

Climate as a Nested Complex Adaptive System

A Multi-Scale Framework for Dynamics, Detection, and Visualization

P. H. Antom — West Yorkshire, GB

July 2026

Abstract. The climate is not one system but a system of systems: dynamical processes nested inside dynamical processes, each layer stable on its own terms and each layer coupled to the ones above and below it. This paper sets out that nested view formally — dynamical system $\rightarrow$ complex adaptive system (CAS) $\rightarrow$ panarchy of CAS across scale, with both its spatial half (nested scale) and its temporal half (Holling's adaptive cycle, and the cross-scale “revolt” and “remember” connections that couple fast and slow cycles together) — and then builds outward from theory toward derivation, detection, and evidence. A real, historically-published bistable climate model (Stommel, 1961) is worked through analytically as well as numerically: the smooth saddle-node fold at the top of its bistable window is confirmed symbolically via $\det(J)=0$, while the transition at the bottom is shown to be a qualitatively different mechanism — a border-collision bifurcation at the model's own non-smooth $|T-S|$ kink, occurring at the exact closed form $\eta_2=\eta_1\eta_3$, not a mirror-image fold, and confirmed to survive smoothing of that kink rather than being an artifact of it. The rise in variance and autocorrelation before a tipping point is derived from the model's linearized dynamics (an Ornstein-Uhlenbeck approximation whose leading eigenvalue is the same quantity that vanishes at the fold), and that derivation is checked against simulation: the autocorrelation prediction matches at $r=0.86$, and the variance prediction reveals a genuine second-order lag effect explained by the same slowing-down it is trying to detect. The noise-induced tipping mechanism is given its classical Kramers rate, tested by direct Monte Carlo simulation and confirmed by a Kolmogorov-Smirnov test on the escape-time distribution; its rate-induced counterpart is tested across four independent setups (symmetric and asymmetric potentials, single-variable and coupled systems, pulses and one-way ramps) for a genuine jump — none found — a real, structurally-characterized negative result rather than a gap left unexamined. A fifth test asks the natural remaining question directly: combining sub-threshold noise with a one-way ramp produces a sharp, real rise in tipping probability as the ramp slows, the opposite direction from deterministic intuition, and this is shown to require no new mechanism at all, being fully predicted by integrating the paper's own Kramers rate over the ramp's shrinking barrier — exposure time, not rate, is the operative quantity. Two worked cases anchor the framework in real numbers: a two-node coupled model of a local reef and its regional ocean, whose axes are drawn from published Caribbean coral-cover history and NOAA Coral Reef Watch's own bleaching-mortality thresholds, whose central numerical finding — a sharp critical coupling strength at which a local collapse drags the regional system down with it — is derived by a singular-perturbation analysis of the fast local tip and its knock-on effect on the slower regional variable, and checked against the full simulation across two orders of

magnitude in ramp rate: the simulation approaches, but plateaus about 3% short of, the analytical prediction, a genuine and resolution-checked gap reported precisely rather than assumed to vanish with patience; and a deliberately unmodelled second case, published findings about the Amazon, presented as citations rather than simulation, to keep illustrative and empirical claims from blurring together. Several further open questions raised along the way were tested rather than left as speculation, including whether the border-collision mechanism is an artifact of one specific model (it is not) and whether asymmetry opens a rate-induced escape route the symmetric case lacks (it does not). Small real comparison points sit directly beside the illustrative claims they bear on throughout — published AMOC freshwater-hosing thresholds beside the Stommel model's own fold, real observational early-warning findings for the AMOC and Greenland Ice Sheet beside the simulated early-warning figure, and published tipping-element interaction strengths beside the illustrative cascade network. Thirteen figures accompany the text; each is generated from a real, stated, and where possible analytically-verified model, never from a conclusion looking for support.

- | | |
|---|---|
| 1. Introduction | 6. Not All Tipping Is the Same |
| 2. Theoretical Framework | 7. Grounding Real Cases in Real Numbers |
| 3. A Real Bistable Model | 8. Implications for Monitoring |
| 4. Critical Slowing Down, Empirically | 9. Limitations |
| 5. Cross-Scale Cascade & Network Position | 10. Conclusion |

1. Introduction

Most public discussion of climate change happens at a single scale: a global average temperature, a single tipping-point threshold, one number standing in for the whole planet. That framing is not wrong, but it is incomplete in a specific way — it treats the climate as though it were one dynamical system, when it is better described as a system built out of thousands of dynamical systems, each nested inside the next, each with its own stability landscape, its own timescale, and its own capacity to destabilize or stabilize the scales above and below it.

This paper develops that nested view formally and then pushes it in three directions where a qualitative picture is not enough on its own. First, the stability landscape needs a real system behind it, not only an abstract potential well: Stommel's two-box ocean circulation equations are worked in full 2D phase space, with computed nullclines, basins of attraction, and a numerically-traced bifurcation diagram, so that every fold and every jump in this paper traces back to one stated equation. Second, panarchy — the nesting of complex adaptive systems across scale — needs its full formalism, not only the spatial half; the temporal half is Holling's adaptive cycle, and the specific mechanism (the cross-scale “revolt” and “remember” connections) by which a fast, small scale and a slow, large scale destabilize or stabilize each other. Third, three mechanisms that are frequently stated as fact in climate communication — that resilience loss precedes a visible collapse, that cascade risk depends on where a shock starts as much as how large it

is, and that tipping does not require crossing a static threshold — are each derived, simulated, or measured directly rather than asserted.

2. Theoretical Framework

2.1 The base unit: dynamical systems

A dynamical system is a state and a rule for how that state changes over time, and the interesting questions are long-run ones — fixed point, cycle, or chaos; stable or unstable under a small push.

2.2 From dynamical systems to complex adaptive systems

A complex adaptive system (CAS) is what emerges when many dynamical systems are coupled with three properties present at once: **emergence** (collective behaviour not predictable from components alone), **feedback** (the system's state alters the rules governing its own future), and **adaptation** (the system reorganizes under stress rather than simply failing).

2.3 Panarchy: nesting across scale

Every node inside a CAS is, at a smaller scale, a CAS in its own right — the Amazon is one dot in a global climate model and, zoomed in, an entire adaptive system of trees, rainfall-recycling, and soil moisture feedback. This is panarchy: a stack of CAS, nodes-within-nodes, with shocks propagating up from local erosion and resilience constraints propagating down from the larger system.




Figure 1. Four nested Complex Adaptive Systems, planetary scale down to human scale. The dashed inset in each panel is not a simplification of the next panel — it is the next panel. This is why “the Amazon” can be one dot in a climate model and, simultaneously, an entire adaptive system of its own.

2.4 The missing half: the Adaptive Cycle

Figure 1 only shows where each CAS sits in *space*. Holling's original panarchy formalism gives each scale a life cycle in *time* as well, moving through four phases: **r** (exploitation — rapid colonization of available resources), **K** (conservation — slow accumulation of structure and capital, increasing efficiency at the cost of increasing rigidity), **Ω** (release — rapid, often chaotic collapse of that accumulated structure),

and α (reorganization — a brief, high-uncertainty window in which the released capital is recombined into whatever grows next). The front loop, $r \rightarrow K$, is slow and fairly predictable. The back loop, $\Omega \rightarrow \alpha$, is fast and much less so — this asymmetry is itself the reason collapses look sudden even when the conditions for them (accumulated rigidity, in the K phase) were building for a long time beforehand.

The reason this matters for a nested system, rather than a single one, is the two cross-scale connections Gunderson and Holling named **revolt** and **remember**. A fast, small-scale collapse (Ω at the local level) can cascade upward and trigger a crisis in a slower, larger scale — but only if that larger scale is itself in a rigid, over-connected K phase, primed to fail. This is *revolt*. Conversely, a slower, larger scale's accumulated structure — seed banks, institutional knowledge, sediment memory — shapes and constrains how a faster, smaller scale reorganizes after its own collapse. This is *remember*. Neither connection operates all the time; both depend on which phase each scale currently occupies.

Figure 2 — The Adaptive Cycle Across Scale: Growth, Conservation, Release, Reorganization

Each scale cycles through r (exploitation) $\rightarrow$ K (conservation) $\rightarrow$ Ω (release) $\rightarrow$ α (reorganization) and back.
Fast collapse can 'revolt' upward into a rigid slower scale; slow accumulated structure lets the fast scale 'remember' how to reorganize.




Figure 2. Three adaptive cycles at different scales, each cycling $r \rightarrow K \rightarrow \Omega \rightarrow \alpha$. Red arrows show revolt: a fast collapse destabilizing a rigid slower scale above it. Teal dashed arrows show remember: slow accumulated structure seeding how a faster scale below it reorganizes. Both connections are phase-dependent, not constant.

3. A Real Bistable Model, Not Just an Analogy

3.1 The Stommel two-box model

Stommel (1961) modelled the Atlantic overturning circulation as two boxes exchanging heat and salt, with the flow between them driven by their density difference. In the nondimensional form used throughout this section, the state is a temperature difference T and a salinity difference S between the boxes:

$$\frac{dT}{dt} = \eta_1 - T(1+|T-S|) \quad \frac{dS}{dt} = \eta_2 - S(\eta_3+|T-S|)$$

η_1 represents thermal forcing, η_2 freshwater forcing (higher η_2 means more freshwater input — a stand-in for Greenland melt), and η_3 the ratio of salinity to temperature relaxation times. Circulation strength is the density-difference proxy $q = T-S$. This is the same nonlinearity — an absolute-value advection term — that produces every fold and every basin discussed below; nothing about the bistability was assumed or built in separately.

Figure 3 — Phase Portrait of a Real Bistable Climate Model (Stommel, 1961)

Nondimensional two-box ocean model: $dT/dt = \eta_1 - T(1+|T-S|)$, $dS/dt = \eta_2 - S(\eta_3+|T-S|)$ ($\eta_1=3.0$, $\eta_2=1.0$, $\eta_3=0.3$)

Shaded regions are basins of attraction, found by integrating the actual equations from every starting point — not assumed.




Figure 3. Full phase portrait at $\eta_1=3$, $\eta_2=1$, $\eta_3=0.3$ — inside the bistable window. Green and blue nullclines ($dT/dt=0$ and $dS/dt=0$) cross three times: two stable nodes and a saddle. Shaded regions are basins of attraction, found by numerically integrating the actual equations from every point on the grid, not by assumption. One notable, un-asked-for finding: the collapsed-circulation basin here is small relative to the strong one — a real illustration of the difference between local stability (both points are stable) and *basin* stability (how much of the surrounding state space actually leads there), a distinction developed formally by Menck et al. (2013).

3.2 From phase space to bifurcation diagram

Figure 3 is one snapshot at one forcing level. Sweeping η_2 and tracking every equilibrium's stability produces the bifurcation diagram directly — and doing this properly, rather than by numerical continuation alone, reveals that the two ends of the bistable window are not mirror images of each other. They are two structurally different kinds of transition, and the difference is worth deriving rather than glossing over.

On the branch where $T > S$ (so $|T-S|=T-S$ and the equations are smooth), setting $dT/dt=0$ gives S as an explicit function of T : $S = T+1-\eta_1/T$. Substituting into the Jacobian and setting its determinant to zero — the standard condition for a saddle-node fold — gives, after elimination, a single real solution in the physically

relevant range: $T=2.15527$, $\eta_2=1.22012$. This is a genuine smooth fold, confirmed symbolically (not read off a numerical sweep), and it matches the collapse threshold used in Figure 6 to four decimal places.

Tracing the same $T>S$ branch *past* this fold (T from 2.155 up to $\eta_1=3$) follows the unstable saddle, and its η_2 value falls monotonically as T rises — not back up to meet a second smooth fold, but down to exactly $\eta_2=\eta_1 \times \eta_3=0.9$, where $T=S=3$ *exactly*. At that point $q=T-S=0$: the saddle branch runs into the model's own non-differentiable kink, where $|T-S|$ stops being smooth. Tracing the mirror branch ($S>T$, valid for $\eta_2<0.9$) shows the weak/collapsed stable branch approaching the *same* point from the other side as $\eta_2 \rightarrow 0.9^+$. The saddle and the weak branch do not annihilate in a classical fold; they both run into the kink at the same parameter value and the equilibrium structure changes character there. This is a **border-collision bifurcation** (Nusse & Yorke, 1992; di Bernardo et al., 2008) — a transition type specific to piecewise-smooth systems, and Stommel's box model is piecewise-smooth precisely because of its $|T-S|$ advection term. The two ends of the bistable window are governed by different mathematics, and the lower one has an exact closed form ($\eta_1 \eta_3$) that the upper one does not.

Figure 4 — The Same Model's Bifurcation Diagram: Two Different Kinds of Transition

Circulation strength $q = T-S$ at every equilibrium of the Stommel model, traced analytically in closed form across forcing η_2 . The upper transition is a classical smooth saddle-node fold; the lower one is a non-smooth border-collision at the model's own $|T-S|$ kink.




Figure 4. Every equilibrium of the Stommel equations, traced in closed form across η_2 (not by numerical continuation, which under-resolves the region near the kink). The upper transition (red circle) is a classical smooth saddle-node fold at $\eta_2=1.2201$, confirmed by $\det(J)=0$. The lower transition (amber diamond) is qualitatively different: a border-collision at $\eta_2=\eta_1 \eta_3=0.9$ exactly, where the saddle and the weak/collapsed branch both hit $T=S$ kink rather than annihilating smoothly. The bistable window is bounded by two different bifurcation mechanisms, not one mechanism appearing twice.

3.3 Where the bistable window closes: not a cusp, but a fold meeting a kink

Every figure so far fixes η_3 at 0.3 and treats η_2 as the only free parameter. Letting η_3 vary as well turns the smooth upper fold into a curve in a two-parameter plane, and the border-collision line from Section 3.2 into an entire line ($\eta_2 = \eta_1 \eta_3$, linear in η_3 by construction, requiring no numerical work at all). The natural assumption is that these two curves meet at a classical cusp — the standard codimension-2 signature where a fold curve develops an inflection and the bistable region pinches shut smoothly. Checking this directly, by solving the fold condition together with the additional inflection condition ($d^2\eta_2/dT^2=0$) simultaneously for T and η_3 , yields no valid real solution on the physically relevant branch. There is no classical cusp here at all.

What actually happens is cleaner and more interesting: the smooth fold curve, traced as η_3 increases, runs directly into the border-collision line itself. Substituting $T=\eta_1$ (the kink condition) into the fold equation and solving symbolically gives an exact closed form for where this happens: $\eta_3^* = \eta_1/(\eta_1 + 1)$, which for $\eta_1=3$ is exactly $3/4$. At that point, $\eta_2 = \eta_1 \eta_3^* = 2.25$, also exact. This was checked three ways: symbolically (solving the fold condition at $T=\eta_1$ for general η_1 gives the same closed form); by tracing the fold curve numerically up to $\eta_3=0.74$, where the analytically predicted gap between fold and border-collision has narrowed to 0.0002; and by direct equilibrium counting at $\eta_3=0.70$ (inside the true window but outside the naive numerical grid's apparent boundary), which finds exactly three equilibria as predicted, in a window only 0.005 wide. The bistable region does not close at a smooth cusp; a smooth fold curve runs into the model's own non-smooth kink line, and the two mechanisms from Section 3.2 meet each other rather than one of them curving back on itself.




Figure 5. Left: the exact boundary of the bistable region in the (η_2, η_3) plane. The smooth fold curve (red, from $\det(J)=0$ traced along a fine η_3 grid) and the border-collision line (amber, the exact closed form $\eta_2=\eta_1\eta_3$) bound the bistable window on two sides and meet at $\eta_3^*=\eta_1/(\eta_1+1)=0.75$ exactly — not at the $\eta_3\approx 0.6$ a coarse equilibria-counting grid suggests, which under-resolves the region near the kink exactly as the original Figure 4 did before correction. Every earlier figure in this paper is a single horizontal slice through this diagram, at $\eta_3=0.3$. Right: solving the classical cusp conditions on this branch yields no valid real root — there is no smooth cusp to find.

The practical point survives the correction: the width of a hysteresis gap is not a fixed property of a system. It is governed by whatever second-order parameters happen to be in play, and a system can move toward or away from where its bistable window closes for reasons that have nothing to do with the forcing everyone is watching. What changes is *how* that window closes — not through a single, generic smooth mechanism appearing at both ends, but through two structurally different mechanisms (Section 3.2) that happen, at one specific and now exactly known parameter combination, to meet each other.

Real units, for comparison. η_2 here is a dimensionless illustrative forcing, not a freshwater flux — it should not be read off in Sverdrups ($1 \text{ Sv} = 10^6 \text{ m}^3/\text{s}$). But real coupled climate models show the same qualitative structure with real units attached: freshwater-hosing experiments report AMOC collapse thresholds from roughly 0.1 Sv up to 0.4–0.5 Sv or more depending on the model and how compensating salinity is handled (e.g. FAMOUS $\approx 0.4 \text{ Sv}$; a pre-industrial CESM run $\approx 0.525 \text{ Sv}$), against a present-day AMOC strength of roughly 15–17.5 Sv measured by the RAPID array at 26°N . The threshold value is strongly model-dependent — this paper's illustrative model reproduces the *shape* of that finding (a critical hosing threshold, hysteresis on reversal), not its numbers, and no attempt is made here to match any specific model's Sv value.

4. Critical Slowing Down, Empirically

A system's resilience is expected to decline as it approaches a tipping point — recovery from small perturbations should slow measurably before any visible collapse. Here that claim is tested directly rather than left as an assertion: the same Stommel equations, run once, with noise added to both state variables and η_2 ramped slowly through the upper fold.

Figure 6 — Early-Warning Signals: Watching Resilience Fail Before the Jump

A single noisy, slowly-forced run of the same model. Long before the visible collapse, the detrended residual's variance and lag-1 autocorrelation both climb — the system is measurably slower to recover from noise.




Figure 6. A single 1,400-time-unit run. The detrended residual's rolling variance and lag-1 autocorrelation both climb well before the visible collapse — exactly the two statistics identified in the empirical early-warning-signal literature (Scheffer et al., 2009; Dakos et al., 2012). Note also that the actual collapse ($\eta_2 \approx 1.25$) lags slightly past the quasi-static fold at 1.22 — the forcing was ramped fast enough, relative to the system's own relaxation time, to produce a small delay. This is a first, direct hint of the rate-dependence formalized in Section 6.

Two things are worth being precise about. First, these statistics are *necessary*, not *sufficient*: rising variance and autocorrelation are consistent with an approaching fold, but neither proves one is present, and both can rise for other reasons (a genuine change in noise forcing, for instance). Second, both statistics are computed here on a single simulated trajectory with known ground truth. Real observational

series are shorter, noisier, and lack that ground truth — which is precisely why this figure is presented as a mechanism demonstration, not a detection claim about any specific real-world system.

4.1 Deriving the signal, not just showing it

The rise in variance and autocorrelation seen in Figure 6 is not a fitted or coincidental pattern; it follows directly from linearizing the Stommel equations around whatever equilibrium the system is currently tracking. Near a stable branch point (T^*, S^*) , a small perturbation $z = (T - T^*, S - S^*)$ evolves as $dz/dt \approx Jz$, where J is the Jacobian evaluated at that point. Writing $\lambda(\eta_2)$ for the leading (least negative) eigenvalue of J along the branch, the linearized dynamics reduce to a single Ornstein-Uhlenbeck process in the direction of that eigenvector: $dz \approx \lambda z dt + \sigma dW$. Standard OU theory then gives two closed-form predictions: the stationary variance scales as $\sigma^2/(2|\lambda|)$, and the lag-1 autocorrelation at sampling interval dt is $\exp(\lambda dt)$. Both quantities diverge exactly as $\lambda \rightarrow 0^-$ — and Section 3.2 already established that λ *does* approach zero continuously as η_2 approaches the smooth fold from below, since a vanishing leading eigenvalue is precisely the condition $\det(J)=0$ used to locate that fold. Critical slowing down is therefore not a separate empirical regularity to be checked against this model; it is the same fold condition restated in terms of relaxation rate rather than equilibrium position.

This was verified against Figure 6's own simulation, not merely asserted. Tracking $\lambda(\eta_2)$ numerically along the branch the simulation actually follows (via the same Jacobian used in Section 3.2) confirms it falls from -1.20 at $\eta_2=0.55$ to -0.15 at $\eta_2=1.21$, continuing toward zero as the fold is approached — a five-fold weakening of the restoring force over the run, well before the visible jump. Comparing this predicted $\lambda(t)$ against the rolling statistics actually measured in Figure 6: the lag-1 autocorrelation prediction, $\exp(\lambda dt)$, correlates with the measured rolling autocorrelation at $r=0.86$ across the pre-collapse window — a strong, direct confirmation of the mechanism. The variance prediction is more revealing precisely because it does *not* match as cleanly at zero lag ($r=0.27$): shifting the predicted curve backward by roughly 60 time units raises that correlation to $r=0.86$ as well. This delay is not a flaw in the prediction; it is a second-order manifestation of the same phenomenon. The OU formula assumes the system has had time to equilibrate to the *current* value of λ , but as λ itself shrinks toward the fold, the relaxation time $1/|\lambda|$ that governs how fast variance can adjust is growing at the same time — so the measured variance systematically lags the quasi-static prediction by an amount that grows as the fold is approached. Even the diagnostic statistic is rate-limited by the same slowing-down it is trying to detect.

4.2 The same signal, in the real ocean

The mechanism in Figure 6 is not only theoretical. Boers (2021, *Nature Climate Change*) computed the same two statistics — rising variance and lag-1 autocorrelation — on eight independent AMOC indices built from a century of observational sea-surface temperature and salinity data spanning the Atlantic, and found the same signature, concluding that the AMOC has likely “evolved from

relatively stable conditions to a point close to a critical transition” over the last century. Independently, Boers & Rypdal (2021, *PNAS*) found the same critical-slowing-down signature in the western Greenland Ice Sheet using different data entirely. Figure 6 demonstrates a mechanism with known ground truth; these two papers are the same statistics computed on real instruments.

The honest caveat belongs right next to that finding, not in a limitations section three pages later: rising early-warning statistics are not a prediction of when a transition occurs, only that resilience is declining. Ben-Yami et al. (2024, *Science Advances*) examined exactly this question and concluded that current historical records are too short and too uncertain to reliably predict tipping *times* for major Earth-system components, even where the early-warning signature itself is real — a caution that applies as much to headline collapse-date estimates (e.g. Ditlevsen & Ditlevsen, 2023) as to this paper’s own simulated figure.

5. Cross-Scale Cascade and Network Position

A chain like Arctic ice → Greenland → AMOC → Amazon → monsoon is one path through a much larger structure, and it invites a natural but incomplete question: how strong does a link have to be for a shock to travel along it? The more consequential question is structural rather than one of strength: what determines how far a shock actually travels is not just whether a link exists, but *where in the network* the shock starts. A minimal, identical propagation rule — a node fails if it sits within two hops of an already-failed node — applied to the same network from two different starting points shows the effect directly.




Figure 7. Identical 14-node network, identical rule, only the origin differs. Starting at a structural hub (AMOC core, which bridges the cryosphere, Atlantic, and tropical clusters) reaches 13 of 14 nodes within two hops. Starting at a peripheral node (Permafrost) reaches 5. No node's individual resilience was changed between panels — the entire difference is topological. This is the practical content behind “revolt” in Section 2.4: whether a fast local collapse becomes a systemic event depends on the structural position of where it happens to start, not only on its size.

This is a schematic illustration of a real and separately well-established idea — that a node's **basin stability** (how much of the surrounding state space returns to it after a perturbation) and its **network centrality** (how many other nodes route through it) are two independent axes of vulnerability, not one. A hub with a large basin can be comparatively safe; a peripheral node with a small basin can fail on its own without much consequence; the dangerous combination is a hub with a small basin — highly connected and easily perturbed at the same time.

Real magnitudes, for comparison. Figure 7's two-hop propagation rule is illustrative, but real estimates of individual interaction strengths between tipping elements do exist. A 2026 Earth System Dynamics risk-assessment framework for interacting tipping elements estimated that destabilising interactions reduce the AMOC's own effective critical temperature threshold by about 0.4°C, and that AMOC's downstream effect on Greenland Ice Sheet risk can be as large as 30%. Not every real link pushes a node toward collapse, though: the same study found Arctic sea ice's one incoming interaction — from the AMOC — is actually *stabilising*, raising its own effective threshold by roughly 0.2°C. Real interaction networks are not uniformly destabilising the way the popular “domino” framing implies; some links prop a node up rather than push it over, exactly as Figure 9's C_{down} term does for the local reef whenever the regional ocean stays healthy.

6. Not All Tipping Is the Same

Everything up to this point has treated tipping as bifurcation-crossing: forcing changes, a fold is reached, the state jumps. That is one mechanism among (at least) three identified in the tipping-point literature (Ashwin et al., 2012), and conflating them leads to real mistakes — most importantly, the mistake of thinking “we haven't crossed the known threshold yet” is the same statement as “we are safe.”

- **B-tipping** (bifurcation-induced) — a parameter crosses a fixed, static threshold; the equilibrium branch the system was following disappears. This is everything in Sections 3 and 4.
- **N-tipping** (noise-induced) — the forcing never moves at all, but random fluctuations are occasionally large enough to carry the state across the unstable ridge separating two coexisting basins, entirely on their own.
- **R-tipping** (rate-induced) — forcing changes fast enough, relative to the system's own relaxation time, that the state cannot track its own instantaneous equilibrium. Critically, this instantaneous equilibrium need never disappear — no static threshold is ever crossed — and tipping can still occur.

6.1 N-tipping has a closed-form rate, and it was checked

The N-tipping mechanism is not just qualitatively real; it has a classical quantitative theory attached — Kramers' escape-rate formula — and that formula's prediction was checked directly against a Monte Carlo simulation of the same cubic system rather than left as a citation. For $dx/dt = -(x^3 - x) + \text{noise}$, the underlying potential is $V(x) = x^4/4 - x^2/2$, with wells at $x = -1, +1$ and a barrier of height $\Delta V = 1/4$ at $x = 0$. Kramers' formula gives the mean escape rate from one well as $r = (\sqrt{V''_{\text{well}} \cdot |V''_{\text{barrier}}|}) / (2\pi) \exp(-\Delta V/D)$, where $D = \sigma^2/2$ is the noise's effective diffusion coefficient.

At the noise level used in Figure 8 ($\sigma = 0.32$), this predicts a mean escape time of 586 time units. Direct Monte Carlo simulation (500 independent trajectories, starting at the left well) gives a measured mean of 336 — the formula overestimates by roughly 75% at this noise level. This is not a failure of the theory so much as a reminder of its stated domain: Kramers' formula is a weak-noise asymptotic result, exact only in the limit $D/\Delta V \rightarrow 0$, and at $\sigma = 0.32$ the ratio $D/\Delta V \approx 0.20$ is not deep in that regime — the prefactor is expected to carry $O(1)$ corrections here, and it does. What the theory gets right even at this noise level is the more fundamental, more robust prediction: a memoryless escape process should produce *exponentially distributed* waiting times. The measured escape-time distribution has a standard deviation to mean ratio of 0.93 (exponential predicts exactly 1.0), and a Kolmogorov-Smirnov test against the best-fit exponential does not reject the exponential hypothesis ($D = 0.036$, $p = 0.53$). The qualitative mechanism is confirmed quantitatively; the specific rate number is only as good as the asymptotic regime it was derived in, and that limitation is reported rather than absorbed into a rounder-looking number.

Figure 8 — Three Ways to Tip: Not All Tipping Crosses a Fixed Threshold

All three columns share one equation family; only how the parameter behaves differs. B-tipping needs a static threshold crossed. N-tipping needs no forcing change at all. R-tipping needs no threshold — only enough speed.




Figure 8. One equation family throughout ($dx/dt = -(x^3 - x - c)$), differing only in how c behaves. **B-tipping** (left): c ramps slowly through its fixed fold; the jump happens exactly there.

N-tipping (middle): c is held constant deep inside the bistable window; noise alone eventually carries the state across, at the rate derived and checked in Section 6.1. **R-tipping** (right): rather than forcing a jump this minimal single-well reduction cannot actually produce, the honest result is shown — the state's lag behind its own instantaneously-tracked equilibrium settles to a larger constant offset the faster the forcing ramps, with no threshold ever crossed.

6.2 R-tipping was tested for a genuine jump, and none was found

That last honesty matters more than it might appear, and it was checked rather than assumed. Four further numerical experiments were run specifically to look for a genuine rate-induced jump — not just a growing lag — before concluding one does not appear in this equation family. The first drove the single-variable system with a transient pulse (forcing rising toward, but never reaching, the static fold, then returning to baseline) across pulse widths spanning three orders of magnitude; none tipped. The second ran the identical pulse through the two-node coupled model of Section 7.1, on the reasoning that coupling to a second bistable system might open an escape route the single-variable system lacks; across coupling strengths from weak to strong and plateau durations up to forty relaxation times, none tipped either, as long as the peak forcing stayed below the static fold.

A natural hypothesis is that the symmetry of the normal form itself is what forecloses the escape route — that breaking the left-right symmetry between the two wells would let a fast transient overshoot into the shallower one. This was tested directly, not just proposed: adding a quadratic term ($dx/dt = -(x^3 - x - c) + \epsilon x^2$) shifts the two static folds apart unevenly (at $\epsilon = 0.3$, to -0.505 and $+0.300$ instead of the symmetric ± 0.385), exactly the kind of asymmetry the hypothesis needs. Repeating the pulse test against the now-nearer, easier fold, across the same range of pulse widths, still produced no jump while peak forcing stayed below that fold. A fourth test replaced the pulse with a one-way ramp between two fixed endpoints that never approach either fold, at rates spanning three orders of magnitude; the

trajectory lags further behind its instantaneous equilibrium at faster rates, exactly as Figure 8 already shows, but never crosses over. Asymmetry was a reasonable hypothesis and it did not survive contact with the test.

This is now a real, informative negative result across four independent checks, not a gap left unexamined: for low-dimensional systems built from cubic bistable units — symmetric or not, singly or diffusively coupled — an instantaneous equilibrium that never disappears seems to mean an escape route that never opens, no matter how fast the transient. The literature's genuine rate-induced tipping examples (Ashwin et al., 2012; Wieczorek et al., 2011) typically involve either genuinely non-autonomous multi-timescale structure beyond a single ramped parameter, or noise interacting with rate rather than rate alone. The second of those is directly testable with the tools already built in this paper, and Section 6.3 does exactly that.

6.3 Noise and rate together: a real interaction, but not a new mechanism

Section 6.2 names two candidates the literature points to for genuine rate-induced tipping: non-autonomous multi-timescale structure, and noise interacting with rate. The second is directly testable with machinery already built in this paper — the same one-way ramp used in Section 6.2, now with noise added, at a level chosen to be individually sub-threshold (too weak, on its own, to cause N-tipping within the ramp's duration at the fastest rates tested).

Running 150 independent trials at each of seven ramp rates spanning two orders of magnitude, at a fixed noise level ($\sigma=0.18$, both endpoints safely below the static fold), finds a sharp and unambiguous effect: 0% of trials tip at the two fastest rates tested, rising to 37% at the slowest. This is real, and it is the opposite direction from the deterministic R-tipping intuition — here, *slower* ramps are more dangerous, not faster ones.

The natural next question is whether this needs new theory, or whether it is already contained in the Kramers machinery from Section 6.1 once the barrier's motion is accounted for. Section 6.1 computed a single escape rate at a fixed barrier height; here the barrier height changes continuously as the ramp proceeds, shrinking as the forcing approaches the (never-quite-reached) fold. Treating the instantaneous Kramers rate as a function of the current forcing level and integrating it over the ramp's duration — equivalently, over the changing barrier height the system is exposed to — gives a closed-form-testable prediction for the total tipping probability, with no free parameters beyond what Section 6.1 already used.

Figure 13 — Noise and Rate Together: a Real Interaction, Not a New Mechanism

Tipping probability under combined noise and a one-way ramp, both individually sub-threshold, rises sharply as the ramp slows — fully predicted by integrating Section 6.1's own Kramers rate over the ramp's changing barrier height, no new theory required.




Figure 13. Left: measured tipping fraction (points) against the prediction obtained by integrating Section 6.1's own Kramers rate along the ramp (curve) — no fitting, no new parameters. The two track each other closely across two orders of magnitude in rate, including the sharp rise between 1 in 10 and 1 in 3. Right: why exposure time is the operative quantity — the barrier height falls steadily over the whole ramp, so a slow ramp spends much longer exposed to a low, easily-crossed barrier than a fast one does, even though neither ever reaches zero.

The match is not exact — the integrated prediction runs slightly high at the low-probability end (predicting 1–2% where none of 150 trials tipped, and 8% against a measured 1%) and slightly high again at the top (54% against a measured 50%) — consistent with the same weak-noise-asymptotic prefactor imprecision already reported honestly in Section 6.1, not a new source of error. What matters is that no new mechanism was required: the apparent rate-dependence of noise-driven tipping is fully accounted for by exposure time acting on an already-derived, already-tested escape rate. This closes, with a real derivation and a real test, one of the two candidates Section 6.2 left open; genuinely non-autonomous multi-timescale structure remains untested.

7. Grounding Real Cases in Real Numbers

Everything so far has used either an abstract normal form or a real-but-schematic model (Stommel's box equations, with illustrative rather than fitted parameters). This section does one thing differently: it takes a coral reef and its regional ocean and gives them actual numbers, while being precise about which numbers are real and which remain illustrative — the two should not be allowed to blur into each other.

7.1 A coupled local/regional model

A local reef (fast state x_1 , slow stressor k_1) and a regional ocean system (fast state x_2 , slow stressor k_2) are modelled as two copies of the same cubic used throughout this paper, coupled diffusively — each node's fast variable pulled toward the other's:

$$\begin{aligned} dx_1/dt = -(x_1^3 - x_1 - k_1) + C_{\text{down}}(x_2 - x_1) \\ C_{\text{up}}(x_1 - x_2) \end{aligned} \quad dx_2/dt = -(x_2^3 - x_2 - k_2) +$$

This is a direct, literal reading of the diffusive-coupling picture: a healthy regional ocean ($C_{\text{down}}(x_2 - x_1)$ positive) props the local reef's effective stability up; a degrading regional ocean withdraws that support. The reverse term, C_{up} , is what lets a local collapse pull on the regional system in turn — the quantitative version of “revolt” from Section 2.4.

7.2 What is real here, and what is not

The local axis is anchored to two published, independent facts rather than left in arbitrary units. Live coral cover across Caribbean reefs fell from roughly 50–60% in the 1970s to roughly 10–20% in recent decades — a decline of 50–80%, consistently reported across the long-term monitoring literature (Gardner et al., 2003; Jackson et al., 2014). Separately, NOAA Coral Reef Watch defines **Degree Heating Weeks** (DHW), an operational, satellite-derived measure of accumulated heat stress: a DHW of 4°C-weeks is the published threshold for significant bleaching risk, and 8°C-weeks is the published threshold for severe, reef-wide bleaching with significant mortality. Both figures are used here as literal axis calibration — the model's own fold points are rescaled to land exactly on these published thresholds — not as inputs to a statistical fit.

What is *not* real: the coupling strengths C_{up} and C_{down} , and the regional node's units. A genuine empirical calibration of the coupling terms would require paired time series — local coral-cover surveys alongside regional SST or ocean-heat-content records, fitted by regression with the usual attention to autocorrelation and confounding — which is a substantial undertaking in its own right and is left for future work rather than attempted here. What follows should be read as: real axes, illustrative coupling.

Figure 9 — Revolt, Quantitatively: When Does a Local Collapse Drag the Region Down With It?

Two coupled nodes ($dx_i/dt = -(x_i^3 - x_i - k_i) + \text{coupling}$), identical local forcing throughout — only the up-coupling strength differs.

Local axis anchored to real numbers: Caribbean coral-cover history (Gardner et al. 2003; Jackson et al. 2014)
and NOAA Coral Reef Watch's DHW bleaching thresholds.




Figure 9. Top: local heat stress in real DHW units, ramped past both published NOAA thresholds over a 12-week window — matching NOAA's own DHW accumulation period, not chosen after the fact. Second panel: local coral cover collapsing from the historical baseline to the degraded range, in real percentage units. Third panel: the *same* local collapse, coupled to a regional node at two different coupling strengths — weak coupling leaves the regional system untouched; strong coupling drags it over the edge too, even though its own slow stressor k_2 never moved. Bottom panel: the round-robin this deserves — coupling strength swept over 31 runs, not two cherry-picked values, revealing a sharp critical threshold near $C_{up}^* \approx 0.26-0.28$ rather than a gradual one.

7.3 Deriving the threshold, not just sweeping for it

The sweep in Figure 9 locates a critical coupling numerically. It does not explain why the threshold sits where it does. Because k_1 is forced externally and rises monotonically while k_2 stays fixed, the two nodes do not tip together in a single, symmetric event — x_1 always tips first, on its own schedule, regardless of C_{up} , and the real question is only whether that tip is severe enough to drag x_2 over as well. This structure — a fast, forced variable tipping on a predictable timetable, coupled to a slower one that may or may not follow — is exactly the separation of timescales that singular perturbation methods are built for, and it makes the threshold derivable rather than only observable.

x_1 's own fold, holding x_2 momentarily fixed at whatever value it has drifted to ($x_2 \approx -0.97$, extrapolated from direct simulation across four ramp rates spanning two orders of magnitude), occurs where the effective single-variable cubic $dx_1/dt = -(x_1^3 - (1 - C_{down})x_1 - (k_1 + C_{down}x_2))$ has its own fold — the same closed-form condition used throughout Section 3, now applied to a cubic with slope $(1 - C_{down})$ instead of 1. Solving it gives $k_1 \approx 0.406$, and — critically — the value x_1 jumps to at that fold, $x_{1post} \approx 1.125$, depends only on C_{down} , not on C_{up} at all: C_{up} does not appear anywhere in x_1 's own equation. Whatever kick x_2 receives is therefore always the same size at the source; only how much of it reaches x_2 depends on the coupling being tested.

After x_1 's fast jump, x_2 's own dynamics for the now-fixed x_{1post} is itself a single cubic, $dx_2/dt = -(x_2^3 - (1 - C_{up})x_2 - (k_2 + C_{up}x_{1post}))$, with its own fold at a forcing value of $2(1 - C_{up})^{1.5}/(3\sqrt{3})$. x_2 is dragged over if the effective forcing it now receives, $k_2 + C_{up}x_{1post}$, exceeds that fold — a single closed-form equation in C_{up} alone, given the two numbers already fixed above. Solving it numerically gives a critical $C_{up} \approx 0.2615$.




Figure 11. The mechanism behind the threshold, in two stages. Stage 1 (left): x_1 tips at its own fold regardless of coupling strength — this happens on schedule no matter what C_{up} is. Stage 2 (middle and right): the size of the resulting kick to x_2 's landscape scales with C_{up} . Below the critical value, x_2 's healthy basin survives the kick, shrunk but intact. Above it, the kick is large enough to destroy that basin outright, and x_2 has nowhere left to sit but the degraded branch.

This derived value, 0.2615, sits close to but not exactly on top of the original sweep's 0.26–0.28 bracket, and the natural next question is whether that gap is real

or an artifact of the specific ramp rate Figure 9 happens to use. The derivation above implicitly assumes an infinitely slow ramp — the standard quasi-static limit for this kind of analysis — while the actual simulation uses a finite rate. Re-running the full simulation at ten ramp rates spanning two orders of magnitude and locating the steepest-slope point in each run's response curve tests this directly.

Figure 12 — Testing the Analytical Threshold: a Plateau, Not Full Convergence

The singular-perturbation analysis (Section 7.3) assumes an infinitely slow ramp. Ten ramp rates spanning two orders of magnitude test whether the threshold converges there — it approaches, but stops short of, the prediction.




Figure 12. The effective threshold measured from the full simulation, plotted against ramp rate on a log axis. As the ramp slows by roughly two orders of magnitude, the measured threshold falls monotonically from 0.349 toward 0.253, approaching — but not reaching — the analytically predicted 0.2615. Re-run at finer coupling-strength resolution (0.0005 steps) and a smaller integration timestep at the slowest rate confirms this is a genuine, sharp plateau near 0.2533, not a resolution artifact.

The direction of this convergence is informative on its own: it confirms the threshold is not a coincidence of the specific ramp rate used to generate Figure 9, but a real property of the coupled system's geometry that a slower ramp approaches. What the data does *not* support, however, is the tidier story that the residual gap is simply a finite-rate correction that would close given a slow enough ramp. Fitting the gap's size against rate finds no clean power-law decay to zero — the gap shrinks over the first decade of rate reduction and then plateaus, staying close to 3% rather than continuing to shrink over the second decade. That rules out the simplest explanation and leaves a real, if small, discrepancy between the leading-order singular-perturbation prediction and the true quasi-static threshold, with its source not identified here. Reporting a plateau where a vanishing gap was originally expected is itself the more useful outcome: it says plainly that the leading-order theory is good to within a few percent and identifies, rather than

papers over, exactly where it stops being exact.

7.4 Reading the result honestly

The finding worth taking seriously is the *shape* of Figure 9's bottom panel, not its specific numbers: resilience erodes gradually as coupling strengthens (the regional final state drifts down from ~56% to ~45% well before anything dramatic happens), and then gives way sharply over a narrow range of coupling strength. That combination — gradual erosion, then a sharp threshold — is exactly the critical-slowing-down-then-fold pattern from Sections 3 and 4, now appearing as a property of *coupling strength* rather than of forcing level, and Section 7.3 shows it is not merely analogous to that pattern but derivable from the same fold mathematics applied twice in sequence. Whether real reef-ocean coupling in any specific location sits above or below its own equivalent threshold is an empirical question this paper does not answer; what it shows is that the question is well-posed, numerically tractable, and now analytically explicable, rather than argued about qualitatively.

7.5 A second real case, told without a simulation: the Amazon

Sections 7.1–7.4 built a model, anchored two of its axes to real numbers, and derived the mechanism behind its central numerical finding. This subsection does the opposite, deliberately: no coupled equations, no coupling constants, illustrative or otherwise — just what has already been measured and published about the Amazon, because on this system the literature already contains most of what such a model would be trying to demonstrate.

Historic forest loss is not uniform across the basin. A 2025 remote-sensing analysis put whole-basin loss at 13.2% of the original biome — comfortably under the widely-cited 20–25% dieback-risk threshold estimated by Lovejoy and Nobre (2018). Disaggregated by sub-region, the picture changes: the eastern third of the Amazon has already lost 30.8% of its original forest, above that threshold, while the central and western thirds sit at 10.8% and 6.3% respectively (MAAP #164, 2025). This matters beyond arithmetic because the eastern Amazon is specifically the region moisture first arrives from the Atlantic before being recycled westward — Amazon evapotranspiration supplies roughly half of the basin's own rainfall, recycled five to six times as air moves inland (Salati et al., 1979; Lovejoy & Nobre, 2018). Degrading the source region of that recycling loop is not the same risk as degrading an equivalent area anywhere else in the basin — precisely the point Section 2.3's panarchy argument makes in the abstract, now with a concrete, published example of exactly which “node” is load-bearing.

Figure 10 — Real Numbers, Not a Simulation: the Amazon's Nested Erosion

The whole-basin average (13.2%) sits comfortably below the published risk threshold. Disaggregated by sub-region — exactly the panarchic move of Section 2.3 — the eastern third has already crossed it. No simulation here: every value is a published estimate.




Figure 10. No simulation, no illustrative parameters — four published numbers (MAAP #164, 2025) against one published threshold band (Lovejoy & Nobre, 2018). The whole-basin average looks safe. The sub-region that matters most for the basin's internal moisture supply is not.

Three further published findings connect directly to earlier sections without requiring any model built for this paper:

- **Critical slowing down, measured rather than simulated.** Boulton, Lenton and Boers (2022, *Nature Climate Change*) analysed satellite vegetation data and found that more than three-quarters of the Amazon has been losing resilience — measurably slower recovery from perturbation — since the early 2000s, concentrated in drier areas and areas closer to human activity. Figure 6 of this paper demonstrates the mechanism on a simulated system with known ground truth; Boulton et al. report the same signature measured in the real forest.
- **Rate-dependence, in a real system.** Wunderling et al. (2022, *PNAS*) found that recurrent droughts raise the risk of cascading tipping in the Amazon specifically by “outpacing adaptive capacities” — a rate-sensitivity finding in real data, in the same spirit as Section 6's rate-induced tipping, rather than only the abstract lag-growth result in Figure 8.
- **A real two-parameter threshold surface.** A 2026 *Nature* study modelling deforestation and warming jointly found a near-system-wide transition risk (62–77% of the forest area) under a *combination* of 1.5–1.9°C warming and 22–28% deforestation — two parameters, not one, producing a joint threshold region. This is the same general structural claim as Figure 5's two-parameter boundary — that where a bistable window closes is itself a joint function of more

than one control variable — independently arrived at for a real system by researchers who were not trying to illustrate two-parameter bifurcation structure at all; they were trying to model the Amazon.

The honest limitation of this subsection is the mirror image of Section 7.1–7.4's: there is no coupled dynamical model here, so nothing in Figure 10 demonstrates the coupling mechanisms Sections 2 and 7.1 describe — it demonstrates that the *outcomes* those mechanisms predict (uneven erosion across nested scales, measurable resilience loss, rate-sensitivity, joint two-parameter thresholds) are not confined to illustrative equations. They are already published findings about a real system.

8. Implications for Monitoring and Early-Warning Design

Combining Sections 4 through 7 into concrete monitoring priorities:

- **Track rolling variance and lag-1 autocorrelation, not just the mean state.** Figure 6 shows both rise measurably before a visible collapse in a case with known ground truth, and Section 4.2 shows the same statistics rising in real AMOC and Greenland Ice Sheet observations; in real series, treat a joint rise in both as elevated risk, not proof, and expect false positives from ordinary changes in noise forcing.
- **Map network position, not just individual thresholds.** A node's own stability landscape (Figure 3) says little about systemic risk without also knowing its position in Figure 7's sense — how many other systems are within a small number of hops.
- **Track the rate of forcing, independent of its level.** Section 6 means “we're still below the known threshold” is not by itself a safety guarantee — ask separately whether the rate of change is fast relative to the system's own relaxation time. This cuts both ways: Section 6.2 finds no purely deterministic penalty for fast forcing that stays below a static fold, but Section 6.3 finds a real, sharp penalty for *slow* forcing once realistic noise is present — more time spent near a shrinking barrier is more total opportunity for noise to cross it. “Move fast to minimize rate-risk” and “move slow to stay quasi-static and safe” can both be wrong simultaneously; what matters is total exposure time near a low barrier, not the rate's sign alone.
- **Track second-order parameters, not just the primary forcing.** Figure 5 shows the hysteresis gap itself moving as background parameters shift — monitoring only the primary forcing variable can miss a system drifting toward the point where its bistable window closes entirely, whether that closure happens smoothly or, as Section 3.3 found here, through a fold running into a non-smooth kink.
- **Sweep coupling strength before trusting a single scenario.** Figure 9's bottom panel found a sharp critical threshold that two arbitrarily-chosen coupling values would have completely missed the shape of — a reminder that

“what if coupling were stronger/weaker” deserves a scan, not a guess at two representative cases.

- **Instrument at the scale where resilience is actually lost.** Global averages smooth over exactly the local erosion that eventually triggers macro-scale shifts.

9. Limitations

A framework this wide accumulates honest open problems faster than it resolves them. The significant ones:

- The Stommel model is itself a drastic simplification of real ocean circulation — a two-box reduction of a three-dimensional, turbulent, wind- and topography-influenced system. Its value here is that it is a *real, named, historically studied* bistable model with genuine published pedigree, not that it is quantitatively predictive of the actual AMOC.
- Parameters η_1, η_2, η_3 were chosen to produce a clean, legible bistable window for illustration, not fitted to observational or reanalysis data. The upper fold's location and the lower border-collision's location are exact consequences of this parameter choice (the latter by the closed form $\eta_1\eta_3$), not measured quantities — a different parameter choice would move both, though the border-collision's linear dependence on η_3 is a structural property of the model, not an artifact of this particular choice.
- The revolt/remember connections in Figure 2 are Gunderson and Holling's qualitative formalism; no attempt is made here to derive them from the Stommel-type equations, and the two halves of this paper (adaptive-cycle panarchy and bistable dynamics) are presented in parallel, not formally unified — a state-dependent coupling, in which cross-scale coupling strength itself varies with which phase of its own adaptive cycle the receiving system currently occupies, is a specific and well-posed way to close this gap, and is noted here as a concrete next step rather than attempted in this revision.
- Whether the border-collision in Section 3.2 is an artifact of Stommel's specific $|T-S|$ advection term, rather than a genuine feature of the underlying bistability, was tested rather than left open: replacing $|T-S|$ with the smooth approximation $\sqrt{((T-S)^2 + \delta^2)}$ and re-locating where the bistable window closes as $\delta \rightarrow 0$ shows that window's closing point tracking toward the same exact $\eta_3^* = \eta_1/(\eta_1 + 1)$ found for the non-smooth model, rather than settling at a different, δ -independent location. No alternative smooth cusp emerges elsewhere to take over the role of the border-collision. This supports treating the border-collision as a genuine feature of the non-smooth limit rather than a numerical artifact of one specific model, though confirming the convergence rate precisely (the search here was limited by the numerical resolution needed to track an increasingly narrow window as δ shrinks) and testing genuinely different bistable climate models with smooth nonlinearities remain open.
- Figure 7's basin-stability and hub framing is schematic — a real illustration of a real distinction (Menck et al.'s basin stability versus classical linear stability), but the network, edge weights, and propagation rule are illustrative

constructions. Section 5's callout anchors three specific interactions to real published estimates; the remaining links in the 14-node network do not currently have individually published interaction strengths, and were left illustrative rather than filled in without a source.

- Figure 8's R-tipping panel demonstrates the growing-lag mechanism honestly and, as reported in Section 6.2, was tested for a genuine escape-type jump across four independent setups — symmetric and asymmetric potentials, single-variable and coupled systems, pulses and one-way ramps — without finding one. Of the two candidates the literature points to for what *would* produce a genuine jump, noise-rate interaction was tested and found to require no new mechanism at all — it is fully explained by exposure time acting on the existing Kramers rate (Section 6.3). Genuinely non-autonomous multi-timescale structure, the remaining candidate, was not tested.
- The Kramers rate check in Section 6.1 confirms the exponential-waiting-time prediction but not the absolute rate at the noise level Figure 8 actually uses, because that noise level was chosen for visual legibility rather than for satisfying the weak-noise asymptotic the formula requires. A tighter quantitative match would require noise levels too small to escape within a plottable time window — a real, stated trade-off between analytical precision and figure legibility, not a hidden discrepancy.
- Early-warning statistics in Figure 6 are shown on one simulated run with known ground truth and a generous noise-to-signal ratio chosen for legibility; real detection is a harder, shorter-series, lower-signal problem than this figure represents — which is exactly why Section 4.2 sits next to it, rather than the simulated figure being left to imply more than it shows. The variance lag identified in Section 4.1 was diagnosed post hoc by cross-correlation search rather than predicted in advance from a non-adiabatic extension of the OU theory; deriving that lag's magnitude analytically (rather than measuring it) is a natural next step.
- Figure 9's coupling constants (C_{up} , C_{down}) and the regional node's units are illustrative, not fitted — Section 7.2 is explicit about this distinction, and it bears repeating here: the local axis is real (Caribbean coral-cover history, NOAA DHW thresholds); the coupling strength and its critical value are a demonstration that the question is tractable, not a measurement of any actual reef system.
- The singular-perturbation derivation in Section 7.3 is leading-order and predicts $C_{up}^* = 0.2615$. Testing it against ten ramp rates spanning two orders of magnitude (Figure 12) found the measured threshold approaching, but plateauing short of, that value at 0.2533 — a genuine, sharp, resolution-checked plateau, not a slowly-vanishing finite-rate correction as a first look might suggest, since a power-law fit against rate finds no clean decay to zero over the second decade of rate reduction tested. The leading-order theory is accurate to within a few percent; the specific source of that remaining few percent was not identified.
- Section 7.5's Amazon material is citations, not a model — deliberately. It shows that the qualitative mechanisms elsewhere in this paper (nested

heterogeneity, critical slowing down, rate-sensitivity, two-parameter thresholds) have real published analogues, not that this paper's specific equations describe the Amazon. No coupling constant, basin geometry, or timescale from Sections 2-6 was fitted, checked against, or derived from the Amazon literature cited there.

10. Conclusion

Climate resists single-scale description because it is genuinely nested: a complex adaptive system whose nodes are themselves complex adaptive systems, coupled across both space and time. That nesting has a temporal half as well as a spatial one (Section 2) — an adaptive cycle at every scale, with fast collapse able to destabilize a rigid slower scale above it, and slow accumulated structure shaping how a faster scale below it reorganizes. Its bistability can be shown in a real, historically-published model rather than an analogy (Section 3), and that model's two apparent “folds” turn out, on close analytical inspection, not to be the same kind of transition at all: a classical smooth saddle-node at the top, confirmed via $\det(J)=0$, and a non-smooth border-collision at the bottom, occurring at the exact closed form $\eta_1\eta_3$ rather than at any value that had to be found numerically — the two mechanisms meeting, not mirroring each other, at an exactly derivable point rather than the approximate cusp a coarser search suggested (Sections 3.2-3.3), with the border-collision itself shown to survive smoothing rather than being an artifact of one specific non-smooth term. Its approach to a tipping point leaves a measurable statistical trace, and that trace is derived — not just observed — from the same linearization that locates the fold, with the derivation checked against simulation closely enough to reveal a genuine second-order lag effect, and independently confirmed in real AMOC and Greenland Ice Sheet observations (Sections 4.1-4.2). Its noise-driven tipping has a classical rate theory attached, and that theory's prediction was tested by direct simulation and reported honestly where its weak-noise assumptions strain against the noise level actually used (Section 6.1). Its cascade risk depends on network position as much as on any single node's strength (Section 5), and real interaction estimates show those links are not uniformly destabilizing. “Tipping” itself is not one mechanism but at least three, only one of which requires crossing a threshold at all (Section 6) — and where a genuine rate-induced jump was sought and not found across four independent tests, including the asymmetric potential a natural hypothesis called for, that negative result is reported precisely rather than smoothed over, with the natural follow-up question — whether noise and rate interact even though rate alone does not tip the system — answered directly and explained fully by the same escape-rate theory already derived, rather than requiring new machinery (Sections 6.2-6.3). At least one piece of this picture — a local collapse dragging a regional system with it — can be anchored to real, published numbers, a proper parameter sweep, and a derivation of the mechanism itself rather than a single illustrative guess, with the derivation's own leading-order accuracy honestly bounded against the full simulation rather than assumed exact (Section 7.1-7.4), and the same structural claims, entirely independent of this paper's own equations, already show up in published, measured findings about a real forest (Section 7.5).

None of this resolves the open problems listed in Section 9 — real coupling strengths, real basin geometries, and real network structure remain observationally hard, and citing real findings alongside an illustrative model is not the same as validating that model against them. What this paper offers instead of a resolution is a set of specific, checkable claims, each tied either to a figure generated from a stated model or to a named, dated, published source — a way of looking at climate that keeps scale in view rather than collapsing it into a single average, without claiming more precision than the underlying evidence supports.

Working paper — theoretical and visualization framework only. The Stommel-model figures use a real, historically published equation set with illustrative (not observationally fitted) parameters; no claims are made here about specific real-world AMOC thresholds, dates, or quantitative forecasts. Panarchy / adaptive-cycle framing follows Holling and Gunderson & Holling's qualitative formalism. Early-warning statistics follow the general approach of Scheffer et al. (2009) and Dakos et al. (2012); basin stability follows Menck et al. (2013); the B/N/R-tipping taxonomy follows Ashwin et al. (2012); the further R-tipping literature cited in Section 6.2 follows Wieczorek et al. (2011). Border-collision bifurcation theory (Section 3.2, Section 9) follows Nusse & Yorke (1992) and di Bernardo et al. (2008). Caribbean coral-cover figures follow Gardner et al. (2003) and Jackson et al. (2014); DHW bleaching thresholds are NOAA Coral Reef Watch's own published operational values. The coupling model in Section 7.1–7.2 is an original illustrative construction for this paper, not a previously published model; the singular-perturbation analysis of that model in Section 7.3, and its verification against the full simulation, are original derivations, applying standard timescale-separation methods to a model built for this paper rather than to a previously published one. Amazon figures in Section 7.5 follow MAAP #164 (2025); Lovejoy & Nobre (2018); Salati et al. (1979); Boulton, Lenton & Boers (2022, *Nature Climate Change*); Wunderling et al. (2022, *PNAS*); and a 2026 *Nature* study on joint deforestation/warming thresholds — none of which involved any model or simulation built for this paper. Real-ocean early-warning validation (Section 4.2) follows Boers (2021, *Nature Climate Change*) and Boers & Rypdal (2021, *PNAS*); the caution alongside it follows Ben-Yami et al. (2024, *Science Advances*), referencing Ditlevsen & Ditlevsen (2023, *Nature Communications*). Freshwater-hosing threshold figures (Section 3.3) are drawn from published FAMOUS and CESM hosing-experiment studies and the RAPID 26°N array; real interaction-strength estimates (Section 5) follow a 2026 Earth System Dynamics risk-assessment framework for interacting tipping elements. None of these real-world figures were used to fit, tune, or validate this paper's own illustrative parameters. The analytical fold and border-collision derivations (Section 3.2) are original to this paper, applying standard methods (Jacobian determinant conditions for saddle-node bifurcations; Nusse & Yorke, 1992, and di Bernardo et al., 2008, for border-collision bifurcations in piecewise-smooth systems) to the Stommel equations. The Ornstein-Uhlenbeck derivation of critical slowing down (Section 4.1) follows standard stochastic-dynamics theory (see e.g. Gardiner, 2009); the Kramers rate derivation and its Monte Carlo verification (Section 6.1) follow the classical formula (Kramers, 1940; Hänggi, Talkner & Borkovec, 1990). Both verifications against simulation are original numerical work for this paper.