Slow Passage and Prepared Hysteresis in Adaptive Spectral Carriers

Jeremy Rodgers

everythingequation.com

DOI: 10.5281/zenodo.23935694

November 2026

Abstract

Stationary bistability does not by itself determine switching under a slowly varying input. We derive slow passage and a prepared hysteresis return for a finite reciprocal adaptive carrier, retaining its complete complex-amplitude and real-memory dynamics. A fixed-fold projection separates one critical direction from seven stable directions and accounts for parameter motion in an extended centre system. Incoming stable-branch preparation selects an Airy solution and an oriented input delay $K\nu^{2/3} + o(\nu^{2/3})$ at sufficiently small fixed local exit levels. For the declared path example the unfitted prefactors are approximately 6.21890 upward and 7.91547 downward. Preparation exactly at the fold instead selects an Airy mixture with pole constant 1.98635, rather than the tracked-branch constant 2.33811. Instantaneous stationary-memory substitution preserves the folds but changes these dynamical coefficients. A strict-loss argument at each fixed fold supplies the outgoing connection to the other minimum; continuous dependence and complete stable-branch tracking then establish a prepared triangular-cycle limit away from the switching inputs. The geometric loop integral has a nonzero limit, whereas physical maintenance work grows as $1/\nu$. Four uninterrupted numerical cycles, localized events, preparation and uniqueness controls, separate modal dynamics and energy ledgers support the predictions. Numerical exit levels are not assigned a certified chart radius, and fixed-rate periodic-orbit attraction remains open.

1 Introduction

A stationary fold specifies where an equilibrium ends, but a driven trajectory need not leave at that input. Near a saddle-node the restoring rate vanishes, and a slowly swept source can remain near the former equilibrium after its static disappearance. The amount of delay depends on the source dynamics, the incoming preparation and the observation used to declare switching. Connecting that local passage to a complete forward-and-reverse return requires a further global capture argument.

We study these questions in the reciprocal adaptive-carrier model underlying the progression from common spectral carriers to evolving and externally sustained states [Rodgers, 2026c,a,f]. The preceding path analysis established a sharp uniqueness threshold, a completely counted bistable regime, full-state stability and nondegenerate saddle-nodes [Rodgers, 2026b]. Here the constitutive law and external effort–flow pairing remain the same. Only the prescribed scalar input amplitude varies in time. The input fixes the spatial mode, phase reference and sweep; reciprocal memory determines the source's evolving carrier and its switching response.

Slow saddle-node passage and centre-manifold reduction are established mathematical tools. Our use of them has three model-specific obligations. The critical coordinate mixes amplitude

and memory, so stationary memory elimination is not a dynamical approximation without an error estimate. Stable directions must be controlled along a passing centre trajectory, rather than by assuming that this trajectory is an attracting equilibrium. Finally, the local exit must return to a distant complete source state before a prepared hysteresis limit can be asserted. We give those bridges explicitly, attribute the standard reduction tools [Carr, 1981, 2006], and derive the Airy normalization from this source [NIST, 2026].

The preparation distinction is quantitative. Tracking the incoming stable branch from a fixed input distance selects the first Airy zero. Starting at the exact fold selects a different mixture and a different leading coefficient. An observation at a fixed local level has the same leading exponent but a finite-threshold displacement. A physical-amplitude crossing or arrival near the other branch is a separate event whose return conditions must be stated. These distinctions prevent a familiar exponent or loop plot from replacing the actual source calculation.

The paper develops a connected argument from bounded-input energy estimates through full-system fold projection, inner matching, outgoing capture and stable tracking. The resulting limit is stated as convergence on compact input intervals away from folds and in finite-L^p input norms, together with an area and work limit. We do not identify an exact vertical transient trace or assert monotone amplitude jumps. Four finite-rate triangular cycles evaluate the predictions in the full source. Their complete states return near the initial branch, but are not assumed to form an exactly closed periodic orbit at a nonzero rate.

The source-to-readout architecture matters throughout. A retained carrier gives linked conditional coherent and dissipative generators, while the freely adapting source retains amplitude, memory and history [Rodgers, 2026d,e]. The delayed response is not determined by a branch plot alone. This construction supplies a finite, externally powered mechanism and explicit observation contracts; it does not establish a universal microscopic law or provide new evidence about atomic, image or rotating-boundary constructions.

2 Time-dependent source, energy and boundedness

2.1 Model and prescribed input

Let $z \in \mathbb{C}^3$ and $m \in \mathbb{R}^2$. Define

$$b_1 = (1,-1,0)^T, \qquad b_2 = (0,1,-1)^T,$$

$$B_e = b_e b_e^T, \quad L = B_1 + B_2, \quad K(m) = \mu I + \sum_{e=1}^{2} \mathrm{e}^{m_e} B_e, \tag{1}$$

$$F(z,m) = \tfrac{1}{2} z^* K(m) z + \tfrac{\kappa}{2}|m|^2, \qquad g_e = \kappa m_e + \tfrac{1}{2}\mathrm{e}^{m_e}|b_e^T z|^2. \tag{2}$$

All parameters $\mu, \kappa, \gamma, \alpha, \eta$ are positive and fixed. With $c = \gamma + i\alpha$ and the normalized highest path mode

$$v_3 = (1,-2,1)^T/\sqrt{6}, \tag{3}$$

the complete equations are

$$\dot{z} = -cK(m)z + a(t)v_3, \qquad \dot{m} = -\eta g(z,m). \tag{4}$$

The input is real in this fixed phase reference. Memory is a free source coordinate, with the reciprocal effort obtained from the same F that generates the amplitude effort. The carrier obeys $K(m) \geq \mu I$ at every finite state.

We use normalized amplitude, energy and time scales A_0, E_0, T_0, and fixed Euclidean metrics on the real amplitude and memory coordinates. All joint norms and Hessians below refer to these normalized coordinates. Before normalization, z has amplitude units A, F energy units E, K units E/A^2, κ units E, m is dimensionless, and γ, α, η have units $A^2/(ET), A^2/(ET), 1/(ET)$, respectively. The amplitude-rate input a has units A/T; its ramp speed $\nu = |\dot{a}|$ has units A/T^2. Baseline edge weights in (1) have been normalized to one.

Fix positive endpoints $a_\ell < a_h$ and define the uninterrupted triangle

$$a_\nu(t) = \begin{cases} a_\ell + \nu t, & 0 \le t \le T_h, \\ a_h - \nu(t - T_h), & T_h \le t \le 2T_h, \end{cases} \qquad T_h = (a_h - a_\ell)/\nu. \tag{5}$$

There are no holds. The input is continuous at reversal; only its slope changes. The principal preparation is the low stationary state at a_ℓ. Nearby complete initial preparations will be identified separately. The numerical experiment uses

$$(\mu, \kappa, \gamma, \alpha, \eta) = (0.1, 1, 0.6, 1.2, 0.8), \quad (a_\ell, a_h) = (2.9, 3.17). \tag{6}$$

2.2 Physical balance and the varying tilt

The actuator supplies an amplitude-rate flow $f = a(t)v_3$ at conjugate effort Kz. Its signed power and exact physical balance are

$$P(t) = \mathrm{Re}[(Kz)^* f(t)], \qquad \dot{F} = -\gamma \|Kz\|^2 - \eta|g|^2 + P. \tag{7}$$

The ideal external supply is outside the state inventory. It may supply or absorb power; no finite reservoir has been introduced. This is an energy-conjugate effort–flow contract [van der Schaft and Jeltsema, 2014].

Let $\ell(t) = f(t)/c$ and $V(x, a) = F(z, m) - \mathrm{Re}[\ell(t)^* z]$. In real coordinates

$$q = (\mathrm{Re}\, z, \mathrm{Im}\, z), \quad x = (q, m), \quad J = \begin{pmatrix} 0 & -I \\ I & 0 \end{pmatrix}, \quad G = \gamma I + \alpha J, \quad D = \mathrm{diag}(G, \eta I_2), \tag{8}$$

the source is $\dot{x} = -D\nabla_x V(x, a(t))$. Differentiating both state and explicit parameter dependence gives, almost everywhere,

$$\dot{V} = -\gamma \|Kz - \ell\|^2 - \eta|g|^2 - \mathrm{Re}[\dot{\ell}^* z]. \tag{9}$$

The final term is essential during a ramp. At the triangular reversal x, F, V remain continuous and (9) holds on both adjacent segments. The constant-input Lyapunov identity cannot be used to claim ramp monotonicity. Physical work is

$$W(t) = \int_0^t P \,\mathrm{d}s, \qquad Q(t) = \int_0^t (\gamma \|Kz\|^2 + \eta|g|^2) \,\mathrm{d}s, \qquad F(t) - F(0) = W(t) - Q(t). \tag{10}$$

The tilted function is not the stored energy of the actuator or source.

Proposition 2.1 (Bounded-input continuation)**.** *For bounded piecewise continuous $f(t)$, the complete source has a global forward solution. If $\|f(t)\| \le f_{\max}$, its amplitude and memories admit bounds independent of ramp speed for fixed initial state and fixed $f_{\max}$.*

Proof. Set $r_z = \|z\|$. Hermiticity and positivity of K give

$$\tfrac{1}{2}\frac{\mathrm{d}}{\mathrm{d}t}\|z\|^2 \le -\gamma\mu \|z\|^2 + f_{\max}\|z\|, \qquad r_z(t) \le r_z(0)\mathrm{e}^{-\gamma\mu t} + \frac{f_{\max}}{\gamma\mu}(1 - \mathrm{e}^{-\gamma\mu t}). \tag{11}$$

The scalar inequality holds in the upper Dini sense at zero. Put $Z = \max(r_z(0), f_{\max}/(\gamma\mu))$ and $M_e = \max(0, m_e(0))$. At $m_e = 0$ its derivative is nonpositive, while on the positive half-line $\dot{m}_e \le -\eta\kappa m_e$. Hence $m_e \le M_e$. Conversely

$$\dot{m}_e \ge -\eta\kappa m_e - \tfrac{\eta}{2}\mathrm{e}^{M_e}\,\|b_e\|^2\,Z^2. \tag{12}$$

Comparison gives

$$m_e(t) \ge m_e(0)\mathrm{e}^{-\eta\kappa t} - \frac{\mathrm{e}^{M_e}\,\|b_e\|^2\,Z^2}{2\kappa}(1-\mathrm{e}^{-\eta\kappa t}). \tag{13}$$

The vector field is smooth in the state and locally Lipschitz uniformly on compact sets with bounded input. These bounds exclude finite-time escape. Piecewise classical, or equivalently absolutely continuous Carathéodory, continuation supplies the solution through the declared reversals. □

For later use, at fixed input (9) loses its last term and is strictly negative away from stationary points. Bounded trajectories then have only stationary omega-limit points: a nonstationary limit point would have a neighborhood of uniformly positive loss, and bounded velocity gives a positive residence time on infinitely many visits, contradicting convergence of V. The omega-limit set is connected as an intersection of compact connected closures of forward tails. If the stationary set is finite, every trajectory converges to one stationary point. We use this fixed-input result only for a reference connecting orbit, following the invariance-principle reasoning of LaSalle [1976].

3 Stationary geometry and complete fold data

We give the path facts needed for the passage and return, including their complete-state content [Rodgers, 2026b]. At fixed z, the memory Hessian is diagonal with entries

$$(H_m)_{ee} = \kappa + \tfrac{1}{2}\mathrm{e}^{m_e}|b_e^T z|^2 > 0. \tag{14}$$

Thus stationary memory is unique:

$$m_e = -s_e, \qquad s_e\mathrm{e}^{s_e} = |b_e^T z|^2/(2\kappa), \qquad s_e = W_0(|b_e^T z|^2/(2\kappa)). \tag{15}$$

The effective stationary potential is

$$V_{\mathrm{eff}}(z,a) = \tfrac{\mu}{2}\,\|z\|^2 + \sum_e \kappa(s_e + s_e^2/2) - \mathrm{Re}[(av_3/c)^* z]. \tag{16}$$

Here W_0 is the principal real Lambert function [Corless et al., 1996]. This minimization is an exact stationary correspondence. It is not used as a memory law during the ramp.

An edge with $w = \mathrm{e}^{-s}$ has tangential effective curvature w and radial curvature $-Q_0(s)$, where

$$Q_0(s) = \mathrm{e}^{-s}\frac{s-1}{s+1}, \quad Q_0'(s) = \mathrm{e}^{-s}\frac{3-s^2}{(1+s)^2}, \quad \chi = (2-\sqrt{3})\mathrm{e}^{-\sqrt{3}} = \max_{s\ge 0} Q_0(s). \tag{17}$$

Together with $v_0 = (1,1,1)^T/\sqrt{3}$ and $v_1 = (1,0,-1)^T/\sqrt{2}$, the vector v_3 is an orthonormal eigenbasis of L, with eigenvalues $0, 1, 3$.

Lemma 3.1 (Complete counted regime)**.** *Suppose $\chi < \mu < 3\chi$ and $a > 0$. Every complete stationary state for $f = av_3$ has*

$$z = \theta R(s)v_3, \quad m = (-s,-s), \quad \theta = |c|/c, \quad R(s) = \sqrt{\tfrac{4\kappa}{3}s\mathrm{e}^s}, \tag{18}$$

$$a(s) = |c|\sqrt{H_3(s)}, \qquad H_3(s) = \tfrac{4\kappa}{3}s\mathrm{e}^s(\mu + 3\mathrm{e}^{-s})^2. \tag{19}$$

There are two folds $s_- \in (1, \sqrt{3})$ and $s_+ > \sqrt{3}$, determined by $Q_0(s_\pm) = \mu/3$. Let

$$a_{\rm up} = a(s_-), \qquad a_{\rm down} = a(s_+), \qquad a_{\rm down} < a_{\rm up}. \tag{20}$$

Inside that open interval the complete stationary set has exactly two stable minima and one index-one saddle. At either endpoint it has exactly two distinct points: a simple saddle-node and the other stable minimum. Outside the closed interval it has one stable equilibrium.

Proof. Write $z = tv_0 + uv_3 + hv_1$ with complex coordinates. The collective part of (16) is $\mu|t|^2/2$, so stationarity forces $t = 0$. Reflection of nodes 1 and 3 reverses h while fixing u. For every fixed u the real two-dimensional h slice is even and has Hessian at least $(\mu - \chi)I > 0$, since $Lv_1 = v_1$. Its only critical point is $h = 0$. Stationary memory is then equal on the two edges. The amplitude equation fixes its phase and gives (18)–(19). This excludes hidden complex or symmetry-broken equilibria.

Direct differentiation yields

$$H_3'(s) = \tfrac{4\kappa}{3}(\mu + 3\mathrm{e}^{-s})\mathrm{e}^{s}(1+s)[\mu - 3Q_0(s)]. \tag{21}$$

The unique maximum of Q_0 gives two simple turning points, a maximum of H_3 at s_- and a minimum at s_+. Since $H_3(0) = 0$ and $H_3(s) \to \infty$, its three monotone pieces give the asserted counts.

For completeness of the stability classification, the complete Hessian is

$$H = \begin{pmatrix} K_R & B \\ B^T & H_m \end{pmatrix}, \quad K_R = \mathrm{diag}(K, K), \quad B_{:e} = \mathrm{e}^{m_e}\,\mathrm{diag}(B_e, B_e)q, \tag{22}$$

with $H_m > 0$ and Schur complement $S = K_R - BH_m^{-1}B^T = \nabla^2 V_{\rm eff}$. Completing the memory square shows that H has the negative and zero counts of S, plus two positive directions. At an equal-edge state the six Schur eigenvalues are

$$\mu,\ \mu - Q_0,\ \mu - 3Q_0; \qquad \mu,\ \mu + w,\ \mu + 3w. \tag{23}$$

For a real symmetric H and $\mathrm{Sym}\, D > 0$, $-DH$ has the negative count of H in the open right half-plane and its positive count in the open left half-plane; zero is semisimple. To verify the bridge, D^{-1} has positive symmetric part. A nonzero imaginary eigenvector would give $(Hv)^* D(Hv) = -i\omega v^* Hv$, impossible by real parts. A generalized zero vector would give $0 = -v_0^* D^{-1} v_0$ for $Hv_0 = 0$, also impossible. Deforming D to I preserves the zero multiplicity and excludes imaginary-axis crossings, hence preserves the counts, including complex algebraic multiplicities. Equation (23) therefore proves the full classification. At a fold only $\mu - 3Q_0$ vanishes, leaving seven decaying modes. □

At either fold denote the complete vector field by

$$X(x, a) = -DY(x, a), \quad Y(x, a) = \nabla F(x) - a\mathfrak{h}, \quad \mathfrak{h} = (\mathrm{Re}(v_3/c), \mathrm{Im}(v_3/c), 0, 0). \tag{24}$$

Write $x_f = x(s_f)$, $a_f = a(s_f)$ and $\mathcal{L}_f = X_x(x_f, a_f) = -DH$. The inherited normalization is

$$r = x_s(s_f), \qquad d_f = r^T D^{-1} r > 0, \qquad l^T = \frac{r^T D^{-1}}{d_f}, \qquad l^T r = 1. \tag{25}$$

Here $r_m = (-1, -1)$; the critical coordinate already contains dynamic memory. Both amplitude quadratures and memory components are retained in r and its complement. The full load and quadratic coefficients are

$$A_f = l^T X_a = \frac{r^T \mathfrak{h}}{d_f} > 0, \qquad B_f = \tfrac{1}{2} l^T X_{xx}[r, r] = -\tfrac{1}{2} A_f a''(s_f). \tag{26}$$

Indeed differentiating $Y(x(s), a(s)) = 0$ twice at $a' = 0$ gives $Hr = 0$ and $r^T Y_{xx}[r, r] = (r^T \mathfrak{h})a''$, while $r^T \mathfrak{h} = R'/|c| > 0$. From (21), $a'' < 0$ at s_- and $a'' > 0$ at s_+. Thus $B_f > 0$ at the upper-input fold and $B_f < 0$ at the lower one. These nonzero coefficients and the simple zero establish the complete saddle-nodes, rather than only scalar turning points.

Fold	a_f	A_f	B_f	K_f
up	3.105013621	0.398391766	0.133395769	6.218901175
down	2.958896746	0.355600158	-0.072477106	7.915472816

Table 1: Source-derived complete fold coefficients for (6), with normalization (25). The delay prefactor is $K_f = \zeta_1/(A_f|B_f|)^{1/3}$. Both folds have one zero and seven stable complete modes; the computed stable spectral gap is 0.06.

4 Fixed-fold reduction with dynamic memory

4.1 Projection and extended nilpotent structure

Fix one fold and define $P_f = r l^T$, $Q_f = I - P_f$. These are complementary oblique projections. Since $\mathcal{L}_f r = 0$ and $l^T \mathcal{L}_f = 0$, the seven-dimensional space $E_s = \ker l^T$ is $\mathcal{L}_f$ invariant; its restriction A_s is Hurwitz. Decompose

$$x - x_f = ry + w, \qquad y = l^T(x - x_f), \quad w \in E_s, \quad \lambda = a - a_f. \tag{27}$$

All basis vectors and the reference state are fixed at the fold. A moving stationary branch or instantaneous eigenbasis is not being substituted.

The input is affine and independent of state, so Taylor expansion gives

$$\begin{aligned} \dot{y} &= A_f \lambda + B_f y^2 + N_y(y, w), \\ \dot{w} &= A_s w + Q_f X_a \lambda + N_w(y, w), \\ N_y &= O(|y|^3 + |y|\,\|w\| + \|w\|^2), \\ N_w &= O(y^2 + |y|\,\|w\| + \|w\|^2). \end{aligned} \tag{28}$$

Since the input is state independent, $X_{xa} = 0$. On a constant-slope ramp leg introduce $\varepsilon = \dot{a}$ as an extended variable:

$$\dot{\lambda} = \varepsilon, \qquad \dot{\varepsilon} = 0. \tag{29}$$

The extended system has three centre coordinates $(y, \lambda, \varepsilon)$ and seven stable ones. Its centre linear part contains $\dot{y} = A_f \lambda$, $\dot{\lambda} = \varepsilon$: it is nilpotent, not a set of three independent semisimple zero modes.

The source is smooth. The finite-dimensional centre-manifold theorem for a centre spectrum of zero real part and a strictly stable complement applies, including parameter augmentation [Carr, 1981, 2006]. Choosing finite regularity C^4, on a sufficiently small chart there is an invariant graph

$$w = \psi(y, \lambda, \varepsilon), \qquad \psi = O(y^2 + |\lambda| + |\varepsilon|). \tag{30}$$

Its linear coefficients satisfy

$$C_\lambda := \psi_\lambda(0) = -A_s^{-1} Q_f X_a, \qquad C_\varepsilon := \psi_\varepsilon(0) = A_s^{-1} C_\lambda. \tag{31}$$

They follow by inserting $\psi = C_\lambda \lambda + C_\varepsilon \varepsilon + \cdots$ into the graph invariance equation. The C_ε term records parameter motion. All inverses here act on E_s, where A_s is invertible. The resulting scalar equation has

$$\dot{y} = A_f \lambda + B_f y^2 + \mathcal{R}, \quad |\mathcal{R}| \le C(|y|^3 + |y\lambda| + \lambda^2 + |\varepsilon|). \tag{32}$$

This graph describes the stable complement of a mixed amplitude–memory critical direction. It does not set either memory to its stationary Lambert-W value.

4.2 Attraction for a passing trajectory

Centre-manifold existence alone is not a license to assume that a passing centre trajectory is a stable equilibrium. The needed estimate follows directly from invariance.

Lemma 4.1 (Finite-chart stable-direction estimate)**.** *As long as a complete trajectory remains in a sufficiently small fold chart, its deviation $e = w - \psi(y, \lambda, \varepsilon)$ from the invariant graph decays exponentially. The centre coordinate may move through and leave the chart; no stability assumption on its equilibrium is required.*

Proof. Subtract the graph invariance equation at the same actual $(y, \lambda, \varepsilon)$ from (28). Parameter derivatives cancel because both systems have (29). Exactly,

$$\begin{aligned} \dot{e} = {} & A_s e + [N_w(y, \psi + e) - N_w(y, \psi)] \\ & - \psi_y [N_y(y, \psi + e) - N_y(y, \psi)]. \end{aligned} \tag{33}$$

The term $B_f y^2$ cancels in the last difference. The derivatives of the nonlinear terms in w vanish at the origin and are $O(\delta)$ on a chart of radius δ. Also $\psi_y(0) = 0$. Therefore $\dot{e} = [A_s + O(\delta)]e + O(\|e\|^2)$. Choose a fixed positive form P_s on E_s solving $A_s^T P_s + P_s A_s = -I$. For a small enough chart,

$$\frac{\mathrm{d}}{\mathrm{d}t}(e^T P_s e) \le -\tfrac{1}{2} \|e\|^2 . \tag{34}$$

Norm equivalence gives $\|e(t)\| \le C \mathrm{e}^{-g(t - t_0)} \|e(t_0)\|$ for the time spent in the chart. Its contribution to the critical equation is an integrable decaying perturbation. This proves precisely the finite-passage estimate used below. □

For an inner-scale trajectory, variation of constants in the stable equation also gives $w = O(\nu^{2/3})$ when $y = O(\nu^{1/3})$, $\lambda = O(\nu^{2/3})$, after stable transients. The seven directions are thus smaller than the critical coordinate on this scale. They have been controlled through the complete dynamics, rather than clamped.

If coordinates were instead centered on a moving stationary branch $x_*(a)$, their equation would contain $-x_*'(a)\dot{a}$. A parameter-dependent eigenbasis would contribute its corresponding connection term. The fixed-fold representation avoids a singular branch derivative in the local calculation, while retaining its effects through (31) and the extended ramp variables. We retain the moving-reference term explicitly when returning to regular stable-branch segments in Section 7.

5 Inner scaling and preparation-dependent passage

5.1 Orientation and Riccati equation

Let $\sigma_f = \operatorname{sign} B_f$, $b_f = |B_f|$, and orient the local variables by

$$u = \sigma_f y, \qquad h = \sigma_f (a - a_f). \tag{35}$$

On the outgoing ramp $\dot{a} = \sigma_f \nu$, $\nu > 0$, we have $\dot{h} = \nu$ and

$$\dot{u} = A_f h + b_f u^2 + O(u^3 + |uh| + h^2 + \nu). \tag{36}$$

The upward delay therefore has $a > a_{\mathrm{up}}$, while the downward delay has $a < a_{\mathrm{down}}$. Define

$$t_\nu = (A_f b_f \nu)^{-1/3}, \quad h_\nu = \frac{\nu^{2/3}}{(A_f b_f)^{1/3}}, \quad u_\nu = \left(\frac{A_f \nu}{b_f^2}\right)^{1/3}. \tag{37}$$

With $T = (t - t_f)/t_\nu = h/h_\nu$ and $U = u/u_\nu$, (36) becomes

$$\frac{\mathrm{d}U}{\mathrm{d}T} = T + U^2 + O(\nu^{1/3}) \tag{38}$$

on bounded (T, U) sets, after stable transients. The identities $\nu t_\nu = h_\nu$ and $b_f u_\nu h_\nu = \nu$ fix the normalization. If u is dimensionless before restoring physical scales, A_f has units $1/A$ and b_f units $1/T$; then t_ν, h_ν, u_ν have the required units of time, input amplitude-rate and one. Positive rescaling of r changes A_f, b_f, u consistently and leaves $A_f b_f$ invariant. The inherited normalization is kept for both the physical projection observable and its reported coefficients.

5.2 Incoming branch tracking selects the Airy solution

For $h < 0$ the attracting stationary branch and its critical restoring rate have

$$u_*(h) = -\sqrt{-A_f h / b_f} + O(h), \qquad \lambda_c(h) = -2\sqrt{A_f b_f (-h)} + O(h). \tag{39}$$

Preparation on the incoming stable branch at a fixed input distance gives ordinary stable tracking until a shrinking fold neighborhood is reached. In that neighborhood the critical error e_c obeys a comparison of the form

$$D^+ |e_c| \le -c_0 \sqrt{-h}\, |e_c| + C \frac{\nu}{\sqrt{-h}} + C |e_c|^2, \tag{40}$$

once fast errors have settled. The forcing includes the motion of the reference branch, whose derivative is $O((-h)^{-1/2})$. The stable block has a uniform gap and is controlled by Lemma 4.1.

For sufficiently large C_0, the barrier $|e_c| \le C_0 \nu / |h|$ is invariant after the initial transient whenever $\nu / |h|^{3/2}$ is small. At its boundary the negative restoring term dominates the forcing and the quadratic error, while the boundary speed $O(\nu^2 / |h|^2)$ is lower order. Choosing an overlap $h = -\nu^p$, $0 < p < 2/3$, gives $e_c = o(\sqrt{-h})$ and $T \to -\infty$. This prices the loss of restoring rate instead of assuming uniform adiabaticity through the fold.

The limiting Riccati equation has the incoming attracting solution

$$U_\infty(T) = \frac{\mathrm{Ai}'(-T)}{\mathrm{Ai}(-T)}, \qquad U_\infty(T) \sim -\sqrt{-T} \quad (T \to -\infty). \tag{41}$$

Indeed $U = -q'/q$ transforms $U' = T + U^2$ into $q'' + Tq = 0$; $q(T) = \mathrm{Ai}(-T)$ has the required incoming asymptotics [NIST, 2026]. The sign in (41) follows from differentiating $\mathrm{Ai}(-T)$ with respect to T.

The order of matching limits matters. At an entrance $T = -M$ with M fixed and large, (40) gives $U = -\sqrt{M} + O(M^{-1}) + o(1)$ as $\nu \to 0$. Riccati sensitivity along the negative branch contracts by $\exp(\int 2U\,\mathrm{d}T)$, of order $\exp(-4M^{3/2}/3)$ between a distant negative entrance and a fixed inner interval. First let $\nu \to 0$ on compact inner intervals, then let $M \to \infty$. This selects (41); it is not a consequence of simply setting $U(0) = 0$.

Theorem 5.1 (Tracked-preparation local delay)**.** *At either complete fold, prepare the source on the attracting incoming branch at a fixed positive input distance before the fold, or in a sufficiently small fixed attracting neighborhood there. Let $\dot{a} = \sigma_f \nu$ through the fold. For every sufficiently small fixed outgoing level $d > 0$, let a_d be the first crossing of*

$$u = \sigma_f l^T (x - x_f) = d. \tag{42}$$

After the incoming and stable-direction transients specified above,

$$\sigma_f (a_d - a_f) = K_f \nu^{2/3} + o(\nu^{2/3}), \qquad K_f = \frac{\zeta_1}{(A_f b_f)^{1/3}}, \quad \zeta_1 = 2.338107410459\ldots, \tag{43}$$

where $\mathrm{Ai}(-\zeta_1) = 0$ is the first Airy zero in positive T.

Proof. On every compact interval $T < \zeta_1$, the matching argument and continuous dependence in (38) give $U \to U_\infty$. Thus $u = O(u_\nu) \to 0$, and a fixed positive d cannot be reached there. This supplies the lower bound for the exit input.

For the upper bound choose a large outgoing value M and the point $T_M < \zeta_1$ where $U_\infty(T_M) = M$. Inner convergence gives the corresponding crossing of $u = Mu_\nu$ at $T \to T_M$. Bootstrap $h \le C_1 h_\nu$ while $Mu_\nu \le u \le d$. Choose d small enough to absorb the cubic remainder into $b_f u^2/4$. The other remainder terms divided by u^2 are controlled by h/u, h^2/u^2 and ν/u^2, which vanish for fixed M as $\nu \to 0$. Since $h > 0$ on this outgoing interval, the $A_f h$ term helps. Hence $\dot{u} \ge b_f u^2/2$, and

$$\Delta t \le \frac{2}{b_f M u_\nu}, \qquad \frac{\Delta h}{h_\nu} \le \frac{2}{M}. \tag{44}$$

This closes the bootstrap with fixed C_1. Therefore the scaled exit input has limsup at most $T_M + 2/M$. Letting $M \to \infty$ gives ζ_1, proving (43). The argument supplies the stated smaller-order remainder, not an explicit full-source finite-rate error constant. □

5.3 Exact-fold and finite-inner preparation

For a finite inner initial condition (T_0, U_0), the quadratic solution instead has

$$U(T) = \frac{\mathrm{Ai}'(-T) + k\,\mathrm{Bi}'(-T)}{\mathrm{Ai}(-T) + k\,\mathrm{Bi}(-T)}, \qquad k = \frac{\mathrm{Ai}'(-T_0) - U_0\,\mathrm{Ai}(-T_0)}{U_0\,\mathrm{Bi}(-T_0) - \mathrm{Bi}'(-T_0)}, \tag{45}$$

when this finite ratio is defined; otherwise the proportional Airy combination is used directly. The first forward zero of its denominator determines passage if that preparation lies on the outgoing passage class.

Prepare the complete source exactly at x_f, $a = a_f$, and start the outward ramp there. Then $T_0 = U_0 = 0$. Its distance from the moving invariant graph is $O(\nu)$ by (31), which is smaller than the inner critical scale. The Airy initial values give $k = 1/\sqrt{3}$ and

$$\mathrm{Ai}(-\zeta_0) + \mathrm{Bi}(-\zeta_0)/\sqrt{3} = 0, \qquad \zeta_0 = 1.986352707430\ldots. \tag{46}$$

The same compact-inner and outgoing-exit argument now gives

$$\sigma_f(a_d - a_f) = \frac{\zeta_0}{(A_f b_f)^{1/3}}\nu^{2/3} + o(\nu^{2/3}). \tag{47}$$

This leading coefficient is about 15 percent smaller. Neither Airy constant is preparation independent.

Critical preparation errors of order u_ν, or input-reference errors of order h_ν, can change the finite scaled initial condition and hence the leading pole constant. Errors smaller than these scales preserve that initial condition if stable errors have settled, for example to $O(\nu^{2/3})$ in the complement. Unsettled larger stable errors require their own induced-critical-error estimate. By contrast, a small fixed perturbation applied before a long stable approach can decay before the inner region is reached. These are different preparation contracts, both tested below.

6 Switching observables and finite-threshold scope

The operational departure observable in (42) is a calibrated linear functional of the complete source. It uses the same fixed l, x_f, σ_f in the reduced and full systems. It is not an amplitude-only measurement or a stationary-memory reconstruction. The numerical levels are $d = 0.25$ and 0.5. No certified analytical chart radius containing either level has been established. Theorem 5.1 applies to sufficiently small fixed local levels; these declared levels provide finite-rate evidence.

For the exactly quadratic reduced equation, the finite-level prediction is explicit. Let $k = 0$ for tracked preparation or $k = 1/\sqrt{3}$ for exact-fold preparation, and let T_d be the first outgoing solution of

$$\frac{\mathrm{Ai}'(-T_d) + k\,\mathrm{Bi}'(-T_d)}{\mathrm{Ai}(-T_d) + k\,\mathrm{Bi}(-T_d)} = d/u_\nu. \tag{48}$$

Then $h_d^{\mathrm{quad}} = h_\nu T_d$. Near the corresponding simple pole ζ, $U(T) \sim 1/(\zeta - T)$, so at fixed d

$$h_d^{\mathrm{quad}} = \frac{\zeta\nu^{2/3}}{(A_f b_f)^{1/3}} - \frac{\nu}{b_f d} + o(\nu). \tag{49}$$

The cubic source terms and stable graph need not preserve this subleading coefficient. We use (48) as an exact comparison within the quadratic equation, not as a universal next-order full-system theorem or a fitted correction.

For two fixed sufficiently small levels $0 < d_1 < d_2$, (44) gives a traversal time at most $2(1/d_1 - 1/d_2)/b_f$, independent of ν. Their input difference is therefore $O(\nu)$, smaller than the leading delay. A physical-amplitude threshold, such as a specified value of $R = |v_3^T z|$, is a different surface. Its passage inherits the same leading delay only when its crossing is linked to local departure by a justified finite-time return, or by a local transverse observable change. That condition is not supplied by a loop plot alone.

Arrival also has several meanings. Entry into a fixed-radius neighborhood of the other stable state can occur a bounded additional time after a fixed local exit, once the connection in Section 7 is established, giving an $O(\nu)$ input correction. On a regular stable-branch segment, a shrinking tube of radius $C\nu$ must be large enough to contain the $O(\nu)$ moving-reference offset. For such a tube, exponential settling gives an additional $O(|\log\nu|)$ time and an $O(\nu|\log\nu|)$ input cost; a smaller tube need not be entered. The executed calculations localize projection exits and check complete-state capture at the reversal and return endpoints. They do not report the first instant of asymptotic arrival or a separate physical-amplitude event.

7 Capture and the prepared quasistatic return

7.1 The fixed-input outgoing connection

The delay theorem is local. It does not, by itself, identify the distant state reached after exit. In this model the complete stationary count and strict constant-input loss supply that bridge.

Lemma 7.1 (Complete connection at a fold)**.** *For $\chi < \mu < 3\chi$, the outgoing local centre orbit at either fixed fold input converges forward to the other stable minimum of the complete source.*

Proof. At $a = a_f$, the local centre equation is $\dot{u} = b_f u^2 + O(u^3) > 0$ for sufficiently small $u > 0$. Its outgoing orbit has backward limit x_f. It follows the complete centre graph, not the stationary-memory curve. At every finite point of this nonstationary orbit, integration of the strict fixed-input loss in (9) from the backward limit gives $V(x, a_f) < V(x_f, a_f)$.

Proposition 2.1 bounds the entire forward orbit. Lemma 3.1 gives precisely two distinct stationary points at the fold input. The fixed-input finite-stationary-set argument in Section 2 forces convergence to one of them. Convergence back to x_f is impossible because the current tilted value is already strictly below its value. The only remaining point is the other stable minimum. No monotonicity of amplitude radius or global planar basin boundary has been assumed. □

At a fixed small local exit $u = d$, Lemma 4.1 and the passage estimate imply that the ramped exit state tends to the corresponding outgoing centre-graph point at $(a, \varepsilon) = (a_f, 0)$. Choose a finite segment of the reference connection in Lemma 7.1 entering a small hyperbolically attracting ball around the other minimum. During this fixed time segment the slowly ramped parameter changes by $o(1)$. Continuous dependence transfers nearby ramp exits into that ball for sufficiently small ν. The ball can be contained in a smooth attracting tube of the other stationary branch for nearby inputs. Constant-input loss is used only to establish this reference connection; no ramp Lyapunov monotonicity is asserted.

7.2 Tracking regular branch segments and the reversal

On a compact stationary-branch interval away from its terminating fold, the complete Jacobian $A(a)$ is uniformly Hurwitz and the branch $x_*(a)$ is smooth. In its moving coordinates $e = x - x_*(a)$,

$$\dot{e} = A(a)e - x_*'(a)\dot{a} + O(\|e\|^2). \tag{50}$$

The reference-motion forcing is retained. Choose the smooth positive matrix $P(a)$ solving

$$A(a)^T P(a) + P(a)A(a) = -I. \tag{51}$$

Its norm, inverse norm and derivative are bounded on this compact interval. For $E = e^T P(a) e$, differentiation gives

$$\dot{E} \le -cE + C\nu\sqrt{E} + O(E^{3/2}), \tag{52}$$

since $\dot{P} = O(\nu)$. A sufficiently small tube is invariant for small ν, and after a decaying transient $\|e\| = O(\nu)$. At the triangle reversal the source state and a are continuous, while $\dot{a}$ changes sign. Equation (50) still has bounded forcing of size $O(\nu)$. We apply the fold chart only on the constant-slope portions near each fold, and this regular tracking argument through the reversal.

Theorem 7.2 (Prepared triangular-cycle limit)**.** *Fix positive parameters with $\chi < \mu < 3\chi$ and endpoints*

$$0 < a_\ell < a_{\mathrm{down}} < a_{\mathrm{up}} < a_h. \tag{53}$$

Initialize at the low equilibrium $x_L(a_\ell)$ and use (5). For sufficiently small ν, the complete source undergoes both local passages of Theorem 5.1, is captured by the other stable branch on each leg, and returns to $x_L(a_\ell)$ with error tending to zero.

Let $x_\nu^{\uparrow}(a)$ and $x_\nu^{\downarrow}(a)$ denote the complete states on the two monotone legs as functions of input. They converge uniformly on compact input subintervals away from the corresponding static fold input to

$$x_0^{\uparrow}(a) = \begin{cases} x_L(a), & a < a_{\mathrm{up}}, \\ x_H(a), & a > a_{\mathrm{up}}, \end{cases}$$
$$x_0^{\downarrow}(a) = \begin{cases} x_L(a), & a < a_{\mathrm{down}}, \\ x_H(a), & a > a_{\mathrm{down}}. \end{cases} \tag{54}$$

The convergence also holds in $L^p([a_\ell, a_h])$ for every finite $p \geq 1$. The switching-input asymptotics at sufficiently small fixed local levels are (43). No conclusion of attraction to a periodic orbit at a fixed nonzero rate is included.

Proof. Incoming full stability and (50) take the initial preparation to the upper fold neighborhood. The matching argument gives its tracked Airy preparation and local delayed exit. Lemma 4.1 identifies the limiting outgoing exit point, and Lemma 7.1 with finite-time continuous dependence supplies capture by the high branch. The attracting tube then carries the source to a_h, through reversal and down to the lower fold neighborhood. The same argument there gives the second delayed exit and capture by the low branch. A final fixed positive input distance before a_ℓ allows stable tracking to remove the transient, proving return to the low state.

For any compact input interval after capture and a fixed distance from the corresponding static fold input, the available tracking time grows like $1/\nu$. Before the first passage the exact stationary preparation supplies the incoming tracking estimate. Equations (50)–(52) give the uniform branch convergence in (54). The global bounded-input estimates are uniform in ν, so dominated convergence gives the finite-L^p statement. The two local input offsets tend to zero with the coefficients already proved; a global trajectory enclosure has not been used. □

A sufficiently small fixed complete perturbation of the initial preparation has the same limiting branch selection after its initial attracting transient. For that variant uniform convergence is asserted away from the initial endpoint as well as the folds; the initial state itself need not equal the branch state. The numerical transverse-preparation test concerns such an early perturbation, not a disturbance applied at the fold.

7.3 The sense of the geometric limit

Theorem 7.2 specifies convergence of the two input-parametrized legs away from folds and in finite-L^p norms. It does not identify every point of a transient trace at the fold, or assert a monotone radius jump. The reference connection can have amplitude overshoot. A bounded fast excursion occupies an input interval whose width tends to zero, so it does not alter the integral limit.

Let $R = |v_3^T z|$ and define the magnitude of the oriented loop integral by

$$\mathcal{A}_\nu = \left| \int_{a_\ell}^{a_h} R_\nu^{\uparrow}(a)\, \mathrm{d}a - \int_{a_\ell}^{a_h} R_\nu^{\downarrow}(a)\, \mathrm{d}a \right|. \tag{55}$$

At a finite rate the state at return need not equal its initial state. A vertical closing segment at $a = a_\ell$ contributes zero to this integral. Thus (55) does not assume an exactly closed source

orbit, nor equal total enclosed area for a self-intersecting trace. It is the stated oriented-integral diagnostic.

Uniform boundedness and (54) give

$$\mathcal{A}_\nu \longrightarrow \mathcal{A}_0 = \int_{a_{\mathrm{down}}}^{a_{\mathrm{up}}} [R_H(a) - R_L(a)]\,\mathrm{d}a > 0. \tag{56}$$

The strict inequality follows from the ordering of the two distinct outer radii. This prepared quasistatic statement is stronger than a finite-rate loop plot, but differs from a theorem about a unique attracting periodic orbit or arbitrary finite-rate basins.

8 Why instantaneous stationary memory changes the delay

The stationary reduction (15) preserves the branch geometry. It need not preserve fold passage. Consider the tempting amplitude-only dynamics obtained by imposing stationary memory at every time:

$$\dot{q} = -G\nabla_q V_{\mathrm{eff}}(q, a). \tag{57}$$

We do not adopt (57) for the experiment. It is an explicit test of the proposed deletion of dynamic memory.

Retain the same branch parameter s and write the complete right vector as $r = (r_q, -1, -1)$. Its full normalization denominator is

$$d_{\mathrm{full}} = r^T D^{-1} r = d_0 + 2/\eta, \qquad d_0 = r_q^T G^{-1} r_q = \frac{\gamma}{|c|^2}\,\|r_q\|^2 > 0. \tag{58}$$

The amplitude-only system has denominator d_0, while its stationary $a''(s_f)$ is unchanged. Applying the same fold calculation therefore gives

$$A_{\mathrm{inst}} = A_f \frac{d_{\mathrm{full}}}{d_0}, \qquad B_{\mathrm{inst}} = B_f \frac{d_{\mathrm{full}}}{d_0}, \qquad K_{\mathrm{inst}} = K_f \left(\frac{d_0}{d_{\mathrm{full}}}\right)^{2/3}. \tag{59}$$

At every finite $\eta > 0$ the two prefactors differ. The deleted contribution $2/\eta$ is the mobility-weighted effect of the two evolving memory coordinates in the mixed critical vector. This failure concerns the requested dynamical observable despite exactly preserved folds. It is not a statement about an independently justified $\eta \to \infty$ singular limit.

Fold	d_0	d_{full}	K_{inst}	K_{full}
up	1.603701	4.103701	3.324136	6.218901
down	7.327314	9.827314	6.508539	7.915473

Table 2: The source-derived memory obstruction for the fixed example. Instantaneous stationary-memory dynamics preserves the static fold inputs but predicts different leading delay coefficients. No such altered source was used in the ramp experiments.

The seven-direction invariant graph in Section 4 has a different status from (57): it is derived from the complete extended vector field, retains parameter-motion corrections, and is accompanied by a stable-error estimate. The independent three-state calculation below is different again: it is an exact symmetry restriction that retains a dynamic memory equation. These reductions preserve different objects and must not be interchanged.

9 Physical work and the evolving carrier

At any stationary point $Kz = av_3/c$ and $g = 0$. The supplied power is therefore

$$P_*(a) = \gamma a^2/|c|^2, \tag{60}$$

independent of which stable branch is occupied. Equal stationary power does not mean equal amplitude, memory, stiffness or response.

Proposition 9.1 (Slow-cycle maintenance work)**.** *Under the exact stationary preparation and hypotheses of Theorem 7.2, over the complete triangular round trip,*

$$\nu W_\nu \longrightarrow C_W = \frac{2\gamma}{|c|^2}\int_{a_\ell}^{a_h} a^2\,\mathrm{d}a, \qquad \nu Q_\nu \longrightarrow C_W, \qquad \Delta F_\nu \longrightarrow 0. \tag{61}$$

The same leading W, Q coefficient holds for nearby fixed preparations with the same branch selection, but their unscaled ΔF need not tend to zero.

Proof. The bounds of Proposition 2.1 are independent of ν, and bound K, g, P as well. On each leg use input as the integration variable, with $\mathrm{d}t = |\mathrm{d}a|/\nu$. At every input away from the two folds the complete state converges to its stationary branch and its power to (60). Dominated convergence gives the first limit, with a factor of two for the two legs. The exact stationary preparation and return give $F(2T_h) \to F(0)$, hence $\Delta F \to 0$. The physical ledger (10) gives the dissipation limit.

For a fixed perturbed initial state, the limiting final stored energy is instead $F(x_L(a_\ell))$, so $\Delta F \to F(x_L(a_\ell)) - F(x(0))$ in general. It remains uniformly bounded and $\nu\Delta F \to 0$, which preserves the leading work and dissipation coefficient. □

Thus slower traversal approaches the static folds while requiring more maintenance work: $W_\nu \sim C_W/\nu$. The finite geometric area in (56) has different units and a different actuator pairing. It is not identified with dissipation or minimum switching work. The prescribed input supplies the spatial mode, phase and traversal. The reciprocal source contributes the carrier change, preparation-dependent delay and branch capture; it does not create an autonomous energy supply.

For a prescribed memory history the conditional amplitude propagator obeys

$$\partial_t U(t,s) = -(\gamma + i\alpha)K(m(t))U(t,s), \qquad U(s,s) = I. \tag{62}$$

Its instantaneous coherent and dissipative generators remain $\alpha K(t)$ and $\gamma K(t)$, with their inherited fixed ratio. A freely adapting perturbation is richer. At a stationary source, with the blocks B, H_m from Section 3, the complete tangent equations are

$$\dot\xi = -GK_R\xi - GB\omega + r_{\mathrm{probe}}, \qquad \dot\omega = -\eta B^T\xi - \eta H_m\omega. \tag{63}$$

Eliminating ω retains its initial condition and the causal kernel $\eta GB\mathrm{e}^{-\eta H_m(t-s)}B^T$. It is not clamping, and does not make complete response a scalar function of the instantaneous carrier. The fold denominator in (58) is one direct manifestation of that retained source dynamics. This connects the result to CSCF's evolving carrier and source-to-readout distinction without promoting the finite constitutive example to an experimentally established universal law.

10 Executed numerical tests and reproducibility

10.1 Frozen experiment and actual event locations

The fixed example is (6), with exact low-branch preparation at $a_\ell = 2.9$. The four declared rates are

$$\nu \in \{10^{-4}, 3\times10^{-5}, 10^{-5}, 3\times10^{-6}\}. \tag{64}$$

Every run traverses $2.9 \to 3.17 \to 2.9$ with immediate reversal and no hold. The model, rates, preparations, thresholds, coefficient formulas, comparison rules and numerical settings were fixed before the ramp calculations. Stationary and square-pulse results supplied development data. All experiments here are synthetic numerical tests.

The primary solver is Radau on eight real source coordinates and four integrated ledgers: dissipation, supplied work, tilted loss, and the explicit term $\int \mathrm{Re}(\dot{\ell}^* z)\,\mathrm{d}t$. It uses the analytic twelve-state Jacobian, relative tolerance 10^{-10} and absolute tolerance 10^{-12}. Each leg is repeated at 10^{-12} and 10^{-14} with half the primary maximum step. The primary cap is

$$\Delta t_{\max} = \min(50,\, 2\times10^{-4}/\nu). \tag{65}$$

This limits input increments while allowing the implicit solver to control stable fast modes. It is not a certified time-step theorem. Switching surfaces (42) are located by dense-output event roots with positive crossing direction. The 501 samples per leg are used for display and ledger reconstruction, not event detection.

ν	d	Δa	$K_f\nu^{2/3}$	Ratio	Quadratic exit
1×10^{-4}	0.25	0.010801	0.013398	0.80616	0.010817
1×10^{-4}	0.50	0.012224	0.013398	0.91235	0.011967
3×10^{-5}	0.25	0.005203	0.006004	0.86659	0.005172
3×10^{-5}	0.50	0.005709	0.006004	0.95080	0.005564
1×10^{-5}	0.25	0.002625	0.002887	0.90942	0.002599
1×10^{-5}	0.50	0.002809	0.002887	0.97323	0.002738
3×10^{-6}	0.25	0.001220	0.001294	0.94290	0.001205
3×10^{-6}	0.50	0.001278	0.001294	0.98788	0.001249

Table 3: Upward input delays at both declared projection levels. Ratio is actual delay divided by the unfitted leading prediction. The last column is the quadratic equation's finite-threshold prediction (48), not a universal full-source next-order law.

ν	d	Δa	$K_f\nu^{2/3}$	Ratio	Quadratic exit
1×10^{-4}	0.25	0.012916	0.017053	0.75738	0.012789
1×10^{-4}	0.50	0.014924	0.017053	0.87513	0.014527
3×10^{-5}	0.25	0.006289	0.007642	0.82295	0.006207
3×10^{-5}	0.50	0.007031	0.007642	0.91996	0.006850
1×10^{-5}	0.25	0.003208	0.003674	0.87316	0.003163
1×10^{-5}	0.50	0.003485	0.003674	0.94860	0.003404
3×10^{-6}	0.25	0.001507	0.001646	0.91525	0.001487
3×10^{-6}	0.50	0.001596	0.001646	0.96939	0.001565

Table 4: Downward oriented delays $a_{\mathrm{down}} - a_d$. Every declared rate and both thresholds are retained. All switches occur below the static lower fold.

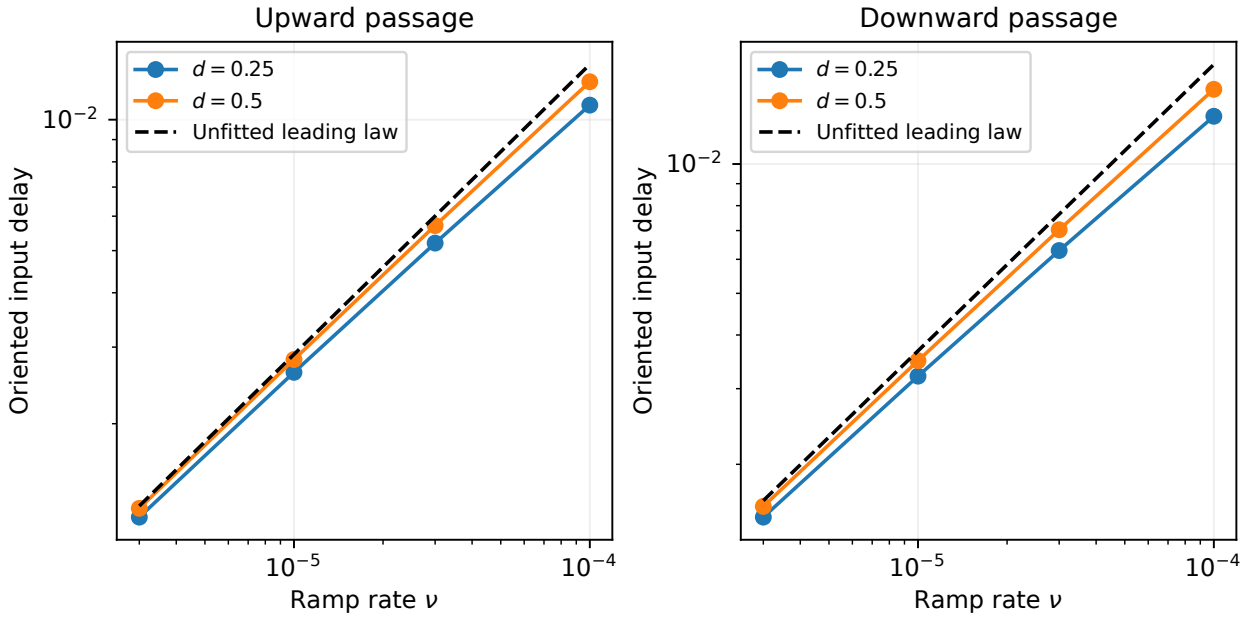


Figure 1: Full-source delays approach the source-derived two-thirds-power law as the ramp slows. The two thresholds retain finite-rate differences. Dashed curves use the fixed coefficients of Table 1, with no fit to these events.

Tables 3–4 show the correct delay direction at both folds. At the slowest rate the actual-to-leading ratios for $d = 0.25$ are about 0.943 upward and 0.915 downward; for $d = 0.5$ they are about 0.988 and 0.969. The difference from the leading value decreases over the declared rate set. Fits of log delay against log rate give exponents approximately 0.613–0.644, depending on fold and threshold. These finite-regime diagnostics reflect unresolved subleading effects; they neither prove the exponent nor supply its prefactor. The asymptotic exponent is derived in Theorem 5.1.

10.2 Preparation, complete-state controls and independent formulation

At $\nu = 3 \times 10^{-6}$, two separate outward runs start at the exact complete fold states and follow their outward slopes for an input distance 0.035. Table 5 compares those events with (47) and the tracked-branch prediction. Their smaller leading delay is not attributed to a change of fold location or memory law.

Fold	d	Δa	Fold-start lead	Ratio	Tracked lead
up	0.25	0.0010298	0.0010990	0.93703	0.0012936
up	0.50	0.0010884	0.0010990	0.99035	0.0012936
down	0.25	0.0012635	0.0013988	0.90331	0.0016465
down	0.50	0.0013536	0.0013988	0.96768	0.0016465

Table 5: Exact-fold preparations at the slowest declared rate. Ratio uses the preparation-specific leading coefficient with ζ_0, while the last column shows the distinct tracked-branch leading value with ζ_1. Both are source predictions, not fitted constants.

The transverse-preparation test adds the following components to the initial low state of the slowest full cycle:

$$\begin{aligned}\delta(\operatorname{Re} z) &= (0.002, -0.001, 0.0015),\\ \delta(\operatorname{Im} z) &= (-0.0015, 0.001, 0.0005),\\ \delta m &= (0.001, -0.0015)\end{aligned} \tag{66}$$

This perturbs collective, reflection-odd, phase and memory coordinates before the long stable approach. Its event changes are numerically unresolved at the demonstrated integration accuracy. This is not immunity to a kick applied inside the fold layer or to noise.

The uniqueness control changes only μ to 0.4 and uses the fastest identical triangle. Here $\mu > 3\chi$, so the effective Hessian bound $\nabla^2 V_{\text{eff}} \geq (\mu - 3\chi)I$ and strict stationary-memory convexity give a unique equilibrium. The complete adaptive control still evolves both memories. The fixed-carrier control instead holds $K_0 = \mu I + L$ fixed and solves the identically driven linear amplitude equation exactly. On each affine input segment, with $k = \mu + 3$, slope $v = \dot{a}$ and initial modal amplitude u_0,

$$u(t) = \frac{a(t)}{ck} - \frac{v}{(ck)^2} + \left(u_0 - \frac{a(0)}{ck} + \frac{v}{(ck)^2}\right) \mathrm{e}^{-ckt}. \tag{67}$$

The small control areas in Table 6 separate finite-rate linear or single-state lag from the bistable cycle.

For an exactly symmetric preparation, ordinary reflection and zero collective amplitude give the invariant subspace $z = uv_3$, $m_1 = m_2 = m$. Its exact dynamic equations are

$$\dot{u} = -c(\mu + 3\mathrm{e}^m)u + a(t), \qquad \dot{m} = -\eta[\kappa m + \tfrac{3}{4}\mathrm{e}^m |u|^2]. \tag{68}$$

A separately coded three-real-coordinate BDF calculation of (68) uses relative tolerance 10^{-11}, absolute tolerance 10^{-13} and half the primary step cap. It retains memory dynamically and compares the four slowest-cycle event inputs with the complete Radau calculation. This exact symmetry restriction is licensed only for the symmetric preparation; it does not replace the full-coordinate perturbation experiment. The two formulations share SciPy and are not a second-platform execution.

Check	Value
Maximum event refinement difference	3.730×10^{-14}
Independent modal BDF difference	3.840×10^{-12}
Full Jacobian finite-difference error	1.042×10^{-9}
Initial perturbation event difference	3.109×10^{-15}
Fixed-carrier sampled area	1.394×10^{-6}
Unique adaptive sampled area	4.569×10^{-5}

Table 6: Actual quality and control diagnostics. The fixed and uniqueness-control areas are magnitudes of sampled oriented integrals. Event refinements, independent modal events and initial-perturbation comparisons have distinct purposes; small values do not certify trajectories by directed arithmetic.

10.3 Continuous cycles, capture and work

Figure 2 shows every declared uninterrupted triangular cycle. Table 7 gives complete-state errors against the stationary high state at reversal and stationary low state at return. These errors quantify finite-rate lag and capture, rather than exact closure at nonzero ν.

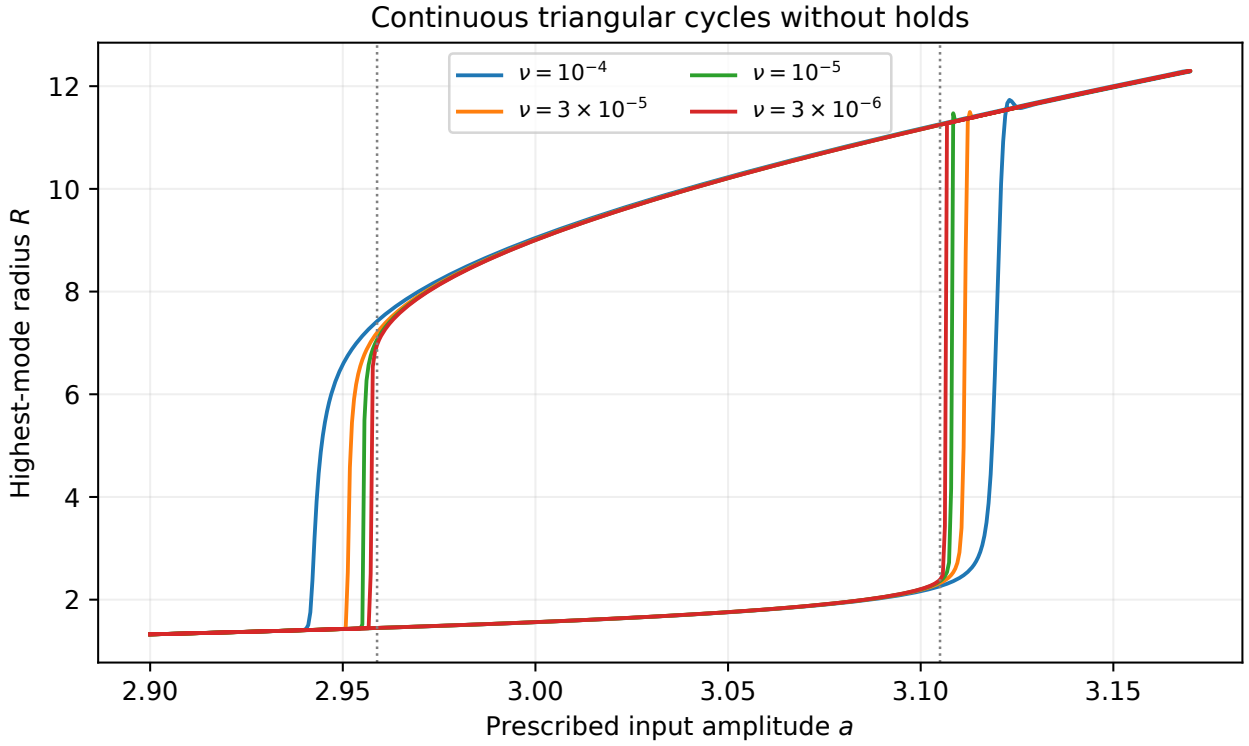


Figure 2: Complete-source forward-and-reverse trajectories at all four rates. Dotted lines are the two static fold inputs. The plotted radius is sampled for display; switching locations in Tables 3–4 are localized independently. Possible transient overshoot is not excluded.

ν	Turn error	Return error	Sampled area	νW
1×10^{-4}	1.296×10^{-2}	8.541×10^{-4}	1.376522	1.65917591
3×10^{-5}	3.880×10^{-3}	2.557×10^{-4}	1.255423	1.65913088
1×10^{-5}	1.293×10^{-3}	8.520×10^{-5}	1.204922	1.65911939
3×10^{-6}	3.877×10^{-4}	2.556×10^{-5}	1.180157	1.65911558

Table 7: Complete-state capture errors and geometric/work diagnostics. Turn and return errors compare with $x_H(3.17)$ and $x_L(2.9)$. Sampled area is (55), evaluated by trapezoids on the declared grid.

ν	W	Q	ΔF	Ledger ratio
1×10^{-4}	16591.75915	16591.75789	1.260×10^{-3}	1.48×10^{-15}
3×10^{-5}	55304.36251	55304.36213	3.772×10^{-4}	2.62×10^{-15}
1×10^{-5}	165911.93921	165911.93909	1.257×10^{-4}	2.65×10^{-15}
3×10^{-6}	553038.52752	553038.52749	3.769×10^{-5}	1.23×10^{-15}

Table 8: Physical supplied work, dissipation and stored-energy change over each complete round trip. The ledger ratio is the maximum per-leg physical or tilted residual divided by $1+|W(t)|+|Q(t)|$. These are actual balances, not an identification of area with energy.

Stationary-root quadrature gives $\mathcal{A}_0 \simeq 1.15880536$ for (56). The sampled finite-rate areas approach it, but their trapezoid error is not a certified area enclosure. The work coefficient in (61) is $C_W = 1.659114$ in the normalized example. The slowest computed νW differs from this coefficient by less than 10^{-6} in relative terms. Figure 3 contrasts these different limits.

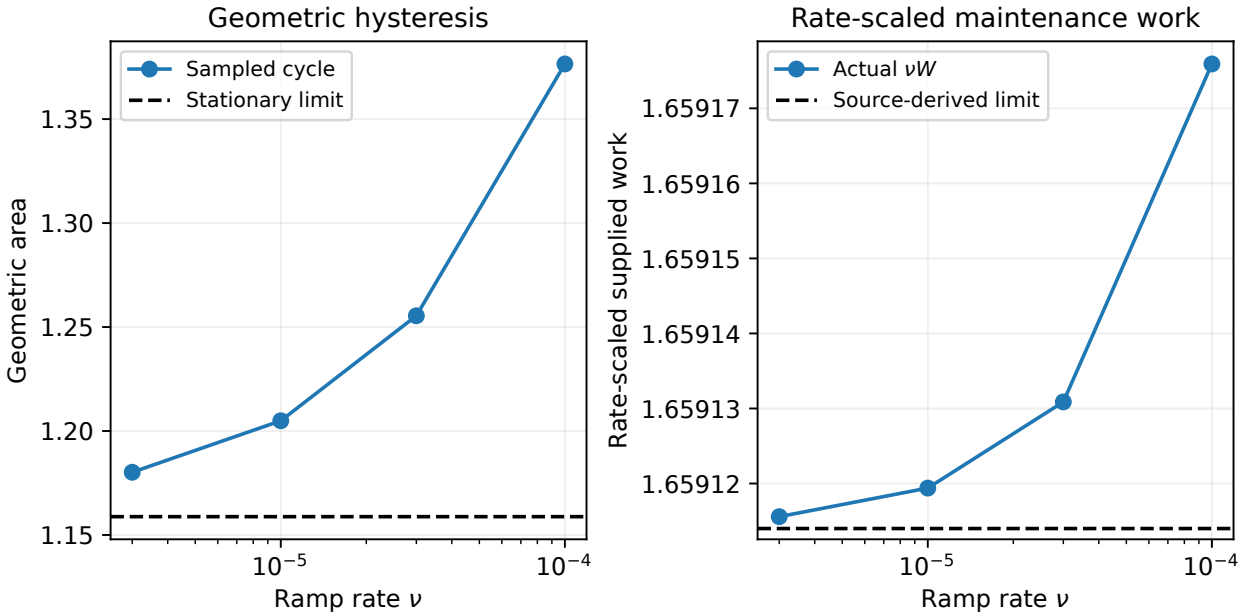


Figure 3: Left: magnitude of the sampled oriented loop integral and stationary-branch limit. Right: rate-scaled supplied work and its source-derived limit. The area remains finite while unscaled maintenance work grows like $1/\nu$. Neither panel is a minimum switching-work claim.

10.4 Verification and correction traceability

The declared criteria include 2×10^{-8} for switching-input refinement and independent modal event differences, 2×10^{-6} for refined endpoint differences, 2×10^{-9} for scaled ledger errors, and 2×10^{-7} for the early-preparation event comparison. Leading-law relative-error criteria are 0.35 at the larger rates and 0.20 at the smallest. These finite-regime criteria were specified before the ramp outputs; they are not explicit error constants in Theorem 5.1. All 113 frozen comparisons passed. Each fold has one zero and seven stable modes, and the analytic source-plus-ledger Jacobian agrees with central finite differences at step 10^{-6}. Reflection equivariance is also checked.

The accompanying public verification archive contains the code, settings, raw and checked outputs, correction records and generated assets. After extraction, run

```
python run_all.py --output ../new_replay
```

using a new empty output directory. The runner re-executes the full source and refinements, separate modal BDF calculation, local Jacobian checks and asset generation. It compares all 24 trajectory arrays at relative tolerance 2×10^{-7} and absolute tolerance 2×10^{-9}, reconstructs whole-cycle physical and tilted balances, and validates own-run tables. Eight checker regressions reject mutations of thresholds, event values, stable counts, labels and pass flags, while accepting scientifically admissible diagnostic variation and irrelevant CSV line endings. Incidental adaptive-step counts are not treated as invariant scientific results. Table and figure generation is bound to that run's inputs and checked outputs.

Four retained raw exact-fold residuals require an explicit reporting correction, detailed in Appendix A. The primary output calculated those differences before replacing the tracked-initialization reference with the exact-fold one. The archived correction recomputes only the four derived differences; switching inputs, prediction fields, leading ratios and the 113 comparison outcomes are unchanged. The original output remains available, and every replay applies the same correction stage. Publication tables use the checked values.

The documented execution uses Windows, Python 3.12.14, NumPy 2.3.5 and SciPy 1.18.1. The public README gives dependency and execution instructions, and `VERIFICATION_RESULTS.json` records the executed checks. The proofs establish existence and asymptotic results under their stated hypotheses; the computations test finite-rate events and energy balances.

11 Discussion

The static folds generate delayed switching through a mixed amplitude–memory critical direction. Its full left/right normalization determines both load transversality and the quadratic rate, and therefore the prefactor of the two-thirds-power law. Preparation selects the appropriate Airy solution. A long stable approach and an exact-fold start are not interchangeable, even when both are described informally as slow driving.

The reduction has two separate warrants. The seven complementary directions follow a local invariant graph with a finite-chart contraction estimate; this retains the parameter-motion terms and does not require a stable centre equilibrium. The later global return uses strict loss at a fixed fold to establish a complete connecting orbit, then finite-time continuity and stable tracking to transfer that connection to the ramp. These mechanisms avoid both instantaneous stationary-memory substitution and the misuse of constant-input monotonicity along a varying input.

The prepared quasistatic conclusion is correspondingly precise. Away from folds, the two input-parametrized legs converge to different stable branches across the coexistence interval. Their local switching offsets have the derived asymptotics, capture occurs on each leg for small rate, and the geometric integral has a positive limit. The argument allows overshoot and does not prescribe an exact transient vertical trace. A stable periodic response at a fixed nonzero rate is a different question, not resolved by one prepared round trip or by its quasistatic limit.

The finite numerical experiment supports the leading prediction across every declared rate and both thresholds. Its controls show how preparation, symmetry and observation enter the comparison. The numerical levels have no certified analytical chart radius, and explicit full-source finite-rate error constants remain open. Higher-order terms, calibrated physical-amplitude events, shrinking-neighborhood arrival, periodic-orbit attraction, noise and physical implementation are focused further problems. None is replaced by a fitted log slope.

Physical maintenance is also distinct from geometric hysteresis. Equal stationary power on different branches yields a simple leading work coefficient, while memory changes the source state and passage time. Slowing the experiment reduces the switching-input displacement but increases accumulated supplied work. Conditional coherent and dissipative linkage remains tied to the same evolving carrier, while complete response retains the richer source and its history. This finite constitutive result extends the programme's maintained-state analysis without asserting a common microscopic origin for unrelated spectral profiles.

A Finite-preparation residual audit

For each exact-fold preparation define the quadratic-comparison residual by

$$E_{\mathrm{quad}} = \Delta a_{\mathrm{actual}} - h_\nu T_d(k = 1/\sqrt{3}). \tag{69}$$

The four raw residual fields instead subtracted $h_\nu T_d(k = 0)$, the tracked-branch comparison, although their stored prediction fields had been updated to the exact-fold preparation. Table 9 retains both differences. The checked output changes no event location, preparation-specific prediction, leading ratio or frozen comparison result. This correction concerns a derived reporting field rather than an altered experiment or tolerance.

Fold	d	Retained raw residual	Checked residual
up	0.25	-1.755594×10^{-4}	1.931485×10^{-5}
up	0.50	-1.604588×10^{-4}	3.418759×10^{-5}
down	0.25	-2.235553×10^{-4}	2.509803×10^{-5}
down	0.50	-2.109648×10^{-4}	3.686897×10^{-5}

Table 9: All four raw reporting differences and their checked exact-fold residuals at $\nu = 3 \times 10^{-6}$. Raw evidence is preserved; the checked values use the preparation-specific quadratic comparison.

The public correction record includes the raw file hash and each old/new value. Its executable correction stage computes the displayed definition from that run's actual events and predictions. The source and comparison criteria that produced the experiment remain unchanged.

References

Jack Carr. *Applications of Centre Manifold Theory*, volume 35 of *Applied Mathematical Sciences*. Springer, 1981. doi: 10.1007/978-1-4612-5929-9. URL https://doi.org/10.1007/978-1-4612-5929-9.

Jack Carr. Center manifold. *Scholarpedia*, 1(12):1826, 2006. doi: 10.4249/scholarpedia.1826. URL https://www.scholarpedia.org/article/Center_manifold.

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.

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.

NIST. Digital library of mathematical functions: Airy and related functions, 2026. URL https://dlmf.nist.gov/9. Sections 9.2, 9.7 and 9.9; accessed November 2026.

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. Sharp uniqueness and bistability in driven adaptive spectral carriers, 2026b. URL https://doi.org/10.5281/zenodo.23932350. Zenodo preprint.

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

Jeremy Rodgers. Source readout non-equivalence: Descent and equivariant reconstruction obstructions, 2026d. 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, 2026e. 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, 2026f. URL https://doi.org/10.5281/zenodo.23927855. Zenodo preprint.

Arjan van der Schaft and Dimitri Jeltsema. Port-Hamiltonian systems theory: An introductory overview. *Foundations and Trends in Systems and Control*, 1(2–3):173–378, 2014. doi: 10.1561/2600000002. URL https://doi.org/10.1561/2600000002.