Spectral response of positive recurrences under memory-filtered balanced coefficient perturbations

W. Schulz

Jul 19, 2026

Abstract

When an AR(1) memory process perturbs a positive two-term recurrence while conserving total coefficient weight, memory systematically accelerates growth. The main theorem gives the exact spectral response and is numerically confirmed; the proof architecture is presented as a skeleton.

A research note. Status: main theorem numerically confirmed (see `verify_theorem.py` and `FINDINGS.md`); proof architecture below is a skeleton, not an audited proof; novelty asserted only as "not located," not established.

Alternative accessible title: When Memory Accelerates Recurrence Growth: A Spectral Response Theorem.

---

1. The model

Two coupled recursions. A driving memory process (AR(1) on the forcing):

$$d_{n+1} = a\, d_n + \varepsilon\, q_n, \qquad 0 \le a < 1,$$

and a positive companion recurrence whose coefficients it perturbs:

$$x_{n+1} = (A + d_n)\, x_n + (B - d_n)\, x_{n-1}, \qquad A, B > 0.$$

The perturbation is special. Written on the coefficient pair,

$$(A, B) \;\longmapsto\; (A + d_n,\; B - d_n),$$

it adds exactly what it subtracts:

$$(A + d_n) + (B - d_n) = A + B \quad \text{(conserved for every } n\text{)}.$$

The environment never changes the *total* coefficient weight; it only transfers influence between the present state x_n and the inherited state x_{n-1}. This is the balanced coefficient transfer, and it is what makes the model distinct from ordinary additive off-diagonal disorder.

Transfer-matrix form

$$\begin{pmatrix} x_{n+1} \\ x_n \end{pmatrix} = M_n \begin{pmatrix} x_n \\ x_{n-1} \end{pmatrix}, \qquad M_n = \begin{pmatrix} A + d_n & B - d_n \\ 1 & 0 \end{pmatrix} = \begin{pmatrix} A & B \\ 1 & 0 \end{pmatrix} + d_n \begin{pmatrix} 1 & -1 \\ 0 & 0 \end{pmatrix}.$$

The perturbation direction is a rank-one, coefficient-conserving update of a companion matrix. Note $\det M_n = -(B - d_n)$ fluctuates with the forcing — the cocycle is *not* naturally $\mathrm{SL}(2, \mathbb{R})$, so the Lyapunov exponent carries a determinant/scale contribution that the usual normalized localization cocycles hide. The derivation below works directly with the positive growth of x_n rather than a determinant-one projective cocycle.

Constants

$$\lambda = \frac{A + \sqrt{A^2 + 4B}}{2} \quad \text{(dominant root of } t^2 = At + B\text{)}, \qquad \beta = \frac{B}{\lambda^2}.$$

At $\varepsilon = 0$ the deterministic recurrence grows like λ^n, so the Lyapunov exponent is $\Lambda(0) = \log \lambda$.

---

2. Main theorem

Assumptions. $q_n = q \circ T^n$ is a bounded, stationary, mean-zero forcing over an ergodic T; σ_q is its spectral (power) measure; and ε is small enough that positivity holds, $A + d_n > 0$ and $B - d_n > 0$.

Claim.

$$\boxed{\Lambda(\varepsilon) = \log \lambda + \varepsilon^2 \int_{-\pi}^{\pi} W_{a,\lambda,\beta}(\omega)\, d\sigma_q(\omega) + O(\varepsilon^3), \qquad \Lambda'(0) = 0,}$$

with the response kernel

$$W_{a,\lambda,\beta}(\omega) = \frac{(\lambda - 1)\left[\, 2\cos\omega + \beta(\lambda + 1) - \lambda + 1 \,\right]}{2\lambda^4 (1 + \beta)(1 + \beta^2 + 2\beta\cos\omega)(1 + a^2 - 2a\cos\omega)}.$$

The linear term vanishes because centered forcing makes the first variation mean-zero (Lemma 6); the leading effect is second order.

The factorization — the real content

$$W_{a,\lambda,\beta}(\omega) = \underbrace{\frac{1}{1 + a^2 - 2a\cos\omega}}_{M_a(\omega)\ \text{— memory filter}} \cdot \underbrace{\frac{(\lambda - 1)\left[\, 2\cos\omega + \beta(\lambda + 1) - \lambda + 1 \,\right]}{2\lambda^4 (1 + \beta)(1 + \beta^2 + 2\beta\cos\omega)}}_{P_{\lambda,\beta}(\omega)\ \text{— projective geometry}}.$$

M_a is exactly the power response of a one-pole (AR(1)) filter; it is everywhere positive and monotone in $|\omega|$. It contains *all* the a-dependence, so the memory parameter cannot move the sign — it only reweights how strongly each frequency is felt.

$P_{\lambda,\beta}$ carries the sign. Its numerator is linear in $\cos\omega$, so it crosses zero at most once on $[0, \pi]$. Hence:

Memory controls strength; recurrence geometry controls sign.

This is sharper than the standard "colored noise changes a Lyapunov exponent": here the *what-frequencies-help* question and the *how-much* question separate cleanly into two factors.

---

3. Exact sub-results

3.1 Frequency transition (fixed recurrence)

For $\lambda > 1$ the response changes sign where the P-numerator vanishes:

$$\boxed{\cos\omega_c = \frac{\lambda - 1 - \beta(\lambda+1)}{2}.}$$

- If $\cos\omega_c \in (-1, 1)$ (Regime I, interior transition): low frequencies $|\omega| < \omega_c$ enhance growth, high frequencies $|\omega| > \omega_c$ suppress it.
- If $\cos\omega_c \notin [-1, 1]$ (Regime II): the kernel has one sign across all physical frequencies.
- $\lambda = 1$ (Regime III, critical): see 3.2.

The threshold is independent of the memory parameter a — confirmed to $< 10^{-3}$ rad in simulation for $a = 0.2$ and $a = 0.6$ (test T4).

3.2 The exact conservation anchor ($A + B = 1$)

$\lambda = 1 \iff A + B = 1$. Because the transfer is balanced,

$$F_d(1) = (A + d) + (B - d) = 1 \quad \text{for every } d,$$

so the ratio $r_n = x_n/x_{n-1} \equiv 1$ is an exact invariant for every admissible forcing sequence — not just to second order. The unit-growth state is protected nonperturbatively:

$$\boxed{A + B = 1 \implies \Lambda \equiv 0 \text{ on the } r \equiv 1 \text{ branch, for all forcing.}}$$

This is *why* the kernel carries the factor $(\lambda - 1)$: at $\lambda = 1$ the whole second-order response vanishes identically.

3.3 Fibonacci specialization ($A = B = 1$)

$$\lambda = \varphi = \tfrac{1+\sqrt{5}}{2}, \qquad \beta = \varphi^{-2}.$$

Sign-change frequency:

$$\boxed{\cos\omega_c = -\frac{1}{2\varphi^2} \approx -0.1910 \implies \omega_c \approx 1.7630 \text{ rad} \approx 101.1°.}$$

So low/medium frequencies accelerate Fibonacci growth (2nd order); fast oscillations suppress it. For spectrally flat (white) forcing the integrated response changes sign at a memory threshold:

$$\boxed{a_c = \frac{1}{3\varphi} \approx 0.2060.}$$

Below a_c the memory is too weak to overcome the harmful high-frequency mass; above it, the low-frequency emphasis of the memory filter makes the net response positive. Both constants are confirmed: the closed-form $\int W\, d\omega$ has its root at 0.20601, and simulated white-noise $C(a)$ crosses zero between $a = 0.1$ ($C < 0$) and $a = 0.35$ ($C > 0$), passing through ≈ 0 at a_c (test T5).

---

4. Proof architecture (skeleton)

Work with the ratio $r_n = x_n/x_{n-1}$, which follows the Riccati/Möbius map

$$r_{n+1} = F_{d_n}(r_n) = (A + d_n) + \frac{B - d_n}{r_n}, \qquad \Lambda = \mathbb{E}[\log r_\infty].$$

- Lemma 1 (bounded stationary memory). $D_\varepsilon(y) = \varepsilon \sum_{j\geq 0} a^j q(T^{-1-j}y)$ solves $D_\varepsilon(Ty) = aD_\varepsilon(y) + \varepsilon q(y)$ with $\|D_\varepsilon\|_\infty \leq \varepsilon\|q\|_\infty/(1-a)$. Controls positivity.
- Lemma 2 (absorbing interval). For $|d| \leq \delta < \min(A, B)$, all positive ratio trajectories enter a common compact interval.
- Lemma 3 (two-step contraction). F_d reverses order, but $F_{d_2} \circ F_{d_1}$ preserves it with derivative $\frac{(B-d_1)(B-d_2)}{((A+d_1)r+B-d_1)^2} < 1$ on the absorbing interval for small δ.
- Lemma 4 (unique invariant graph). The two-step graph transform is a contraction $\Rightarrow$ unique R_ε with $R_\varepsilon(Ty) = F_{D_\varepsilon(y)}(R_\varepsilon(y))$; every positive initial ratio synchronizes exponentially.
- Lemma 5 (smoothness). Uniform contraction + analyticity of F_d away from $r = 0 \Rightarrow \varepsilon \mapsto R_\varepsilon$ is C^3 near 0.
- Lemma 6 (first variation). Expanding $R_\varepsilon = \lambda + \varepsilon W + \varepsilon^2 V + O(\varepsilon^3)$, $W \circ T = -\beta W + \gamma Z$ with $\gamma = 1 - 1/\lambda$; centered forcing gives $\mathbb{E}[W] = 0$, hence $\Lambda'(0) = 0$.
- Lemma 7 (second variation). $V \circ T = -\beta V + ZW/\lambda^2 + BW^2/\lambda^3$; solving the linear filters for Z, W and applying the spectral theorem for stationary processes yields the kernel $W_{a,\lambda,\beta}$.

Caveat that must be stated in the theorem. The constants are *not* uniform as $a \to 1^-$: since $\|D_\varepsilon\|_\infty \leq \varepsilon\|q\|_\infty/(1-a)$, the admissible perturbation scales like $|\varepsilon| \lesssim 1 - a$. The theorem must read "for every fixed $a < 1$ there exists $\varepsilon_0(a) > 0$," not one ε_0 for all $a \in [0, 1)$.

---

5. Relation to known work (threat assessment)

Established, not claimed here: - random / weakly-disordered transfer matrices and perturbative Lyapunov exponents; - correlated diagonal and off-diagonal disorder; correlation-induced mobility edges; - random Fibonacci recurrences and their Riccati ratio maps; - regularity/response of Lyapunov exponents under perturbation; - general linear-response machinery for skew products with contracting fibers.

Not located in targeted searches ($\Rightarrow$ *candidate* novelty, not proven): the specific balanced-transfer model $(A + d_n, B - d_n)$ with an external AR(1) memory parameter, its closed-form second-order kernel, the

clean $M_a \cdot P$ factorization (sign independent of memory), and the Fibonacci constants $\cos \omega_c = -1/2\varphi^2$, $a_c = 1/3\varphi$.

Verdict after literature review (deflated). The real check was done and the general result is established prior art. The form *"weakly/randomly perturbed 2×2 recurrence ⇒ $\Lambda(\varepsilon)=\Lambda_0+\varepsilon^2\cdot$(spectral integral)$+O(\varepsilon^3)$, $\Lambda'(0)=0$, kernel integrated against the forcing's spectral density"* is owned by weak-noise Lyapunov expansions (Arnold, Papanicolaou & Wihstutz, *SIAM J. Appl. Math.* 46 (1986) 427–450; and "products of random 2×2 matrices close to the identity," arXiv:1208.6430) and by weak-disorder Lyapunov/localization theory (Thouless-type γ, correlated/colored disorder feeding a power spectrum into the kernel). The AR(1) memory `M_a` is just a Lorentzian noise spectrum plugged into that known formula. The specific balanced-transfer direction and the golden-ratio constants appear unpublished *in this exact form*, but they fall out mechanically — the equivalent of evaluating a known general integral for a particular integrand. Honest classification: a new *worked example* inside a fully established framework, not a new theorem and not a new field. This supersedes the earlier "original special theorem" self-assessment. Full citations and options in `FINDINGS-spectral-response.md`.

6. What is verified vs. open

Verified numerically (`verify_theorem.py`, T1–T5 all PASS): the baseline $\log \varphi$; clean ε^2 scaling with $\Lambda'(0) = 0$; kernel shape match to $k \cdot W_a(\omega)$ at shape-correlation 1.0000 for two memory values; a-independence of ω_c; the flat-forcing threshold $a_c = 1/3\varphi$ (both formula-self-consistent and simulation-confirmed).

Open: audited proof of Lemmas 1–7; the n-th-order generalization (transfer across many past states); nonlinear updates; and the practical applications in `docs/applications-directions.md`.

Published by NeuroGenesys (AW3 Technology, Inc.). Canonical version: https://neurogenesys.si/research/spectral-response