

# Non-normal eigenvector amplification in multidimensional Kesten processes

Virgile Troude  and Didier Sornette \*

*Institute of Risk Analysis, Prediction and Management (Risks-X), Academy for Advanced Interdisciplinary Sciences,  
Southern University of Science and Technology, Shenzhen, China*



(Received 13 October 2025; accepted 3 May 2026; published 13 July 2026)

Heavy-tailed fluctuations and power law statistics pervade physics, finance, and economics, yet their origin is often ascribed to systems poised near criticality. Here, we show that such behavior can emerge far from instability through a universal mechanism of *non-normal eigenvector amplification* in multidimensional Kesten processes  $\mathbf{x}_{t+1} = \mathbf{A}_t \mathbf{x}_t + \boldsymbol{\eta}_t$ , where  $\mathbf{A}_t$  are random interaction  $N \times N$  matrices and  $\boldsymbol{\eta}_t$  represents external inputs, capturing the evolving interdependence among  $N$  coupled components. Even when each random multiplicative matrix is spectrally stable, nonorthogonal eigenvectors generate transient growth that renormalizes the Lyapunov exponent and decreases the tail exponent, producing stationary power laws without eigenvalues crossing the stability boundary. We derive explicit relations, first in the form of bounds and subsequently as approximations, linking the Lyapunov exponent  $\gamma \sim \gamma_0 + \ln \kappa$  and the tail index  $\alpha \sim -2\gamma/\sigma_\kappa^2$  to the statistics of the condition number with mean  $\kappa$  and standard deviation  $\sigma_\kappa$ , confirmed by numerical simulations. This framework offers a unifying geometric perspective that helps interpret diverse phenomena, including polymer stretching in turbulence, magnetic-field amplification in dynamos, volatility clustering, and wealth inequality in financial systems. Non-normal interactions provide a collective route to scale-free behavior in globally stable systems, defining a universality class where multiplicative feedback and transient amplification generate critical-like statistics without spectral criticality.

DOI: [10.1103/ztdv-7t5j](https://doi.org/10.1103/ztdv-7t5j)

## I. INTRODUCTION

Heavy-tailed statistics pervade complex systems, appearing in phenomena as diverse as turbulent flows, financial markets, ecological dynamics, and biological growth. A widely used framework for generating such statistics is the class of Kesten-type stochastic recursions [1], where multiplicative growth combined with additive noise yields stationary power law tails. These *Kesten processes* form a unifying framework for scale-free behavior across disciplines, from disordered transport and self-organized criticality [2–4] to population dynamics [5], macroeconomic fluctuations [6–9], and wealth inequality [10,11].

Traditionally, the origin of power laws in such models is traced to episodes of *spectral supercriticality*, when eigenvalues of the random multiplicative operator temporarily exceed the stability boundary. Heavy tails are then viewed as the signature of rare supercritical bursts moderated by global mean reversion. Yet, this explanation overlooks a distinct geometric aspect of linear dynamics: Even when all eigenvalues indicate stability, nonorthogonal eigenvectors can transiently amplify fluctuations.

The  $L_2$ -norm formulation underlying Kesten theory is mathematically complete in that it characterizes instability

and power law tails through products of random matrices. However, it does not by itself elucidate the physical mechanisms responsible for the amplification that gives rise to power law distributions. The purpose of the present work is precisely to open this black box: We distinguish between growth driven by intermittent spectral supercriticality, and growth arising from a fundamentally geometric mechanism involving non-normality, rotation, and switching between modes.

*Non-normality* has long been recognized in hydrodynamic stability theory, where it accounts for transient energy growth in linearly stable shear flows [12–16]. Non-normal operators can exhibit responses several orders of magnitude stronger than predicted by eigenvalue analysis alone [17], and stochastic forcing can sustain persistent variability through such amplification mechanisms [18,19]. More broadly, non-normal amplification has emerged as a generic organizing principle of complex systems [20], governing phenomena ranging from turbulence [14,15] to neural and ecological networks [21,22].

In this paper, we identify a general mechanism linking non-normality to the statistics of multiplicative stochastic processes, arising naturally within the Kesten framework. We show that, even in the absence of spectral instability, *non-normal eigenvector amplification* can generically produce power law stationary distributions. When random matrices are non-normal (i.e., not unitarily diagonalizable), transient growth emerges from the coupling of nearly aligned eigenvectors, enabling intermittent bursts of amplification even when all eigenvalues lie strictly within the unit circle. This non-normal rotational effect reveals a distinct and generic pathway to heavy-tailed statistics, independent of classical spectral criticality.

\*Contact author: dsornette@ethz.ch

We then develop a quantitative framework connecting the Lyapunov exponent  $\gamma$  and the tail exponent  $\alpha$  to the statistical properties of eigenvector geometry, as captured by the matrix condition number  $\kappa$ . Specifically, we first derive rigorous bounds, and then explicit approximations, showing that the effective Lyapunov exponent is renormalized by the statistics of  $\kappa$ , while the tail exponent is correspondingly reduced. These relations make explicit how transient amplification, driven by eigenvector nonorthogonality, directly shapes the emergence of stationary heavy-tailed statistics.

It is important to distinguish the mechanism identified here from a trivial diagonal amplification mechanism. For instance, consider matrices with random diagonal entries of the form  $\text{diag}(z_t, z_t^{-1})$ . As soon as  $|z_t| \neq 1$ , one of the two diagonal elements exceeds unity, so any heavy-tailed behavior or positive Lyapunov contribution follows directly from intermittent supercriticality along a coordinate axis. This is *not* the regime studied in the present work. Our focus is instead on situations in which the spectrum remains strictly stable, while transient growth is generated by eigenvector nonorthogonality and by the repeated reinjection of the dynamics into amplifying directions through rotations or mode switching. This distinction is central: The key point is not merely that the operator norm can become large, but that heavy tails and effective destabilization can arise *without* any eigenvalue crossing the stability boundary. In this sense, the operator norm alone is not an explanation. It indicates *that* amplification is possible, but not *how* it is dynamically sustained. The mechanism we identify is precisely this: Constructive amplification emerges from the interplay of non-normal geometry and reinjection via rotations or mode switching.

The implications of this finding are broad. Non-normal amplification provides a form of *universality* in the emergence of power laws: Fat tails may arise not only from rare excursions into instability but also from generic geometrical properties of high-dimensional random operators. We illustrate this mechanism across multiple domains. In turbulent polymer solutions, it elucidates how intermittent stretching events emerge from transient alignments of the velocity-gradient eigenvectors. In magnetohydrodynamics, it clarifies the kinematic stage of the small-scale turbulent dynamo, where magnetic-field amplification follows the same geometric principle. In financial systems, it provides a natural multidimensional generalization of Kesten-generalized autoregressive conditional heteroskedasticity (GARCH)-type dynamics that rationalizes the near-universal inverse-cubic law ( $\alpha \approx 3$ ) observed across assets and markets without invoking exogenous heavy-tailed shocks.

By revealing non-normal eigenvector amplification as a fundamental and universal mechanism for the emergence of heavy tails, our work unifies previously disconnected phenomena, from fluid instabilities and dynamo action to financial volatility clustering, under a single theoretical framework. It thereby extends the classical theory of Kesten processes and situates non-normality as a central organizing principle in the dynamics and statistics of complex systems.

The paper is organized as follows. Section II introduces the general mathematical framework. Section III examines the two-dimensional case, which provides analytical tractability and offers key insights into the underlying mechanism.

Section IV extends the quantitative analysis to the general  $N$ -dimensional setting. Section V illustrates the relevance of the results through two physical examples (polymer stretching in turbulent flows and the small-scale turbulent dynamo) and four applications in finance and economics (volatility models of the autoregressive conditional heteroskedasticity (ARCH) type, cointegration models, factor models, and wealth inequality dynamics). Section VI summarizes the main findings and outlines perspectives for future work.

## II. MATHEMATICAL FRAMEWORK

Throughout this section, we separate two conceptually different routes to heavy-tailed distributions. The first is the classical one, in which stochastic excursions of the spectrum across the unit circle generate intermittent supercritical growth. The second, which is the focus of this paper, occurs when the spectrum remains stable but the geometry of the eigenvectors produces transient amplification through non-normality. Our goal is to make this distinction precise, first through general bounds and then through explicit approximations.

We begin by formalizing the multidimensional Kesten process and recalling classical results. This section provides the necessary preliminaries before turning to the role of non-normality in driving critical behavior.

This section serves two complementary purposes. First, it recalls the compact matrix-product formulation in which the  $L_2$  norm encapsulates the Kesten mechanism. Second, it clarifies why this formulation, while mathematically complete, remains physically incomplete unless the dynamical origin of norm growth is explicitly identified, namely, the interplay of eigenvector geometry, reinjection, and mode switching that underlies transient amplification in our framework.

### A. Multidimensional Kesten process

We consider a generalized  $N$ -dimensional Kesten process  $(\mathbf{x}_t)_{t \geq 0}$  defined by

$$\mathbf{x}_{t+1} = \mathbf{A}_t \mathbf{x}_t + \sqrt{2\delta} \boldsymbol{\eta}_t, \quad \boldsymbol{\eta}_t \stackrel{i.i.d.}{\sim} \mathcal{N}(0, \mathbf{I}), \quad (1)$$

where the random matrices  $(\mathbf{A}_t)$  are assumed to be independent and identically distributed (i.i.d.).

Following the framework introduced by Furstenberg and Kesten [23], we define the product norm

$$\pi_t = \|\Pi_t\|, \quad \Pi_t = \mathbf{A}_t \Pi_{t-1}, \quad \Pi_0 = \mathbf{I}, \quad (2)$$

where  $\|\cdot\|$  denotes the matrix norm, taken here to be the  $L_2$  norm. The associated Lyapunov exponent is then

$$\gamma := \lim_{t \rightarrow \infty} \frac{1}{t} \mathbb{E}[\ln \pi_t]. \quad (3)$$

A classical sufficient condition for the existence of a unique stationary solution (and for the recursion to be geometrically ergodic) is the “average contraction” condition  $\gamma < 0$ , together with the nondegeneracy condition  $\mathbb{P}[\boldsymbol{\eta} = 0] < 1$ . In this case, the stationary forward solution is given by

$$\mathbf{x} \stackrel{d}{=} \lim_{T \rightarrow \infty} \sqrt{2\delta} \sum_{t=0}^T \Pi_t \boldsymbol{\eta}_{T-t}. \quad (4)$$

The intuition is that when  $\gamma < 0$ , the products  $\Pi_t$  decay exponentially fast to zero. Conversely, if  $\gamma > 0$ , the process diverges exponentially and no stationary distribution exists.

A second fundamental result concerns the tail behavior of the stationary solution. Under the following assumptions,

- (1) Global stability:  $\gamma < 0$ .
- (2) Local instability:  $\mathbb{P}[\|\mathbf{A}_t\| > 1] > 0$ .
- (3) A mild nonarithmetic (nonslattice) condition on the distribution of  $\ln \|\mathbf{A}_t\|$ .
- (4) Defining the moment generating function of the random matrix product  $\Pi_t$  through

$$\phi(\alpha) := \lim_{t \rightarrow \infty} \frac{1}{t} \ln \mathbb{E}[\|\Pi_t\|^\alpha], \quad (5)$$

there exists an  $\alpha > 0$  such that  $\phi(\alpha) = 0$  and  $\mathbb{E}[\|\eta_t\|^\alpha] < \infty$ .

Then, the stationary distribution exhibits a power law tail of the form

$$\mathbb{P}[\|\mathbf{x}\| > x] \sim Cx^{-\alpha}, \quad x \rightarrow \infty. \quad (6)$$

Moreover, the stationary solution has finite moments of order  $p$  if and only if  $p < \alpha$ , that is,

$$\mathbb{E}[\|\mathbf{x}\|^p] < \infty \quad \text{if } p < \alpha, \quad (7)$$

and diverges for  $p \geq \alpha$ . This result constitutes the  $N$ -dimensional extension of the classical scalar Kesten-Goldie theorem.

### B. One-dimensional illustration

In the one-dimensional setting, the process reduces to

$$x_{t+1} = \rho_t x_t + \sqrt{2\delta} \eta_t, \quad \eta_t \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1), \quad (8)$$

where  $(\rho_t)$  is an i.i.d. sequence of random multipliers. In this case, the Lyapunov exponent (3) simplifies to

$$\gamma = \mathbb{E}[\ln \rho_t]. \quad (9)$$

To obtain explicit numerical results, we must specify the distribution of  $\rho_t$ . A convenient choice is to assume  $\ln \rho_t \sim \mathcal{N}(\ln \rho, \sigma_\rho^2)$ , so that  $\ln \rho_t$  follows a normal distribution with mean parameter  $\ln \rho$ . Here,  $\ln \rho$  being the expectation of  $\ln \rho_t$  directly defines the Lyapunov exponent  $\gamma = \ln \rho$ .

Under this assumption, the power law exponent  $\alpha$  is determined by the equation

$$\phi(\alpha) := \ln \mathbb{E}[\rho_t^\alpha] = 0. \quad (10)$$

Evaluating this expression yields

$$\phi(\alpha) = \gamma\alpha + \frac{1}{2}\sigma_\rho^2\alpha^2 = 0 \implies \alpha = -2\frac{\gamma}{\sigma_\rho^2}. \quad (11)$$

Therefore, as long as  $\gamma < 0$  (ensuring average contraction), the system remains stable. However, if the variance of the multiplicative noise satisfies  $\sigma_\rho^2 > |\gamma|$ , the stationary distribution exists, but its variance is infinite. This illustrates how heavy-tailed behavior can arise even in a stable one-dimensional Kesten process.

In the one-dimensional Kesten process (8), the emergence of power law tails can be understood in terms of spectral criticality. In standard dynamical systems language, a system becomes critical when  $\rho$  approaches the unit root,

i.e.,  $\rho \rightarrow \pm 1$ . However, in a Kesten process, the stochastic fluctuations of  $\rho_t$  imply that the process may occasionally cross the unit root, since  $0 < \mathbb{P}[|\rho_t| > 1] < 1$ . The power law (6) arises from the interplay between two opposing exponential processes [10,24,25]. Transient episodes where  $|\rho_t| > 1$  cause multiplicative (and thus exponential) amplification of  $x_t$  over their duration  $n$ . However, such supercritical bursts occur with an exponentially decaying probability as  $n$  increases. The balance between the exponential growth of  $x_t$  during these rare intervals and their exponentially decreasing frequency produces an emergent power law tail. In essence, the heavy tail reflects the statistics of transient supercritical phases that exponentially amplify  $x_t$  but whose occurrence becomes exponentially rarer with increasing duration.

### C. Normal Kesten processes

In the  $N$ -dimensional case (1), the same intuition applies: Power law tails arise from fluctuations between subcritical and supercritical regimes. To make this precise, consider the decomposition  $\mathbf{A}_t = \mathbf{P}_t \mathbf{\Lambda}_t \mathbf{P}_t^{-1}$ , where  $\mathbf{P}_t$  is a unitary matrix, i.e.,  $\mathbf{P}_t^{-1} = \mathbf{P}_t^\dagger$ , and  $\mathbf{\Lambda}_t = \text{Diag}(\lambda_{i,t} \mid i = 1, \dots, N)$ , with the eigenvalues  $\lambda_{i,t}$ .

For a product of matrices, we have the inequality

$$\pi_t \leq \prod_{s=1}^t \|\mathbf{A}_s\| = \prod_{s=1}^t \rho_s, \quad (12)$$

because  $\mathbf{P}_t$  are unitary matrices, each individual matrix  $\mathbf{A}_t$  is normal, and so  $\rho_s = \|\mathbf{A}_s\|$  is the spectral radius of each individual matrix  $\mathbf{A}_s$ .

Defining

$$X_t = \ln \pi_t \quad \text{and} \quad X_t^0 = \sum_{s=1}^t \ln \rho_s, \quad (13)$$

we have the inequality  $X_t \leq X_t^0$ . From the *central limit theorem* (CLT),  $X_t^0 \sim \mathcal{N}(t \ln \rho, t \sigma_\rho^2)$  where  $\sigma_\rho^2 = \text{Var}[\ln \rho_s]$ .

This provides an upper bound for the Lyapunov exponent:

$$\gamma \leq \gamma_0 := \ln \rho. \quad (14)$$

Thus, in the normal case, the scalar results (11) and (12) provide an upper bound for the Lyapunov exponent  $\gamma$ .

The submultiplicativity of the operator norm (12) implies the upper bound for  $\phi(\alpha)$  (5),

$$\phi(\alpha) \leq \phi_0(\alpha), \quad (15)$$

where

$$\phi_0(\alpha) = \lim_{t \rightarrow \infty} \frac{1}{t} \ln \mathbb{E}[\exp(\alpha X_t^0)]. \quad (16)$$

Since  $\phi(\alpha)$  is convex, if  $\alpha$  is the solution of  $\phi(\alpha) = 0$ , then any  $\alpha_0 > 0$  with  $\phi_0(\alpha_0) = 0$  satisfies  $\alpha_0 < \alpha$ , where

$$\alpha_0 = -2\frac{\gamma_0}{\sigma_\rho^2}, \quad (17)$$

which provides a lower bound on the true tail exponent. The detailed proof of this result is as follows. The two functions

$\phi, \phi_0 : [0, \infty) \rightarrow \mathbb{R}$  are convex functions with

$$\phi(0) = \phi_0(0) = 0, \quad \phi'(0) = \gamma < 0, \quad \phi'_0(0) = \gamma_0 \geq \gamma, \quad (18)$$

and for all  $\alpha \geq 0$  we have the pointwise bound (15) with (16). By convexity and  $\phi'(0) < 0$ , both  $\phi$  and  $\phi_0$  admit a *unique* zero in  $(0, \infty)$ , which we denote by

$$\phi(\alpha) = 0 \text{ at } \alpha > 0, \quad \phi_0(\alpha_0) = 0 \text{ at } \alpha_0 > 0. \quad (19)$$

Since  $\phi_0(\alpha_0) = 0$  and  $\phi(\alpha) \leq \phi_0(\alpha)$  for all  $\alpha \geq 0$ , we have

$$\phi(\alpha_0) \leq \phi_0(\alpha_0) = 0. \quad (20)$$

Because  $\phi$  is convex with  $\phi(0) = 0$  and  $\phi'(0) < 0$ ,  $\phi$  is strictly negative on a right neighborhood of 0, and then (by convexity) crosses 0 exactly once at  $\alpha > 0$ . The inequality  $\phi(\alpha_0) \leq 0$  means that  $\alpha_0$  lies at or to the left of the zero of  $\phi$ , hence  $\alpha_0 \leq \alpha$ . Moreover, if either  $\gamma_0 > \gamma$  or the bound is strict for some  $\alpha > 0$  (which is the case under the normal-versus-general comparison), then  $\phi(\alpha_0) < 0$  and the crossing of  $\phi$  must occur strictly to the right of  $\alpha_0$ , yielding  $\alpha_0 < \alpha$ .

In summary, for normal random matrices  $\mathbf{A}_t$ , the one-dimensional Kesten process aligned with the leading eigendirection provides an extremal benchmark: It yields the largest possible Lyapunov exponent and the smallest possible tail exponent. In this precise sense, it represents the worst-case scenario for stability (maximal growth) and heavy tailedness (maximal tail weight) among multidimensional Kesten processes with normal matrices.

This also clarifies the role of mode switching in the normal case. When the eigenbasis is orthogonal, changes of direction do not produce constructive amplification: They merely redistribute growth across independent modes and typically reduce the effective persistence of the dominant one. Thus, in normal systems, rotations and switching do not create an amplification mechanism beyond the usual spectral one.

#### D. Non-normal Kesten processes

A matrix is called *non-normal* if it cannot be diagonalized in a unitary basis, that is, if  $\mathbf{P}_t^{-1} \neq \mathbf{P}_t^\dagger$  in the decomposition  $\mathbf{A}_t = \mathbf{P}_t \mathbf{\Lambda}_t \mathbf{P}_t^{-1}$ . In this case, by the classical bound for diagonalizable matrices (see, e.g., Ref. [26]), with equality (and  $\kappa = 1$ ) when the eigenbasis is orthogonal, the matrix norm satisfies the inequality

$$\|\mathbf{A}_t\| \leq \rho_t \kappa_t, \quad \kappa_t = \|\mathbf{P}_t\| \|\mathbf{P}_t^{-1}\|, \quad (21)$$

where  $\kappa_t$  denotes the condition number of the eigenbasis transformation matrix  $\mathbf{P}_t$ , called *the degree of non-normality*, and  $\rho_t$  is the spectral radius of  $\mathbf{A}_t$ . In this section, the only relevant quantities are the variance of the log-spectral radius and log-degree of non-normality. Since the matrices  $(\mathbf{A}_t)$  are i.i.d., the condition numbers  $(\kappa_t)$  are also i.i.d.

Taking products and using submultiplicativity of the norm, the upper bound in Eq. (13) becomes

$$X_t^0 = \sum_{s=1}^t \ln(\rho_s \kappa_s), \quad (22)$$

so that  $\ln \pi_t = X_t \leq X_t^0$ , and by the CLT  $X_t^0 \sim \mathcal{N}(t\gamma_0, t\sigma_0^2)$ , allowing us to modify the upper bounds for the Lyapunov

exponent as

$$\gamma \leq \gamma_0 = \ln \rho + \ln \kappa, \quad \ln \kappa = \mathbb{E}[\ln \kappa_t] \geq 0. \quad (23)$$

Therefore, the non-normality of the matrix increases the upper bound on the Lyapunov exponent. Similarly, we can extract a lower bound for the tail exponent as

$$\alpha \geq \alpha_0 = -2 \frac{\gamma_0}{\sigma_0^2}, \quad (24)$$

with

$$\sigma_0^2 = \sigma_\rho^2 + \sigma_\kappa^2 + 2 \text{Cov}[\ln \rho_t, \ln \kappa_t], \quad (25)$$

where  $\sigma_\rho^2 = \text{Var}[\ln \rho_s]$  and  $\sigma_\kappa^2 = \text{Var}[\ln \kappa_s]$ . Therefore, the lower bound for the power law exponent is decreased by the presence of non-normality.

In summary, non-normality of the matrices  $\mathbf{A}_t$  raises the upper bound on the Lyapunov exponent and lowers the lower bound on the tail exponent. Consequently, it enlarges the parameter region in which the system approaches instability ( $\gamma \geq 0$ ), while simultaneously promoting heavier tails in the stationary distribution.

However, the key point is that the factor  $\kappa_t$  should not be viewed as a purely formal correction in a norm inequality. Its dynamical impact depends on whether successive amplifying directions combine constructively over time. Non-normality becomes physically effective only when rotation or mode switching repeatedly reinjects the dynamics into these directions; otherwise, large instantaneous values of  $\kappa_t$  need not translate into an increased Lyapunov exponent.

*Remark: Non-normality versus diagonal intermittency.* The bound  $\|\mathbf{A}_t\| \leq \rho_t \kappa_t$  should not be conflated with mechanisms based on diagonal amplification or eigenvalues exceeding unity. If  $\rho_t > 1$  occurs with positive probability, heavy tails may already arise from the classical Kesten mechanism of intermittent spectral supercriticality. The genuinely distinct regime emphasized here is that in which  $\rho_t < 1$  almost surely, yet fluctuations of  $\kappa_t$ , together with the associated eigenvector geometry, generate transient growth, enhanced effective Lyapunov exponents, and heavier stationary tails.

#### E. Synthesis

We have introduced Kesten processes, a class of stochastic recursions driven by random multiplicative dynamics and additive noise. In the one-dimensional setting, such processes exhibit stationary distributions with heavy-tailed power law behavior, provided the Lyapunov exponent  $\gamma$  is negative (ensuring average contraction), but the local dynamics occasionally exceed the unit root. The tail exponent  $\alpha$  is then determined by the balance between contraction and fluctuations, capturing how moments of the stationary solution may diverge.

Extending this picture to higher dimensions reveals the key mechanism is unchanged: Criticality emerges when random fluctuations intermittently push the system across the spectral stability boundary. For normal random matrices, the Lyapunov exponent and the power law exponent can be bounded using the one-dimensional theory, linking criticality directly to the eigenvalue spectrum. This corresponds to the familiar notion of *spectral criticality* in natural systems, where

behavior changes qualitatively as eigenvalues approach the unit circle.

However, when the random matrices are non-normal, a distinct phenomenon arises. Because the condition number of the eigenbasis amplifies fluctuations, non-normality increases the effective Lyapunov exponent and decreases the tail exponent, thereby enlarging the region of parameter space where the system behaves as if it were critical. In other words, even when the spectrum indicates stability, non-normality can drive the process closer to true criticality and produce heavier-tailed stationary distributions. This highlights non-normality as a key mechanism that can mimic or enhance critical behavior in stochastic dynamical systems.

Throughout the derivation, we have assumed the standard Kesten conditions. Beyond these, the sole additional assumption is the finiteness of the second moments of the log-spectral radius and the log-condition number, enabling the application of the central limit theorem.

In the next section, we turn to explicit examples to illustrate how the theoretical bounds derived here manifest in practice, and how non-normality alters the effective stability and tail properties of Kesten processes.

### III. TWO-DIMENSIONAL ILLUSTRATION

The general theory developed in the previous section applies to arbitrary dimension, but it can be difficult to develop intuition. To gain more insight, we now focus on the two-dimensional case. This setting is simple enough to allow explicit calculations, while still capturing the essential role of non-normality in modifying stability and tail behavior.

We proceed step by step: first by analyzing a tractable special form of random matrices (where two-step products become diagonal), then by studying how the Lyapunov exponent and the tail exponent can be computed exactly, and finally by comparing these results with the general bounds derived earlier. This two-dimensional illustration serves both as a consistency check and as a bridge toward more complex applications presented later.

#### A. A first example

To gain intuition about the effect of non-normality in higher dimensions, let us start with a two-dimensional example. We consider matrices of the form

$$\mathbf{A}_t = \rho_t \begin{pmatrix} 0 & z_t \\ z_t^{-1} & 0 \end{pmatrix}, \quad z_t = \frac{s_{1,t}}{s_{2,t}}, \quad (26)$$

where  $\rho_t$  sets the spectral radius and  $s_{1,t}, s_{2,t}$  are the singular values of the eigenbasis transformation matrix  $\mathbf{P}_t$ .

Our goal is to compute explicitly the Lyapunov exponent  $\gamma$  and the tail exponent  $\alpha$ . We begin with the product norm  $\pi_t$  defined in Eq. (2). A key simplification arises if we look at two-step products:

$$\begin{aligned} \mathbf{A}_t^{(2)} &= \mathbf{A}_{2t} \mathbf{A}_{2t-1} = \rho_{2t} \rho_{2t-1} \begin{pmatrix} z_t^{(2)} & 0 \\ 0 & (z_t^{(2)})^{-1} \end{pmatrix}, \\ z_t^{(2)} &= \frac{z_{2t}}{z_{2t-1}}. \end{aligned} \quad (27)$$

Thus, every two-step product is diagonal, which makes the analysis tractable as it reduces to the one-dimensional setup. The product norm over  $2t$  steps reads

$$\begin{aligned} \pi_{2t} &= \left\| \prod_{s=1}^t \mathbf{A}_s^{(2)} \right\| = \left( \prod_{s=1}^{2t} |\rho_s| \right) \max(Z_t, Z_t^{-1}), \\ Z_t &= \prod_{s=1}^t z_s^{(2)}. \end{aligned} \quad (28)$$

Taking logarithms and using monotonicity of  $\ln(\cdot)$  gives

$$\ln \pi_{2t} = \sum_{s=1}^{2t} \ln |\rho_s| + |\ln Z_t|. \quad (29)$$

By the law of large numbers, the Lyapunov exponent is

$$\gamma = \ln \rho + \frac{1}{2} |\mathbb{E}[\ln z_1^{(2)}]|, \quad \ln \rho := \mathbb{E}[\ln |\rho_1|]. \quad (30)$$

However, since  $z_t^{(2)} = z_{2t}/z_{2t-1}$  and  $z_i$  are i.i.d., we have  $\mathbb{E}[\ln z_1^{(2)}] = \mathbb{E}[\ln z_2 - \ln z_1] = 0$ . Hence, the Lyapunov exponent simplifies to

$$\gamma = \ln \rho. \quad (31)$$

In other words, in this specific construction, the non-normality cancels out on average, and the Lyapunov exponent is determined solely by the spectral radius.

The tail behavior, however, is more subtle as it depends on the distribution of  $\ln z_t^{(2)}$ . If  $\text{Var}[\ln z_t^{(2)}] < \infty$ , then by the central limit theorem,

$$\ln Z_t \sim \mathcal{N}(0, 2t\sigma_z^2), \quad \sigma_z^2 := \text{Var}[\ln z_1]. \quad (32)$$

At the same time,

$$\sum_{s=1}^{2t} \ln |\rho_s| \sim \mathcal{N}(2t \ln \rho, 2t\sigma_\rho^2), \quad \sigma_\rho^2 := \text{Var}[\ln |\rho_1|]. \quad (33)$$

Therefore, the logarithm of the product norm is approximately Gaussian:

$$\ln \pi_{2t} \sim \mathcal{N}(2t \ln \rho, 2t\sigma_0^2), \quad (34)$$

where

$$\sigma_0^2 = \sigma_\rho^2 + \sigma_z^2 + 2\text{Cov}[\ln |\rho_1|, \ln z_1], \quad (35)$$

which accounts for the covariance term since  $\rho_t$  and  $z_t$  may be correlated. It follows that the tail exponent is given by the solution of  $\phi(\alpha) = 0$  with  $\phi(\alpha)$  given by Eq. (5), which yields

$$\alpha = -\frac{2\gamma}{\sigma_0^2}. \quad (36)$$

Thus, for non-normal matrices (26), non-normality does not affect stability (since  $\gamma$  remains unchanged), but it does decrease the tail exponent and therefore increases the heaviness of the stationary distribution. This result shows that the geometry of the eigenvectors by itself can produce and enhance fat-tailed stationary distributions. Remarkably, even when transient supercritical excursions are absent ( $\sigma_\rho = 0$ ), the non-normal amplification mechanism alone gives rise to power law tails characterized by  $\alpha = -\frac{2\gamma}{\sigma_z^2}$ .

This point is conceptually important. Although the  $L_2$  norm of  $\mathbf{A}_t$  depends explicitly on  $z_t$  and thus on the degree of

non-normality, this dependence alone does *not* increase the Lyapunov exponent in the present two-dimensional setting; it only lowers the tail exponent. Hence, the mere appearance of  $\kappa_t$  in the norm is not, by itself, the mechanism responsible for destabilization: Constructive accumulation requires an additional reinjection mechanism. At the same time, this solvable example highlights the limitation of a purely norm-based interpretation: Non-normality can broaden the stationary distribution without increasing the long-run growth rate. To affect the Lyapunov exponent itself, one must understand how amplifying directions are dynamically coupled across time.

Finally, we observe that the log-condition number is

$$\ln \kappa_t = |\ln z_t|, \quad (37)$$

which implies  $\sigma_z^2 \geq \sigma_\kappa^2$ . Hence, the explicit solutions for the Lyapunov exponent (31) and the tail exponent (36) are consistent with the general bounds derived in Eqs. (23) and (24).

### B. General case

In the previous subsection, we have considered a special  $2 \times 2$  construction for which the Lyapunov exponent turned out to be independent of non-normality. To understand the situation more generally, we refer to Ref. [20], which showed that any non-normal  $2 \times 2$  matrix can be written in the form

$$\mathbf{A} = \frac{1}{2} \mathbf{U} \begin{pmatrix} \lambda_1 + \lambda_2 & z(\lambda_1 - \lambda_2) \\ z^{-1}(\lambda_1 - \lambda_2) & \lambda_1 + \lambda_2 \end{pmatrix} \mathbf{U}^\dagger, \quad (38)$$

where  $\lambda_1, \lambda_2$  are the eigenvalues of  $\mathbf{A}$ ,  $z = s_1/s_2$  is the ratio of the singular values of the eigenbasis transformation matrix  $\mathbf{P}$ , and  $\mathbf{U}$  is a unitary matrix.

For real  $2 \times 2$  matrices, the eigenvalues are either both real or appear as complex conjugates. In both cases, the spectral radius is given by

$$\begin{aligned} \rho &= \max(|\lambda_1|, |\lambda_2|) \quad (\text{real case}), \\ \rho &= |\lambda_1| = |\lambda_2| \quad (\text{complex conjugate case}). \end{aligned} \quad (39)$$

(1) *Case 1: Real eigenvalues.* When the eigenvalues are real, it is convenient to introduce the relative spectral distance

$$\delta = \frac{|\lambda_1| - |\lambda_2|}{|\lambda_1| + |\lambda_2|} \in [0, 1], \quad (40)$$

which quantifies how far apart the eigenvalues are compared to their average size. With this notation, the matrix can be expressed as

$$\mathbf{A} = \rho \mathbf{U} \begin{pmatrix} \delta & (1 - \delta)z \\ (1 - \delta)z^{-1} & \delta \end{pmatrix} \mathbf{U}^\dagger. \quad (41)$$

(2) *Case 2: Complex conjugate eigenvalues.* If the eigenvalues are complex conjugates, one can instead define

$$\delta = \frac{\text{Re}(\lambda_1)}{|\lambda_1|} \in [-1, 1], \quad (42)$$

which captures the ratio between the real and modulus parts of the eigenvalue. The matrix then takes the form

$$\mathbf{A} = \rho \mathbf{U} \begin{pmatrix} \delta & i\sqrt{1 - \delta^2}z \\ i\sqrt{1 - \delta^2}z^{-1} & \delta \end{pmatrix} \mathbf{U}^\dagger. \quad (43)$$

In both situations, the parameters  $\rho$ ,  $\delta$ , and  $z$  jointly control the stability and tail properties of the corresponding Kesten process. The parameter  $\rho$  governs the spectral radius (and thus the baseline stability condition),  $\delta$  measures the distance to spectral degeneracy, and  $z$  encodes the non-normality of the eigenbasis.

It is particularly informative to study the limiting regimes  $\delta \approx 0$  (close to spectral degeneracy, where non-normality plays the strongest role) and  $1 - \delta \approx 0$  (where the system is nearly normal). In the next subsections, we analyze how these different regimes affect the Lyapunov exponent and the tail exponent.

### C. Approximation for large non-normality

For the general two-dimensional cases (41) and (43), let us study how non-normality modifies the Lyapunov and tail exponents. To make analytical progress, we consider the regime where the parameter  $\delta$  is small (i.e., close to spectral degeneracy). In this setting, perturbation theory provides a useful approximation by expanding the dynamics to first order in  $\delta$ .

(1) *Case 1: Real eigenvalues.* When the eigenvalues are real, the matrix can be expressed as

$$\mathbf{A}_t = \rho_t (A_{0,t} + \delta(\mathbf{I} - A_{0,t})), \quad A_{0,t} = \begin{pmatrix} 0 & z_t \\ z_t^{-1} & 0 \end{pmatrix}. \quad (44)$$

Expanding the product over  $2t$  steps gives

$$\begin{aligned} \Pi_{2t} &= \prod_{s=1}^{2t} \mathbf{A}_s = \left( \prod_{s=1}^{2t} \rho_s \right) \prod_{s=1}^{2t} [A_{0,s} + \delta(\mathbf{I} - A_{0,s})] \\ &= \left( \prod_{s=1}^{2t} \rho_s \right) \left[ \prod_{s=1}^{2t} A_{0,s} + \delta \sum_{r=1}^{2t} \left( \prod_{s=1}^{r-1} A_{0,s} \right) (\mathbf{I} - A_{0,r}) \right. \\ &\quad \left. \times \left( \prod_{s=r+1}^{2t} A_{0,s} \right) \right] + \mathcal{O}(\delta^2). \end{aligned} \quad (45)$$

After simplification, this reduces to

$$\Pi_{2t} = \left( \prod_{s=1}^{2t} \rho_s \right) \left( \prod_{s=1}^{2t} A_{0,s} \right) [\mathbf{I} - 2t\delta(\mathbf{I} - \mathbf{B}_t)] + \mathcal{O}(\delta^2), \quad (46)$$

where

$$\begin{aligned} \mathbf{B}_t &= \begin{pmatrix} 0 & k_{2t}^+ \\ k_t^- & 0 \end{pmatrix}, \\ k_t^+ &= \frac{1}{t} \sum_{s=1}^t z_s^{(-1)^s}, \quad k_t^- = \frac{1}{2t} \sum_{s=1}^{2t} z_s^{(-1)^{s+1}}. \end{aligned} \quad (47)$$

We note that the product of the  $A_{0,s}$  terms simplifies to a diagonal matrix. In particular, defining

$$\mathbf{Z}_t := \prod_{s=1}^t z_s, \quad (48)$$

we have

$$\prod_{s=1}^{2t} A_{0,s} = \text{Diag}(\mathbf{Z}_t, \mathbf{Z}_t^{-1}). \quad (49)$$

We thus obtain

$$\ln \pi_{2t} = \sum_{s=1}^{2t} \ln |\rho_s| + |\ln Z_t| - 2t\delta + \mathcal{O}(\delta^2). \quad (50)$$

Therefore, in the large- $t$  limit,

$$\ln \pi_{2t} \sim \mathcal{N}(2t(\ln \rho - \delta), 2t\sigma_0^2) + \mathcal{O}(\delta^2). \quad (51)$$

This implies the approximations

$$\gamma = \ln \rho - \delta + \mathcal{O}(\delta^2), \quad \alpha = -\frac{2\gamma}{\sigma_0^2} + \mathcal{O}(\delta^2). \quad (52)$$

Thus, to first order in  $\delta$ , the Lyapunov exponent decreases linearly with the relative spectral distance (40), while the tail exponent retains the same form as in the  $\delta = 0$  case. The correction to the tail exponent  $\alpha$  appears only at second order in  $\delta$ .

(2) *Case 2: Complex conjugate eigenvalues.* When the eigenvalues are complex conjugates, the expansion reads

$$\mathbf{A}_t = i\rho_t(A_{0,t} - i\delta\mathbf{I}) + \mathcal{O}(\delta^2). \quad (53)$$

The product then becomes

$$\begin{aligned} \Pi_{2t} &= \prod_{s=1}^{2t} \mathbf{A}_s = i^{2t} \left( \prod_{s=1}^{2t} \rho_s \right) \prod_{s=1}^{2t} (A_{0,s} - i\delta\mathbf{I}) + \mathcal{O}(\delta^2) \\ &= i^{2t} \left( \prod_{s=1}^{2t} \rho_s \right) \left( \prod_{s=1}^{2t} A_{0,s} \right) [\mathbf{I} - 2i\delta t \mathbf{B}_t] + \mathcal{O}(\delta^2). \end{aligned} \quad (54)$$

Consequently, the logarithm of the product norm is

$$\ln \pi_{2t} = \sum_{s=1}^{2t} \ln |\rho_s| + |\ln Z_t| + \mathcal{O}(\delta^2). \quad (55)$$

In this case, the first-order  $\delta$  correction cancels, so the Lyapunov and tail exponents are unchanged at order  $\delta$ :

$$\gamma = \ln \rho + \mathcal{O}(\delta^2), \quad \alpha = -\frac{2\gamma}{\sigma_0^2} + \mathcal{O}(\delta^2). \quad (56)$$

To summarize, when  $\delta$  is small, perturbation theory shows that

(1) for real eigenvalues, non-normality decreases the Lyapunov exponent linearly in the degree  $\delta$  of spectrum degeneracy, slightly pushing the system closer to instability;

(2) for complex conjugate eigenvalues, the first-order correction vanishes, and non-normality affects the exponents only at order  $\delta^2$ .

These results connect smoothly with the exactly solvable case  $\delta = 0$  and illustrate how the influence of non-normality depends sensitively on the spectral structure of the underlying matrices.

#### D. Non-normal reinjection via rotation

So far, we have mainly considered the case in which the rotation matrix  $\mathbf{U}$  in Eq. (38) is fixed. In that setting, non-normality affects the distribution of growth rates, but its impact on the Lyapunov and tail exponents remains limited. In the more general case, however,  $\mathbf{U}_t$  may itself be stochastic and time dependent. This introduces an additional mechanism: Rotations can continually *reinject the dynamics* into the

direction of maximal amplification. As a result, non-normality can significantly enhance trajectory growth and modify both the Lyapunov exponent and the tail exponent.

This provides the missing dynamical ingredient absent from the previous example. Rotation does not merely act as a change of coordinates: In a non-normal system, it can repeatedly redirect the state into reactive directions, thereby converting transient amplification into cumulative growth. Through this mechanism, non-normality begins to influence the Lyapunov exponent itself, and not only the tail exponent.

To illustrate this phenomenon, let us focus on the simplest case of real eigenvalues with  $\delta = 0$ . The matrices then take the form

$$\begin{aligned} \mathbf{A}_t &= \rho_t \mathbf{U}(\theta_t) \begin{pmatrix} 0 & z_t \\ z_t^{-1} & 0 \end{pmatrix} \mathbf{U}(\theta_t)^\dagger, \\ \mathbf{U}(\theta) &= \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}. \end{aligned} \quad (57)$$

A useful identity is

$$\mathbf{U}(\theta)^\dagger \mathbf{U}(\phi) = \cos(\theta - \phi) \mathbf{I} + \sin(\theta - \phi) \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}, \quad (58)$$

which will allow us to keep track of how rotations couple across time.

The product of two consecutive matrices

$$\mathbf{A}_t^{(2)} = \mathbf{A}_{2t} \mathbf{A}_{2t-1} \quad (59)$$

can be obtained explicitly using the identity (58) to obtain

$$\begin{aligned} \mathbf{A}_t^{(2)} &= \rho_{2t} \rho_{2t-1} \mathbf{U}(\theta_{2t}) [\cos(\theta_{2t} - \theta_{2t-1}) \mathbf{D}_t^{(2)} \\ &\quad + \sin(\theta_{2t} - \theta_{2t-1}) \mathbf{K}_t^{(2)}] \mathbf{U}(\theta_{2t-1})^\dagger, \end{aligned} \quad (60)$$

$$\text{where } \mathbf{D}_t^{(2)} = \begin{pmatrix} z_t^{(2)} & 0 \\ 0 & (z_t^{(2)})^{-1} \end{pmatrix}, \quad z_t^{(2)} = \frac{z_{2t}}{z_{2t-1}}, \quad (61)$$

$$\text{and } \mathbf{K}_t^{(2)} = \begin{pmatrix} 0 & -k^{(2)} \\ 1/k^{(2)} & 0 \end{pmatrix}, \quad k^{(2)} = z_{2t} z_{2t-1}. \quad (62)$$

Unlike the purely diagonal case studied before, the product  $\mathbf{A}_t^{(2)}$  is now a combination of a diagonal contribution  $\mathbf{D}_t^{(2)}$  (which tends to average out non-normality) and an anti-diagonal contribution  $\mathbf{K}_t^{(2)}$  (which can reinforce it). This constructive “reinjection” is the key mechanism.

The impact of the anti-diagonal contribution  $\mathbf{K}_t^{(2)}$  can be quantified by noting the following. Let us consider a matrix

$$\begin{aligned} \mathbf{A}_t &= \rho_t \mathbf{U}(\theta_t) \mathbf{B}(z_t) \mathbf{U}(\theta_t)^\dagger, \\ \mathbf{B}(z) &:= \begin{pmatrix} 0 & z \\ z^{-1} & 0 \end{pmatrix}, \quad \mathbf{U}(\theta) = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}. \end{aligned} \quad (63)$$

Consider a two-step block  $\mathbf{A}_s^{(2)} := \mathbf{A}_{2s} \mathbf{A}_{2s-1}$ . Using unitarity,

$$\mathbf{A}_s^{(2)} = \rho_{2s} \rho_{2s-1} \mathbf{U}(\theta_{2s}) [\mathbf{B}(z_{2s}) \mathbf{R}_{2s,2s-1} \mathbf{B}(z_{2s-1})] \mathbf{U}(\theta_{2s-1})^\dagger, \quad (64)$$

where

$$\begin{aligned} \mathbf{R}_{2s,2s-1} &= \mathbf{U}(\theta_{2s})^\dagger \mathbf{U}(\theta_{2s-1}) \\ &= \cos(\theta_{2s} - \theta_{2s-1}) \mathbf{I} + \sin(\theta_{2s} - \theta_{2s-1}) \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}, \end{aligned} \quad (65)$$

This allows us to obtain the two identities (checked by direct multiplication):

$$\begin{aligned} \mathbf{D}^{(2)}(z, w) &:= \mathbf{B}(z) \mathbf{I} \mathbf{B}(w) = \begin{pmatrix} z/w & 0 \\ 0 & (z/w)^{-1} \end{pmatrix}, \\ \mathbf{K}^{(2)}(z, w) &:= \mathbf{B}(z) \mathbf{Q} \mathbf{B}(w) = \begin{pmatrix} 0 & -zw \\ (zw)^{-1} & 0 \end{pmatrix}. \end{aligned} \quad (66)$$

Hence,

$$\begin{aligned} \mathbf{B}(z_{2s}) \mathbf{R}_{2s,2s-1} \mathbf{B}(z_{2s-1}) &= \cos(\theta_{2s} - \theta_{2s-1}) \mathbf{D}_s^{(2)} \\ &\quad + \sin(\theta_{2s} - \theta_{2s-1}) \mathbf{K}_s^{(2)}. \end{aligned} \quad (67)$$

Operator norms are invariant under unitary rotations before or after a matrix, so

$$\begin{aligned} \|\mathbf{A}_s^{(2)}\| &= |\rho_{2s} \rho_{2s-1}| \left\| \cos(\theta_{2s} - \theta_{2s-1}) \mathbf{D}_s^{(2)} \right. \\ &\quad \left. + \sin(\theta_{2s} - \theta_{2s-1}) \mathbf{K}_s^{(2)} \right\|. \end{aligned} \quad (68)$$

In the non-normal (antidiagonal)-dominated regime, we keep the  $\mathbf{K}^{(2)}$  term,

$$\begin{aligned} \|\mathbf{A}_s^{(2)}\| &\approx |\rho_{2s} \rho_{2s-1}| |\sin(\theta_{2s} - \theta_{2s-1})| \|\mathbf{K}_s^{(2)}\| \\ &= |\rho_{2s} \rho_{2s-1}| |\sin(\theta_{2s} - \theta_{2s-1})| \max(k_s^{(2)}, (k_s^{(2)})^{-1}). \end{aligned} \quad (69)$$

Multiplying the  $t$  two-step blocks and using submultiplicativity,

$$\begin{aligned} \pi_{2t} &= \left\| \prod_{s=1}^{2t} A_s \right\| = \left\| \prod_{s=1}^t A_s^{(2)} \right\| \\ &\approx \left( \prod_{s=1}^{2t} |\rho_s| \right) \left( \prod_{s=1}^t |\sin(\theta_{2s} - \theta_{2s-1})| \right) \\ &\quad \times \max(K_t, K_t^{-1}), \end{aligned}$$

$$\text{with } K_t := \prod_{s=1}^t k_s^{(2)} = \prod_{s=1}^{2t} z_s. \quad (70)$$

Taking logarithms and writing  $\ln \kappa_s := |\ln z_s|$  so that  $\sum_{s=1}^t \ln \max(k_s^{(2)}, (k_s^{(2)})^{-1}) = \sum_{s=1}^{2t} \ln \kappa_s$ , we obtain

$$\begin{aligned} \ln \pi_{2t} &\approx \sum_{s=1}^{2t} \ln |\rho_s| + \sum_{s=1}^t \ln |\sin(\theta_{2s} - \theta_{2s-1})| \\ &\quad + \sum_{s=1}^{2t} \ln \kappa_s. \end{aligned} \quad (71)$$

By the CLT, assuming all variances exist, this behaves as

$$\ln \pi_{2t} \sim \mathcal{N}(2t(\ln \rho + \ln \kappa - \mu_\theta), 2t\sigma_0^2), \quad (72)$$

with

$$\sum_{s=1}^t \ln |\sin(\theta_{2s} - \theta_{2s-1})| \sim \mathcal{N}(-2t\mu_\theta, t\sigma_\theta^2), \quad (73)$$

and where  $\sigma_0^2$  accounts for the variance of all the processes and their covariances, as in the previous sections.

We thus obtain the following approximations for the exponents:

$$\gamma \approx \ln \rho + \ln \kappa - \mu_\theta, \quad \alpha \approx -\frac{2\gamma}{\sigma_0^2}. \quad (74)$$

This shows explicitly how random rotations can strongly amplify growth: If  $\mu_\theta = \sigma_\theta = 0$ , then at each step the rotation reinjects the system perfectly along the most expanding direction [corresponding to  $|\sin(\theta_{2s} - \theta_{2s-1})| = 1$ ]. In this idealized situation, the Lyapunov exponent nearly attains its upper bound and the tail exponent its lower bound. In practice, stochastic fluctuations in  $\theta_t$  reduce this effect, but the mechanism highlights how non-normality can push the system much closer to instability even when it is spectrally subcritical.

### E. Numerical application

To illustrate the theoretical results derived in the previous sections, we now turn to a numerical study of the two-dimensional Kesten process. For this purpose, we use the generic random matrix (41)

$$\mathbf{A}_t = \rho_t \mathbf{U}(\theta_t) \begin{pmatrix} \delta & (1-\delta)z_t \\ (1-\delta)z_t^{-1} & \delta \end{pmatrix} \mathbf{U}(\theta_t)^\dagger, \quad (75)$$

where  $\rho_t$  controls the spectral radius,  $z_t$  characterizes non-normal variations,  $\delta$  quantifies spectral imbalance, and  $\mathbf{U}(\theta_t)$  is a rotation matrix. Throughout this section, we assume that  $\delta$  is constant,  $\rho_t$  and  $z_t$  are log-normally distributed [ $\ln \rho_t \sim \mathcal{N}(\rho, \sigma_\rho^2)$  and  $\ln z_t \sim \mathcal{N}(0, \sigma_z^2)$ ], and  $\theta_t$  is either uniformly distributed on  $[0, 2\pi]$  or fixed to zero. Therefore, in order to disentangle the respective roles of the spectral radius, non-normal shear, and rotation, we treat these different degrees of freedom as independent of one another. The validity and implications of this independence assumption are nontrivial, and are discussed in greater detail when extending the results to the general  $N$ -dimensional case.

Let us define the (log-)condition number by  $\ln \kappa_t := |\ln z_t|$ . Let  $X := \ln z_t \sim \mathcal{N}(0, \sigma_z^2)$  and  $Y := |X|$  (a half-normal variable). The pdf of  $Y$  is

$$f_Y(y) = \frac{2}{\sqrt{2\pi} \sigma_z} e^{-y^2/(2\sigma_z^2)}, \quad y \geq 0. \quad (76)$$

*Mean.* Using  $E[Y] = \int_0^\infty y f_Y(y) dy$ ,

$$\begin{aligned} E[Y] &= \frac{2}{\sqrt{2\pi} \sigma_z} \int_0^\infty y e^{-y^2/(2\sigma_z^2)} dy \\ &= \frac{2}{\sqrt{2\pi} \sigma_z} [-\sigma_z^2 e^{-y^2/(2\sigma_z^2)}]_0^\infty \\ &= \sigma_z \sqrt{\frac{2}{\pi}}. \end{aligned} \quad (77)$$

*Second moment.* Since  $Y^2 = X^2$  and  $E[X^2] = \sigma_z^2$  for a centered normal,  $E[Y^2] = E[X^2] = \sigma_z^2$ .




FIG. 1. Simulation of the two-dimensional Kesten process (1), with matrices  $A_t$  defined by Eq. (75). Parameters are set to  $\delta = 0$  and  $\ln \rho = -1$  (hence,  $\sigma_\rho = 0$ ). Left column: no rotation ( $\theta_t = 0$ ). Middle column: uniform rotation [ $\theta_t \sim \mathcal{U}(0, 2\pi)$ ]. Each row corresponds to a different level of non-normal variability, from top to bottom:  $\sigma_z = 0.5, 1, 1.5$ . The y axis reports the crossing mean  $\langle x \rangle$  between the two components of the system defined by expression (83). The right panel shows the complementary cumulative distribution function (CCDF) of the absolute mean, normalized so that the fifth-largest observation equals 1. Solid lines correspond to the no-rotation case, while dashed lines correspond to uniform rotations. The numerical procedures used to estimate the Lyapunov and tail exponents are detailed in the Appendix.

Variance. Hence,

$$\begin{aligned} \text{Var}(Y) &= E[Y^2] - (E[Y])^2 = \sigma_z^2 - \sigma_z^2 \frac{2}{\pi} \\ &= \left(1 - \frac{2}{\pi}\right) \sigma_z^2. \end{aligned} \quad (78)$$

Identifying  $Y = \ln \kappa_t$ , we obtain

$$\begin{aligned} \ln \kappa &= E[\ln \kappa_t] = \sigma_z \sqrt{\frac{2}{\pi}}, \\ \sigma_\kappa^2 &= \text{Var}(\ln \kappa_t) = \left(1 - \frac{2}{\pi}\right) \sigma_z^2. \end{aligned} \quad (79)$$

For uniformly distributed rotations,  $\theta_t \sim \mathcal{U}(0, 2\pi)$ , the differences  $\theta_t - \theta_{t+1}$  are also i.i.d. uniform on  $[0, 2\pi]$ . This implies

$$\mu_\theta = \ln 2, \quad \sigma_\theta^2 = \frac{\pi^2}{12}. \quad (80)$$

Hence, the system is governed by four hyperparameters:  $\rho$ ,  $\sigma_\rho$ ,  $\sigma_z$ , and  $\delta$ .

We focus on two limiting scenarios:

(1) *Case 1: No rotation.*  $\theta_t = 0$ . In this case, the theoretical predictions reduce to

$$\begin{aligned} \gamma &= \ln \rho - \delta + \mathcal{O}(\delta^2), \\ \alpha &= -\frac{2\gamma}{\sigma_\rho^2 + \sigma_z^2} + \mathcal{O}(\delta^2). \end{aligned} \quad (81)$$

(2) *Case 2: Uniform rotation.*  $\theta_t \sim \mathcal{U}(0, 2\pi)$ . Here, the Lyapunov and tail exponents become

$$\begin{aligned} \gamma &\approx \ln \rho + \sqrt{\frac{2}{\pi}} \sigma_z - \ln 2 - \delta + \mathcal{O}(\delta^2), \\ \alpha &\approx -\frac{2\gamma}{\sigma_\rho^2 + \left(1 - \frac{2}{\pi}\right) \sigma_z^2 + \frac{\pi^2}{12}} + \mathcal{O}(\delta^2). \end{aligned} \quad (82)$$

Figure 1 shows simulations for the case  $\delta = 0$  with  $\ln \rho_t = -1$  fixed and for increasing non-normal variability  $\sigma_z = 0.5, 1, 1.5$ . We measure both the Lyapunov exponent  $\gamma$  and the tail exponent  $\alpha$  (see the Appendix for methodology). We compare the no-rotation and uniform-rotation cases and show in the first two columns the temporal evolution of the crossing mean  $\langle x \rangle$  between the two components of the system. The crossing mean  $\langle x \rangle$  is defined as follows. Let  $x_t = (x_t^{(1)}, x_t^{(2)})^\top$  denote the two components of the two-dimensional Kesten process. Since the dynamics alternates between the two coordinates, a convenient scalar observable is obtained by averaging them at their crossings. We define

$$\langle x_t \rangle = \frac{1}{2} (x_t^{(1)} + x_t^{(2)}), \quad (83)$$

measured at the instants  $t$ , where  $x_t^{(1)}$  and  $x_t^{(2)}$  cross each other (i.e., when  $x_t^{(1)} = x_t^{(2)}$  or  $x_t^{(1)} = -x_t^{(2)}$ ). This “crossing mean” provides a single representative trajectory that summarizes the joint growth of the two components.

For *no rotation*, increasing  $\sigma_z$  decreases the tail exponent while leaving the Lyapunov exponent near  $\gamma = -1$ , as predicted (81). For *uniform rotation*,  $\gamma$  increases while  $\alpha$  decreases with  $\sigma_z$  (82). In particular,  $\gamma$  vanishes at a critical value of  $\sigma_z$  given by  $\sigma_c \approx 2.1$ , indicating the onset of instability.

Figure 2 further explores these trends by plotting  $\gamma$  and  $\alpha$  as functions of  $\sigma_z$  for the case of a constant spectral radius  $\ln \rho_t = -1$ , and  $\ln \rho_t \sim \mathcal{N}(-1, 1/9)$ . Here, the theoretical tail exponent for  $\delta = 0$ , no rotation, and  $\sigma_z = 0$  is  $\alpha = 18$ . The numerical results confirm that, with no rotation, the Lyapunov exponent remains constant, while with uniform rotation  $\gamma$  grows approximately linearly with  $\sigma_z$ , with slope  $\sqrt{2/\pi}$  as predicted. However, a finite-size “transition region” appears for small  $\sigma_z$ , where the linear theory does not apply.




FIG. 2. Study of the two-dimensional Kesten process (1), with matrices  $\mathbf{A}_t$  defined by Eq. (75). The system is simulated with  $\ln \rho_t = -1$  (top) and  $\ln \rho_t \sim \mathcal{N}(-1, 1/9)$  (bottom), while varying the non-normal variance  $\sigma_z$  and the imbalance parameter  $\delta$ , both with uniform rotations (+) and without rotations ( $\cdot$ ). Theoretical predictions without rotations (solid lines) and with rotations (dashed lines) are taken from Eqs. (81) and (82), respectively. We considered the case without any imbalance, i.e.,  $\delta = 0$  (black) and with imbalance, i.e.,  $\delta = 0.1$  (blue). Left panel: Lyapunov exponent  $\gamma$  as a function of  $\sigma_z$ . Right panel: Corresponding tail exponent  $\alpha$  (see the Appendix for details of the numerical estimation). We estimate the tail exponent over ten independent ensembles, report the ensemble average for each data point, and use the corresponding standard deviation as the error bar. We focus on the range  $\alpha \in [0, 7]$ . Missing points correspond to cases where the tail exponent estimate diverges, typically when  $\alpha$  approaches zero and/or when the Lyapunov exponent becomes positive.

For the tail exponent, the measured values follow the theoretical trend but do not exactly match for small non-normality, where the measured  $\alpha$  tends to saturate. This discrepancy is expected: Estimating  $\alpha$  from finite samples is notoriously difficult, especially when  $\alpha$  is large or close to unity, and may yield spurious or divergent results.

Finally, we investigate the effect of nonzero  $\delta$ . We observe that, with uniform rotation, the system behaves as predicted, with a shift given by  $-\delta$ , but in the absence of rotation, we also observed that the system exhibits a growth in the Lyapunov exponent due to the non-normal behavior, revealing an additional destabilizing mechanism occurring at higher order power in  $\delta$ .

### F. Physical meaning and interpretation of rotation

In two-dimensional systems, the rotation matrix  $\mathbf{U}(\theta)$  plays a key role in enabling non-normality to enhance the Lyapunov exponent. This rotation may be stochastic or deterministic, for instance, arising from Perron-type instabilities along a limit cycle [27–29]. Its presence can be made explicit in various dynamical settings.

*Turbulence and polymer stretching.* In the context of turbulence, random rotations arise naturally from the temporal evolution of the velocity-gradient tensor along Lagrangian trajectories. Even in statistically isotropic and homogeneous flows, the strain and vorticity components continuously

rotate the local principal stretching directions. As a result, material elements (or polymer end-to-end vectors) are repeatedly reoriented relative to the instantaneous eigenbasis of the velocity-gradient tensor. This mechanism is well documented in studies of polymer stretching in random and turbulent flows, where the dynamics of the polymer conformation tensor is driven by a sequence of random velocity-gradient matrices whose eigenvectors rotate in time. Such rotations are essential for maintaining a stationary distribution by reinjecting probability mass into directions that can subsequently experience amplification. See, for example, Refs. [30–32].

More broadly, random reorientation of stretching directions is a generic feature of Lagrangian chaos and has long been emphasized in the theory of turbulent mixing and stretching [33,34]. In this setting, the rotation-induced reinjection provides a minimal stochastic representation of these physical processes.

*Financial systems and portfolio dynamics.* In financial contexts, analogous rotation mechanisms arise through changes in the effective basis of risk factors or asset correlations. Time-dependent correlations, portfolio rebalancing, and regime shifts continuously rotate the directions along which multiplicative growth or decay acts. In stochastic portfolio models and generalized multiplicative processes, this leads to repeated redistribution of wealth across components, preventing collapse onto a single dominant direction and enabling stationary heavy-tailed distributions. Such mechanisms are discussed, for instance, in the literature on random growth with redistribution and time-dependent correlations [25,35].

*Interpretation in the present framework.* In our model, random rotations should therefore be understood as an effective representation of these ubiquitous physical mechanisms that continuously reshuffle the state vector relative to the instantaneous eigenbasis of the multiplicative dynamics. Their role is not to generate amplification by themselves, but to ensure ergodicity and reinjection, allowing rare but strong amplification events driven by non-normal eigenvector geometry to contribute to stationary heavy-tailed statistics.

### G. Synthesis

In this section, we have explored the two-dimensional Kesten process as a tractable yet sufficiently rich setting to understand the interplay between spectral properties, non-normality, and critical behavior. Starting from an exactly solvable case, we showed that the Lyapunov exponent depends only on the spectral radius, while the tail exponent is rescaled by the variance of non-normal fluctuations. Through perturbative expansions, we extended these results to strong non-normal matrices and demonstrated that rotation dynamics can either suppress or reinforce instabilities by reinjecting trajectories into expanding directions. Finally, numerical simulations confirmed the theoretical predictions, while also highlighting subtle finite-size effects and nonlinear corrections when spectral imbalance is present.

The two-dimensional case thus provides a clear picture: non-normality does not necessarily alter the stability criterion (through the Lyapunov exponent), but it can significantly reduce the tail exponent, effectively pushing the system into heavy-tailed regimes. When combined with stochastic

rotations and/or spectral imbalance (via the parameter  $\delta$ ), non-normality can even shift the system toward true criticality, despite being spectrally subcritical.

In the next section, we move beyond the two-dimensional setting and investigate the generalization to  $N$ -dimensional Kesten processes. This will allow us to understand how non-normality scales with dimension, how collective effects emerge, and how the interplay of Lyapunov and tail exponents shapes the onset of critical phenomena in high-dimensional stochastic systems.

## IV. GENERALIZED $N$ -DIMENSIONAL CASE

We now extend the analysis to the case of large dimension  $N \gg 1$ . The aim is to derive approximate expressions for the Lyapunov exponent and the tail exponent, and to identify how non-normality interacts with dimensionality to push the system closer to criticality. To do so, we build on the spectral and singular value decompositions of the random matrices  $\mathbf{A}_t$  and apply tools from random matrix theory and extreme value theory (EVT).

We begin with the decomposition

$$\mathbf{A}_t = \mathbf{U}_t \boldsymbol{\Sigma}_t \mathbf{V}_t^\dagger \mathbf{A}_t \mathbf{V}_t \boldsymbol{\Sigma}_t^{-1} \mathbf{U}_t^\dagger, \quad (84)$$

where

$\boldsymbol{\Lambda}_t = \text{Diag}(\lambda_{1,t}, \dots, \lambda_{N,t})$  contains the eigenvalues,  $\mathbf{P}_t = \mathbf{U}_t \boldsymbol{\Sigma}_t \mathbf{V}_t^\dagger$  is the eigenbasis transformation matrix,  $\boldsymbol{\Sigma}_t = \text{Diag}(s_{1,t}, \dots, s_{N,t})$  the singular values of  $\mathbf{P}_t$ ,  $\mathbf{U}_t$  and  $\mathbf{V}_t$  are unitary matrices.

We assume that all second moments exist, so we can apply the CLT.

### A. Dominant approximation via condition number

This subsection develops, in full detail, a compact and constructive approximation. The key idea is that the dominant contribution of each matrix in the product can be captured through the leading singular directions associated with the eigenbasis transformation. This representation makes explicit how both the condition number and the reinjection geometry enter into the growth of the product norm. In this way, the approximation provides a concrete underpinning for the earlier bounds and shows how the increase in the Lyapunov exponent arises dynamically, rather than merely as a consequence of an inequality.

The general  $N$ -dimensional setting further clarifies that large non-normality alone is not sufficient. What matters is whether the most amplifying singular directions at successive times can couple constructively. If these directions remain dynamically misaligned, the amplification does not accumulate. In contrast, when rotation and mode switching repeatedly align the output of one step with the reactive direction of the next, non-normality becomes an effective mechanism for enhanced Lyapunov growth and heavier tails.

Starting from the singular value decomposition (SVD)-based decomposition (84), we write the singular-vector expansions

$$\boldsymbol{\Sigma}_t = \sum_{i=1}^N s_{i,t} \hat{\mathbf{u}}_{i,t} \hat{\mathbf{v}}_{i,t}^\dagger, \quad \boldsymbol{\Sigma}_t^{-1} = \sum_{i=1}^N s_{i,t}^{-1} \hat{\mathbf{v}}_{i,t} \hat{\mathbf{u}}_{i,t}^\dagger, \quad (85)$$

so that

$$\begin{aligned}\Sigma_t \mathbf{V}_t^\dagger \Lambda_t \mathbf{V}_t \Sigma_t^{-1} &= \sum_{i,j=1}^N \frac{s_{i,t}}{s_{j,t}} \hat{\mathbf{u}}_{i,t} (\hat{\mathbf{v}}_{i,t}^\dagger \mathbf{V}_t^\dagger \Lambda_t \mathbf{V}_t \hat{\mathbf{v}}_{j,t}) \hat{\mathbf{u}}_{j,t}^\dagger \\ &= \sum_{i,j=1}^N \frac{s_{i,t}}{s_{j,t}} \lambda_{ij,t} \hat{\mathbf{u}}_{i,t} \hat{\mathbf{u}}_{j,t}^\dagger,\end{aligned}\quad (86)$$

where we denote

$$\lambda_{ij,t} := (\mathbf{V}_t^\dagger \Lambda_t \mathbf{V}_t)_{ij}. \quad (87)$$

Let  $s_{\max,t}$  and  $s_{\min,t}$  be the extreme (largest and smallest) singular values of  $\Sigma_t$ , with corresponding singular vectors  $\hat{\mathbf{u}}_{\max,t}$ ,  $\hat{\mathbf{u}}_{\min,t}$ . Then, the condition number  $\kappa_t = \|\mathbf{P}_t\| \|\mathbf{P}_t^{-1}\|$  (21) is given by  $\kappa_t := s_{\max,t}/s_{\min,t}$ . Then, in the double sum above, the coefficient  $s_{i,t}/s_{j,t}$  is maximized at  $(i, j) = (\max, \min)$  and equals  $\kappa_t$ , while for any other pair  $(i, j)$  one has  $s_{i,t}/s_{j,t} \leq \kappa_t$  with strict inequality unless  $(i, j) = (\max, \min)$ . If we moreover assume independence and i.i.d. in time of  $\{\lambda_{ij,t}\}$  and that these entries do not concentrate pathologically on subdominant  $(i, j)$  pairs, the rank-one term with  $(i, j) = (\max, \min)$  dominates

$$\Sigma_t \mathbf{V}_t^\dagger \Lambda_t \mathbf{V}_t \Sigma_t^{-1} \approx \kappa_t \lambda_{\max,\min,t} \hat{\mathbf{u}}_{\max,t} \hat{\mathbf{u}}_{\min,t}^\dagger. \quad (88)$$

Left and right multiplication by the unitary  $\mathbf{U}_t$  leaves the operator norm unchanged, so inserting back into  $\mathbf{A}_t$  yields the dominant-direction (rank-one) approximation

$$\mathbf{A}_t \approx \kappa_t \lambda_t \hat{\mathbf{u}}_{\max,t} \hat{\mathbf{u}}_{\min,t}^\dagger, \quad \lambda_t := (\mathbf{V}_t \Lambda_t \mathbf{V}_t^\dagger)_{\max,\min}. \quad (89)$$

The approximation (89) is valid in the *ill-conditioned, large- $N$  (rank-one-dominance)* regime:

(1)  $\kappa_t \gg 1$  so that the ratio  $s_{i,t}/s_{j,t}$  is sharply peaked at  $(i, j) = (\max, \min)$ ;

(2) the spectral mixing  $\mathbf{V}_t^\dagger \Lambda_t \mathbf{V}_t$  has entries  $\lambda_{ij,t}$  with bounded moments and no anomalous amplification on subdominant pairs  $(i, j) \neq (\max, \min)$ .

Under these conditions (typical at large  $N$ , where extreme singular values separate), the outer-product term  $\hat{\mathbf{u}}_{\max,t} \hat{\mathbf{u}}_{\min,t}^\dagger$  provides the leading contribution to  $\mathbf{A}_t$ , and Eq. (89) controls the growth in the product  $\Pi_t$  used later.

Under this approximation, the product of matrices becomes

$$\Pi_t \approx \left( \prod_{s=1}^t \kappa_s \lambda_s \right) \left( \prod_{s=1}^{t-1} \hat{\mathbf{u}}_{\min,s} \cdot \hat{\mathbf{u}}_{\max,s+1} \right) \hat{\mathbf{u}}_{\max,1} \hat{\mathbf{u}}_{\min,t}^\dagger, \quad (90)$$

so that the product norm satisfies

$$\ln \pi_t \approx X_t, \quad (91)$$

$$X_t = \sum_{s=1}^t \ln |\lambda_s| + \sum_{s=1}^t \ln \kappa_s + \sum_{s=1}^{t-1} \ln |\hat{\mathbf{u}}_{\min,s} \cdot \hat{\mathbf{u}}_{\max,s+1}|. \quad (92)$$

By the CLT, in the limit  $t \rightarrow \infty$ , the distribution of  $X_t$  is approximately Gaussian,

$$\begin{aligned}X_t &\sim \mathcal{N}(t\gamma_0, t\sigma_0^2), \\ \gamma_0 &= \ln \rho + \ln \kappa - \mu_\theta, \\ \sigma_0^2 &= \sigma_\rho^2 + \sigma_\kappa^2 + \sigma_\theta^2 + \text{correlation terms}, \\ \ln \rho &:= \mathbb{E}[\ln |\lambda_1|], \quad \sigma_\rho^2 := \text{Var}[\ln |\lambda_1|],\end{aligned}$$

$$\ln \kappa := \mathbb{E}[\ln \kappa_1], \quad \sigma_\kappa^2 := \text{Var}[\ln \kappa_1],$$

$$\mu_\theta := -\mathbb{E}[\ln |\hat{\mathbf{u}}_{\min,1} \cdot \hat{\mathbf{u}}_{\max,2}|],$$

$$\sigma_\theta^2 := \text{Var}[\ln |\hat{\mathbf{u}}_{\min,1} \cdot \hat{\mathbf{u}}_{\max,2}|], \quad (93)$$

where the term  $\mu_\theta$  makes transparent what the norm bound alone cannot say. The effective growth is controlled not only by the instantaneous degree of non-normality, but also by the ability of the dynamics to *reinject* amplification from one step to the next. In this sense,  $\mu_\theta$  provides the missing physical interpretation of when a large condition number actually matters.

Therefore, the Lyapunov and tail exponents can be approximated by

$$\gamma \approx \gamma_0 \quad \text{and} \quad \alpha \approx \alpha_0 = -2 \frac{\gamma_0}{\sigma_0^2}. \quad (94)$$

Thus, the Lyapunov exponent depends on three contributions: the eigenvalue statistics, the condition number of  $\mathbf{P}_t$  and its statistics, and the reinjection term due to random orientations.

### B. Eigenvalue-eigenvector independence assumption

In the previous section, we derived an approximation for the Lyapunov and tail exponents, where the only additional assumption is the existence of the second moments of all the processes considered. In this section, we present generic results based on the Ginibre ensemble, which will later support qualitative arguments justifying the assumption that the eigenvectors and eigenvalues of the system can be treated as independent when the spectral radius remains well below unity (i.e., far from the critical boundary).

Let  $\mathbf{A} \in \mathbb{R}^{N \times N}$  denote a (real) Ginibre matrix [36] with i.i.d. entries  $(\mathbf{A})_{ij} \sim \mathcal{N}(0, 1/\sqrt{N})$ , so that the empirical spectral radius is  $O(1)$  and the eigenvalue cloud obeys the circular law (with its support being approximately the unit disk). We denote the eigenvalue-eigenvector decomposition (when diagonalizable)  $\mathbf{A} = \mathbf{P} \Lambda \mathbf{P}^{-1}$ ,  $\Lambda = \text{Diag}(\lambda_1, \dots, \lambda_N)$ , with (right) eigenvectors  $\mathbf{r}_i$  (columns of  $\mathbf{P}$ ) and left eigenvectors  $\ell_i^\dagger$  (rows of  $\mathbf{P}^{-1}$ ), normalized so that  $\ell_i^\dagger \mathbf{r}_i = 1$ . The local (or eigenpair) condition/overlap is defined as

$$\mathcal{O}_i = \|\mathbf{r}_i\| \|\ell_i\|, \quad (95)$$

and the global eigenbasis condition number is

$$\kappa(P_N) = \|\mathbf{P}_N\| \|\mathbf{P}_N^{-1}\| \gtrsim \max_{1 \leq i \leq N} \mathcal{O}_i. \quad (96)$$

Below, we summarize the key statistical properties of the ensemble that are essential for the subsequent analysis. In particular, the spectral and overlap properties are standard in the random matrix literature on non-Hermitian matrices [37,38].

*Circular law and bulk regime.* As  $N \rightarrow \infty$ , the eigenvalues fill the unit disk (circular law). Throughout this work, we consider eigenvalues conditioned to lie well inside the spectral boundary, i.e.,  $|\lambda_i| < 1 - \varepsilon$  for some fixed  $\varepsilon > 0$ . This is the “far from spectral criticality” regime: Spectral radii are bounded away from unity and the leading eigenvalue magnitudes do not approach instability thresholds.

*Typical eigenvector nonorthogonality (overlaps).* For the complex Ginibre ensemble, the diagonal overlap (for an

eigenvalue conditioned at  $z$  in the bulk) has expectation of order

$$\mathbb{E}[\mathcal{O}_i | \lambda_i = z] \asymp \frac{N}{1 - |z|^2} \quad (|z| < 1). \quad (97)$$

Thus, overlaps scale linearly with  $N$  (and blow up as  $|z| \rightarrow 1$ ). This is the Chalker-Mehlig scaling, later made rigorous by subsequent authors; it is the principal reason non-normal effects grow with system size in the bulk. Intuitively, the factor  $1/(1 - |z|^2)$  encodes the reduced local spectral stability near the boundary, while the factor  $N$  is a combinatorial/density effect of the many degrees of freedom [37,38].

*Heavy tails and rare large overlaps.* Overlaps have heavy-tailed fluctuations: The distribution of  $\mathcal{O}_i$  (properly rescaled) admits heavy tails/inverse-Gamma type limits in the bulk. Consequently, the maximum overlap [hence  $\kappa(\mathbf{P})$ ] is dominated by rare large events (near-colliding eigenvalues, tight local clusters), not just by the typical value above. This heavy-tailed character is important when computing extreme rare-event contributions to Lyapunov/tail exponents.

*Local spacing dependence.* When an eigenvalue  $\lambda_i$  has a nearby eigenvalue at distance  $\Delta$ , the corresponding overlaps can scale like  $(N\Delta^2)^{-1}$ : Small nearest-neighbor gaps produce large overlaps. In practice, this means that configurations with small eigenvalue gaps (even if rare) strongly increase  $\kappa(\mathbf{P})$ . Thus, the geometry of the eigenvalue cloud (positions and spacings) directly controls eigenvector conditioning.

In several derivations of Kesten dynamics throughout this paper, we assume that the statistical properties of the eigenbasis (singular values and eigenvectors) are approximately independent of those of the eigenvalue set. This approximation is well controlled when the system is *far from criticality*, that is, when all eigenvalues satisfy  $|\lambda_i| < 1 - \varepsilon$  uniformly and no systematically small eigenvalue gaps occur. This assumption can be justified on several theoretical and probabilistic grounds, supported by heuristic reasoning.

*Biunitary invariance.* For the complex Ginibre ensemble, the distribution of matrices is exactly biunitarily invariant for any  $N$ , meaning that  $A \stackrel{d}{=} UAV$  for all  $U, V \in \text{U}(N)$ . This invariance ensures that the singular vectors are uniformly distributed and statistically independent of the singular values. However, this bi-invariance does *not* imply that the eigenvectors and eigenvalues are independent: At finite  $N$ , their joint density involves correlations through the nonorthogonality matrix  $G$  [37]. In the large- $N$  limit, these correlations vanish in the bulk of the spectrum, so that eigenvectors and eigenvalues become asymptotically independent. For non-Gaussian i.i.d. ensembles, the same asymptotic factorization holds by universality, but independence is no longer exact at finite  $N$ .

*Bulk decorrelation.* For Ginibre matrices, the eigenvectors (left/right) are, at leading order and in distribution, approximately isotropic objects whose directional statistics are weakly dependent on the macroscopic eigenvalue location provided the location is in the bulk (i.e., at distance  $> \varepsilon$  from the spectral edge). Thus, conditioning on  $\lambda_i = z$  with  $|z| < 1 - \varepsilon$  modifies eigenvector statistics only mildly at  $O(1)$  level, while the typical overlaps scale like  $N$ , which is a separation of scales that allows approximate factorization in many averaged calculations.

*Absence of near resonances.* The main failure mode of independence is clustering/near collisions of eigenvalues: If two eigenvalues are atypically close, eigenvectors associated with them are strongly coupled and independence breaks down. By conditioning away from small spacings, i.e., assuming a minimal spacing  $\Delta_{\min}$  not smaller than  $O(N^{-1/2+\eta})$  for some  $\eta > 0$ , these rare pathological configurations are excluded and the independence approximation becomes valid for the dominant contributions that we study.

*Controlled error in Kesten averages.* In the Kesten-product formulas, the eigenvalue contribution enters multiplicatively via  $\prod_t |\lambda_{i,t}|$  while non-normal amplification enters via multiplicative factors related to  $\kappa(P_t)$ . When eigenvalues stay uniformly subcritical ( $|\lambda| < 1 - \varepsilon$ ), the product of spectral moduli decays and the leading corrections coming from eigenvector-eigenvalue correlations manifest as finite  $O(1)$  modifications to Lyapunov or tail exponents; they do not change the scaling exponents unless the system approaches spectral criticality (where small changes can flip the sign of the Lyapunov exponent). This is precisely why assuming independence is safe in the far from criticality regime used in our main results.

The independence assumption provides accurate leading-order predictions for the Lyapunov and tail exponents, as well as for the scaling of non-normal amplification, while rare violations can be treated separately as extreme events.

We presented some generic results associated with the Ginibre ensemble, allowing us to argue later in the paper that, when the matrix spectral radius is far from unity, the eigenbasis and eigenvalues of the random matrices may be treated as approximately independent.

### C. Scaling of the contribution with dimension $N$

For each of the parameters in Eq. (93), we can characterize how it scales with the dimension  $N$  of the system by using the law of large numbers, the central limit theorem, and EVT, and considering that the rotations, non-normal shears, and spectrum are independent. We now discuss each contribution in turn.

Let us define the average spectral contribution

$$\ln \rho := \mathbb{E}[\ln |\lambda_{i,t}|], \quad (98)$$

and decompose each eigenvalue as

$$\lambda_{i,t} = e^{\ln \rho} \lambda_{i,t}^0, \quad \mathbb{E}[\ln |\lambda_{i,t}^0|] = 0. \quad (99)$$

Hence,

$$\ln |\lambda_t| = \ln \rho + \ln \left| \sum_k \lambda_{k,t}^0 \hat{\mathbf{v}}_{ik} \hat{\mathbf{v}}_{jk}^* \right|, \quad (100)$$

where  $(i, j) = (i_{\max}, i_{\min})$  are the indices associated with the largest and smallest singular values. To obtain Eq. (100), we have used  $\lambda_t := (\mathbf{V}_t \mathbf{A}_t \mathbf{V}_t^\dagger)_{\max, \min}$  as defined in Eq. (89).

The coefficients  $\hat{\mathbf{v}}_{ik}$  and  $\hat{\mathbf{v}}_{jk}$  are entries of two independent rows. The vectors  $\hat{\mathbf{v}}_i$  and  $\hat{\mathbf{v}}_j$  can be treated as independent isotropic random vectors on the unit sphere, so their overlaps have mean zero and variance  $O(1/N)$ . Because  $\lambda_{k,t}^0$  are i.i.d.

with finite variance, the central limit theorem implies that

$$\sum_k \lambda_{k,t}^0 \hat{\mathbf{v}}_{ik} \hat{\mathbf{v}}_{jk}^* \xrightarrow{d} \mathcal{N}\left(0, \frac{\sigma_0^2}{N}\right), \quad N \gg 1, \quad (101)$$

where  $\sigma_0$  is the standard deviation of the fluctuations of the normalized eigenvalues  $\lambda_{i,t}^0$  introduced in Eq. (99). Intuitively, the random overlaps  $\hat{\mathbf{v}}_{ik} \hat{\mathbf{v}}_{jk}^*$  behave like  $O(1/\sqrt{N})$  Gaussian coefficients, so the whole sum is a Gaussian random variable of amplitude  $\sim 1/\sqrt{N}$ .

If  $X \sim \mathcal{N}(0, \sigma_0^2/N)$ , then  $|X|$  has a folded normal distribution. The statistics of  $\ln |X|$  are well known:

$$\mathbb{E}[\ln |X|] = \ln \sigma_0 - \frac{1}{2} \ln N - \frac{g + \ln 2}{2}, \quad \text{Var}[\ln |X|] = \frac{\pi^2}{8}, \quad (102)$$

where  $g$  is Euler's constant. This follows from the fact that  $X/\sigma_0\sqrt{N}$  is standard normal, and the distribution of  $\ln |Z|$  for  $Z \sim \mathcal{N}(0, 1)$  is explicitly known.

Combining these results, we obtain

$$\ln \lambda = \ln \rho + \ln \sigma_0 - \frac{g + \ln 2}{2} - \frac{1}{2} \ln N, \quad \sigma_\lambda^2 = \frac{\pi^2}{8}. \quad (103)$$

Thus, the eigenvalue contribution to the Lyapunov exponent decreases as  $-(1/2) \ln N$  in high dimensions, with Gaussian fluctuations of universal variance  $\pi^2/8$ .

We now analyze how the condition number

$$\kappa_t = \frac{s_{\max,t}}{s_{\min,t}} \quad (104)$$

scales with the dimension  $N$ . Taking logarithms, we have

$$\ln \kappa_t = \max_i \ln s_{i,t} - \min_i \ln s_{i,t}. \quad (105)$$

We assume that the singular values are log-normally distributed, i.e.,

$$\ln s_{i,t} \sim \mathcal{N}(0, \sigma^2), \quad i = 1, \dots, N \quad (106)$$

independently across  $i$ . Hence,  $\ln s_{i,t}$  are  $N$  i.i.d. Gaussian random variables. From EVT, the maximum of  $N$  i.i.d.  $\mathcal{N}(0, \sigma^2)$  variables satisfies

$$\max_i \ln s_{i,t} \approx \sigma \sqrt{2 \ln N} + \frac{\sigma}{\sqrt{2 \ln N}} G^+, \quad (107)$$

where  $G^+$  converges in distribution to a standard Gumbel variable. Similarly, the minimum satisfies

$$\min_i \ln s_{i,t} \approx -\sigma \sqrt{2 \ln N} + \frac{\sigma}{\sqrt{2 \ln N}} G^-, \quad (108)$$

with  $G^-$  an independent Gumbel variable. Taking the difference, we obtain

$$\begin{aligned} \ln \kappa_t &= \max_i \ln s_{i,t} - \min_i \ln s_{i,t} \\ &\approx 2\sigma \sqrt{2 \ln N} + \frac{\sigma}{\sqrt{2 \ln N}} (G^+ - G^-). \end{aligned} \quad (109)$$

Thus, the leading-order scaling of  $\ln \kappa_t$  grows like  $\sqrt{\ln N}$ . The large  $N$  dependence of  $\ln \kappa_t$  is therefore

$$\ln \kappa \approx 2\sigma \sqrt{2 \ln N}, \quad (110)$$

while the variance is controlled by the fluctuations of the Gumbel distribution. Since  $\text{Var}(G^+ - G^-) = \pi^2/3$ , we find

$$\sigma_\kappa^2 \approx \frac{\pi^2 \sigma^2}{6 \ln N}. \quad (111)$$

This shows that the condition number diverges slowly as the dimension  $N$  increases: The mean of  $\ln \kappa_t$  grows like  $\sqrt{\ln N}$ , while its fluctuations shrink like  $1/\sqrt{\ln N}$ . Hence, for large  $N$ , the instability induced by non-normality (measured through  $\kappa_t$ ) becomes progressively more pronounced but also more concentrated around its typical value.

The last contribution in Eq. (93) comes from the scalar products  $\hat{\mathbf{u}}_{\min,t} \cdot \hat{\mathbf{u}}_{\max,t+1}$  between eigenvectors associated with the most contracting and expanding directions at successive times. Assuming isotropy of the eigenbasis, both  $\hat{\mathbf{u}}_{\max,t}$  and  $\hat{\mathbf{u}}_{\min,t}$  can be treated as independent and uniformly distributed on the unit sphere in  $\mathbb{R}^N$ .

For two independent isotropic unit vectors in  $\mathbb{R}^N$ , the dot product satisfies

$$\sqrt{N} \hat{\mathbf{u}}_{\min,t} \cdot \hat{\mathbf{u}}_{\max,t} \xrightarrow{d} u_t \sim \mathcal{N}(0, 1), \quad N \gg 1. \quad (112)$$

Thus, the overlap is typically of order  $1/\sqrt{N}$ . Taking logarithms, we obtain

$$\ln |\hat{\mathbf{u}}_{\min,t} \cdot \hat{\mathbf{u}}_{\max,t}| = -\frac{1}{2} \ln N + \ln |u_t|, \quad (113)$$

where  $u_t$  is standard Gaussian. Therefore, the reinjection term is a deterministic shift  $-\frac{1}{2} \ln N$  plus Gaussian-logarithmic fluctuations, for which it is well known that for  $u_t \sim \mathcal{N}(0, 1)$ ,

$$\mathbb{E}[\ln |u_t|] = -\frac{g + \ln 2}{2}, \quad \text{Var}[\ln |u_t|] = \frac{\pi^2}{8}, \quad (114)$$

where  $g$  is Euler's constant. Combining these contributions, the reinjection term has

$$\mu_\theta = \frac{1}{2} \ln N + g + \frac{\ln 2}{2}, \quad \sigma_\theta^2 = \frac{\pi^2}{8}. \quad (115)$$

The overlap between expanding and contracting directions shrinks like  $N^{-1/2}$ , which produces a negative correction to the Lyapunov exponent. At the same time, the variance remains constant ( $\pi^2/8$ ), reflecting universal logarithmic fluctuations associated with Gaussian overlaps. Hence, reinjection acts as a dimensional penalty on stability, counterbalancing the growth induced by the condition number.

#### D. Asymptotic Lyapunov and tail exponents

Putting everything together, the Lyapunov exponent is approximated as

$$\gamma \approx \ln \rho + \ln \sigma_0 - \ln N - \ln 2 - g + 2\sigma \sqrt{2 \ln N}. \quad (116)$$

Hence, for large  $N$ , there exists a critical non-normal standard deviation  $\sigma_c(N)$  of the singular values defined in Eq. (106) such that  $\gamma \approx 0$ :

$$\sigma_c(N) = \frac{1}{2\sqrt{2}} \sqrt{\ln N} \left[ 1 + \frac{\ln 2 + g - \ln \rho - \ln \sigma_0}{\ln N} \right]. \quad (117)$$

From the CLT decomposition  $\alpha \approx -\frac{2\gamma}{\sigma_\rho^2 + \sigma_\kappa^2 + \sigma_\theta^2}$  derived in Eq. (74), the variance entering the denominator of Eq. (74) is

the sum of the three contributions:

$$\sigma_\rho^2 + \sigma_\kappa^2 + \sigma_\theta^2 \approx \underbrace{\sigma_\rho^2}_{\text{fixed}} + \underbrace{\left(3 + \frac{2\sigma^2}{\ln N}\right)}_{\text{EVT range of Gaussians}} + \underbrace{\frac{\pi^2}{8}}_{\text{from Eq. (115)}}. \quad (115)$$

Collecting constants and expanding consistently to first order around  $\sigma = \sigma_c(N)$  leads to the compact near-critical form

$$\alpha \approx \frac{\sqrt{2 \ln N}}{\frac{48}{\pi^2}} \frac{\sigma_c(N) - \sigma}{3 + \frac{2\sigma^2}{\ln N}} + o\left(\frac{1}{\sqrt{\ln N}}\right), \quad (118)$$

Thus, the tail exponent decreases *linearly* with the distance to criticality, with a slope that grows like  $\sqrt{2 \ln N}$  and a denominator collecting the variance contributions of spectrum, condition number, and reinjection. These results show that, in high dimensions, the combined effect of spectral fluctuations, condition numbers, and random reinjection induces a delicate balance between stability and criticality. In particular,

(1) the critical non-normal variance  $\sigma_c(N)$  grows like  $\sqrt{\ln N}$ , implying that larger systems tolerate more heterogeneity before destabilizing when  $\gamma \approx 0$ ;

(2) the tail exponent  $\alpha$  decreases as  $\sigma$  approaches  $\sigma_c(N)$ , confirming that high-dimensional non-normality drives heavy-tailed fluctuations even when the system is spectrally stable.

This highlights how non-normal amplification, coupled with dimensionality, pushes the system closer to criticality in the sense of diverging moments and vanishing tail exponents.

### E. Singular-value perturbations under random fluctuations

In the previous section, we focused on the case  $\mathbb{E}[\ln \kappa_1] \gg 1$ , which allowed us to approximate the Lyapunov exponent using the dominant singular value (condition number) only. Here, we consider the opposite regime, where the deviations from normality are small and can be treated perturbatively. To this end, we write

$$\ln s_{i,t} = \sigma \epsilon_{i,t}, \quad \mathbb{E}[\epsilon_{i,t}] = 0, \quad \text{Var}[\epsilon_{i,t}] = 1, \quad (119)$$

with  $\sigma \ll 1$  controlling the amplitude of the non-normal perturbation.

Using the matrix decomposition (84), we can expand

$$\mathbf{A}_t = \mathbf{A}_t^0 + \sigma \mathbf{A}_t^1 + \mathcal{O}(\sigma^2),$$

$$\text{where } \mathbf{A}_t^0 = \mathbf{U}_t \mathbf{V}_t^\dagger \mathbf{A}_t \mathbf{V}_t \mathbf{U}_t^\dagger,$$

$$\mathbf{A}_t^1 = \mathbf{U}_t \mathbf{E}_t \mathbf{V}_t^\dagger \mathbf{A}_t \mathbf{V}_t \mathbf{U}_t^\dagger - \mathbf{U}_t \mathbf{V}_t^\dagger \mathbf{A}_t \mathbf{E}_t \mathbf{V}_t \mathbf{U}_t^\dagger, \quad (120)$$

and  $\mathbf{E}_t = \text{Diag}(\epsilon_{1,t}, \dots, \epsilon_{N,t})$ , so that  $\mathbf{A}_t^0$  is Hermitian (symmetric) and  $\mathbf{A}_t^1$  is anti-Hermitian (antisymmetric). This allows us to expand the product

$$\mathbf{\Pi}_t = \mathbf{\Pi}_t^0 + \sigma \mathbf{\Pi}_t^1 + \mathcal{O}(\sigma^2),$$

$$\text{where } \mathbf{\Pi}_t^0 = \prod_{s=1}^t \mathbf{A}_s^0,$$

$$\mathbf{\Pi}_t^1 = \sum_{s=1}^t \left( \prod_{u=1}^{s-1} \mathbf{A}_u^0 \right) \mathbf{A}_s^1 \left( \prod_{u=s+1}^t \mathbf{A}_u^0 \right). \quad (121)$$

Since  $\mathbf{\Pi}_t^0$  is Hermitian and  $\sigma \ll 1$ , the  $L_2$  norm of the product can be expanded via standard eigenvalue perturbation theory as

$$\|\mathbf{\Pi}_t\| \approx \|\mathbf{\Pi}_t^0\| \left[ 1 + \sigma \frac{\hat{\mathbf{v}}_{\max}^\dagger \mathbf{\Pi}_t^1 \hat{\mathbf{v}}_{\max}}{\|\mathbf{\Pi}_t^0\|} + \sigma^2 \frac{1}{\|\mathbf{\Pi}_t^0\|} \sum_{k \neq \max} \frac{|\hat{\mathbf{v}}_k^\dagger \mathbf{\Pi}_t^1 \hat{\mathbf{v}}_{\max}|^2}{\|\mathbf{\Pi}_t^0\| - \lambda_k} \right], \quad (122)$$

where  $\hat{\mathbf{v}}_{\max}$  is the eigenvector of  $\mathbf{\Pi}_t^0$  associated with the largest eigenvalue  $\|\mathbf{\Pi}_t^0\|$ , and  $\hat{\mathbf{v}}_k$  and  $\lambda_k$  denote the other eigenvectors and eigenvalues of  $\mathbf{\Pi}_t^0$ . We note that the first-order correction  $\mathbf{\Pi}_t^1$  in the expansion (121) contains exactly one occurrence of the perturbation  $\mathbf{A}_t^s$ , while all other factors are the unperturbed Hermitian matrices  $\mathbf{A}_u^0$ . Because  $\mathbf{A}_t^s$  is anti-Hermitian, its eigenvalues are purely imaginary, so it generates infinitesimal unitary rotations. This can be seen by writing, for any vector  $v$ ,

$$\frac{d}{d\epsilon} \|(I + \epsilon \mathbf{A}_t^s) v\|^2 \Big|_{\epsilon=0} = 2 v^\dagger \frac{1}{2} (\mathbf{A}_t^s + \mathbf{A}_t^{s\dagger}) v = 0, \quad (123)$$

since  $\mathbf{A}_t^s + \mathbf{A}_t^{s\dagger} = 0$ . Therefore, the action of  $\mathbf{A}_t^s$  changes only the *direction* of  $v$ , not its norm, at first order in  $\epsilon$ . In other words,  $\mathbf{A}_t^s$  acts as an infinitesimal generator of rotations in the complex vector space.

Because  $\mathbf{A}_t^s$  has zero mean,  $\mathbb{E}[\mathbf{A}_t^s] = 0$ , the first-order correction  $\mathbf{\Pi}_t^1$  also averages to zero over the ensemble of random realizations. The reason is that each  $\mathbf{\Pi}_t^1$  contains a single insertion of  $\mathbf{A}_t^s$ , which upon averaging gives

$$\mathbb{E}[\mathbf{\Pi}_t^1] = \sum_{s=1}^t \mathbb{E} \left[ \left( \prod_{u=1}^{s-1} \mathbf{A}_u^0 \right) \mathbf{A}_s^1 \left( \prod_{u=s+1}^t \mathbf{A}_u^0 \right) \right] = 0, \quad (124)$$

since  $\mathbb{E}[\mathbf{A}_s^1] = 0$  by construction. Thus, any contribution to the norm at order  $\sigma$  cancels on average.

Geometrically, the anti-Hermitian perturbation introduces only infinitesimal rotations of the eigenbasis of  $\mathbf{\Pi}_t^0$ , which conserve the  $L_2$  norm. Hence, the norm  $\|\mathbf{\Pi}_t\|$  is invariant at first order, and corrections to the Lyapunov exponent arise only at order  $\sigma^2$ . The physical interpretation is that anti-Hermitian fluctuations produce random rotational flows in matrix space that redistribute amplitudes among directions without altering the total norm to first order, so that only quadratic (variance-driven) effects contribute to the long-term growth rate.

Moreover, since  $\mathbf{\Pi}_t^1$  is a sum of random matrices with finite variance, by the central limit theorem, its matrix elements tend to a Gaussian distribution with zero mean for large  $t$ . Consequently, the logarithm of the norm can be approximated as

$$\ln \pi_t \sim \mathcal{N}(t(\gamma_0 + \sigma^2 c_1), t(\sigma_0^2 + c_2 \sigma^2)), \quad t \rightarrow \infty, \quad (125)$$

where  $\gamma_0$  is the Lyapunov exponent of the purely normal system ( $\sigma = 0$ ),  $\sigma_0^2 t$  is the variance of  $\ln \pi_t$  in the normal system, and  $c_1, c_2 \geq 0$  are constants independent of  $\sigma$  determined by the variance of the perturbation.

Hence, for weak non-normality, the Lyapunov exponent and the associated tail exponent  $\alpha$  admit the expansions

$$\gamma \approx \gamma_0 + c_1 \sigma^2, \quad \alpha \approx -2 \frac{\gamma}{\sigma_0^2 + c_2 \sigma^2}. \quad (126)$$

We see explicitly that, in this regime, the correction to the Lyapunov exponent is quadratic in  $\sigma$ , rather than linear, in agreement with the intuition that anti-Hermitian perturbations contribute only at second order to the growth rate.

### F. Results of numerical simulations

We numerically investigate the Lyapunov exponent  $\gamma$  for the  $N$ -dimensional Kesten process (1), when the random interaction matrices  $\mathbf{A}_t$  are generated from the spectral-singular decomposition (84). The purpose of these numerical experiments is twofold: (1) to illustrate the qualitative effect of non-normality on  $\gamma$  and  $\alpha$ , and (2) to test the large- $N$  scaling expressions (116) and (118).

Concretely, the eigenvalues are distributed according to  $\ln \lambda_{i,t} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(-1, 1/5)$ . We model the singular values of the eigenbasis transform by  $\ln s_{i,t} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2)$ . The unitary matrix  $\mathbf{V}_t$  is sampled uniformly at each step. We consider two alternatives for  $\mathbf{U}_t$ : either  $\mathbf{U}_t = \mathbf{I}$  for all  $t$  (“no rotation”) or  $\mathbf{U}_t$  is sampled uniformly at each step (“uniform rotation”). For each parameter set we estimate  $\gamma$  by the standard QR decomposition-based reorthonormalization method described in the Appendix.

Figure 3 displays the estimated Lyapunov exponent as a function of the singular-value dispersion  $\sigma$  for different system sizes. The main empirical findings are as follows:

(1) When the system is *normal* ( $\sigma = 0$ ), the measured Lyapunov exponent is essentially constant as a function of  $N$ . This matches the expectation for normal matrices, where the spectral radius controls  $\gamma$  and no  $\ln N$  dependence is observed.

(2) Introducing non-normality ( $\sigma > 0$ ) produces an approximate linear increase of  $\gamma$  with  $\sigma$ . The slope of this linear contribution is larger in the “uniform-rotation” case than in the “no-rotation” case.

(3) The predicted negative  $-\frac{1}{2} \ln N$  contribution from eigenvector overlaps (derived in the asymptotic analysis) is not visible in these experiments:  $\gamma$  does not systematically decrease with  $N$ . Instead, the dominant finite- $N$  effect is well described by  $\gamma \approx \gamma_0 + \mathbb{E}[\ln \kappa]$ , with  $\gamma_0$  being the Lyapunov exponent of the normal system, i.e., for  $\sigma = 0$ , and  $\mathbb{E}[\ln \kappa]$  growing approximately like  $\sqrt{\ln N}$  as theory predicts for  $\ln \kappa$ .

The difference between the analytic  $-\ln N/2$  correction and the numerical outcome likely stems from selection and correlation effects that were not fully captured in the heuristic asymptotic derivation. In particular,

(1) when  $\mathbf{U}_t = \mathbf{I}$  (no random rotation), the reinjection into the most-amplified direction occurs only when the previous mode projects on the non-normal direction. Under isotropy, this happens on average roughly  $t/N$  times over a time horizon  $t$ , which reduces the effective contribution of  $\sum_s \ln \kappa_s$  by a factor  $\sim 1/N$ . This explains why the slope in the “no-rotation” case is smaller;

(2) when  $\mathbf{U}_t$  is uniformly random, rotations frequently reorient modes such that the non-normal amplification is applied

more often; hence the larger slope in the “uniform-rotation” case;

(3) the absence of an observed  $-\frac{1}{2} \ln N$  term suggests that the cross-term responsible for this contribution (roughly, sums of the form  $\sum_k \lambda_k^0 v_{ik} v_{jk}^*$  evaluated at extremal indices) does not have the small- $O(1/\sqrt{N})$  variance assumed in the derivation; conditioning on extremal singular-value indices or correlations between  $\mathbf{V}_t$  and  $\mathbf{A}_t$  can increase the variance of these sums to  $O(1)$ , removing the  $-\frac{1}{2} \ln N$  penalty in practice.

For intermediate values of the expected logarithm of the condition number, the Lyapunov exponent is no longer described by the dominant approximation but rather by the perturbative regime of weak non-normality, in which it grows quadratically with  $\sigma$ . This quadratic dependence is recovered in the inset of Fig. 3.

Figure 4 shows that the Lyapunov exponent increases, while the tail exponent decreases, with the system dimension  $N$ . The magnitude of  $\sigma$  modulates this trend: larger  $\sigma$  values enhance, whereas smaller  $\sigma$  values attenuate, the rate at which these exponents vary with increasing  $N$ .

The simulations confirm several qualitative predictions: (1) non-normality increases the effective Lyapunov exponent, (2) random rotations amplify this effect, and (3) the leading contribution from the log-condition number behaves as predicted by extreme-value arguments.

### V. APPLICATIONS TO PHYSICS, FINANCE, AND ECONOMICS

The theoretical developments presented above, though derived in an abstract mathematical framework, provide a versatile foundation for understanding a wide spectrum of real-world phenomena. The generality of the stochastic Kesten framework, extended here to encompass non-normal multiplicative dynamics, makes it directly relevant to numerous domains. In particular, systems characterized by multiplicative feedback, stochastic reinforcement, or transient amplification of fluctuations, which are ubiquitous features of physical, financial, and economic processes, naturally fall within the scope of this theory.

These examples should not be viewed as isolated case studies. Rather, they illustrate, in concrete settings, the general mechanism developed above: Heavy tails and instability are governed not only by spectral fluctuations, but by whether non-normal amplification can be constructively reinjected through rotations, switching, or evolving interaction patterns. In this sense, the applications provide physical realizations of the general theory and help clarify how the underlying mechanism operates in practice.

In physical systems, the same formal structure naturally emerges in the dynamics of polymer stretching in turbulent flows and in the kinematic stage of the small-scale turbulent dynamo. In both cases, the local velocity-gradient tensor plays the role of a random non-normal operator: The transient alignment of its eigenvectors leads to intermittent amplification of either the polymer extension or the magnetic field, precisely as predicted by the non-normal Kesten mechanism developed here. These examples demonstrate that the framework is not limited to abstract or economic settings, but captures




FIG. 3. Lyapunov exponent  $\gamma$  for the Kesten process (1) with matrices  $\mathbf{A}_t$  generated via the decomposition (84) as a function of  $\sigma\sqrt{\ln N}$ , testing the scaling predictions: (1) For small  $\mathbb{E}[\ln \kappa_t] \sim \sigma\sqrt{\ln N}$ , the Lyapunov exponent scales as  $\gamma \sim \text{Var}[\ln \kappa_t] \propto \sigma^2 \ln N$  (126) (as presented in the inset); (2) for large  $\mathbb{E}[\ln \kappa_t]$ , the Lyapunov exponent scales as  $\gamma \sim \mathbb{E}[\ln \kappa_t]$  (93). For the top panels, we generated  $\mathbf{V}_t^\dagger \mathbf{A}_t \mathbf{V}_t$  once at the beginning and kept it fixed throughout the simulation, ensuring that the “normal” part of the matrix is fixed. For the bottom panels, we again used Eq. (84) but generated the matrix  $\mathbf{V}_t^\dagger \mathbf{A}_t \mathbf{V}_t$  at each step. We set  $\ln \lambda_{i,t} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(-1, 1/9)$  and  $\ln s_{i,t} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2)$ . Each curve shows  $\gamma$  as a function of  $\sigma$  for different dimensions  $N$ . The “no-rotation” cases correspond to fixing  $\mathbf{U}_t = \mathbf{I}$ ; the “uniform-rotation” cases correspond to sampling  $\mathbf{U}_t$  uniformly at each time step. To measure  $\gamma$ , we use the QR-based Lyapunov estimator described in the Appendix.

fundamental amplification processes in fluid and magnetohydrodynamic turbulence.

Then, we illustrate the relevance of our formalism in four complementary applications in finance and economics. The first application concerns volatility models of ARCH type, where the non-normal formulation introduces mechanisms of asymmetric feedback between returns and volatility. We then

turn to the theory of cointegration and vector autoregressive (VAR) models, which naturally map to stochastic Kesten processes when cross-asset interaction matrices are reestimated dynamically. Next, we reinterpret factor models as dimensional reductions of high-dimensional stochastic systems, showing how crises can be associated with effective factor collapse and proximity to non-normal criticality. Finally, we




FIG. 4. Lyapunov exponent  $\gamma$  (left panels) and power law tail exponent  $\alpha$  (right panels) for the Kesten process (1) with matrices  $\mathbf{A}_t$  generated via the decomposition (84), as a function of the system dimension  $N$ . For the top panels, we generated  $\mathbf{V}_t^\dagger \mathbf{A}_t \mathbf{V}_t$  once at the beginning and kept it fixed throughout the simulation, ensuring that the “normal” part of the matrix remained constant. In the bottom panels, we again used Eq. (84) and generated the matrix  $\mathbf{V}_t^\dagger \mathbf{A}_t \mathbf{V}_t$  at each step. We set  $\ln \lambda_{i,t} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(-1, 1/9)$  (eigenvalues) and  $\ln s_{i,t} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2)$  (singular values). Each curve shows  $\gamma$  as a function of  $N$  for different values of  $\sigma$ . The “no-rotation” cases correspond to fixing  $\mathbf{U}_t = \mathbf{I}$ ; the “uniform-rotation” cases correspond to sampling  $\mathbf{U}_t$  uniformly at each time step. To measure  $\gamma$ , we use the QR-based Lyapunov estimator described in the Appendix. Both the Lyapunov and tail exponents are measured over ten realizations, and the error bars on the tail exponent are obtained by measuring the standard deviation of the tail exponent across realizations.

extend the framework to the problem of wealth inequality, connecting our mechanism to established economic theories of capital accumulation and highlighting how endogenous heavy tails can emerge in interacting-agent economies.

Taken together, these examples show that the mathematical results derived above transcend their abstract origin and apply directly to diverse physical and socio-economic systems. They demonstrate that non-normal eigenvector amplification constitutes a generic mechanism for heavy-tailed fluctuations in systems driven by multiplicative interactions, from

turbulent flows and magnetic dynamos to markets and wealth distributions.

#### A. Polymer stretching in turbulent flows as an instance of eigenvector amplification

The coil-stretch transition of single polymer molecules in random or turbulent flows is a paradigmatic example of multiplicative stochastic amplification balanced by nonlinear relaxation. In such flows, the local velocity gradients act

as a sequence of random matrices that repeatedly stretch and reorient the polymer, while entropic elasticity provides a reinjection mechanism that prevents collapse and ensures a stationary distribution of extensions. This competition between random multiplicative amplification and nonlinear restoring forces gives rise to a power law distribution of end-to-end polymer lengths [30,39].

Following the formulation of Refs. [30,39], the stochastic evolution of the polymer's end-to-end vector  $\mathbf{R}$  obeys

$$\dot{\mathbf{R}} = (\nabla \mathbf{v}) \mathbf{R} - \frac{1}{\tau} \mathbf{R}, \quad (127)$$

where  $\nabla \mathbf{v}$  is the local velocity-gradient tensor and  $\tau$  the polymer relaxation time. Writing  $\mathbf{R} = R \mathbf{n}$ , with  $\mathbf{n}$  a unit orientation vector, the dynamics separates into

$$\dot{R} = R z - \frac{R}{\tau}, \quad \dot{\mathbf{n}} = (\nabla \mathbf{v}) \mathbf{n} - z \mathbf{n}, \quad (128)$$

where  $z(t) = \mathbf{n}^\top (\nabla \mathbf{v}) \mathbf{n}$  is the instantaneous stretching rate. The logarithmic extension  $r(t) = \ln(R/R_0)$  then evolves as

$$r(t) = \int_0^t \left[ z(t') - \frac{1}{\tau} \right] dt'. \quad (129)$$

Using large-deviation theory for the time-integrated process  $z(t)$ , Refs. [30,39] derived an exponential tail for  $r$ , which translates into a power law distribution for  $R$ :

$$P(R) \sim \frac{1}{R^{1-2a}}, \quad (130)$$

with the exponent  $a$  determined by the statistics of finite-time Lyapunov exponents of the random velocity gradient. The transition at  $a = 0$  corresponds to the classical coil-stretch threshold, where the mean stretching rate (the flow Lyapunov exponent  $\lambda$ ) equals the inverse relaxation time  $1/\tau$ .

*Interpretation in terms of non-normal amplification.* The above derivation describes  $z(t)$  as a stochastic process driven by the local strain field, but its origin can be understood more generally within the framework of non-normal random stochastic multiplicative dynamics. The instantaneous velocity-gradient tensor  $\nabla \mathbf{v}$  is generically non-normal: its eigenvectors are nonorthogonal, and hence transient growth of  $\|\mathbf{R}\|$  can occur even if all eigenvalues have negative real parts. The polymer extension  $\mathbf{R}$  therefore experiences random episodes of *eigenvector alignment*, which are moments when  $\mathbf{n}$  becomes nearly collinear with a subset of nonorthogonal stretching directions of  $\nabla \mathbf{v}$ . During such episodes, constructive interference among these directions leads to transient amplification of  $\|\mathbf{R}\|$  that can far exceed the exponential growth expected from the eigenvalues alone. These intermittent amplification events correspond precisely to the large-deviation fluctuations that dominate the power law tail of  $P(R)$ .

In this interpretation, the “coil-stretch transition” is a stochastic bifurcation between two regimes of eigenvector dynamics: (1) A dissipative regime where random rotations dominate and alignment is quickly lost, keeping the polymer coiled, and (2) an amplification regime where alignment persists long enough for transient growth to overcome elastic relaxation. The critical condition  $\lambda = 1/\tau$  thus separates a state of rotational diffusion from one dominated by coherent transient amplification. This viewpoint naturally extends the large-deviation theory of Balkovsky-Fouxon-Lebedev [39]

and Chertkov [30] by relating it to the spectral geometry of non-normal operators.

*Connection with multiplicative stochastic processes.* Equation (127) can be rewritten as a multiplicative update

$$\mathbf{R}(t + \Delta t) = \left[ \mathbf{I} + \Delta t \left( \nabla \mathbf{v}_t - \frac{1}{\tau} \mathbf{I} \right) \right] \mathbf{R}(t), \quad (131)$$

revealing its analogy with a Kesten-type process driven by random matrices  $\mathbf{A}_t = \mathbf{I} + \Delta t \nabla \mathbf{v}_t$ . In this discrete form, the polymer extension corresponds to the norm of a product of non-normal random matrices, and the power law distribution of  $R$  arises from the same multiplicative reinjection mechanism that governs heavy-tailed statistics in other systems. From this perspective, the velocity-gradient tensor plays the role of a non-normal amplification operator, while the elastic relaxation  $-\mathbf{R}/\tau$  provides a restoring term that maintains statistical stationarity.

*Physical significance.* The non-normal eigenvector amplification picture clarifies why extreme polymer stretching events are intermittent and burstlike. Even in statistically stationary turbulence, regions of high strain correspond to transiently coherent structures of the velocity gradient, in which eigenvectors become strongly aligned. These rare but powerful episodes drive the tail of the extension distribution and are the microscopic counterparts of finite-time Lyapunov fluctuations responsible for the heavy-tailed statistics predicted by large-deviation theory. Thus, polymer stretching in turbulence provides a vivid physical realization of stochastic amplification through non-normal matrix dynamics.

## B. Small-scale turbulent dynamo as eigenvector amplification

A closely related structure governs the small-scale turbulent dynamo, where the magnetic field is repeatedly stretched and rotated by fluctuating velocity gradients, while resistive diffusion ensures saturation [40,41]. In the *kinematic regime*, when the magnetic field is still dynamically weak and does not influence the flow, the induction equation

$$\partial_t \mathbf{B} = (\mathbf{B} \cdot \nabla) \mathbf{u} - \mathbf{B}(\nabla \cdot \mathbf{u}) + \eta \nabla^2 \mathbf{B} \quad (132)$$

reduces along Lagrangian trajectories to

$$\dot{\mathbf{B}} = (\nabla \mathbf{u}) \mathbf{B} + \eta \nabla^2 \mathbf{B}. \quad (133)$$

Neglecting diffusion above the resistive cutoff, the magnitude  $B = |\mathbf{B}|$  and direction  $\mathbf{b} = \mathbf{B}/B$  evolve as

$$\dot{s} = \frac{d}{dt} \ln B = \mathbf{b}^\top S \mathbf{b}, \quad \dot{\mathbf{b}} = \Omega \mathbf{b} + (S - \mathbf{b} \mathbf{b}^\top S) \mathbf{b}, \quad (134)$$

where

$$S = \frac{1}{2}(\nabla \mathbf{u} + (\nabla \mathbf{u})^\top) \quad (135)$$

is the rate-of-strain tensor and

$$\Omega = \frac{1}{2}(\nabla \mathbf{u} - (\nabla \mathbf{u})^\top) \quad (136)$$

the vorticity tensor. The growth of  $\ln B$  thus depends on the instantaneous projection  $\mathbf{b}^\top S \mathbf{b}$ , while the orientation  $\mathbf{b}$  is rotated and tilted by  $\Omega$  and by the off-diagonal action of  $S$ . In this kinematic stage, the mechanism of *non-normal eigenvector amplification* becomes transparent: The non-normality

of  $\nabla \mathbf{u}$  implies nonorthogonal and time-dependent eigenvectors, so transient alignments between  $\mathbf{b}$  and nearly colinear stretching directions can produce amplification bursts far exceeding the growth predicted by eigenvalues alone. Random sequences of such alignments form a multiplicative matrix process for  $\mathbf{B}$ , while resistive diffusion  $\eta \nabla^2 \mathbf{B}$  provides the reinjection and ultimate saturation at small scales. This linear regime corresponds precisely to the multidimensional Kesten setting. At later times, once the magnetic energy becomes comparable to the kinetic energy of viscous-scale eddies, nonlinear feedback through the Lorentz force modifies the velocity gradients and limits further growth, leading to the slow resistive-scale evolution described by Schekochihin *et al.* [41]. Thus, the eigenvector-amplification mechanism governs the kinematic stage and seeds the transition to the nonlinear saturated dynamo.

### C. Multidimensional non-normal ARCH models

One of the earliest and most influential applications of multiplicative processes in economics and finance is the ARCH model, introduced by Engle in his seminal paper [42]. The ARCH framework was developed to account for the empirical observation that financial time series exhibit volatility clustering, i.e., periods of high variability tend to be followed by periods of high variability, and periods of calm are followed by calm.

The original ARCH( $q$ ) model specifies that a return process  $\{r_t\}$  is given by

$$r_t = \sqrt{v_t} \varepsilon_t, \quad \varepsilon_t \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1), \quad (137)$$

with the conditional variance evolving as

$$v_t = v_0 + \sum_{i=1}^q a_i r_{t-i}^2, \quad v_0 > 0, \quad a_i \geq 0. \quad (138)$$

Here,  $v_t$  captures the time-varying variance, driven by past squared returns. The key innovation of the ARCH process is that volatility is no longer constant, but an evolving random variable reacting endogenously to shocks.

It was quickly noticed that ARCH processes naturally generate unconditional distributions for returns that are power law distributed. This comes directly from the multiplicative structure: Returns are Gaussian conditional on  $v_t$ , but since  $v_t$  itself is random, the marginal distribution becomes a mixture of Gaussians, yielding power law distributions [43].

The asymptotic tail exponent  $\alpha$  of an ARCH-type process can be analyzed through Kesten's model [1], by considering the simple ARCH(1) case, the squared return recursion is

$$r_t^2 = \varepsilon_t^2 [v_0 + a_1 r_{t-1}^2], \quad (139)$$

which is a Kesten process in the variable  $r_t^2$ , where the multiplicative and additive noise are mutually correlated since the same noise term  $\varepsilon_t^2$  drives them both. Comparing with definition (8), we have  $\rho_t = a_1 \varepsilon_t^2$  and  $\sqrt{2\delta} \eta_t = v_0 \varepsilon_t^2$ . We note that the classical tail theory of Kesten processes leading to the power law tail (6) does not require  $\rho_t$  and  $\eta_t$  to be independent of each other; it only requires the i.i.d. driving sequence to be independent of the past state.

Empirical studies of financial data typically find power law tails with exponent  $\alpha \approx 3$ , sometimes referred to as the

"inverse-cubic law" [44]. However, standard ARCH and even its popular generalization, GARCH [45], tend to generate larger exponents (often  $\alpha > 4$ ), corresponding to tails that are "too light" compared to real data [46].

This discrepancy has been labeled the *tail exponent puzzle* in the econometrics and econophysics literature. It finds a natural resolution once the analysis is extended from scalar to multidimensional Kesten dynamics. In real markets, returns are not governed by isolated variance processes but by the joint interactions of many correlated assets (stocks, indices, and portfolios) that together form a high-dimensional multiplicative system. As the system dimension  $N$  increases, non-normal eigenvector amplification lowers the effective tail exponent through two complementary effects: (1) an increase of the Lyapunov exponent  $\gamma$ , driven by  $\mathbb{E}[\ln \kappa] \propto \ln N$ , and (2) an enhanced variance of  $\ln \kappa$ , which decreases  $\alpha$  as  $1/\sqrt{\ln N}$ . Together, these collective effects provide a natural mechanism for the emergence of heavier tails across financial assets. Even when all eigenvalues remain subcritical, enlarging  $N$  increases the condition number of the eigenbasis, amplifying transient growth and reducing  $\alpha$  from the unrealistically high values predicted by scalar GARCH models to empirically consistent ones. Importantly, because this amplification arises from the non-normal structure of cross-asset interactions rather than from asset-specific parameters, the same effective exponent naturally emerges across diverse markets and instruments, as expressed by

$$\mathbb{P}[\mathbf{n} \cdot \mathbf{x}_t > x_n] \sim x_n^{-\alpha}, \quad x_n \rightarrow \infty, \quad (140)$$

where  $\mathbf{n} \cdot \mathbf{x}_t$  denotes any projection (the same asymptotics hold for the  $L_2$  norm).

For completeness, we note that several alternative approaches have been proposed to reconcile ARCH-type models with the empirically observed tail exponents.

*Introducing Student- $t$  innovations:* By replacing Gaussian shocks  $\varepsilon_t$  with Student- $t$  innovations, the model can directly match heavier tails [46]. However, this shifts the responsibility for tails away from the volatility dynamics, and is sometimes criticized as *ad hoc*.

*Long-memory extensions:* fractionally integrated generalized autoregressive conditional heteroskedasticity (FIGARCH) and related models [47] incorporate fractional integration, allowing volatility to decay slowly and reproduce empirically observed scaling laws.

*Nonlinear GARCH models:* Threshold GARCH, exponential generalized autoregressive conditional heteroskedasticity (EGARCH), and others introduce asymmetry and nonlinear effects, which can modify the distributional properties and bring tail exponents closer to empirical ones.

*Stochastic volatility models:* In continuous time, models with latent stochastic volatility can also reproduce power-law-like tails under certain parametrizations [48].

*Quadratic Hawkes model:* The Quadratic Hawkes model [49] adds a weak nonlinear feedback term to a linear self-exciting Hawkes process, capturing the Zumbach effect where aligned past returns enhance future volatility. It is written in terms of the QHawkes intensity as  $\lambda_t = \lambda_\infty + H_t + n_Z Z_t^2$  with  $H_t = \int \phi(t-s) dN_s$ ,  $Z_t = \int k(t-s) dP_s$ . This quadratic term  $n_Z Z_t^2$  produces multiplicative volatility diffusion, leading to

stationary power law tails with exponent  $\alpha = 1 + \frac{1}{n_Z(1+a^*)}$ . The exponent  $a^*$  is given by  $a^* \approx n_H/(1 - n_H)$ , where  $n_H$  is the branching ratio of the linear Hawkes part, in the limit where the quadratic (trend) feedback relaxes faster than the linear Hawkes memory, so volatility dynamics are dominated by short-term self-excitation. The obtained exponent is  $\alpha \approx 3-4$  even for small nonlinearity  $n_Z \sim 0.05-0.1$  [49]. The Quadratic Hawkes model also breaks time-reversal symmetry, thereby reproducing both the fat tails and volatility asymmetry observed empirically.

#### D. Stochastic cointegration as the remnant of global stability under non-normal amplification

A central concept in empirical finance and econometrics is the theory of *cointegration*, introduced by Engle and Granger [50] and later generalized by Johansen [51]. The idea is that, even if individual asset prices are nonstationary, often modeled as unit-root processes such as geometric Brownian motions, certain linear combinations of them may remain stationary. This framework captures the coexistence of long-run equilibrium relations between assets and short-run deviations, forming the empirical backbone of error-correction models and market efficiency tests.

Classical cointegration theory thus begins with the assumption that individual log-prices are intrinsically nonstationary and then searches for exceptional linear combinations that restore stationarity. These combinations are interpreted as manifestations of deep equilibrium relationships between the nonstationary variables, but they often appear empirically fragile and require delicate statistical tuning to identify. In this traditional perspective, cointegration represents an *exceptional alignment* among nonstationary processes, namely, an imposed condition rather than a natural dynamical property of the system.

Formally, let  $\mathbf{x}_t \in \mathbb{R}^N$  denote the vector of asset log-returns at time  $t$ . A standard representation is the VAR model:

$$\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \boldsymbol{\eta}, \quad \boldsymbol{\eta} \sim \mathcal{N}(0, \mathbf{I}_N), \quad (141)$$

where  $\mathbf{A}$  is the  $N \times N$  interaction matrix. Cointegration arises when  $\mathbf{A}$  has unit roots while certain linear projections of  $\mathbf{x}_t$  remain stationary.

In practice, financial interactions between assets (arbitrage, risk premia, liquidity spillovers) are not fixed but time varying. This motivates replacing the constant matrix  $\mathbf{A}$  by a random sequence  $\{\mathbf{A}_t\}$ , reestimated or updated as market information arrives:

$$\mathbf{x}_{t+1} = \mathbf{A}_t \mathbf{x}_t + \boldsymbol{\eta}, \quad (142)$$

which is exactly the form of the multidimensional Kesten process (1). Here, the random matrix  $\mathbf{A}_t$  captures the stochastic and possibly asymmetric nature of cross-asset feedback, while maintaining market efficiency, by considering  $\mathbb{E}[\mathbf{A}_t] = 0$ . Unlike the classical VAR case, where  $\mathbf{A}$  is constant and symmetric assumptions are usually imposed for tractability, our formulation allows  $\mathbf{A}_t$  to be random, nonsymmetric, and non-normal, which is arguably more realistic in financial systems.

The key insight is to begin with a spectrally stable system, where every random matrix  $\mathbf{A}_t$  has all eigenvalues inside the unit circle [ $\rho(\mathbf{A}_t) < 1$ ], ensuring stability in

the absence of non-normal effects. Yet, non-normal interactions between components can still induce positive Lyapunov exponents, thereby creating a limited number of unstable, nonstationary directions despite the local spectral stability of each  $\mathbf{A}_t$ .

In the absence of non-normal effects, all directions of  $\mathbf{x}_t$  would be mean reverting, and the system would converge to a stationary distribution. However, when the matrices  $\mathbf{A}_t$  are non-normal and noncommuting, transient amplification can raise the leading Lyapunov exponent  $\lambda_1$  above zero, even though every instantaneous spectrum remains stable. This mechanism creates a *partial loss of stability*: A subset of directions (those associated with  $\lambda_i > 0$  or  $\lambda_i \approx 0$ ) become nonstationary or weakly trending, while the remaining directions with  $\lambda_i < 0$  stay stationary.

*Inverted cointegration logic.* This spectral reordering naturally generates a structure that mirrors the classical notion of cointegration but with inverted causality. Instead of searching for stationary combinations among nonstationary prices, we start from a fully stationary system and let non-normal amplification make a few combinations nonstationary. The surviving stable directions form the *cointegrating subspace* automatically—the residual imprint of global stability under partial instability. In this sense, cointegration is not an artificial equilibrium constraint, but the generic outcome of a system whose Lyapunov spectrum straddles zero:

$$\lambda_1 > \lambda_2 > \dots > \lambda_r > 0 > \lambda_{r+1} > \dots > \lambda_N.$$

Such partially unstable systems will almost necessarily exhibit stationary linear combinations of  $x_t$ , defined by the left Oseledets vectors associated with the negative Lyapunov exponents. The subspaces  $E_i(\omega)$  and their duals  $E_i^*(\omega)$  are the right and left Oseledets spaces associated with the Lyapunov exponents  $\lambda_i$ , as established by the multiplicative ergodic theorem [52,53]. The corresponding vectors (covariant Lyapunov vectors) span the dynamically invariant directions of the random cocycle and its adjoint [54,55].

Hence, cointegration becomes the *norm* for high-dimensional financial systems that are globally stable in expectation but locally destabilized by non-normal interactions. This interpretation calls for a systematic analysis of the *Lyapunov spectrum* of random non-normal VAR processes. One should derive conditions on the distribution of  $\mathbf{A}_t$ , for instance, on the moments, the variance of off-diagonal elements, or the expected  $\log \kappa(\mathbf{A}_t)$ , which determine how many Lyapunov exponents can cross zero as non-normality increases. A key step is to characterize the geometry of the corresponding invariant subspaces, showing how the left Oseledets vectors associated with  $\lambda_i < 0$  define cointegrating combinations and how their orientation fluctuates with the realizations of  $\mathbf{A}_t$ . Finally, one must quantify how power law tails emerge in the stable block, where  $\Pr(\|\mathbf{A}_t|_E\| > 1) > 0$ , and how the tail exponent  $\alpha$  depends on the statistics of non-normal amplification.

Empirically, this perspective can be tested with rolling VAR estimations on multivariate financial time series. One first estimates sequences of interaction matrices  $\hat{\mathbf{A}}_t$  and computes their instantaneous condition numbers  $\kappa(\hat{\mathbf{A}}_t)$  to quantify non-normality. Next, the empirical Lyapunov spectrum

$\hat{\lambda}_1 \geq \dots \geq \hat{\lambda}_N$  can be evaluated using QR-based orthogonalization on products of  $\hat{\mathbf{A}}_t$ . The cointegration rank is then identified as  $N - r + 1 = \#\{i : \hat{\lambda}_i < 0\}$ , and the corresponding left Oseledets vectors are extracted as empirical cointegrating directions. Finally, one can test for Kesten-type heavy tails in  $\mathbf{b}^\top \mathbf{x}_t$  for  $\mathbf{b}$  in this subspace, verifying that the estimated power law exponents  $\hat{\alpha}$  are consistent with the theoretical predictions.

This stochastic formulation reframes cointegration as a dynamic property of systems near pseudocriticality: Global stability survives in most directions, but non-normal coupling creates transient instabilities that translate into persistent trends in a few directions. As a result, the coexistence of stationary and nonstationary modes, often treated as a statistical curiosity in econometrics, emerges here as an intrinsic and robust feature of high-dimensional multiplicative dynamics. Cointegration thus appears not as an exceptional equilibrium condition, but as the natural *signature of partial stability* in a non-normal stochastic system.

### E. Factor models and dimensional reduction in financial systems

In the previous section, we highlighted the connection between the cointegration formalism and the  $N$ -dimensional Kesten process. Both the derivation and the numerical illustrations showed that, as the system dimension increases, the power law exponent of the stationary distribution decreases and the Lyapunov exponent increases.

This behavior can be traced back to results from EVT, which explain how the expected growth of the condition number scales with  $N$ . Since financial markets consist of thousands of interacting assets, one might therefore expect that, in such high-dimensional settings, the system would be generically unstable and characterized by power laws with exponents arbitrarily close to zero. Yet, this is not what empirical data show. One piece of resolution lies in the fact that the effective number of degrees of freedom in financial markets is much smaller than the raw number of assets. In practice, a handful of dominant factors govern the dynamics, inducing a strong dimensional reduction in the system.

A widely used framework to formalize this idea is the *factor model* [56,57]. The central assumption is that the returns of  $N$  assets, collected in the vector  $\mathbf{r}_t \in \mathbb{R}^N$ , can be represented through a reduced set of  $K \ll N$  latent factors:

$$\mathbf{r}_t = \mathbf{B}\mathbf{f}_t + \mathbf{e}_t, \quad (143)$$

where  $\mathbf{f}_t \in \mathbb{R}^K$  denotes the factor realizations,  $\mathbf{B} \in \mathbb{R}^{N \times K}$  is the loading matrix, and  $\mathbf{e}_t$  contains idiosyncratic components with weak cross-sectional correlation. Within this framework, the return covariance structure simplifies to

$$\Sigma_r = \mathbf{B}\Sigma_f\mathbf{B}^\top + \Sigma_e, \quad (144)$$

making estimation tractable even when  $N$  is very large. Standard implementations emphasize symmetric covariance matrices and typically neglect asymmetric interactions, as the main objective is to capture comovements rather than directional feedback loops.

Factor models are not only a theoretical tool but are routinely employed in industry and policy institutions. Banks, hedge funds, and central banks rely on factor-based and

VAR-type models for risk management, stress testing, and monetary policy analysis [58,59]. In practice, large-scale factor models are often implemented as VAR systems, where the dynamic interaction among factors (or asset groups) is explicitly modeled through a stochastic matrix of coefficients. Since a VAR process is mathematically equivalent to a linear recursion, the time-varying structure of the VAR matrix can be mapped naturally to a Kesten process. This perspective highlights that industry-standard tools already rely implicitly on a stochastic-matrix formalism, and our framework extends this by allowing for asymmetry and non-normal amplification in the dynamics.

In the stochastic Kesten representation, factor dynamics are governed by random matrices  $\mathbf{A}_t$ :

$$\mathbf{f}_{t+1} = \mathbf{A}_t\mathbf{f}_t + \boldsymbol{\eta}, \quad (145)$$

where  $\mathbf{A}_t$  may be nonsymmetric and non-normal. From our  $N$ -dimensional analysis, we found that the critical standard deviation of the logarithm of the singular values scales as  $\sqrt{\ln N}$  (117), i.e.,  $\sigma_c(N) \propto \sqrt{\ln N}$ . This means that as the dimensionality of the system grows, stronger heterogeneity in singular values is required to approach criticality ( $\gamma \approx 0$ ) and hence to observe heavy-tailed fluctuations. However, these results only hold in the “highly” non-normal regime. We observed that, for intermediate non-normality, increasing the system dimension enhances the Lyapunov exponent and reduces the tail exponent of the distribution. Low dimensions, e.g.,  $N = 3$ , might be sufficient to give rise to the inverse-cubic law for the tail exponent of return distributions [44]; and to push the Lyapunov exponents toward zero.

Factor models achieve an implicit dimensional reduction: The dynamics of returns unfold not in the full  $N$ -dimensional space, but within a low-dimensional subspace of dimension  $K \ll N$ . It is within this reduced subspace that collective effects emerge, such as power law tails and market-wide instabilities.

Empirical studies have documented that, during bubbles and crashes, the effective dimensionality of financial markets tends to collapse onto a single dominant factor, often interpreted as a “market mode” [60,61]. This means that the covariance spectrum becomes concentrated, and nearly all assets move together.

This interpretation offers several strengths. It provides a unified perspective by linking classical econometric models, such as factor structures and covariance estimation, to dynamical systems with stochastic matrices. It also explains crisis phenomena by naturally accounting for why crises are characterized by both increased comovement and extreme fluctuations: During such episodes, factor collapse reduces the system’s effective dimensionality, leading to dynamics governed by a low-dimensional Kesten process.

At the same time, there are important limitations. A key interpretational challenge arises because factor models are statistical constructs, and connecting them directly to the underlying feedback matrices  $\mathbf{A}_t$  requires assumptions that may be difficult to test. The framework also entails a neglect of microstructure, as real markets include frictions, strategic interactions, and institutional behaviors not captured by linear Kesten-type recursions. Finally, there remains an empirical controversy: Although factor collapse is widely observed, it

is still debated whether it represents a cause or merely a symptom of financial crises.

Overall, embedding factor models within the Kesten process framework suggests that dimensional reduction is not only a statistical convenience but also a source of systemic fragility. In particular, the observed market tendency to align along a single factor during turbulent periods can be reinterpreted as a mechanism driving the system toward non-normal criticality, thus providing another application of our formalism to financial instability.

### F. Wealth inequality in an $N$ -agent Kesten process

We now sketch how the  $N$ -dimensional Kesten framework can be applied to model wealth inequality. Two distinct literatures motivate this link. On the one hand, econophysicists have long modeled wealth dynamics by multiplicative random processes and interacting-agent models [4,35,62–64], showing how power law tails and wealth condensation emerge from such dynamics. On the other hand, political economy arguments [65] emphasize systematic differences between returns to capital ( $r$ ) and income growth ( $g$ ) as drivers of inequality: When the rate of return on capital ( $r$ ) exceeds the growth rate of the economy ( $g$ ), capital owners accumulate relatively faster than wages, generating an inexorable rise in capital share and inequality. More recent political-economic diagnostics (often labeled “technofeudalism” in contemporary debates) stress technological/platform rents as mechanisms of concentration [66,67].

Consider an economy of  $N$  agents. Let  $w_{i,t}$  denote agent  $i$ ’s wealth at time  $t$ . A very general linearized interaction model that captures trade, investment returns, and redistribution is

$$\mathbf{w}_{t+1} = \mathbf{A}_t \mathbf{w}_t + \boldsymbol{\eta}_t, \quad (146)$$

where  $\mathbf{w}_t = (w_{1,t}, \dots, w_{N,t})^\top$ ,  $\mathbf{A}_t \in \mathbb{R}^{N \times N}$  encodes how wealth flows, amplifies or transfers between agents (via returns, trades, credit, dividends, rents, etc.), and  $\boldsymbol{\eta}_t$  is a source vector representing exogenous inflows (income, wages, natural resource rents, subsidy, or capital extraction from the “ground”). Equation (146) is precisely an  $N$ -dimensional Kesten-type recursion: Multiplicative (matrix) dynamics plus additive inputs. Under mild assumptions on  $\{\mathbf{A}_t, \boldsymbol{\eta}_t\}$ , the stationary distribution of  $\mathbf{w}_t$  (when it exists) develops Pareto tails.

Several economically meaningful mechanisms map naturally to the entries of  $\mathbf{A}_t$  and to  $\boldsymbol{\eta}_t$ :

(1)  $\mathbf{A}_t$  diagonal entries can represent idiosyncratic capital returns (multiplicative growth of own wealth), while off-diagonal entries represent transfers, credit exposure, network spillovers, profit sharing, or predatory extraction.

(2) Non-normality (large condition numbers, nonorthogonal eigenvectors) corresponds to amplification directions: shocks aligned with certain transiently amplifying modes are boosted far more than classical spectral analysis would predict.

(3)  $\boldsymbol{\eta}_t$  represents capital entering the system (wages, profits extracted from the “ground,” fiscal transfers). The spatial pattern of  $\boldsymbol{\eta}_t$  (who receives the capital) substantially affects inequality dynamics.

This representation makes clear that asymmetry in interactions (nonsymmetric  $\mathbf{A}_t$ )—e.g., rent extraction by platforms, monopolistic pass throughs, or strongly asymmetric network positions—is a plausible, and empirically relevant, source of wealth concentration. Interacting-agent models [35] show how multiplicative stochastic exchange with weak redistribution produces Pareto tails and, in some regimes, “wealth condensation” (a small fraction of agents holding most of the mass).

A common (and intuitive) route from multiplicative dynamics to Zipf/Pareto laws is to assume a form of stationarity or ergodicity that links cross-sectional and temporal distributions. Concretely,

(1) if an agent’s wealth evolves multiplicatively and additive shocks are small relative to multiplicative fluctuations, Kesten-type results show the stationary tail of  $w$  is Pareto:  $\mathbb{P}(w > x) \sim Cx^{-\alpha}$ ;

(2) if one posits as a modeling assumption that the cross-sectional distribution at any time is statistically similar to the empirical time distribution of a typical agent’s wealth path, then the observed cross section inherits the same Pareto tail (this is a form of equivalence between ensemble and time averages used in many econophysics arguments).

This yields the ubiquitous observation of Zipf-like tails for top wealth. Note that this is an additional modeling assumption (ensemble/time equivalence), which needs careful empirical validation, but it is commonly invoked to justify why single-agent multiplicative rules produce cross-sectional Pareto tails.

Our Kesten-based mechanism, formulated in Eq. (146), differs in several important ways. It is inherently stochastic and network-driven, as inequality arises from random multipliers and asymmetric interactions among agents. Even when the average return is small or the overall economy is contracting, stochastic amplification through  $\mathbf{A}_t$  can generate extremely large outcomes for a few agents. It also incorporates sudden reconfiguration and non-normality: Network switches, which are rapid changes in  $\mathbf{A}_t$ , or temporary alignments of shocks with amplifying directions, can cause abrupt and large deviations for certain nodes, providing a natural explanation for episodes of rapid wealth concentration such as bubbles and crashes. Finally, it spans a broad range of regimes: While classical models account for long-run secular trends, the Kesten approach captures heavy tails and large fluctuations over short to medium horizons, and can also reproduce persistent concentration when the dynamics repeatedly favor the same agents.

An important implication of our  $N$ -dimensional analysis is that the threshold variance of the log-singular values needed to push the Lyapunov exponent to zero scales like  $\sqrt{\ln N}$ . Equivalently, lower effective dimension makes the system more fragile: If the effective number of independent channels of exchange or diversification collapses (e.g., due to a dominant platform, a single “winner” technology, or a coordinated market momentum), the economy becomes more susceptible to non-normal amplification and to the emergence of extreme wealth concentration even for moderate levels of heterogeneity.

This observation connects to empirical claims about crisis dynamics and “factor collapse” in markets: When economic

relations become dominated by a single factor or platform, the effective interaction dimension drops and the system becomes more prone to large endogenous inequalities (a “winner-takes-most” dynamic). In the context of modern digital rents (the “technofeudalism” diagnostic), platform dominance can reduce effective competition and diversify channels, thus mechanically lowering the barrier for concentration.

Because the Kesten mechanism is driven by multiplicative products of random matrices, a few extreme realizations of the product can create huge outliers (very large wealth) even when the average growth is nonpositive. That is, heavy tails and “condensation” are driven by multiplicative amplification events rather than by steady positive drift. Consequently, a contracting or stagnating economy can still display very large inequality if the interaction structure occasionally aligns to amplify particular nodes. This complements the long-run drift view and helps explain episodic and severe concentration events (technology windfalls, platform monopolies, rent extraction) without requiring sustained  $r > g$ .

The Kesten-network interpretation implies several empirical signatures. First, it predicts a heavy-tailed wealth distribution characterized by a Pareto exponent that may vary over time and respond to measures of network non-normality, such as condition numbers or the concentration of singular vectors. Second, it anticipates a correlation between concentration and effective dimension: periods of low effective dimension, indicated, for example, by a dominant eigenvalue in the cross-asset or cross-firm covariance, should coincide with greater upper-tail concentration of wealth or firm size. Third, it suggests the presence of transient amplification events, where abrupt wealth jumps of particular agents align with measurable reconfigurations of  $\mathbf{A}_t$ , such as mergers, platform launches, regulatory arbitrage, or major product successes. Testing these predictions will require (1) time-resolved network or flow estimates to construct  $\mathbf{A}_t$ , (2) careful inference of effective dimension or singular-value dispersion, and (3) detailed panel data on individual wealth or firm size.

## VI. CONCLUSION

We have developed a general analysis of multidimensional Kesten processes [1] to understand the emergence of power law statistics in globally stable high-dimensional stochastic systems. Our key result is that heavy-tailed stationary distributions can arise even in the absence of spectral instability, through a collective mechanism rooted in the geometry of non-normal interactions.

At the theoretical level, we derived explicit scaling relations linking the Lyapunov exponent  $\gamma$  and the tail exponent  $\alpha$  to statistical properties of the eigenvector geometry, quantified by the condition number  $\kappa$  of the random multiplicative matrices. These relations,

$$\gamma \sim \gamma_0 + \ln \kappa, \quad \alpha \sim -2\gamma/\sigma_\kappa^2, \quad (147)$$

where  $\gamma_0$  is the Lyapunov exponent in the normal limit and we assumed stability (spectrally stable eigenvalues inside the unit circle), demonstrate that non-normality renormalizes the effective stability and tail behavior of the system. In this picture, transient eigenvector alignment produces bursts of amplification that mimic the effects of criticality, giving rise to stationary power laws even

when all eigenvalues lie strictly within the unit circle. Numerical simulations confirm these analytical predictions, showing that the heavy-tailed regime and enhanced Lyapunov growth are accurately captured by the statistics of  $\ln \kappa$  alone.

Beyond our mathematical formulation, we have illustrated the broad applicability of this framework across physics, finance, and economics. In fluid dynamics, non-normal amplification underlies the transient stretching of polymers in turbulent flows and the kinematic stage of small-scale magnetic dynamos, where geometric alignment of the velocity-gradient eigenvectors produces intermittent energy bursts. In finance and economics, the same mechanism explains fat tails and volatility clustering in ARCH and VAR models, as well as inequality amplification in interacting-agent economies. Across all these systems, nonorthogonal interactions act collectively to amplify fluctuations, producing power laws as a generic outcome of high-dimensional multiplicative dynamics.

This unified perspective highlights that heavy tails and large deviations need not signal proximity to a critical point. Instead, they can emerge far from criticality, as a collective consequence of non-normal geometry and transient amplification. In this sense, non-normal Kesten processes define a universality class for stochastic systems: One where global stability coexists with local bursts of instability and where multiplicative feedback and eigenvector alignment generate critical-like statistics without spectral criticality.

From a broader standpoint, these results provide a bridge between distinct research traditions, from the theory of non-modal hydrodynamic instabilities [12,13,18] to econometric models of stochastic volatility and macroeconomic growth. They suggest that the same underlying principles govern amplification, intermittency, and heavy tails across a wide spectrum of natural and socioeconomic systems. Whether the system describes turbulent eddies, magnetic fields, or market returns, the governing equations share the same algebraic backbone: a stochastic sequence of non-normal transformations whose geometry organizes the emergence of collective fluctuations.

More generally, our results show that the matrix-product and  $L_2$ -norm formulation of Kesten theory, while mathematically synthetic and powerful, does not by itself identify the dynamical source of amplification. The present work complements this formalism by showing when growth is produced by spectral intermittency, and when it is instead generated by non-normal geometry, reinjection, and mode switching. This distinction is essential for the physical interpretation of multiplicative systems across disciplines.

Ultimately, the framework introduced here reframes heavy-tailed statistics as a manifestation of cooperative dynamics rather than of criticality. It explains how complex systems can remain globally stable while exhibiting persistent power law fluctuations, linking the mathematics of non-normal amplification to the phenomenology of turbulence, dynamo action, and economic inequality. By unifying these domains under a single theoretical mechanism, our results point to a broader principle: Collective power law behavior is not an exception arising near critical points, but a natural feature of high-dimensional systems governed by multiplicative and non-normal interactions.

## ACKNOWLEDGMENTS

D.S. was partially supported by the National Natural Science Foundation of China (Grants No. T2350710802 and No. U2039202), Shenzhen Science and Technology Innovation Commission Project (Grants No. GJHZ20210705141805017 and No. K23405006), and the Center for Computational Science and Engineering at Southern University of Science and Technology.

## DATA AVAILABILITY

The numerical data generated for the figures are available from the authors upon reasonable request.

## APPENDIX: NUMERICAL TOOLS

### 1. Numerical estimation of Lyapunov exponents

In this Appendix, we briefly describe the numerical method used to estimate the Lyapunov exponents of products of random matrices, which is based on the classical *QR reorthonormalization* technique introduced by Benettin *et al.* [68,69]. This procedure is significantly more stable and accurate than a naive “brute force” approach.

*Brute force estimation.* A straightforward way to approximate the largest Lyapunov exponent  $\gamma$  is to iterate

$$\mathbf{\Pi}_t = \prod_{s=1}^t \mathbf{A}_s, \quad \gamma \approx \frac{1}{t} \ln \|\mathbf{\Pi}_t\|, \quad (\text{A1})$$

for large  $t$ . However, this method is numerically unstable: The norm  $\|\mathbf{\Pi}_t\|$  grows (or decays) exponentially fast, leading to rapid loss of numerical precision. In addition, if one is interested in the *full spectrum* of exponents (not just the largest), the brute force method is not applicable, since it only tracks the growth rate along a single direction.

*QR reorthonormalization (Benettin method).* The Benettin method overcomes these issues by evolving an orthonormal basis of vectors instead of a single one. At each iteration, the random matrix  $\mathbf{A}_t$  is applied to the current basis, and the resulting set of vectors is reorthonormalized via a QR decomposition:

$$\mathbf{A}_t \mathbf{Q}_{t-1} = \mathbf{Q}_t \mathbf{R}_t, \quad (\text{A2})$$

where  $\mathbf{Q}_t$  is an orthogonal matrix and  $\mathbf{R}_t$  is upper triangular. The diagonal elements of  $\mathbf{R}_t$  encode the local stretching factors of the basis vectors. By accumulating the logarithms of these diagonal entries over time, we obtain estimates of the Lyapunov exponents:

$$\gamma_i = \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t \ln |(R_s)_{ii}|, \quad i = 1, \dots, N. \quad (\text{A3})$$

This procedure has several advantages:

(1) It maintains numerical stability, since the vectors are continuously reorthonormalized and do not collapse into the dominant eigendirection.

(2) It provides the *entire Lyapunov spectrum*, not just the maximal exponent.

(3) Averaging over multiple independent realizations (ensembles) further reduces statistical noise and yields estimates of standard errors.

*Practical considerations.* In practice, the algorithm proceeds as follows:

(1) Initialize with an orthonormal basis  $\mathbf{Q}_0 = I$ .

(2) For each time step, multiply the basis by a random matrix  $\mathbf{A}_t$  and apply QR decomposition to reorthonormalize.

(3) Accumulate the logarithms of the diagonal entries of  $\mathbf{R}_t$ , and normalize by the total time horizon  $t_{\max}$  to estimate the exponents.

(4) Optionally, repeat over multiple ensembles to improve statistical accuracy and compute error bars.

This method is therefore more reliable than brute force multiplication, particularly in high-dimensional settings or when long time horizons are required. Throughout our numerical analysis, we employ this QR reorthonormalization scheme to compute the Lyapunov spectrum of Kesten processes, and we refer the reader to this Appendix for methodological details.

### 2. Numerical estimation of the tail exponent

In this Appendix, we describe the methodology used to estimate the tail exponent  $\alpha$  of heavy-tailed distributions. Naively, one might attempt to fit a straight line in log-log coordinates to the survival function

$$\mathbb{P}[X > x] \sim Cx^{-\alpha}, \quad (\text{A4})$$

and take the slope as an estimate of  $\alpha$ . However, such linear regression on doubly logarithmic plots is statistically inefficient and strongly biased: The choice of fitting range is arbitrary, correlations between points are ignored, and finite-sample fluctuations often dominate the visual slope. For these reasons, modern approaches rely instead on maximum-likelihood estimation (MLE) combined with goodness-of-fit diagnostics. In our work, we adopt the Clauset-Shalizi-Newman (CSN) method [70], which has become a standard in the analysis of heavy-tailed data.

*Pareto maximum-likelihood estimator.* Suppose we fix a threshold  $x_{\min}$  above which the data are assumed to follow a Pareto law. Given  $N$  observations  $\{x_i \geq x_{\min}\}$ , the log likelihood of a Pareto tail is maximized by

$$\hat{\alpha}(x_{\min}) = \frac{N}{\sum_{i=1}^N \ln \left( \frac{x_i}{x_{\min}} \right)}. \quad (\text{A5})$$

This estimator is unbiased in the limit  $N \rightarrow \infty$ , and is statistically more efficient than slope-fitting in log-log space.

*Choosing the threshold.* A central difficulty is the choice of  $x_{\min}$ : too low and the data deviate from a pure power law, too high and too few points remain in the tail. The CSN method selects  $x_{\min}$  by minimizing the Kolmogorov-Smirnov (KS) distance between the empirical cumulative distribution function (CDF) of the tail and the fitted Pareto CDF. This balances goodness-of-fit with statistical power.

*Goodness-of-fit and uncertainty.* To assess whether the Pareto model is plausible, the CSN method uses a bootstrap

procedure: Synthetic datasets are generated from the fitted model and refitted in the same way, yielding a distribution of KS statistics. The  $p$  value is the fraction of synthetic datasets with KS statistic at least as large as the empirical one. This provides a rigorous test of the hypothesis “the tail follows a power law.” In addition, confidence intervals for  $\alpha$  can be obtained from the bootstrap samples.

*Error on the estimate.* Following Ref. [70], we know that the error on the fit is given by

$$\text{Std}(\hat{\alpha}) = \frac{\hat{\alpha}}{\sqrt{n_{\text{tail}}}} + \mathcal{O}\left(\frac{1}{n_{\text{tail}}}\right), \quad (\text{A6})$$

where  $n_{\text{tail}}$  is the number of observations in the tail (i.e., those exceeding  $x_{\min}$ ). In our case, we decided to use multiple realizations  $\hat{\alpha}$  across different simulations, and we used the average estimation among the estimated  $\hat{\alpha}$ . Therefore, rather than using the error on a single realization, we measured the standard deviation among all the observations.

*Alternative estimators.* Another widely used approach is the Hill estimator, which directly uses the top  $k$  order

statistics:

$$\hat{\alpha}_{\text{Hill}}(k) = \frac{k}{\sum_{i=1}^k \ln\left(\frac{x_{(i)}}{x_{(k)}}\right)}, \quad (\text{A7})$$

where  $x_{(i)}$  denotes the  $i$ th largest observation. While useful for exploratory analysis, it requires careful tuning of  $k$  and can be sensitive to sample variability. In contrast, the CSN method automates the threshold selection and provides statistical tests.

*Summary.* In summary, our estimation procedure consists of the following:

- (1) Sorting the data and considering candidate thresholds  $x_{\min}$ .
- (2) For each threshold, estimating  $\alpha$  by MLE and computing the KS statistic.
- (3) Selecting the  $x_{\min}$  that minimizes the KS distance.
- (4) Validating the fit and estimating confidence intervals via bootstrap resampling.

This methodology avoids the pitfalls of linear regression in log-log space, yields statistically principled estimates of the tail exponent, and allows us to assess both parameter uncertainty and goodness of fit. In the numerical analysis, all reported values of  $\alpha$  are obtained using this CSN framework.

- 
- [1] H. Kesten, Random difference equations and renewal theory for products of random matrices, *Acta Math.* **131**, 207 (1973).
  - [2] C. de Calan, J. M. Luck, T. M. Nieuwenhuizen, and D. Petritis, On the distribution of a random variable occurring in 1D disordered systems, *J. Phys. A: Math. Gen.* **18**, 501 (1985).
  - [3] H. Takayasu, A. H. Sato, and M. Takayasu, Stable infinite variance fluctuations in randomly amplified Langevin systems, *Phys. Rev. Lett.* **79**, 966 (1997).
  - [4] A. Blank and S. Solomon, Power laws in cities population, financial markets and internet sites (scaling in systems with a variable number of components), *Physica A* **287**, 279 (2000).
  - [5] A. Statman, M. Kaufman, A. Minerbi, N. E. Ziv, and N. Brenner, Synaptic size dynamics as an effectively stochastic process, *PLoS Comput. Biol.* **10**, e1003846 (2014).
  - [6] T. Lux and D. Sornette, On rational bubbles and fat tails, *J. Money Credit Bank.* **34**, 589 (2002).
  - [7] T. Lux and S. Alfarano, Financial power laws: Empirical evidence, models, and mechanisms, *Chaos, Solitons & Fractals* **88**, 3 (2016).
  - [8] X. Gabaix, Power laws in economics and finance, *Annu. Rev. Econ.* **1**, 255 (2009).
  - [9] E. Derman, K. S. Park, and W. Whitt, A stochastic-difference-equation model for hedge-fund returns, *Quant. Financ.* **10**, 701 (2010).
  - [10] D. Sornette, Linear stochastic dynamics with nonlinear fractal properties, *Physica A* **250**, 295 (1998).
  - [11] J. Benhabib, A. Bisin, and M. Luo, Wealth distribution and social mobility in the US a quantitative approach, *Am. Econ. Rev.* **109**, 1623 (2019).
  - [12] L. N. Trefethen, A. E. Trefethen, S. C. Reddy, and T. A. Driscoll, Hydrodynamic stability without eigenvalues, *Science* **261**, 578 (1993).
  - [13] B. F. Farrell and P. J. Ioannou, Generalized stability theory. Part I: Autonomous operators, *J. Atmos. Sci.* **53**, 2025 (1996).
  - [14] J.-M. Chomaz, Global instabilities in spatially developing flows: Non-normality and nonlinearity, *Annu. Rev. Fluid Mech.* **37**, 357 (2005).
  - [15] P. J. Schmid, Nonmodal stability theory, *Annu. Rev. Fluid Mech.* **39**, 129 (2007).
  - [16] S. C. Reddy and D. S. Henningson, Energy growth in viscous channel flows, *J. Fluid Mech.* **252**, 209 (1993).
  - [17] M. Embree and L. N. Trefethen, *Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators* (Princeton University Press, Princeton, 2005).
  - [18] B. F. Farrell and P. J. Ioannou, Generalized stability theory. Part II: Non-autonomous operators, *J. Atmos. Sci.* **53**, 2041 (1996).
  - [19] B. F. Farrell and P. J. Ioannou, Structural stability of turbulent jets, *J. Atmos. Sci.* **60**, 2101 (2003).
  - [20] V. Troude and D. Sornette, Unifying framework for amplification mechanisms: Criticality, resonance and non-normality, *Phys. Rev. Res.* **7**, L042048 (2025).
  - [21] S. Allesina and S. Tang, The stability–complexity relationship at age 40: A random matrix perspective, *Popul. Ecol.* **57**, 63 (2015).
  - [22] M. Asllani, R. Lambiotte, and T. Carletti, Structure and dynamical behavior of non-normal networks, *Sci. Adv.* **4**, eaau9403 (2018).
  - [23] H. Furstenberg and H. Kesten, Products of random matrices, *Ann. Math. Statist.* **31**, 457 (1960).
  - [24] D. Sornette, Multiplicative processes and power laws, *Phys. Rev. E* **57**, 4811 (1998).
  - [25] D. Sornette, *Critical Phenomena in Natural Sciences: Chaos, Fractals, Selforganization and Disorder: Concepts and Tools*,

- Springer Series in Synergetics*, 2nd ed. (Springer, Berlin, Heidelberg, 2006).
- [26] R. A. Horn and C. R. Johnson, *Matrix Analysis*, 2nd ed. (Cambridge University Press, Cambridge, 2013).
- [27] O. Perron, Die stabilitätsfrage bei differentialgleichungen, *Math. Z.* **32**, 703 (1930).
- [28] G. A. Leonov and N. V. Kuznetsov, Time-varying linearization and the Perron effects, *Int. J. Bifurcation Chaos* **17**, 1079 (2007).
- [29] I. N. Sergeev, Definition and properties of Perron stability of differential systems, in *Proceedings of the International Workshop QUALITDE—2019* (Tbilisi, Georgia, 2019), p. 162.
- [30] M. Chertkov, Polymer stretching by turbulence, *Phys. Rev. Lett.* **84**, 4761 (2000).
- [31] A. Celani, A. Puliafito, and D. Vincenzi, Dynamical slowdown of polymers in laminar and random flows, *Phys. Rev. Lett.* **97**, 118301 (2006).
- [32] S. Musacchio, V. Steinberg, and D. Vincenzi, Polymer stretching in laminar and random flows: Entropic characterization, *Phys. Rev. Fluids* **8**, 053301 (2023).
- [33] J. M. Ottino, *The Kinematics of Mixing: Stretching, Chaos, and Transport* (Cambridge University Press, Cambridge, England, 1989).
- [34] G. Falkovich, K. Gawędzki, and M. Vergassola, Particles and fields in fluid turbulence, *Rev. Mod. Phys.* **73**, 913 (2001).
- [35] J.-P. Bouchaud and M. Mézard, Wealth condensation in a simple model of economy, *Physica A* **282**, 536 (2000).
- [36] J. Ginibre, Statistical ensembles of complex, quaternion, and real matrices, *J. Math. Phys.* **6**, 440 (1965).
- [37] J. T. Chalker and B. Mehlh, Eigenvector statistics in non-Hermitian random matrix ensembles, *Phys. Rev. Lett.* **81**, 3367 (1998).
- [38] P. Bourgade and G. Dubach, The distribution of overlaps between eigenvectors of Ginibre matrices, *Probab. Theory Relat. Fields* **177**, 397 (2020).
- [39] E. Balkovsky, A. Fouxon, and V. Lebedev, Turbulent dynamics of polymer solutions, *Phys. Rev. Lett.* **84**, 4765 (2000).
- [40] A. P. Kazantsev, Enhancement of a magnetic field by a conducting fluid, *Sov. Phys. JETP* **26**, 1031 (1968).
- [41] A. A. Schekochihin, S. C. Cowley, G. W. Hammett, J. L. Maron, and J. C. McWilliams, A model of nonlinear evolution and saturation of the turbulent MHD dynamo, *New J. Phys.* **4**, 84 (2002).
- [42] R. Engle, Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation, *Econometrica* **50**, 987 (1982).
- [43] T. Mikosch and C. Stărică, ARCH and GARCH models: Properties, processes and examples, *Encycl. Actuar. Sci.* **1**, 93 (2004).
- [44] P. Gopikrishnan, V. Plerou, L. Amaral, M. Meyer, and H. Stanley, Scaling of the distribution of fluctuations of financial market indices, *Phys. Rev. E* **60**, 5305 (1999).
- [45] T. Bollerslev, Generalized autoregressive conditional heteroskedasticity, *J. Econom.* **31**, 307 (1986).
- [46] T. Bollerslev, A conditionally heteroskedastic time series model for speculative prices and rates of return, *Rev. Econ. Stat.* **69**, 542 (1987).
- [47] R. Baillie, T. Bollerslev, and H. Mikkelsen, Fractionally integrated generalized autoregressive conditional heteroskedasticity, *J. Econom.* **74**, 3 (1996).
- [48] S. Heston, A closed-form solution for options with stochastic volatility with applications to bond and currency options, *Rev. Financ. Stud.* **6**, 327 (1993).
- [49] P. Blanc, J. Donier, and J.-P. Bouchaud, Quadratic Hawkes processes for financial prices, *Quant. Financ.* **17**, 171 (2017).
- [50] R. F. Engle and C. W. Granger, Co-integration and error correction: Representation, estimation, and testing, *Econometrica* **55**, 251 (1987).
- [51] S. Johansen, Statistical analysis of cointegration vectors, *J. Econ. Dyn. Control* **12**, 231 (1988).
- [52] V. I. Oseledets, A multiplicative ergodic theorem: Lyapunov characteristic numbers for dynamical systems, *Trans. Mosc. Math. Soc.* **19**, 197 (1968).
- [53] L. Arnold, *Random Dynamical Systems* (Springer, Berlin, 1998).
- [54] F. Ginelli, P. Poggi, A. Turchi, H. Chaté, R. Livi, and A. Politi, Characterizing dynamics with covariant Lyapunov vectors, *Phys. Rev. Lett.* **99**, 130601 (2007).
- [55] C. L. Wolfe and R. M. Samelson, An efficient method for recovering Lyapunov vectors from singular vectors, *Tellus A: Dyn. Meteorol. Oceanogr.* **59**, 355 (2007).
- [56] S. A. Ross, The arbitrage theory of capital asset pricing, *J. Econ. Theory* **13**, 341 (1976).
- [57] E. F. Fama and K. R. French, Common risk factors in the returns on stocks and bonds, *J. Financ. Econ.* **33**, 3 (1993).
- [58] J. H. Stock and M. W. Watson, Vector autoregressions, *J. Econ. Perspect.* **15**, 101 (2001).
- [59] C. A. Sims, Macroeconomics and reality, *Econometrica* **48**, 1 (1980).
- [60] L. Laloux, P. Cizeau, J.-P. Bouchaud, and M. Potters, Noise dressing of financial correlation matrices, *Phys. Rev. Lett.* **83**, 1467 (1999).
- [61] V. Plerou, P. Gopikrishnan, B. Rosenow, L. A. Amaral, and H. E. Stanley, Random matrix approach to cross correlations in financial data, *Phys. Rev. E* **65**, 066126 (2002).
- [62] M. Levy and S. Solomon, New evidence for the power-law distribution of wealth, *Physica A* **242**, 90 (1997).
- [63] S. Solomon and P. Richmond, Power laws of wealth, market order volumes and market returns, *Physica A* **299**, 188 (2001).
- [64] S. Solomon and P. Richmond, Stable power laws in variable economies; Lotka–Volterra implies Pareto–Zipf, *Eur. Phys. J. B* **27**, 257 (2002).
- [65] T. Piketty, *Capital in the Twenty-First Century* (Harvard University Press, Cambridge, MA, 2014).
- [66] Y. Varoufakis, *Technofeudalism: What Killed Capitalism* (Melville House, Brooklyn, NY, 2024).
- [67] C. Durand, *How Silicon Valley Unleashed Techno-Feudalism* (Verso, London, 2024).
- [68] G. Benettin, L. Galgani, A. Giorgilli, and J.-M. Strelcyn, Lyapunov characteristic exponents for smooth dynamical systems and for Hamiltonian systems; a method for computing all of them. Part 1: Theory, *Meccanica* **15**, 9 (1980).
- [69] G. Benettin, L. Galgani, A. Giorgilli, and J.-M. Strelcyn, Lyapunov characteristic exponents for smooth dynamical systems and for Hamiltonian systems; a method for computing all of them. Part 2: Numerical application, *Meccanica* **15**, 21 (1980).
- [70] A. Clauset, C. R. Shalizi, and M. E. J. Newman, Power-law distributions in empirical data, *SIAM Rev.* **51**, 661 (2009).