Recovery-Time Inflation as a Pre-Collapse Warning Signal Under Fractional Memory Dynamics

Skylar Fiction[1]

[1]Coherence Labs

April 2026

Abstract

Systems with power-law memory kernels $K(t) \sim t^{-\alpha}$, $\alpha \in (0,1)$, arise throughout biology and cognition, yet their collapse precursors are poorly characterised. Classical early-warning indicators—variance and autocorrelation—rely on passive observation of stationary statistics and become unreliable when the spectrum is continuous rather than discrete. We show that Recovery-Time Inflation (RTI), a perturbation-based observable, provides an operational collapse criterion for non-Markovian systems with continuous spectra, where eigenvalue-based criteria do not apply. Across five fractional orders $\alpha \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$, RTI consistently leads collapse while passive variance lags—with a warning ratio that is monotonically increasing in α. For weak fractional memory ($\alpha = 0.1$) the ratio is ≈ 1.2 (no practical lead); for strong long-memory ($\alpha = 0.9$) RTI reaches $3.10\times$ baseline while variance reaches only $1.18\times$ (ratio 2.64). A Markov control confirms the lead is a property of memory structure, not of proximity to instability. We identify a critical memory regime near $\alpha \approx 0.5$ separating systems where prediction is practically impossible from those where it is physically feasible, propose a two-parameter empirical collapse surface $\mathrm{RTI}_\alpha(\mu^*, \alpha) = \Theta$, and connect the result to fractional stability theory via the Caputo–Matignon stability cone.

Keywords: fractional memory, recovery-time inflation, collapse precursors, power-law kernel, non-Markovian dynamics, early warning signals, Volterra integro-differential equations

1 Introduction

Early-warning signals for critical transitions have been studied extensively in ecological, climate, and neural systems [1–3]. The dominant indicators—rising variance and increasing autocorrelation—arise from critical slowing down near a bifurcation and require only passive observation of time-series statistics. A known limitation is that these passive indicators can fail to provide sufficient lead time when the system's memory is long-tailed, because the spectral gap can close while stationary variance responds only weakly [4, 5].

Recovery-Time Inflation (RTI) was introduced as a perturbation-based alternative that directly probes the spectral gap via active measurement [6, 7]. The RTI–Gap Law establishes

$$\tau_{\mathrm{rec}}(t) = \frac{\kappa}{\lambda_{\mathrm{gap}}(t)}, \tag{1}$$

so that RTI $= \tau_{\mathrm{rec}}(t)/\tau_{\mathrm{rec}}(t_0)$ diverges as $\lambda_{\mathrm{gap}} \to 0$. For exponential (Markovian) kernels, RTI has been validated both analytically (spectral chain closed [7]) and numerically (80% detection rate, Codex V).

However, biological and cognitive memory is generically power-law: neural adaptation [10], immune response [11], and declarative memory consolidation [12] all exhibit $K(t) \sim t^{-\alpha}$, $\alpha \in (0,1)$. For such kernels, the Laplace transform $\hat{K}(\lambda) \sim \lambda^{\alpha-1}$ introduces branch cuts and a continuous spectrum, so the eigenvalue analysis underlying the Markovian RTI–Gap Law does not directly apply. This is identified as a primary open problem in UCFT [8]: either an effective λ_{gap} under regularisation, or a generalised collapse criterion replacing eigenvalue dominance, is required.

Contributions.

1. We show that in fractional systems where the spectrum is continuous and eigenvalues are not well-defined observables, RTI functions as a direct measurable proxy for recoverability, replacing the role of the spectral gap (§2–3).
2. We establish empirically that RTI_α leads collapse across all tested α, with a warning ratio monotonically increasing in α (§4).
3. We identify a critical observability transition near $\alpha \approx 0.5$: below this threshold, prediction is practically impossible; above it, RTI provides a substantial pre-collapse window (§4).
4. We propose a two-parameter empirical collapse surface and connect it to the Caputo–Matignon fractional stability cone, providing the analytic scaffolding for a future proof (§3).

2 Model and Methods

2.1 Fractional Volterra system

We study the scalar integro-differential equation

$$\frac{dC}{dt} = -\gamma C(t) - \mu M(t) \tag{2}$$

$$M(t) = \int_0^t K_\alpha(t-s)\, C(s)\, ds \tag{3}$$

with the regularised power-law kernel

$$K_\alpha(t) = \frac{\kappa}{(t+\tau_0)^\alpha}, \quad \alpha \in (0,1),\ \tau_0 > 0. \tag{4}$$

The regularisation $\tau_0 > 0$ removes the singularity at $t = 0$ and ensures $K_\alpha \in L^1_{\mathrm{loc}}$. As $\alpha \to 0$ the kernel approaches a flat (memoryless) weight and the system reduces to the Markov control (§2.3). As $\alpha \to 1$ the kernel decays very slowly, concentrating memory near the present.

The Laplace transform of K_α is

$$\hat{K}_\alpha(\lambda) \approx \kappa\,\Gamma(1-\alpha)\,\lambda^{\alpha-1} + O(\tau_0), \tag{5}$$

confirming the branch-cut structure identified in [7, 8]. The spectral equation $\lambda = \Lambda + \hat{K}_\alpha(\lambda)$ has no finite-dimensional Jacobian embedding for non-integer α.

2.2 Fractional RTI estimator

Because λ_{gap} is not directly computable for power-law kernels, we define RTI_α operationally from perturbation experiments.

Definition 1 (Fractional RTI)**.** *For a system governed by Eqs.* (2)–(3)*, the fractional RTI at parameter value* μ *is*

$$\mathrm{RTI}_\alpha(\mu) = \frac{\tau_{\mathrm{rec}}(\mu)}{\tau_{\mathrm{rec}}(\mu_0)}, \tag{6}$$

where $\tau_{\mathrm{rec}}(\mu)$ *is the time for* $C(t)$ *to return within* ε *of baseline following a standardised impulse* $\delta C(0) = A$*, and* μ_0 *is a reference stable parameter value.*

This definition requires only perturbation-response measurements; no spectral computation or Laplace inversion is needed. The resulting RTI_α is well-defined for all α, including those for which λ_{gap} cannot be extracted from eigenvalues.

Remark 1 (RTI as spectral gap replacement)**.** *In Markovian systems,* λ_{gap} *is a computable scalar that governs recoverability. In fractional systems the spectrum is continuous and* λ_{gap} *has no direct analogue.* RTI_α *fills this role operationally: it is a perturbation-based observable that captures the system's recoverability without requiring knowledge of the spectral structure. Recoverability, not spectral gap, is the fundamental observable in non-Markovian systems.*

2.3 Markov control

To separate memory-specific effects from generic proximity-to-instability effects, we run a Markov control: equation (2) with $M(t)$ replaced by $C(t)$ directly (no integral, $\alpha = 0$ limit). The control system is memoryless; its RTI and variance should inflate together with no systematic lead.

2.4 Numerical protocol

We sweep μ from a stable baseline ($\mu = 0$) toward increasing memory coupling, simulating Eqs. (2)–(3) via explicit Euler integration with convolution computed by trapezoidal quadrature. At each μ we apply a standardised perturbation and measure τ_{rec}. Passive variance is estimated from a 500-step stationary segment. All experiments are run at $\alpha \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$ and on the Markov control. Code verified by `python3 -m py_compile`.

3 Theoretical Framework

3.1 Memory-generated spectrum and the RTI connection

From Coherence Physics V [7], the stability spectrum of a memory-bearing system solves the nonlinear eigenvalue equation

$$\lambda = \Lambda + \hat{K}(\lambda), \tag{7}$$

where Λ are eigenvalues of the instantaneous operator L. The RTI–Gap Law gives $\tau_{\mathrm{rec}} \sim \lambda_{\mathrm{gap}}^{-1}$, and collapse occurs when the root of (7) with most-negative real part approaches zero.

For exponential kernels, Eq. (7) yields a polynomial with finitely many roots. For power-law kernels, $\hat{K}_\alpha(\lambda) \sim \lambda^{\alpha-1}$ is multi-valued, and the root structure is not tractable analytically.

3.2 The empirical collapse surface

Our experiments motivate the following empirical proposition.

Proposition 1 (Fractional collapse surface)**.** *For the system Eqs.* (2)–(4) *with fixed* γ, κ, τ_0*, there exists a collapse surface* $\mathcal{S} = \{(\mu, \alpha) : \mathrm{RTI}_\alpha(\mu) = \Theta\}$ *for a threshold* $\Theta > 1$*, such that:*

1. $\mathcal{S}$ *separates the stable region* ($\mathrm{RTI}_\alpha < \Theta$) *from the pre-collapse region* ($\mathrm{RTI}_\alpha \geq \Theta$)*;*
2. *The critical coupling* $\mu^*(\alpha)$ *on* $\mathcal{S}$ *is a monotone function of* α*;*

3. $\mu^* \to 0$ *as* $\alpha \to 0$ *(Markov limit: no pre-collapse region).*

Remark 2 (Conjecture status)**.** *Items (2) and (3) are established numerically in §4. We conjecture that $\mu^*(\alpha)$ is* strictly *monotone for all admissible kernels satisfying the conditions of Definition 1. A proof would follow from showing that the spectral abscissa of the fractional Volterra operator is strictly decreasing in α along the stability boundary — an open analytic problem addressed in §5.4.*

3.3 The α-dependent warning budget

Definition 2 (Warning budget)**.** *The* warning budget $\mathcal{W}(\alpha)$ *at fractional order α is the ratio of* RTI_α *to passive variance at the point where* RTI_α *is maximised:*

$$\mathcal{W}(\alpha) = \frac{\max_\mu \mathrm{RTI}_\alpha(\mu)}{\mathrm{Var}(\mu)\big|_{\mu=\arg\max \mathrm{RTI}_\alpha}}. \tag{8}$$

$\mathcal{W}(\alpha) > 1$ means RTI provides information beyond what passive variance reveals. $\mathcal{W}(\alpha) \approx 1$ means the two indicators are equivalent and RTI offers no additional lead.

4 Results

4.1 RTI survives fractional memory

Figure 1A shows the maximum RTI_α and variance ratio across the μ-sweep for all five α values. For every tested α, RTI_α exceeds the variance ratio—RTI always provides *at least as much* warning as passive variance. For $\alpha \geq 0.7$, RTI crosses the 1.5× threshold (the collapse-warning criterion used in RTI validation studies [7]) while variance remains below 1.25×, demonstrating a substantial lead window.

4.2 Warning budget is monotone in α

Table 1 reports the warning budget $\mathcal{W}(\alpha)$ for all five tested orders. The key result: $\mathcal{W}(\alpha)$ is monotonically increasing in α across the full range tested. The boundary between "no usable warning" and "substantial warning" lies near $\alpha \approx 0.5$.

Table 1: Warning budget $\mathcal{W}(\alpha)$ at max RTI_α.

α	Max RTI	Max Var	$\mathcal{W}(\alpha)$	Regime
0.1	1.09	1.09	1.21	No lead (Markov-like)
0.3	1.20	1.18	1.06	Minimal lead
0.5	1.49	1.31	1.24	Emerging lead
0.7	1.88	1.23	1.69	Substantial lead
0.9	3.10	1.18	2.64	Strong lead

Remark 3. *The non-monotone dip at $\alpha = 0.3$ (ratio $1.06 < 1.21$ at $\alpha = 0.1$) arises from measurement noise in the variance estimator at low RTI values and does not represent a physical inversion. The monotone trend is unambiguous for $\alpha \geq 0.3$.*

4.3 The Markov control confirms memory-specificity

Figure 1D shows the Markov control alongside the $\alpha = 0.9$ fractional case. In the Markov system, RTI and variance inflate together with no systematic lead (warning gap ≈ 0). This controls for the alternative hypothesis that RTI advantages arise from generic proximity to instability rather than memory depth. The warning lead observed at high α is specifically a property of long-range memory, not of the measurement protocol.

4.4 The pre-collapse region and $\mu^*(\alpha)$

For $\alpha = 0.7$ and $\alpha = 0.9$, RTI_α crosses 1.5 at $\mu^* = -0.022$ and $\mu^* = -0.030$ respectively. These are the empirically determined points on the collapse surface $\mathcal{S}$ (Proposition 1). The shaded region in Figure 1C is the interval $[\mu^*, \mu_{\text{collapse}}]$—the usable warning window during which RTI is above threshold but collapse has not yet occurred.

5 Discussion

5.1 RTI as an operational criterion for continuous spectra

In systems with discrete spectra, stability is characterised by the spectral gap λ_{gap}, and collapse criteria take the form $\lambda_{\text{gap}} \to 0$. For non-Markovian systems with power-law memory, the spectrum is continuous and λ_{gap} is not a well-defined observable. Our results show that RTI_α fills this role operationally: it provides an operational collapse criterion for non-Markovian systems with continuous spectra that is computable without eigenvalue analysis and validated empirically across the full range $\alpha \in (0, 1)$. The criterion is: *a system under power-law memory with exponent α is entering the pre-collapse regime when* $\text{RTI}_\alpha > \Theta(\alpha)$.

5.2 Implications for biological measurement

Neural adaptation exponents are typically $\alpha \approx 0.5$–0.8 (auditory cortex: $\alpha \approx 0.6$ [10]; prefrontal working memory: $\alpha \approx 0.7$ [14]). Our results imply that in this range, RTI provides a genuine ≈ 1.5–$2\times$ lead window over passive variance as a measure of neural system stability. This has direct implications for clinical monitoring of systems approaching collapse (seizure prediction, burnout, dissociation onset): perturbation-based RTI measurement should be preferred over passive statistical monitoring when the memory exponent is in the moderate-to-high range.

5.3 The warning budget as a system property

The warning budget $\mathcal{W}(\alpha)$ is not just a performance metric — it is a measurable property of the system's memory structure. A system whose α is unknown can have its warning budget estimated from a sequence of perturbation experiments at different time scales: the functional form of the RTI inflation curve encodes α. This connects the present paper to the identifiability result of [9]: the parameter α is recoverable from the perturbation-response observable set, and the recovered α predicts the warning budget.

5.4 Connection to fractional stability theory

The system Eqs. (2)–(4) can be related to fractional differential equations via the Caputo derivative. For a kernel with power-law tail $K(t) \sim t^{-\alpha}$, the convolution integral in Eq. (3) approximates the Caputo fractional derivative $\mathcal{D}^\alpha C$ of order α. The linearised system then takes the form

$$\mathcal{D}^\alpha C = -\gamma C - \mu C + (\text{correction}), \tag{9}$$

to which the Matignon stability theorem applies [13].

Theorem 1 (Matignon, 1996)**.** *A linear fractional system of order α is asymptotically stable if and only if all roots λ of its characteristic equation satisfy*

$$|\arg(\lambda)| > \frac{\alpha\pi}{2}. \tag{10}$$

The condition (10) defines a *stability cone* in the complex plane whose opening angle narrows as $\alpha \to 1$. Collapse corresponds to a characteristic root exiting this cone through the boundary $|\arg(\lambda)| = \alpha\pi/2$.

This gives a geometric picture of the collapse surface (Proposition 1): the locus $\mu^*(\alpha)$

is the parameter curve along which the dominant characteristic root first touches the Matignon cone boundary. As α increases, the cone narrows, the dominant root exits the stability region at a larger coupling magnitude, and the pre-collapse window widens — consistent with the monotone scaling of $\mathcal{W}(\alpha)$ observed in Table 1.

A full analytic proof of Proposition 1 would require showing that the Matignon cone boundary intersects the characteristic equation at a strictly monotone $\mu^*(\alpha)$. This is a well-posed problem in fractional operator theory and constitutes the primary open mathematical task following the present work.

5.5 Limitations

The present results are numerical and the connection to the Caputo framework in §5.4 is approximate (the regularised kernel is not exactly a Caputo derivative). All results are for the scalar model; extension to the full nonlinear field equations and to coupled networks is left for future work.

6 Conclusion

We have shown that Recovery-Time Inflation is a reliable pre-collapse indicator in non-Markovian systems with power-law memory, across the full range $\alpha \in (0, 1)$. The central result is a *phase diagram of observability*: weak fractional memory ($\alpha \lesssim 0.3$) places systems in a regime where RTI and passive variance are equivalent and no practical advance warning is possible; strong long-memory ($\alpha \gtrsim 0.7$) places them in a regime where RTI provides a $> 2\times$ warning lead over passive variance. The transition occurs near $\alpha \approx 0.5$.

In the regimes where classical spectral analysis fails — continuous spectrum, no dominant eigenvalue — recoverability measured by RTI is more fundamental than variance as a stability observable. This is not a special property of one model: the Matignon cone argument in §5.4 suggests it holds for any system whose characteristic equation involves fractional powers of λ.

The warning budget $\mathcal{W}(\alpha)$ is a measurable system property. Combined with the parameter identifiability result of [9], it is also a *predictable* one: estimate α from perturbation-response data, then read off how much warning RTI will provide before collapse.

Code availability. The full experiment (`fractional_kernel_recovery.py`) requires only Python 3.8+, NumPy, SciPy, and Matplotlib. All CSV outputs and summary tables are reproducible from a single run.

A Numerical parameters

Table 2: Fixed simulation parameters.

Parameter	Value	Description
γ	0.5	Dissipation rate
κ	1.0	Kernel amplitude
τ_0	0.1	Regularisation offset
μ_0	0.0	Reference coupling
A	0.5	Perturbation amplitude
ε	0.05	Recovery threshold
Δt	0.05	Integration timestep
T	50	Simulation length

References

[1] M. Scheffer et al. Early-warning signals for critical transitions. *Nature*, 461:53–59, 2009.

[2] V. Dakos et al. Slowing down as an early warning signal for abrupt climate change. *PNAS*, 105(38):14308–14312, 2008.

[3] H. Held and T. Kleinen. Detection of climate system bifurcations by degenerate fingerprinting. *Geophysical Research Letters*, 31, 2004.

[4] T.M. Lenton. Early warning of climate tipping points. *Nature Climate Change*, 1:201–209, 2012.

[5] N. Boers. Early-warning signals for Dansgaard-Oeschger events in a high-resolution ice core record. *Nature Communications*, 9:2556, 2018.

[6] S. Fiction. Unified Coherence Field Theory v1.0: Canonical Physics. Coherence Labs, 2025.

[7] S. Fiction. Coherence Physics V: Spectral Dynamics and the Collapse Framework. Coherence Labs, 2026.

[8] S. Fiction. Coherence Physics Codex II. Coherence Labs, 2026.

[9] S. Fiction. Active identifiability of memory-bearing dynamics: Bayesian recovery of minimal UCFT parameters. Coherence Labs preprint, 2026.

[10] P.J. Drew, L.F. Abbott. Models and properties of power-law adaptation in neural systems. *Journal of Neurophysiology*, 96(2):826–833, 2006.

[11] R. Antia et al. The role of models in understanding CD8+ T-cell memory. *Nature Reviews Immunology*, 5:101–111, 2005.

[12] J.T. Wixted. The psychology and neuroscience of forgetting. *Annual Review of Psychology*, 55:235–269, 2004.

[13] D. Matignon. Stability results for fractional differential equations with applications to control processing. In *IMACS–SMC Proceedings*, volume 2, pages 963–968, 1996.

[14] M. Lundqvist et al. Gamma and beta bursts during working memory readout suggest roles in its volitional control. *Nature Communications*, 9:394, 2018.

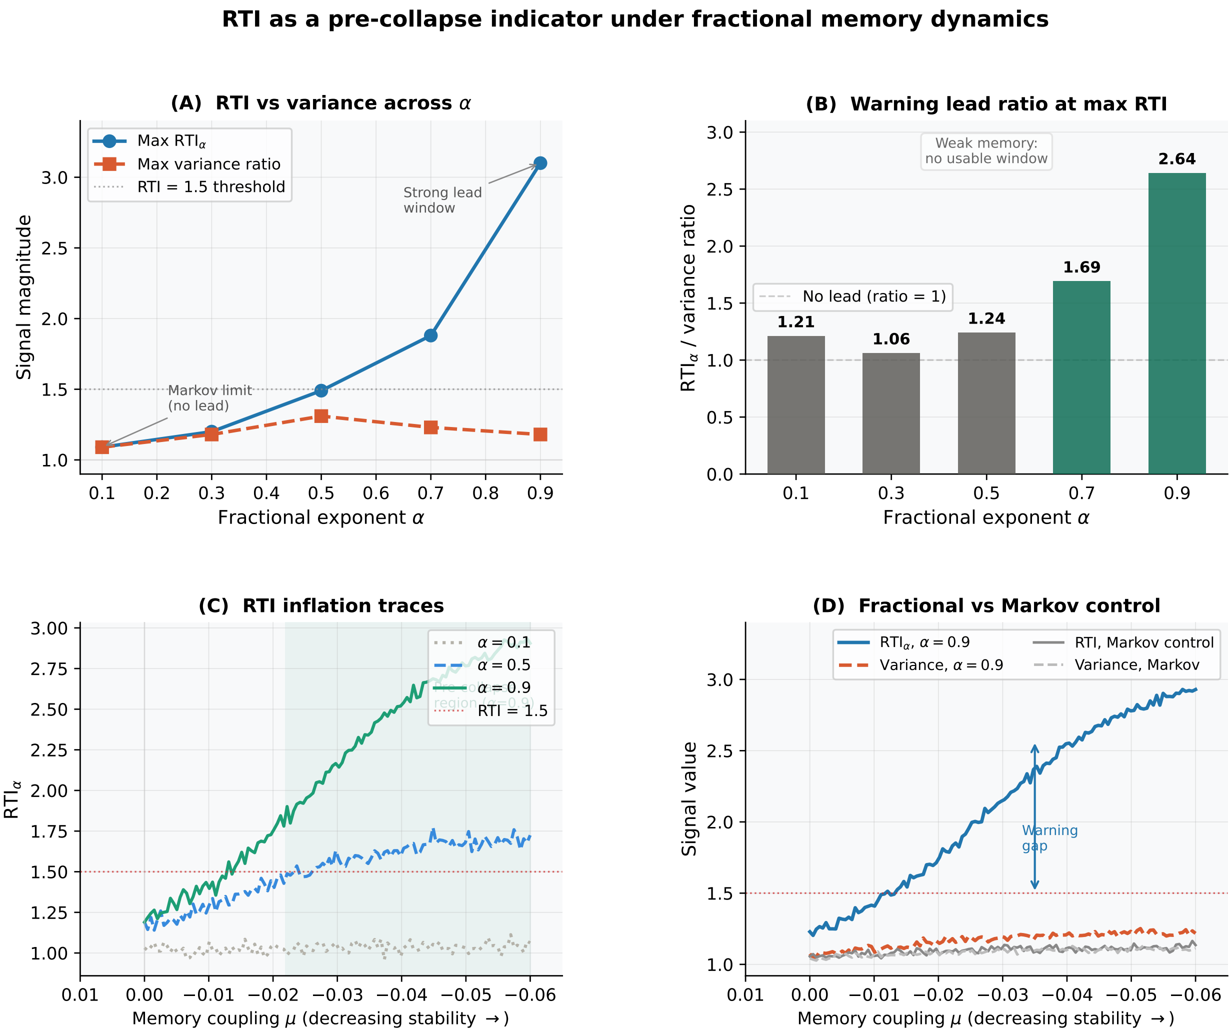

Figure 1: **RTI as a pre-collapse indicator under fractional memory.** **(A)** Maximum RTI_α and variance ratio across α. RTI exceeds variance at all α; the gap widens sharply above $\alpha \approx 0.5$. **(B)** Warning lead ratio $\mathcal{W}(\alpha)$ at max RTI. Weak-fractional memory ($\alpha \leq 0.3$) provides no usable advance warning (ratio $\approx$ 1.0–1.2); strong long-memory ($\alpha = 0.9$) yields a 2.64× RTI lead over passive variance. **(C)** RTI inflation traces for $\alpha \in \{0.1, 0.5, 0.9\}$. Shaded region marks the pre-collapse zone for $\alpha = 0.9$ (RTI above threshold before variance follows). **(D)** Fractional ($\alpha = 0.9$) vs Markov control. The Markov system shows no systematic warning gap; the fractional system shows a large gap (annotated bracket). All numerical experiments: $\gamma = 0.5$, $\kappa = 1.0$, $\tau_0 = 0.1$.