 (Sharp path threshold and optimal universal coefficient)**.** *On the unchanged three-node path, every constant complex input has one globally asymptotically stable equilibrium if and only if $\mu \geq 3\chi$. Hence χ is the smallest coefficient in a universal sufficient condition of the form $\mu \geq C\,\|L\|$ for this model class.*

Proof. Since $\|L\| = 3$, sufficiency is Theorem 3.1. For every $0 < \mu < 3\chi$, choose a highest-mode load strictly between the two folds. Both outer scalar roots satisfy $Q < \mu/3$, hence both are fully stable by (32) and Lemma 4.1. This proves failure of all-input uniqueness even when $\mu \leq \chi$, where additional equilibria have not been completely counted. Any proposed coefficient $C < \chi$ admits a positive $\mu < 3\chi$ with $\mu \geq 3C$, and this path supplies a counterexample. □

This separates a sufficient theorem for every network, optimality of its coefficient over networks, an exact all-input threshold on this particular path, and uniqueness for one specified input. It is not a necessary graph-norm test for every individual network or every individual drive.

At fixed physical a, reactive mobility changes the scalar load. In the completely counted parameter range the three-state condition is exactly

$$H_3(s_+) < \frac{a^2}{\gamma^2 + \alpha^2} < H_3(s_-). \tag{35}$$

Equivalently $a^2/H_3(s_-)-\gamma^2 < \alpha^2 < a^2/H_3(s_+)-\gamma^2$, intersected with $\alpha > 0$. At a positive endpoint the load derivative is $-2\alpha a^2/(\gamma^2+\alpha^2)^2 \neq 0$, so varying this constant family parameter can pass through a fold. This does not contradict fixed-H inertia invariance. The parameter η does not enter the stationary equations.

6 Saddle-nodes and complete modal cubics

6.1 Full dynamical nondegeneracy

The scalar turning points alone do not prove dynamical saddle-nodes. We verify the usual simple-center, transversality and quadratic conditions in the complete source [Kuznetsov, 2004]. Work in the range of Theorem 5.1, with fixed c, and define

$$R(s) = \sqrt{\tfrac{4\kappa}{3} s\mathrm{e}^s}, \quad z(s) = \frac{|c|}{c} R(s) v_3, \quad m(s) = (-s, -s), \quad a(s) = |c|\sqrt{H_3(s)}. \tag{36}$$

Let $x(s)$ denote this real complete state and

$$Y(x,a) = \nabla F(x) - a\mathfrak{h}, \quad \mathfrak{h} = (\mathrm{Re}(v_3/c), \mathrm{Im}(v_3/c), 0, 0), \quad X(x,a) = -DY(x,a). \tag{37}$$

The ordering $-D$ on the left of the stationary gradient is essential.

At a fold $a' = 0$, put $r = x'(s)$. Differentiating $Y(x(s), a(s)) = 0$ gives $Hr = 0$. Theorem 5.1 proves that this kernel is one dimensional. Since

$$r^T\mathfrak{h} = \frac{R'(s)}{|c|} > 0, \qquad d = r^T D^{-1} r > 0, \tag{38}$$

the row vector $l^T = r^T D^{-1}/d$ is a normalized left zero eigenvector of $A = -DH$: $l^T A = 0$ and $l^T r = 1$. The right vector need not be unit length; this normalization is the convention for the reported coefficients. The load coefficient is

$$l^T X_a = \frac{r^T\mathfrak{h}}{d} > 0. \tag{39}$$

Twice differentiating the stationary identity at the fold and left multiplying by r^T yields

$$r^T Y_{xx}[r,r] = (r^T\mathfrak{h})a'', \quad \frac{1}{2} l^T X_{xx}[r,r] = -\frac{(r^T\mathfrak{h})a''}{2d}. \tag{40}$$

Here $a'' = |c|H_3''/(2\sqrt{H_3})$ because $H_3' = 0$. Direct differentiation gives

$$H_3''(s) = \frac{4\kappa}{3}\left[\mu^2 \mathrm{e}^s(s+2) + 9\mathrm{e}^{-s}(s-2)\right]. \tag{41}$$

Alternatively (29) at a fold makes its sign opposite to Q'. Thus $a'' < 0$ at s_- and $a'' > 0$ at s_+, and neither vanishes. Equations (39)–(40), with the simple zero and seven decaying modes, establish both complete saddle-nodes. The quadratic coefficient has positive sign at the upper-input fold and negative sign at the lower-input fold in this coordinate convention.

At $\mu = 3\chi$, $s = \sqrt{3}$, the two scalar turning points meet. The effective potential is still strictly convex, and the stationary state at the critical input is unique, but the complete Jacobian has a zero eigenvalue. Linear exponential stability and a general positive local exponential estimate cannot be asserted there. A cusp classification would require additional higher-order and unfolding conditions; it is not inferred from the meeting of scalar folds.

6.2 An independent all-coordinate block formulation

Rotate the common equilibrium phase to make z real, use the three orthonormal node modes (26), and use even and odd memory combinations $(\delta m_1 \pm \delta m_2)/\sqrt{2}$. These are orthogonal transformations of the complete eight-dimensional real source. The collective amplitude contributes exactly

$$p = -\mu(\gamma \pm i\alpha). \tag{42}$$

For each $\lambda = 1, 3$, the remaining Hessian and mobility block can be written, after choosing the sign of the memory basis vector, as

$$H_\lambda = \begin{pmatrix} k & 0 & \beta \\ 0 & k & 0 \\ \beta & 0 & h \end{pmatrix}, \qquad D_\lambda = \begin{pmatrix} \gamma & -\alpha & 0 \\ \alpha & \gamma & 0 \\ 0 & 0 & \eta \end{pmatrix}, \tag{43}$$

where

$$k = \mu + \lambda w, \quad h = \kappa(1+s), \quad \beta^2 = 2\kappa\lambda s w, \quad w = \mathrm{e}^{-s}. \tag{44}$$

For a source mode v_j, $j \in \{1, 3\}$, the squared coupling is $w^2 R^2 \sum_e (b_e^T v_j)^2 (b_e^T v_\lambda)^2$. The sum equals $j\lambda/2$, and $R^2 = (4\kappa/j)s\mathrm{e}^s$, giving the displayed β^2. For the highest-mode source the $\lambda = 3$ block couples to even memory and $\lambda = 1$ to odd memory; the low-mode assignment is reversed. Each block retains radial amplitude, phase quadrature and one memory coordinate. Their characteristic polynomials are

$$\begin{aligned} P_\lambda(p) &= (p^2 + 2\gamma k p + |c|^2 k^2)(p + \eta h) - \eta\beta^2(\gamma p + |c|^2 k) \\ &= p^3 + A_2 p^2 + A_1 p + A_0, \\ A_2 &= 2\gamma k + \eta h, \\ A_1 &= |c|^2 k^2 + \eta\gamma(2hk - \beta^2), \\ A_0 &= \eta|c|^2 k(hk - \beta^2). \end{aligned} \tag{45}$$

The decisive Schur curvature is

$$\rho = k - \beta^2/h = \mu - \lambda Q(s). \tag{46}$$

Moreover $2hk - \beta^2 = 2\kappa[(1+s)\mu + \lambda w] > 0$ and

$$\begin{aligned} A_2 A_1 - A_0 = {} & 2\gamma|c|^2 k^3 + \eta|c|^2 k\beta^2 \\ & + \eta\gamma(2\gamma k + \eta h)(2hk - \beta^2) > 0. \end{aligned} \tag{47}$$

Thus $A_2, A_1 > 0$. If $\rho > 0$, all cubic roots lie strictly left. For completeness, hold $A_2, A_1 > 0$ fixed and vary the constant coefficient over $0 < A_0 < A_2 A_1$. At $A_0 = 0$ there is one simple zero and a stable quadratic; a small positive A_0 moves zero left with derivative $-1/A_1$. A nonzero imaginary root during this polynomial continuation would require $p = i\omega$, $\omega^2 = A_1$ and $A_0 = A_2 A_1$, excluded by (47). If $\rho = 0$, zero is simple and the other two roots decay. If $\rho < 0$, $P_\lambda(0) < 0$ and $P'_\lambda(p) > 0$ for $p \geq 0$, giving exactly one positive real root. The remaining quadratic has positive product and negative sum, hence roots with negative real parts. This also supplies an elementary cubic stability argument independent of the general inertia proof.

The two cubics and collective pair account for all eight eigenvalues with multiplicity. They prevent a transverse or collective omission from being hidden by a favorable scalar branch plot. Numerical polynomial roots and direct full-matrix eigenvalues test these different formulations, while sharing the numerical eigensolver backend.

7 Full-state basin regions and powered switching

7.1 Cylinder barrier in the complete state

Fix the three-equilibrium regime of Theorem 5.1. Write V_M for the middle stationary value and $R_M = |v_3^T z_M|$. Denote the low and high equilibria by x_L, x_H.

Theorem 7.1 (Sufficient complete basin criterion)**.** *Every complete initial state satisfying*

$$V(x) < V_M, \quad |v_3^T z| < R_M \tag{48}$$

converges to x_L. Every complete initial state satisfying

$$V(x) < V_M, \quad |v_3^T z| > R_M \tag{49}$$

converges to x_H. These criteria allow arbitrary collective, reflection-odd, phase and memory disturbances subject to the displayed inequalities.

Proof. For every full state, $V(z, m) \geq V_{\text{eff}}(z)$. At fixed highest-mode coordinate u, the collective quadratic is minimized at $t = 0$ and the strictly convex, reflection-even slice in the odd-mode coordinate at $h = 0$. At fixed radius $R = |u|$ the forcing term is minimized by aligning u with a/c. Writing $\theta = |c|/c$, on the entire cylinder $|v_3^T z| = R_M$ we obtain

$$V(z, m) \geq V_{\text{eff}}(uv_3) \geq V_{\text{eff}}(\theta R_M v_3) = V_M. \tag{50}$$

The middle term retains the actual phase of u; the last term minimizes that phase. Each inequality applies to the complete state, including all memory and transverse coordinates.

The strict sublevel $V < V_M$ is forward invariant. Continuity prevents its trajectories from crossing the cylinder (50). Proposition 2.1 gives convergence to one of the three equilibria. Convergence to the middle one is excluded by the strict initial energy inequality and monotonicity. Of the remaining two, only x_L has radius below R_M and only x_H has radius above it. This proves both assertions. □

Both minima lie strictly below V_M. To check this without relying on a plot, restrict the effective potential to the phase-aligned highest-mode ray and call it $W(R)$. Then

$$W'(R) = (\mu + 3\mathrm{e}^{-s(R)})R - a/|c|. \tag{51}$$

Its sign through the three simple roots is successively negative, positive, negative, positive. Integrating between adjacent roots proves $V_L < V_M$ and $V_H < V_M$. Hence the sufficient regions include neighborhoods of both minima. The two open attraction basins and the global stable set of the saddle partition state space, since every trajectory converges to one of these three points. Equations (48)–(49) do not give the full geometry of that stable set or decide membership for every state above V_M.

7.2 Pulse experiments and energy resources

Changing the input across a fold and returning it can place the source in the other attraction basin. This is a dynamical claim about a prescribed finite schedule, rather than a consequence of the stationary diagram alone. The executed schedules are listed in Table 1. They are two independent runs: the upward experiment starts at x_L, the downward experiment at x_H. They are not asserted to be one continuously executed cycle.

Pulse	Initial	a_{pulse}	T_{pulse}	a_{return}	T_{return}
up	low	3.3	300	3.0	600
down	high	2.7	300	3.0	600

Table 1: Prescribed square pulses, followed by return to the identical base input $f = 3v_3$. Time and amplitude use the normalized units of Section 2.

The physical ledger on each complete run is

$$W_{\text{in}}(T) = \int_0^T P\,\mathrm{d}t, \quad Q_{\text{diss}}(T) = \int_0^T (\gamma\,\|Kz\|^2 + \eta|g|^2)\,\mathrm{d}t, \quad \Delta F = W_{\text{in}} - Q_{\text{diss}}. \tag{52}$$

At each input jump the state and F remain continuous, while V changes by the input-dependent linear term in (6). Equation (8) applies segment by segment; no monotonicity across the change of tilt is used. In both executed runs the final state satisfies the strict base-input barrier criterion for its target, as well as a small target-state error. The computations are floating numerical trajectories with convergence checks, not directed interval-flow certificates. They do not establish a minimum switching work or a dynamic quasistatic hysteresis limit.

At any stationary state $Kz = f/c$ and $g = 0$, so

$$P_* = \gamma\frac{\|f\|^2}{|c|^2} > 0 \quad (f \neq 0). \tag{53}$$

For a trajectory converging at fixed input, $P(t) \to P_*$, whence $W_{\text{in}}(T) = P_*T + o(T)$. The stored energy converges and the accumulated dissipation has the same leading growth. Both stable memories are therefore maintained states of an externally powered open system. An unrefilled finite energy source cannot sustain either indefinitely under this contract.

8 Observable response and source-to-readout information

8.1 Conditional carrier versus freely responding memory

For an imposed memory history $m(t)$ the amplitude propagator satisfies

$$\partial_t U(t,s) = -cK(m(t))U(t,s), \quad U(s,s) = I, \quad z(t) = U(t,0)z(0) + \int_0^t U(t,s)f(s)\,\mathrm{d}s. \tag{54}$$

Its instantaneous dissipative and coherent generators are γK and αK, linked by the fixed ratio α/γ. This conditional linkage is not a statement that the complete freely adapting Jacobian is simply $-cK$.

At a stationary source, a small real amplitude-rate probe $r(t)$ gives the complete linear response

$$\dot{\xi} = -GK_R\xi - GB\nu + r, \qquad \dot{\nu} = -\eta B^T\xi - \eta H_m\nu. \tag{55}$$

Eliminating memory in this linear dynamical equation is exact only with its initial condition and causal history retained:

$$\begin{aligned}\dot{\xi}(t) = &-GK_R\xi(t) - GB\mathrm{e}^{-\eta H_m t}\nu(0) + r(t)\\ &+\eta\int_0^t GB\mathrm{e}^{-\eta H_m(t-s)}B^T\xi(s)\,\mathrm{d}s.\end{aligned} \tag{56}$$

In particular the matrix order in the kernel cannot be commuted. With zero initial perturbations the transfer inverse is

$$\chi_{\text{full}}(p)^{-1} = pI + GK_R - \eta GB(pI + \eta H_m)^{-1}B^T. \tag{57}$$

Clamping memory instead gives $\chi_{\text{clamp}}(p) = (pI + GK_R)^{-1}$. At either minimum $H > 0$ and A is Hurwitz. Since $H_m > 0$, block Schur inversion of $pI - A$ shows that (57) is invertible throughout $\operatorname{Re} p \geq 0$. The resulting full response is stable and causal. At zero frequency its operator is

$$\chi_{\text{full}}(0) = S^{-1}G^{-1}, \qquad \chi_{\text{clamp}}(0) = K_R^{-1}G^{-1}. \tag{58}$$

These factors have the indicated order, since the static equation is $GS\xi = r$.

At the middle saddle S is nonsingular, so $S^{-1}G^{-1}$ is still the algebraic derivative of the stationary branch with respect to a probe. It is not a stable long-time susceptibility. The complete Jacobian has one simple positive real eigenvalue. Its right and left eigenvectors must both have nonzero amplitude components: a zero amplitude component would force that positive eigenvalue to belong to $-\eta H_m$, which is negative definite. Hence the pole is visible in the amplitude-to-amplitude transfer operator, and generic free perturbations grow away. At a fold S is singular and even this algebraic static inverse fails. The clamped amplitude dynamics remain stable at all three stationary states because $K > 0$, which does not establish full source stability.

8.2 A precise coarse-readout obstruction

On an incidence graph with a collective mode, the smallest eigenvalue of K is exactly μ. Since $|c|G^{-1}$ is orthogonal in the real amplitude metric,

$$\|\chi_{\text{clamp}}(0)\| = \frac{1}{|c|\mu}, \qquad \left\|S^{-1}G^{-1}\right\| = \frac{1}{|c| \min_j |\lambda_j(S)|} \tag{59}$$

when S is nonsingular. Equal clamped norms can therefore conceal different carriers and contrast responses. Table 3 gives the numerical example; the middle entry is explicitly algebraic.

Define the coarse present summary

$$\mathcal{O}(x_*) = (\bar{z}_*, P_*, \|\chi_{\text{clamp}}(0)\|). \tag{60}$$

Both stable states in the example have $\mathcal{O} = (0, 3, 1/(|c|\mu))$, but different $\|\chi_{\text{full}}(0)\|$. The operators must differ, hence some common small phase-sensitive amplitude-rate probe has different stationary linear responses when memory is free. Thus this coarse summary cannot predict every admitted subsequent free-memory amplitude probe from both source states. The continuation class includes calibrated additive probes, phase-sensitive amplitude access, and a distinction between clamped and free memory. Without those operations no stronger nonclosure claim is inferred merely from the equal summary.

This is a concrete source-to-readout information loss in the sense of Rodgers [2026c,d]. It is not an assertion that amplitude data are intrinsically insufficient. A complete calibrated stationary measurement of z determines each memory through (11). The word stationary and the calibration are essential: arbitrary instantaneous $z(t)$ does not determine an independently evolving $m(t)$, and reconstructing memory at a steady state does not remove the kernel (56) during a free perturbation.

The stable states differ in amplitude, stored energy, s, edge weight e^{-s} and full response. Their conditional coherent and dissipative edge stiffnesses change together. Their common positive supplied power follows from the actuator norm, not from equality of stored state or response. These distinctions connect maintained adaptive carriers to CSCF's linked generators while retaining the richer source needed to predict observations.

9 Fixed example and numerical evidence

9.1 Analytic design and stationary values

All computations use the unchanged path and the fixed normalized parameters

$$(\mu, \kappa, \gamma, \alpha, \eta) = (0.1, 1, 0.6, 1.2, 0.8), \qquad f = 3v_3. \tag{61}$$

Here $\chi \simeq 0.04740589436$, so $\chi < \mu < 3\chi$ and the global strong-convexity margin $\mu - 3\chi$ is negative. The load is exactly $a^2/|c|^2 = 5$. The branch law becomes

$$H_3(s) = \tfrac{1}{75} s\mathrm{e}^s + \tfrac{4}{5} s + 12 s\mathrm{e}^{-s}. \tag{62}$$

Exact rational Taylor bounds, using order 40 and a geometric tail bound, give $H_3(0) < 5 < H_3(1)$, $H_3(3) < 5 < H_3(4)$; in particular $20.08 < \mathrm{e}^3 < 20.09$ suffices for the signed third-point comparison. With the proved two-turn shape, the brackets $(0, 1)$, $(1, 3)$ and $(3, 4)$ isolate exactly the three roots. These are exact scalar inequalities, not interval certificates for complete matrix spectra or trajectories.

Branch	s	R	F	V	n_+
low	0.811988	1.561583	2.405227	-1.086579	0
middle	2.088707	4.742040	9.664461	-0.939063	1
high	3.006220	9.000611	19.100345	-1.025633	0

Table 2: Complete stationary states from the actual calculation. $R = \|z_*\|$, $m_1 = m_2 = -s$, and n_+ counts complete Jacobian eigenvalues in the open right half-plane. Both outer states have eight decaying modes; the middle state has seven decaying modes and one unstable mode.

Branch	e^{-s}	P_*	$\|\chi_{\mathrm{clamp}}\|$	$\|\chi_{\mathrm{full}}\|$
low	0.443975	3.000000	7.453560	7.453560
middle	0.123847	3.000000	7.453560	24.074130
high	0.049478	3.000000	7.453560	29.039222

Table 3: Stored carrier weight, stationary power and static norms. The middle full entry is the norm of an algebraic stationary derivative, not a stable susceptibility. Both outer full entries describe stable full-feedback response. All three means are zero.

Fold	s	a	$l^T X_a$	$\frac{1}{2} l^T X_{xx}[r, r]$
Upper input	1.269082	3.105014	0.398392	0.133396
Lower input	2.585036	2.958897	0.355600	-0.072477

Table 4: Fold inputs and complete dynamical coefficients, with $r = x'(s)$ and $l^T r = 1$ as in Section 6. The upper input corresponds to the smaller fold coordinate. Both coefficients are bounded away from zero in the executed numerical checks.

Table 2 quantifies simultaneous stable states at identical input. The larger carrier amplitude is accompanied by more negative memory, weaker edge stiffness and greater physical stored energy.

The higher-amplitude minimum is not the lower tilted minimum. Table 3 exhibits equal power and clamped norm with different full response. Table 4 places the base input strictly between the two saddle-nodes. Figure 1 combines the analytic branch curve with the executed full-state basin tests.

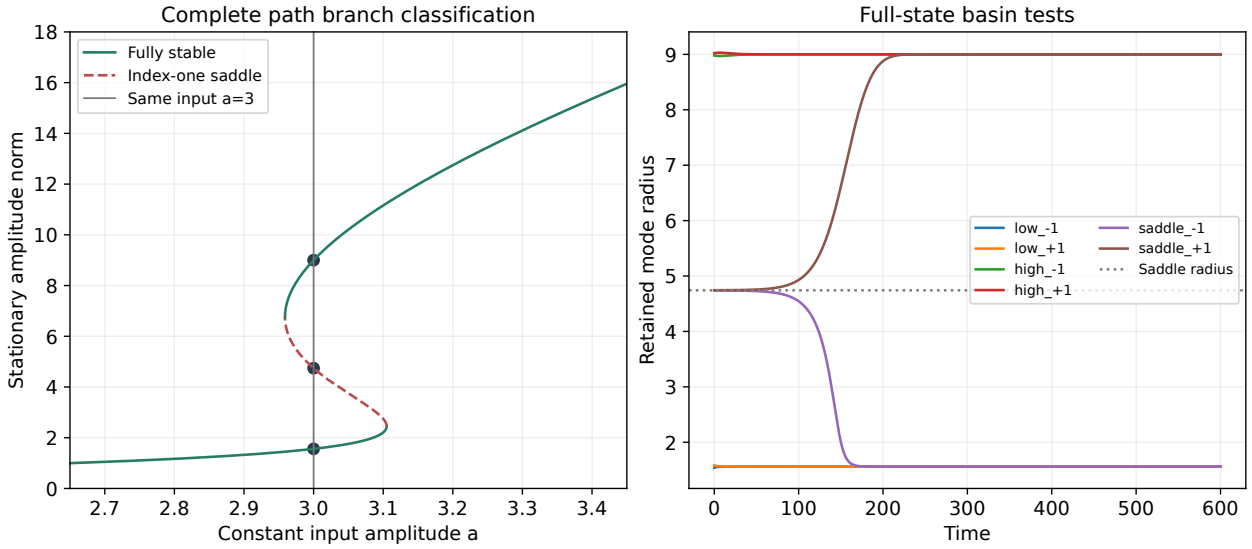


Figure 1: Left: exact branch formulas sampled for display, with the proved full stability classification and the common base input. Right: numerical complete-state preparations approach the two stable radii. The dotted saddle radius is a barrier only when combined with the tilted-energy condition; it is not a global basin separator by radius alone.

9.2 Preparations and matched controls

The real coordinate order is $(\operatorname{Re} z_1, \operatorname{Re} z_2, \operatorname{Re} z_3, \operatorname{Im} z_1, \operatorname{Im} z_2, \operatorname{Im} z_3, m_1, m_2)$. Four preparations use the low or high stationary state plus or minus

$$d = (0.02, -0.01, 0.015, -0.015, 0.01, 0.005, 0.01, -0.015). \tag{63}$$

These perturb collective, reflection-odd, phase and memory coordinates. Two more use the middle state plus or minus 0.001 times its normalized real unstable eigenvector, oriented toward increasing R. Every preparation satisfies the full barrier test on its respective side. Table 5 reports its initial margin, terminal state error and physical ledger residual. These results test full-state dynamics; they are not constrained-sector integrations.

The matched fixed-carrier control sets $m = 0$ permanently and uses $K_0 = \mu I + L$, retaining the same input and starting from the amplitude of either unperturbed low or high adaptive equilibrium. Its exact matrix-exponential solution converges to the same $z_*^{\text{fixed}} = f/(c(\mu+3))$, with norm $\sqrt{5}/3.1$, irrespective of which adaptive equilibrium supplied its initial amplitude. This removes reciprocal feedback while preserving the drive. A second adaptive control raises only μ to $0.4 > 3\chi$, uses the same f and starts from the main high equilibrium plus d. It converges to the unique equilibrium guaranteed by Theorem 3.1. Table 6 includes numerical control values and the separate real-gradient integration comparison.

Preparation	Target	$V_M - V(0)$	Final error	Ledger error
low_-1	low	1.466e-01	3.48e-14	3.02e-11
low_+1	low	1.466e-01	3.53e-14	2.89e-11
high_-1	high	8.575e-02	1.04e-13	8.07e-11
high_+1	high	8.574e-02	1.06e-13	4.18e-11
saddle_-1	low	1.217e-08	2.80e-14	1.94e-09
saddle_+1	high	1.216e-08	1.63e-13	1.11e-09

Table 5: Six full-state basin preparations, evolved to $T = 600$. The final error is the Euclidean complete-state distance from the analytically identified target. The ledger column is the maximum absolute physical energy-balance residual in the sampled run. Displayed diagnostic digits are not acceptance thresholds.

Control	Quantity	Value
Fixed carrier	Stored energy	0.806452
Fixed carrier	Stationary power	3
Unique adaptive	μ	0.4
Unique adaptive	Final error	2.75885e-15
Real-gradient RK4	Difference	3.41061e-13

Table 6: Matched control outputs and the independent real-gradient RK4 comparison. The fixed-carrier power equals that of both adaptive stationary states; its stored energy and amplitude differ. The RK4 difference compares state and physical ledgers over its separately specified $T = 60$ run.

9.3 Switches and their physical ledgers

Table 7 and Figure 2 report the schedules of Table 1. Each pulse runs for 300 time units and the identical base input is then restored for 600. At the final time the upward experiment is in the high-state strict barrier region, and the downward experiment in the low-state strict barrier region. Target errors are below 4×10^{-13} in the retained calculations, and both physical ledger residuals are below 2×10^{-9}. The large cumulative work mostly maintains the driven source over a long interval; it is not the tilted barrier height.

Pulse	Work	Dissipation	ΔF	Final margin
up	2900.664693	2883.969575	16.695118	0.086570
down	2520.086406	2536.781523	-16.695118	0.147515

Table 7: Actual pulse work and dissipation over $T = 900$. Final margin is $V_M - V(x(T))$ at the restored base input. The opposite signs of ΔF reflect different stored endpoint energies, while both runs require positive net supplied work.

9.4 Methods, independent formulations and reproducibility

The primary integration is DOP853 with relative tolerance 10^{-10}, absolute tolerance 10^{-12} and separate solves at pulse jumps. Each prescribed run is repeated at 10^{-12} and 10^{-14}. A source-

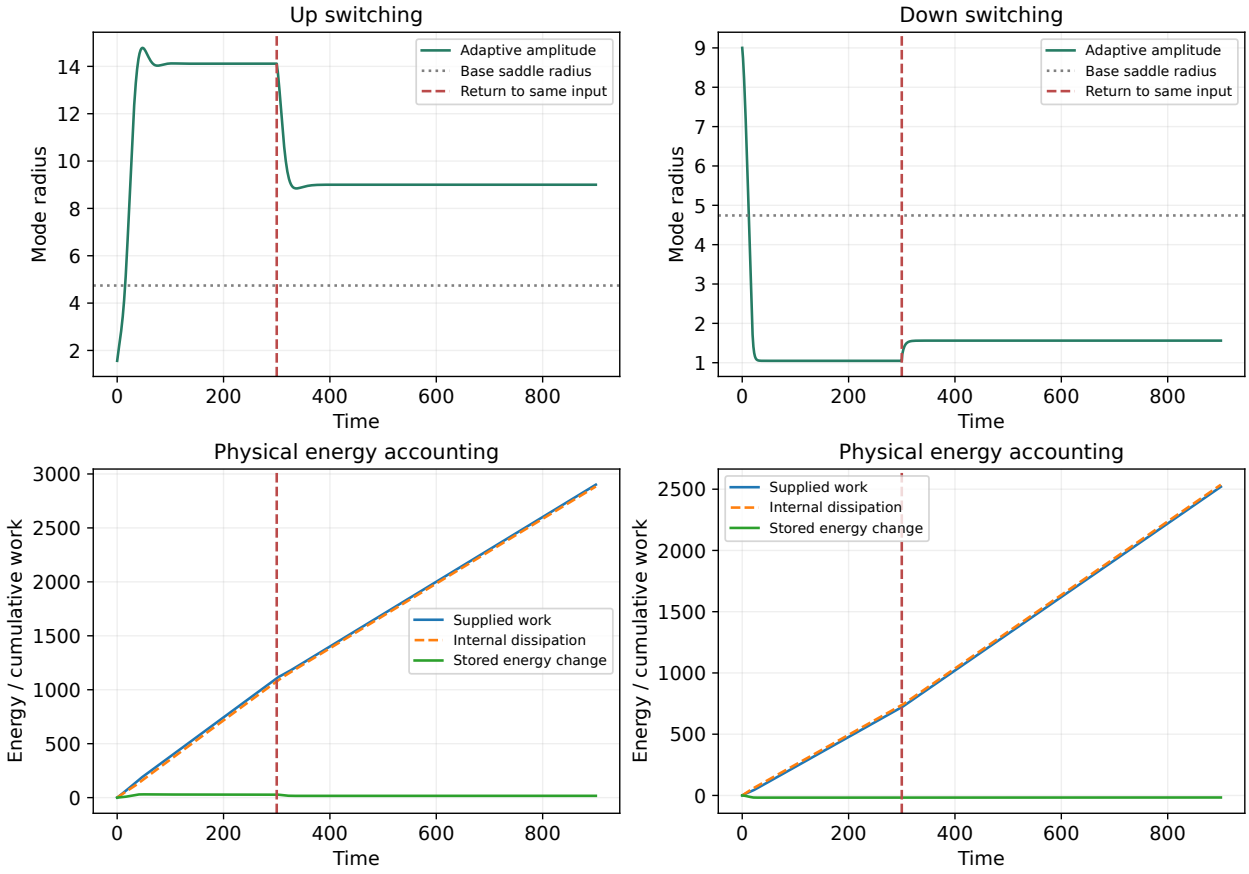


Figure 2: Separate upward and downward switching experiments with physical work ledgers. Dashed vertical lines mark return to the same base input. The displayed difference between cumulative work and dissipation equals the stored-energy change within the numerical ledger tolerance. No slowly ramped input is executed.

derived maximum step is

$$\Delta t_{\max} = \frac{1}{8|c|\,[\mu + \exp(M_0)\,\|L\|]}, \qquad M_0 = \max(0, \max_e m_e(0)). \tag{64}$$

The bound follows because $\dot{m}_e \le -\eta\kappa m_e$ when $m_e \ge 0$, so $m_e(t) \le \max(0, m_e(0))$ and $K(t) \le \mu I + \exp(M_0)L$. It controls the carrier scale; it does not turn an adaptive solver into a rigorous enclosure. Basin trajectories have 1201 stored times over $T = 600$; pulse segments have 601 each with the common endpoint retained once. Physical supplied-work and dissipation ledgers are integrated as extra variables.

Scalar roots use bracketing with absolute tolerance 10^{-13}. Direct eight-coordinate stationary solves use the full Hessian Jacobian and are compared with the scalar constructions. The complete Schur spectrum is checked against (32); central finite differences of the complete vector field use step 10^{-5}, and fold second differences use 10^{-4}. Preserved scientific ceilings include 10^{-9} for stationary residuals and Schur-spectrum differences, 2×10^{-6} for Jacobian finite differences, 2×10^{-5} for fold coefficient differences, and 2×10^{-7} for trajectory convergence and ledger errors. Fold transversality and quadratic magnitudes must exceed 10^{-6}; the zero-eigenvalue tolerance is 10^{-7}. These criteria are in the package specification, rather than chosen from incidental small final digits.

The scalar check separately evaluates exponential laws with 70-digit Decimal arithmetic and 220 bisections, without SciPy or a Lambert-W routine. Fraction arithmetic supplies exact order-40 scalar inequalities, including the rejected low-mode candidate. A separately implemented real-gradient RK4 evolves eight source coordinates and two physical ledgers for $T = 60$, step 0.002, and agrees with a tight DOP853 solution well inside 2×10^{-7}. This does not independently re-integrate the full switching schedules. Six modal fixtures compare the two cubic spectra plus collective pair with complete matrices, including stationary branches, folds and the retained negative candidate. Their maximum observed spectral difference is below 6×10^{-15}, but the two eigenvalue routes share NumPy's backend. Fixed-H mobility fixtures use $\alpha = 0.12, 1.2, 12$ with $f = c\ell$ adjusted to keep the Hessian unchanged. The endpoint fixture has one zero and no unstable eigenvalue.

The retained specification contains 114 numerical conditions. The public runner re-executes them, compares all 20 saved trajectory arrays with relative tolerance 2×10^{-8} and absolute tolerance 2×10^{-10}, checks six cubic fixtures, and verifies tables generated from its own run. Eight checker regressions reject changes to scientific thresholds, states, counts, labels and pass flags, while accepting sub-ceiling diagnostic variation and irrelevant CSV line endings. Eigenvalue comparisons match multisets rather than depending on numerical ordering. A separate saved-identity audit checks the closed fixed-carrier endpoint, energy and response identities; it is a post-result algebraic audit, not a new holdout.

Every numerical table and figure in this paper is generated from the retained actual outputs and declared analytic plotting formulas. The public package contains inputs, specifications, source code, dependencies, reference outputs, hashes, own-run table generation and clean-extraction instructions. Mathematical proofs establish universal and complete-count statements; exact scalar inequalities certify only their scalar scope; floating matrices and trajectories evaluate the worked example; separate formulations and same-code replay provide different forms of checking. No fresh dependency installation, directed trajectory certificate or physical experiment is claimed.

10 Discussion and limits

The sharp threshold is a balance between the baseline stiffness and the most negative stationary edge curvature. Memory weakens an excited edge, but its quadratic cost limits the resulting effective curvature loss. Above and at $\mu = \chi \|L\|$ the effective potential remains strictly convex, so a constant input selects one complete stationary source. Below the path threshold, weakening can create two stable operating states separated by a saddle at the same input. The complete Hessian, not the positive conditional carrier alone, determines which branch survives all perturbations.

The highest path mode makes two distinct reductions compatible. Equal edge magnitudes give an explicit scalar branch, while strict convexity in the odd mode excludes hidden symmetry-broken equilibria when $\mu > \chi$. Lifting through the stationary memory map and the mobility lemma restores the full dynamics. The lower-mode candidate shows why this return matters: sector stability can coexist with a transverse unstable direction. For $\mu \leq \chi$ the outer highest-mode states still prove nonuniqueness, but the argument does not enumerate every possible additional equilibrium. Broader networks below the universal bound require further geometry; the norm criterion is not presented as their individual necessary threshold.

The folds provide genuine complete-system saddle-nodes under the stated hypotheses. Reactive mobility cannot change the inertia at fixed Hessian, yet it changes load at fixed physical input and can move the system through these folds. The distinction has operational consequences: a comparison must specify whether it matches actuator input or stationary load. Likewise a memory clamp changes the response experiment. Conditional linkage of coherent and dissipative generators survives, but full feedback adds the causal kernel and can greatly change static gain.

The full-state barrier supplies explicit sufficient basin regions, not a global formula for the saddle stable set. Numerical square pulses demonstrate switching with a return to the same input and an explicit physical energy ledger. They do not prove a quasistatic limit, minimum work, noisy retention time or experimental realizability. A physical implementation would need a calibrated reciprocal constitutive law, controlled input, accessible amplitude phase and a justified method for comparing clamped and free memory. These requirements are additional resources, not consequences of a formal operator identity.

The result extends the programme from a common carrier, through evolving and externally maintained carriers, to a sharp distinction between unique and multistable maintained states. Its source-to-readout consequence is equally specific: identical coarse present observables can fail to determine a declared future response, while richer calibrated stationary data can restore the hidden memory. Continuous external power sustains both alternatives. This finite model supplies a mathematically complete mechanism in its counted regime and a reproducible test case, without establishing a shared physical origin for otherwise different spectral domains.

References

Robert M. Corless, Gaston H. Gonnet, D. E. G. Hare, David J. Jeffrey, and Donald E. Knuth. On the Lambert W function. *Advances in Computational Mathematics*, 5:329–359, 1996. doi: 10.1007/BF02124750. URL https://doi.org/10.1007/BF02124750.

Yuri A. Kuznetsov. *Elements of Applied Bifurcation Theory*, volume 112 of *Applied Mathematical Sciences*. Springer, 3 edition, 2004. doi: 10.1007/978-1-4757-3978-7. URL https://doi.org/10.1007/978-1-4757-3978-7.

J. P. LaSalle. *The Stability of Dynamical Systems*, volume 25 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. SIAM, 1976. doi: 10.1137/1.9781611970432. URL https://doi.org/10.1137/1.9781611970432.

Jeremy Rodgers. Adaptive spectral carriers: Time ordering, reciprocal memory, and linked coherent–dissipative response, 2026a. URL https://doi.org/10.5281/zenodo.23924603. Zenodo preprint.

Jeremy Rodgers. The canonical spectral-curvature field: A unified operator origin for coherence, dissipation, and spectral budgets, 2026b. URL https://doi.org/10.5281/zenodo.23899458. Zenodo preprint.

Jeremy Rodgers. Source readout non-equivalence: Descent and equivariant reconstruction obstructions, 2026c. URL https://doi.org/10.5281/zenodo.21369967. Zenodo preprint.

Jeremy Rodgers. Observable quotients and exact projected dynamics: Closure, memory, minimal dynamical completion, and effective field operators, 2026d. URL https://doi.org/10.5281/zenodo.21371251. Zenodo preprint.

Jeremy Rodgers. Sustained adaptive spectral states: Global attraction, nonlinear gain, and linked coherent–dissipative response, 2026e. URL https://doi.org/10.5281/zenodo.23927855. Zenodo preprint.

Arjan van der Schaft and D. Jeltsema. *Port-Hamiltonian Systems Theory: An Introductory Overview*, volume 1 of *Foundations and Trends in Systems and Control*. Now Publishers, 2014. doi: 10.1561/2600000002. URL https://doi.org/10.1561/2600000002. Nos. 2–3.