

LYAPUNOV EXPONENTS FROM RANDOM FIBONACCI
SEQUENCES TO THE LORENZ EQUATIONS

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Divakar Viswanath

August 1998

© Divakar Viswanath 1998

ALL RIGHTS RESERVED

LYAPUNOV EXPONENTS FROM RANDOM FIBONACCI SEQUENCES TO
THE LORENZ EQUATIONS

Divakar Viswanath, Ph.D.

Cornell University 1998

Lyapunov exponents give a way to capture the central features of chaos and of stability in both deterministic and stochastic systems using just a few real numbers. However, exact analytic determination of Lyapunov exponents is rarely possible, and as we will show, even an accurate numerical computation is not a trivial task.

One of the principal results of this thesis is about *random Fibonacci sequences*. Random Fibonacci sequences are defined by $t_1 = t_2 = 1$ and $t_n = \pm t_{n-1} \pm t_{n-2}$ for $n > 2$, where each \pm sign is independent and either $+$ or $-$ with probability $1/2$. Using Stern-Brocot sequences, we prove that

$$\sqrt[n]{|t_n|} \rightarrow 1.13198824\dots \quad \text{as } n \rightarrow \infty$$

with probability 1.

Other contributions of this thesis include formulas for condition numbers of random triangular matrices and an accurate computation of the Lyapunov exponents of the Lorenz equations.

Biographical Sketch

Divakar Viswanath was born in India 27 years ago. The first three acts of his life were played out in Bangalore, a city that he greatly loves; in Mumbai, that sea of humanity; and then in sometimes white, sometimes grey, sometimes green Ithaca, a place that at first seemed a bit uninhabited to him. He awaits the future with hope and suspense. It must have been the summer of 1986 when the slim book *Inequalities* by P.P. Korovkin fell into his hands. Too proud to give up, he gasped his way through the book and won a faint appreciation of a beautiful inequality sometimes attributed to Lyapunov:

$$\left(\frac{x_1^\alpha + \cdots + x_n^\alpha}{n}\right)^{1/\alpha} \leq \left(\frac{x_1^\beta + \cdots + x_n^\beta}{n}\right)^{1/\beta}$$

for $x_i > 0$ and $\alpha < \beta$. He notes with pleasure that this inequality plays a role in this thesis 12 eventful years later.

Acknowledgements

Grateful thanks are due to Nick Trefethen. I have benefited immensely from his maturity, his generosity, and his exceptional scientific taste. Nick has directed me with intelligence, patience, and interest, and in the process helped cure me of several delusions about research.

Discussions with Robert Strichartz about various aspects of this thesis were very helpful. I thank him for much encouragement.

Steve Vavasis graciously agreed to read this thesis at the last minute.

My meetings with Persi Diaconis, though few in number, were invaluable; during occasional conversations with him, I got a sense of what it means to “internalize” a piece of math.

Table of Contents

1	Introduction to Lyapunov Exponents	1
1.1	Classical Stability Theory of Differential Equations	1
1.2	Chaos and Sensitive Dependence on Initial Conditions	4
1.3	Linear and Nonlinear Stochastic Systems	7
1.4	Random Matrix Products	8
1.5	Issues in Computing Lyapunov Exponents	9
1.6	Contributions of this Thesis	10
2	Random Fibonacci Sequences and the Number 1.13198824...¹	12
2.1	Introduction	12
2.2	Furstenberg's Formula	16
2.3	The Stern-Brocot Tree and Construction of the Invariant Measure ν_f	19
2.4	$e^{\gamma_f} = 1.13198824\dots$	25
2.5	Appendix : Rounding Error Analysis	29
3	A Brief Overview of Random Matrix Products	37
3.1	Existence of Lyapunov Exponents and the Theorem of Furstenberg and Kesten	37
3.2	Furstenberg's Theorem	39
3.3	Computability of Lyapunov Exponents	41
4	Random Recurrences and Condition Numbers of Random Triangular Matrices²	44
4.1	Introduction	44
4.2	Inverse of a Random Triangular Matrix	48
4.3	Moments of \mathbf{T}_k	50
4.4	Exponential Growth of κ_n	51
4.5	Unit Triangular Matrices	54
4.6	A Comment on the Stability of Gaussian Elimination	55
4.7	Complex Matrices	56

¹This chapter is adapted from [112].

²This chapter is adapted from [114].

4.8	Matrices with Entries from Stable Distributions	59
4.9	A Conjecture about Random Recurrences	61
4.10	Summary	62
5	Lyapunov Exponents of the Lorenz Equations ³	65
5.1	Introduction	65
5.2	Computing Lyapunov Exponents	68
5.3	Errors in Computed Lyapunov Exponents	70
5.4	The Lorenz Equations	73
5.5	Reliability of Numerical Evidence for Chaos	73
	References	78

³This chapter is adapted from [113].

List of Tables

5.1	Lyapunov exponents of the Lorenz equations	66
5.2	Estimates of the top Lyapunov exponent of the Lorenz equations . .	74
5.3	Estimates of the top Lyapunov exponent of the Lorenz equations . .	75

List of Figures

2.1	Computer generated random Fibonacci sequences	13
2.2	Invariance condition for the random walk induced by the random Fibonacci recurrence	17
2.3	The Stern-Brocot tree	20
2.4	Invariant measure for the random walk induced by the random Fibonacci recurrence	23
2.5	Illustration of Y. Peres's theorem	28
4.1	Empirical distribution functions of $\kappa_n^{1/n}$	45
4.2	Inversion of random triangular matrices	48
4.3	Illustration of Theorem 4.9	62
4.4	Another illustration of Theorem 4.9	63
5.1	Chaos in the Lorenz equations	67
5.2	Higher derivatives of $x(t)$ with respect to $x(0)$	69
5.3	Error in computed Lyapunov exponents vs. length of integration	71
5.4	Error in computed Lyapunov exponents vs. time step	72
5.5	Chaotic transients in coupled Ginsburg-Landau oscillators	75

Chapter 1

Introduction to Lyapunov

Exponents

Lyapunov exponents were first introduced to study stability of differential equations. They are also useful for describing the sensitive dependence on initial conditions of dynamical systems and the asymptotic behaviour of random matrix products. This introduction will present in a nontechnical way the diverse uses of Lyapunov exponents, mainly those related to stability and sensitive dependence on initial conditions; existing theory and its deficiencies; obstacles to computing Lyapunov exponents, which include long transients; and a unifying idea — in every case Lyapunov exponents can be defined by considering matrix products $M_n \dots M_1$, $n \geq 1$, where the matrices $M_n \in R^{d,d}$ are generated according to some rule. Let us mention at the outset the great debt that this introduction owes to the excellent article by Arnold and Wihstutz [9]. This chapter concludes with a summary of the contributions of this thesis.

1.1 Classical Stability Theory of Differential Equations

Ordinary differential equations of the form

$$\frac{dx(t)}{dt} = F(t, x), \quad x(0) = x_0, \quad (1.1)$$

where $x(t) \in R^d$ (here and in the rest of this chapter), are pervasive in the mathematical sciences. Their uses range from the description of planetary motion in physics to the evolution of chemical processes in chemistry to the workings of a

living cell in biology. For example, if $x(t)$ gives the coordinates of the planets in the solar system at time t , classical mechanics allows us to describe the evolution of the solar system using an equation that can be written in the form (1.1). Given the imprecision of all measurements and observations, it is necessary to ask if a small change in the initial condition x_0 can cause a much bigger change in the solution $x(t)$ for large times t .

This question is directly addressed by two concepts of stability introduced by Lyapunov in his famous memoir on differential equations [80, 1892]. A solution $x(t)$ of (1.1) is stable if for every $\epsilon > 0$, there exists a $\delta > 0$ such that any solution $x'(t)$ of (1.1) satisfying

$$|x(0) - x'(0)| < \delta$$

also satisfies

$$|x(t) - x'(t)| < \epsilon,$$

for all times $t \geq 0$. If in addition $|x(t) - x'(t)| \rightarrow 0$ as $t \rightarrow \infty$, $x(t)$ is said to be *asymptotically stable*. As Lyapunov himself pointed out, given a solution $x(t)$ of (1.1), it is convenient to write down a differential equation for a perturbation $\delta(t)$ such that $x(t) + \delta(t)$ is a solution of (1.1) with $x(0) = x_0 + \delta(0)$. Pretending now that the perturbation $\delta(t)$ is $x(t)$, this leads us to consider first differential equations of the form

$$\frac{dx(t)}{dt} = A(t)x + f(t, x), \quad x(0) = x_0, \quad (1.2)$$

where $A(t) \in R^{d,d}$ is bounded and $\|f(t, x)\| = o(\|x\|)$ as $x \rightarrow 0$ uniformly over finite intervals of t , and then the stability of the solution $x(t) \equiv 0$ for $t \geq 0$ of (1.2) [24, chapter 13] [95].

The stability of (1.2) is relatively easy to decide when $A(t)$ is either constant or periodic; except for some borderline cases, the eigenvalues and Floquet numbers will determine the stability of (1.2). More generally, when $A(t)$ is neither constant nor periodic, Lyapunov exponents can be defined for

$$\frac{dx(t)}{dt} = A(t)x, \quad x(0) = x_0, \quad (1.3)$$

with the intention of understanding the stability of (1.2); examining the linearization (1.3) suffices to understand the stability of the nonlinear equation (1.2) under some conditions. Call the solution of (1.3) $x(t; x_0)$. The Lyapunov exponent of (1.3) is defined by

$$\gamma(x_0) = \limsup_{t \rightarrow \infty} \frac{\log \|x(t; x_0)\|}{t}.$$

The number $\gamma(x_0)$ can vary with x_0 , but it can take at most d different values. This is because for scalars $\alpha, \beta \neq 0$, $\gamma(\alpha x_0) = \gamma(x_0)$ and $\gamma(\alpha x_0 + \beta x_1) = \max(\gamma(x_0), \gamma(x_1))$

if $\gamma(x_0) \neq \gamma(x_1)$. The d values $\gamma_1 \geq \gamma_2 \geq \dots \geq \gamma_d$ are the Lyapunov spectrum of (1.3). It is possible to choose a nested sequence of subspaces $0 \subsetneq V_d \subsetneq \dots \subsetneq V_1 = R^d$ such that the Lyapunov exponent $\gamma(x_0)$ for any x_0 in $V_k - V_{k+1}$ is γ_k [9] [95].

The Lyapunov exponents γ_i can also be defined using a matrix product. The solution $x(t)$ of the linearized equation (1.3) at $t > s$ can be written as $x(t) = M(t, s)x(s)$, where $M(t, s) \in R^{d,d}$. Obviously, $M(t, r) = M(t, s)M(s, r)$ if $t > s > r$, and

$$\gamma(x_0) = \limsup_{t \rightarrow \infty} \frac{\log \|M(t, 0)x_0\|}{t}.$$

The relevance of matrix products to Lyapunov exponents is clear right away since $M(t, 0) = M(t, t-1)M(t-1, t-2) \dots M(1, 0)$ when t is an integer.

It is immediate from the definitions that the linearized equation (1.3) is asymptotically stable if $\gamma_1 < 0$. But the stability of (1.3) does not always imply the stability of the nonlinear equation (1.2). For example, the solution $x(t) \equiv 0$ of

$$\frac{dx(t)}{dt} = x^2$$

is obviously unstable even though $x(t) \equiv 0$ is a stable solution of the linearized equation $\frac{dx(t)}{dt} = 0$. The Lyapunov exponent of this linearized equation is 0.

So when does the stability of (1.3) imply the stability of (1.2)? The answer is contained in the following theorem, which is Theorem 33 in Chapter IX of Sansone and Conti [95].

Theorem 1.1. *Let $\|M(t, s)\| < Ne^{-\gamma(t-s)}$ for $0 \leq s \leq t$ for some $N > 0$, $\gamma > 0$. If $\|f(t, x)\| = o(\|x\|)$ as $\|x\| \rightarrow 0$ uniformly on finite intervals of t , $t \geq 0$, then the solution $X(t) \equiv 0$ of (1.2) is asymptotically stable.*

This result goes back to Lyapunov. Lyapunov's assumption about $f(t, x)$ was very restrictive. The result attained its present form after the work of several researchers including Perron, Petrovski, and Bellman. The inequality $\gamma_1 < 0$ alone does not imply the conditions of Theorem 1.1; $\gamma_1 < 0$ gives a bound on $\|M(t, 0)\|$ in the limit $t \rightarrow \infty$ but not on $\|M(t, s)\|$ with finite s and t .

Let us introduce the important notion of regularity. The system (1.3) is said to be regular if

$$\sum_{i=1}^d \gamma_i = \liminf_{t \rightarrow \infty} \frac{\log |\det M(t)|}{t}.$$

According to Arnold and Wihstutz [9], a stronger version of regularity together with $\gamma_1 < 0$ implies the asymptotic stability of (1.2).

Below is an example of an irregular system due to Lyapunov:

$$\begin{aligned} \frac{dx_1(t)}{dt} &= (\cos \log t)x_1 + (\sin \log t)x_2 \\ \frac{dx_2(t)}{dt} &= (\sin \log t)x_1 + (\cos \log t)x_2. \end{aligned}$$

The sum of the Lyapunov exponents for this system is 2. But the right hand side in the definition of regularity in the previous paragraph evaluates to $-\sqrt{2}$. There are examples, due to Persidskii, Vinograd, and others, whose Lyapunov exponents are discontinuous with respect to small changes in $A(t)$. For references see [95]. After Lyapunov's basic work, Perron's contribution [93] has been influential.

1.2 Chaos and Sensitive Dependence on Initial Conditions

Numerical methods and the arrival of computers made it clear that systems of the form

$$\frac{dx(t)}{dt} = F(x) \tag{1.4}$$

can have unexpectedly complicated dynamics as $t \rightarrow \infty$. The trajectories $x(t)$ do not always have to converge to a limit point or a limit cycle or some other geometrically simple object asymptotically. In chaotic systems, the trajectory $x(t)$, $t \geq 0$, depends sensitively on the initial conditions $x(0)$: a small perturbation of $x(0)$ gives a trajectory that initially diverges exponentially from $x(t)$ and becomes totally uncorrelated with $x(t)$ even though all trajectories are bounded as $t \rightarrow \infty$. But amazingly, almost every trajectory not only converges to the same set called a strange attractor but is dense in it. The strange attractors are usually fractal sets of exotic geometry. Classical tools like Fourier analysis, the theory of special functions, and numerous other analytic techniques do not provide a means for understanding chaos. Although topological methods have been more successful, the versatility and usefulness of numerical methods have not been matched by any kind of mathematics that is divorced from the computer. Many examples of chaos are described in [32], [58], [102].

The definitions of stability given in the previous section are not enough to understand chaos. But a different definition in the same framework can be used to quantify the sensitive dependence on initial conditions. Consider (1.4) with $x(0) = x_0$. Since it is sensitive dependence on initial conditions that interests us, consider the Jacobian $\frac{\partial x(t)}{\partial x(0)}$ and call it $M(t; x_0)$. As before, $M(t+s; x_0) = M(t; x(s))M(s; x_0)$. The d Lyapunov exponents at x_0 $\gamma_1(x_0) \geq \dots \geq \gamma_d(x_0)$ are the possible values of

$$\limsup_{t \rightarrow \infty} \frac{\log \|M(t; x_0)v\|}{t},$$

for $v \in R^d$. The multiplicities of $\gamma_i(x_0)$ are decided using a nested sequence of subspaces as in the previous section. It is easy to see (just differentiate (1.4) with respect to x_0) that $M(t; x_0)$ satisfies the linearized equation

$$\frac{dM(t; x_0)}{dt} = \frac{\partial F(x(t))}{\partial x},$$

with $M(0; x_0)$ being the identity.

Regularity is again an important issue, and is defined as in the previous section. If the matrices $M(t; x_0)$ form a regular system, then

$$\gamma_i(x_0) = \lim_{t \rightarrow \infty} \frac{\log \sigma_i(M(t; x_0))}{t}, \quad 1 \leq i \leq d,$$

where σ_i is the i th singular value, with σ_1 being the 2-norm. Further, the right singular vectors also converge as $t \rightarrow \infty$. We will call these $v_i(x_0)$ ignoring some minor technicalities related to coincidence of Lyapunov exponents. The vector $v_i(x_0)$ is called a stretching direction at x_0 if $\gamma_i(x_0) > 0$, because small perturbations to x_0 in the direction $v_i(x_0)$ are “stretched” exponentially as t increases.

A notable fact about chaotic dynamical systems is that $\gamma_i(x_0)$, though not $v_i(x_0)$, are independent of where x_0 is chosen on a given strange attractor. Thus we can speak of the Lyapunov exponents γ_i of a strange attractor, and even of a chaotic system if it has only one strange attractor. This independence from x_0 is because trajectories are dense in the strange attractor. Chaos means $\gamma_1 > 0$.

Unfortunately, regularity is very hard to prove for chaotic systems that arise in practice. But the numerical evidence for regularity of chaotic systems like the Lorenz equations is very convincing (as we show in Chapter 5). A proof of regularity for a wide class of dynamical systems would be a great advance, but may not have any implications for numerical methods to compute Lyapunov exponents; it is unlikely that such a proof will show us how to compute 8 digits of the Lyapunov exponents for the Lorenz equations, and not just 4 digits as in Chapter 5.

The question of regularity is addressed by ergodic theory from a different direction. But first we note that it suffices to consider discrete dynamical systems. Nothing changes if we work with the map $x_n \rightarrow x(T; x_n)$, where $T > 0$ is fixed and $x(T; x_n)$ gives the trajectory of $\frac{dx(t)}{dt} = F(x(t))$ with $x(0) = x_n$ at $t = T$. Consider a measure μ on the strange attractor which is invariant and ergodic under some discrete map $x_n \rightarrow x_{n+1}$. Such measures always exist with some compactness assumption, but as we will soon see there may be too many of them. Under such a measure, the sequence of matrices $M(T; x_0), M(T; x_1), M(T; x_2), \dots$ are stationary and ergodic. Osseledac’s theorem [89] [25] guarantees that the $\gamma_i(x)$ and $v_i(x)$ exist, that the matrix products are regular and that $\gamma_i(x)$ is a constant γ_i almost surely with respect to μ . Thus it guarantees regularity with respect to an invariant, ergodic measure.

The problem with this approach to regularity is that there may be too many invariant measures. Consider the map $x_{n+1} = 2x_n \bmod 1$ over $[0, 1)$. This map moves out the leftmost bit in the binary representation of x_n to get x_{n+1} . Using the famous extension theorem of Kolmogorov (upon which modern probability theory is founded), we can construct a measure μ_p on $[0, 1)$ such that, for x picked from μ_p , the bits of x are independent random variables with $p(0) = p$, $0 \leq p \leq 1$. All these measures μ_p are invariant and ergodic. Which μ_p to work with? In this case it does not matter — the Lyapunov exponent will be 2 for all μ_p . But in general different invariant measures can correspond to different Lyapunov exponents.

A natural choice for an invariant measure is the Sinai-Ruelle-Bowen (SRB) measure. It is required to be smooth, or absolutely continuous with respect to the Lebesgue measure, in the stretching directions $v_i(x)$ at every x . This is a natural requirement. Imagine a square with some measure on it. If it is stretched into a rectangle with a long length and a short breadth, the measure obviously becomes smoother in the stretching direction. For the dynamical system $x_{n+1} = 2x_n \pmod 1$ on $[0, 1)$ the Lebesgue measure $\mu_{1/2}$ is the only SRB measure. But, unfortunately, SRB measures are not known to exist in great generality. In every case where the existence of an SRB measure has been proven, the proof seems to involve a detailed knowledge of the dynamics [11] [17]. Further, SRB measures may not always exist, and some physical measures may not be SRB [37].

We might ask what ergodic theory does for deterministic dynamical systems. At best, it can give a way around complications like cycles trapped inside strange attractors, which impede topological approaches, by making statements with respect to an invariant measure. But at the moment, it seems to us that most of the complications of the dynamics are wrapped inside the invariant measure.

Regardless of theoretical difficulties, there are strong pragmatic reasons for working with Lyapunov exponents. Numerical computations can give us sufficient confidence in their existence. Besides, one can work with Lyapunov exponents even in high dimensions, as we show in Chapter 5. Other concepts in the theory of dynamical systems like invariant manifolds, Poincaré maps and Smale horseshoes are not as easy to work with in high dimensions. It is definitely easier to compute Lyapunov exponents to investigate chaos in the solar system than it is to look for a Smale horseshoe or an invariant manifold.

Lyapunov exponents are related to the dimension of the attractor. A small perturbation δx to x_0 will increase exponentially at the rate γ_1 if δx has a component along $v_1(x_0)$. Similarly, the volume of a small k dimensional volume at x_0 will increase exponentially at the rate $\gamma_1 + \dots + \gamma_k$, when it is translated forward in time by the dynamical system, if it has a component along the space spanned by $v_1(x_0), \dots, v_k(x_0)$. Thus if $\gamma_1 + \dots + \gamma_k > 0$ one may expect the attractor to be at least k dimensional, and if $\gamma_1 + \dots + \gamma_k < 0$ the attractor may be expected to have dimension less than k . This has been proven [67]. Kaplan and Yorke [68] have gone further and conjectured that if $\gamma_1 + \dots + \gamma_k \geq 0$ but $\gamma_1 + \dots + \gamma_{k+1} < 0$, then “generically” the Hausdorff dimension of an SRB measure is given by

$$k + \frac{\gamma_1 + \dots + \gamma_k}{|\gamma_{k+1}|}.$$

Since it relates concepts that are quite hard to access analytically or axiomatically, this is probably not a conjecture that can be easily settled; but we do not count ourselves among the believers. For a survey of ergodic theory and various notions of dimension, see [37].

Lyapunov exponents can be computed even for partial differential equations. The theory here is beset with even graver difficulties, but the pragmatic reasons for computing Lyapunov exponents are just the same. Manneville [83] has computed

the top few Lyapunov exponents for the Kuramoto-Sivashinsky equation, which is often used as a model partial differential equation that becomes turbulent. The Kuramoto-Sivashinsky equation is given by

$$\partial_t \varphi + \partial_{xx}^2 \varphi + \partial_{xxxx}^4 \varphi + 2\varphi \partial_x \varphi = 0.$$

Typical boundary conditions are $\varphi = \partial_x \varphi = 0$ at $x = 0$ and $x = L$. The propagation of a small perturbation to the initial conditions is governed by

$$\partial_t \psi + \partial_{xx}^2 \psi + \partial_{xxxx}^4 \psi + 2\psi \partial_x \varphi + 2\varphi \partial_x \psi = 0.$$

The top Lyapunov exponent defined by

$$\gamma_1 = \lim_{t \rightarrow \infty} \frac{\log \|\psi(x, t)\|}{t}$$

seems to be about 0.078 for $L = 50$. Lyapunov exponents could conceivably be computed to investigate chaos or turbulence in other partial differential equations like the Navier-Stokes equations of fluid mechanics. Indeed, we will argue in the last section of this thesis that computing Lyapunov exponents carefully in such situations may be a way to avoid mistaking transients for asymptotic behaviour.

1.3 Linear and Nonlinear Stochastic Systems

The stochastic analog of the linear system (1.3) is easy to state:

$$\frac{dx(t, \omega)}{dt} = A(t, \omega)x(t, \omega), \quad x(0, \omega) = x_0.$$

Here $A(t, \omega)$ is a random matrix and ω is a sample point in a probability space. Now the Lyapunov exponent

$$\lambda(x_0, \omega) = \limsup_{t \rightarrow \infty} \frac{\log \|x(t, \omega)\|}{t}$$

can depend on both x_0 and ω . The ergodic theory of linear stochastic systems such as this has been worked out in [8]. For a concrete example, consider the damped harmonic oscillator

$$\begin{aligned} \frac{d^2 u(t)}{dt^2} + 2\xi \frac{du(t)}{dt} + k^2(1 + \eta(t))u(t) &= 0, \\ \frac{du(0)}{dt}, u(0) \text{ are given, } t \geq 0, k > 0, 0 \leq \xi < 1, \end{aligned}$$

with $\eta(t)$ being the noise term. Assume $\eta(t)$ to be a random telegraph process; more specifically, assume $\eta(t)$ has two states $\{-b, +b\}$ with the Markov transition

function $\frac{1}{2} \begin{pmatrix} 1+e^{-2\lambda t} & 1-e^{-2\lambda t} \\ 1-e^{-2\lambda t} & 1+e^{-2\lambda t} \end{pmatrix}$. With this assumption, the expected transition time for $\eta(t)$ is λ . Loparo and Blankenship [78] show that

$$\lim_{t \rightarrow \infty} \frac{\log \left\| \begin{pmatrix} u'(t) \\ u(t) \end{pmatrix} \right\|}{t} = \frac{k^2 \lambda b^2}{8(\lambda^2 + k^2)}$$

with probability 1 and with the damping constant $\xi = 0$. Their result for $\xi > 0$ is too cumbersome to state here. They point out that identical results were obtained earlier by Pastur and Feldman [91] using a formal calculation. Hařminskii [62] is a basic contribution to the theory of linear stochastic systems.

The stochastic analogue of the nonlinear dynamical system (1.4) is

$$\frac{dx(t, \omega)}{dt} = F(x, \eta(t, \omega)), \quad x(0, \omega) = x_0,$$

where $\eta(t, \omega)$ is a stochastic process, possibly noise, and ω is a sample point in a probability space. As before, the Lyapunov exponents can be defined by considering either the Jacobian $\frac{\partial x(t, \omega)}{\partial x_0}$ or the linearization

$$\frac{dM(t, \omega)}{dt} = \frac{\partial F(x, \eta(t, \omega))}{\partial x}, \quad M(0, \omega) = I,$$

where I is the identity matrix. Theory for these stochastic nonlinear systems is more developed than theory for deterministic dynamical systems; see the two volumes [81] [82]. But no exact analytic calculation of Lyapunov exponents seems to be available. Relating Lyapunov exponents to Itô integrals, which are widely used in stochastic dynamics, might hold interesting possibilities.

The definitions of the Lyapunov exponents in the stochastic situation can be phrased using matrix products, with the only change from the deterministic situation being the presence of the sample point ω . The definition of regularity and our remarks about its significance carry over from the deterministic case. Furthermore, as before, Lyapunov exponents are related to stability and sensitive dependence on initial conditions.

1.4 Random Matrix Products

We have seen that matrix products and Lyapunov exponents provide a framework for formulating notions of stability and sensitive dependence on initial conditions for linear or nonlinear, stochastic or deterministic systems. Bellman [10, 1952] initiated the study of random matrix products $M_n \dots M_1$, where each M_i is independent and identically distributed over $R^{d,d}$. Thanks mainly to Furstenberg's remarkable work [46, 1963]¹, the theory in this situation is very well developed. For example, as we

¹Gian-Carlo Rota begins his review of this paper (Mathematical Reviews:29 #648) with the words "This is a profound memoir."

will show in Chapter 3, there are simple algorithms for bounding the Lyapunov exponents in this setting. The advanced state of the theory for random matrix products is a peculiar situation because deterministic matrix products that govern sensitive dependence on initial conditions are barely understood; it is as if the strong law of large numbers were well understood without a satisfactory theory of convergence of infinite series.

The elements of the theory of random matrix products are carefully explained in the beautiful monograph by Bougerol [16]. The basic result about Lyapunov exponents,

$$\lim_{n \rightarrow \infty} \frac{\log \|M_n \dots M_1\|}{n} = \gamma_1 \quad \text{with probability 1,}$$

was proved by Furstenberg and Kesten [47]. Furstenberg [46] presented conditions for strict separation between the Lyapunov exponents $\gamma_1 \geq \dots \geq \gamma_d$. Hašminskii's work on the stability of linear stochastic systems [62] was influenced by Furstenberg's paper. Carverhill [21] has generalized some results of Furstenberg to nonlinear stochastic systems. We outline the theory of random matrix products from an elementary point of view in Chapter 3.

Random matrix products have been used in image compression [33], demography [111], statistical physics [28] and other areas. The applications, or possible applications, of Lyapunov exponents of dynamical systems are even more diverse [81] [82]. However, the significance of Lyapunov exponents goes beyond their current use. As for other basic concepts like eigenvalues, the simplicity and immediacy of the definitions are compelling. The widespread use of eigenvalues today is a consequence of both a careful understanding of the theory (which goes back to the 19th century) and the easy availability of software implementing highly effective algorithms (MATLAB, LAPACK, EISPACK). Advances in both theory and algorithms are perhaps necessary to elevate the use of Lyapunov exponents to a comparable level.

1.5 Issues in Computing Lyapunov Exponents

How efficient and reliable can algorithms to compute Lyapunov exponents be? Algorithms to compute eigenvalues of matrices are remarkably efficient; Trefethen and Bau [108] explain that often computing eigenvalues in floating point arithmetic is as easy (or as hard) as solving a linear system. Though Lyapunov exponents in a sense generalize eigenvalues, this comparison is misleading. The rates of stretching on the surface of a strange attractor vary a lot — in other words most attractors are not uniformly hyperbolic — and Lyapunov exponents average these stretching rates. Thus a more appropriate comparison is to the integration of highly non-smooth functions in multiple dimensions. This comparison leads us to think that the efficiency of Monte Carlo integration is perhaps what algorithms to compute Lyapunov exponents should aim for. As we explain in Chapter 3, Le Page's cen-

tral limit theorem guarantees this level of efficiency, independent of dimension, for random matrix products.

Reliability is a trickier issue. Lyapunov exponents are defined in the asymptotic regime $t \rightarrow \infty$. Since all computations are finite, there is a possibility that the transients persist all the way through the computation. How to identify transients? This is a formidable problem. Chapter 5 states some rules of thumb, but this problem is mostly unresolved. We suspect that the number of instances in the published literature where a long chaotic transient has been mistaken for asymptotic chaos is not small. In fact, when A is a non-normal matrix, the simple linear system $\frac{dx(t)}{dt} = Ax$ can have transient behaviour that is quite different from its asymptotic behaviour. In this fully linear situation, Trefethen [106] [107] has generalized eigenvalues to pseudospectra to understand the transient phenomena. Trefethen has also pointed out numerous instances where it is the transient and not the asymptotic regime that is of physical interest. As the notion of pseudospectra makes clear, eigenvalues are especially hard to compute accurately when the asymptotic behaviour of the linear system differs from the transient behaviour. This is only natural, because from an analytic point of view eigenvalues, like Lyapunov exponents, are asymptotic concepts.

1.6 Contributions of this Thesis

The list of contributions below also serves as a summary of the rest of this thesis.

- In Chapter 2, we exactly determine the Lyapunov exponent of random Fibonacci sequences. Random Fibonacci sequences are defined by the recurrence $t_1 = t_2 = 1$ and $t_n = \pm t_{n-1} \pm t_{n-2}$ for $n > 2$, where each \pm is independent and either $+$ or $-$ with probability $1/2$. We prove that

$$\sqrt[n]{|t_n|} \rightarrow 1.13198824\dots$$

with probability 1. The number $\log 1.13198824\dots$ is the top Lyapunov exponent of the random matrix product which picks one of the four matrices $\begin{pmatrix} 0 & 1 \\ \pm 1 & \pm 1 \end{pmatrix}$ with probability $1/4$. We determine the invariant measure exactly for this random matrix product using Stern-Brocot sequences and then use Furstenberg's formula to compute the constant. This computation is validated by rigorous rounding error analysis, and then repeated on two entirely different computer systems. Random Fibonacci recurrence is a rare, and in our opinion the most natural, example of a random matrix product whose top Lyapunov exponent can be exactly determined.

- Chapter 3 is a brief overview of the theory of random matrix products.

- Chapter 4 is about random recurrences of the form

$$\begin{aligned}
 t_1 &= 1/\alpha_{11} \\
 t_2 &= (\alpha_{21}t_1)/\alpha_{22} \\
 t_3 &= (\alpha_{31}t_1 + \alpha_{32}t_2)/\alpha_{33} \\
 &\vdots \\
 t_n &= (\alpha_{n1}t_1 + \cdots + \alpha_{n,n-1}t_{n-1})/\alpha_{nn},
 \end{aligned}$$

where α_{ij} , $i > j$, are all independent and identically distributed, and α_{ii} are all either independently drawn from the same distribution or fixed at 1. We show for a variety of distributions — the real and complex normal distributions, the Cauchy distribution and others — that

$$\sqrt[n]{|t_n|} \rightarrow C \quad \text{as } n \rightarrow \infty$$

with probability 1 and for a constant $C > 1$ that we determine exactly. The formulas for C resemble Furstenberg's formula for random matrix products. For example, when $\alpha_{ii} = 1$ and α_{ij} , $i \neq j$, are normal variables with mean 0 and variance σ ,

$$C = \exp\left(\frac{1}{2\sqrt{2\pi\sigma^2}} \int_{-\infty}^{\infty} \log(1+x^2)e^{-x^2/2\sigma^2} dx\right).$$

Recurrences like the one above are related to condition numbers of random triangular matrices. Our conclusion is that random triangular matrices are exponentially ill-conditioned. This fact has somewhat tangential implications for the numerical stability analysis of Gaussian elimination.

- Chapter 5 examines the accurate computation of Lyapunov exponents of continuous, deterministic dynamical systems. We demonstrate the importance of long integrations, short time steps, and higher order discretizations for the Lorenz equations, which form a low dimensional system, and for coupled Ginsburg-Landau oscillators, which form a high dimensional system. We accurately compute 4 digits of the Lyapunov exponents of the Lorenz equations with three standard parameter choices. The important issue of identifying long chaotic transients to be just transients is discussed.

Chapters 2, 4, and 5 are based on [112], [114], and [113], respectively.

Chapter 2

Random Fibonacci Sequences and the Number 1.13198824... ¹

2.1 Introduction

The Fibonacci numbers defined by $f_1 = f_2 = 1$ and $f_n = f_{n-1} + f_{n-2}$ for $n > 2$ are widely known. It is equally well-known that $|f_n|$ increases exponentially with n at the rate $(1 + \sqrt{5})/2$. Consider random Fibonacci sequences defined by the random recurrence $t_1 = 1$, $t_2 = 1$, and for $n > 2$, $t_n = \pm t_{n-1} \pm t_{n-2}$, where each \pm sign is independent and either $+$ or $-$ with probability $1/2$. Do the random Fibonacci sequences level off because of the subtractions? Or do the random Fibonacci sequences increase exponentially with n like the Fibonacci sequence? If so, at what rate? The answer to these questions brings Stern-Brocot sequences, a beautiful way to divide the real number line that was first discovered in the 19th century, and fractals and random matrix products into play. The final answer is obtained from a computer calculation, raising questions about computer assisted theorems and proofs.

Below are three possible runs of the random Fibonacci recurrence:

1, 1, -2, -3, -1, 4, -3, 7, -4, 11, -15, 4, -19, 23, -4, ...

1, 1, 2, 3, 5, 8, 13, 21, 34, 55, 89, 134, 223, 357, 580, ...

1, 1, -2, 1, 1, -2, 1, 1, -2, 1, 1, -2, 1, 1, -2, ...

The first of the runs was randomly generated on a computer. The second is the familiar Fibonacci sequence. The last is a sequence that remains bounded as $n \rightarrow \infty$; but such runs with no exponential growth occur with probability 0. For longer, typical runs see Figure 2.1. Numerical experiments in Figure 2.1 illustrate our result (Theorem 2.6) that

$$\sqrt[n]{|t_n|} \rightarrow 1.13198824\dots \text{ as } n \rightarrow \infty$$

¹This chapter is adapted from [112].




Figure 2.1: (a) A semilog plot of $|t_n|$ vs. n for a computer generated random Fibonacci sequence t_n showing a clear exponential trend. The dashed line is 1.132^n . (b) Plot of $\sqrt[n]{|t_n|}$ vs. n . As n increases to a million, $\sqrt[n]{|t_n|}$ seems to settle down to a constant close to 1.132.

with probability 1. Thus $1.13198824\dots$ gives the exponential rate of increase of $|t_n|$ with n for random Fibonacci sequences just as the golden ratio $(1 + \sqrt{5})/2$ gives the exponential rate of increase of the Fibonacci numbers.

For the random Fibonacci recurrence $t_n = \pm t_{n-1} \pm t_{n-2}$ as well as the recurrence $t_n = \pm t_{n-1} + t_{n-2}$ with each \pm independent and $+$ or $-$ with probability $1/2$, $|t_n|$ is either $|t_{n-1}| + |t_{n-2}|$ or $||t_{n-1}| - |t_{n-2}||$ with probability $1/2$. As our interest is in $|t_n|$ vs. n as $n \rightarrow \infty$, we restrict focus to $t_n = \pm t_{n-1} + t_{n-2}$ and call it the random Fibonacci recurrence. As a result, the presentation becomes briefer, especially in Section 2.3.

The next step is to rewrite the random Fibonacci recurrence using matrices. In matrix form the random Fibonacci recurrence is $\begin{pmatrix} t_{n-1} \\ t_n \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ 1 & \pm 1 \end{pmatrix} \begin{pmatrix} t_{n-2} \\ t_{n-1} \end{pmatrix}$, with one of the two matrices

$$A = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}, \quad B = \begin{pmatrix} 0 & 1 \\ 1 & -1 \end{pmatrix} \quad (2.1)$$

picked independently with probability $1/2$ at each step. Let μ_f denote the distribution that picks A or B with probability $1/2$. Then the random matrix M_n chosen at the n th step is μ_f -distributed and independent of M_i for $i \neq n$. Moreover,

$$\begin{pmatrix} t_{n-1} \\ t_n \end{pmatrix} = M_{n-2} \dots M_1 \begin{pmatrix} 1 \\ 1 \end{pmatrix},$$

where $M_{n-2} \dots M_1$ is a product of independent, identically distributed random matrices.

Known results from the theory of random matrix products imply that

$$\frac{\log \|M_n \dots M_1\|}{n} \rightarrow \gamma_f \text{ as } n \rightarrow \infty, \quad (2.2)$$

$$\sqrt[n]{|t_n|} \rightarrow e^{\gamma_f} \text{ as } n \rightarrow \infty, \quad (2.3)$$

for some constant γ_f with probability 1 [16, p. 11, p. 157]. Our aim is to determine γ_f or e^{γ_f} exactly. Theorem 2.6 realizes this aim by showing that $e^{\gamma_f} = 1.13198824\dots$. The limit in (2.3) is the same γ_f for any norm over 2-dimensional matrices because all norms over a finite dimensional vector space are equivalent. In the rest of this chapter, all norms are 2-norms, and all matrices and vectors are real and 2-dimensional except when stated otherwise. Thus, in this chapter for a vector x , $\|x\|$ is its Euclidean length in the real plane, and for a matrix M , $\|M\| = \sup_{x \neq 0} \frac{\|Mx\|}{\|x\|}$.

The limit (2.2) for M_i independent but identically distributed over d -dimensional matrices has been a central concern of the theory of random matrix products. Furstenberg and Kesten [47, 1960] showed that the limit exists under very general conditions. When it exists, the limit is usually denoted by γ and called the upper Lyapunov exponent. Furstenberg [46, 1963] showed that when the normalizing condition $|\det M_i| = 1$ holds, as it does for μ_f , “usually” $\gamma > 0$. Furstenberg’s theorem implies, for example, that $\gamma_f > 0$, and hence, that $|t_n|$ increases exponentially with n with probability 1. The basic theory of random matrix products is explained from an elementary point of view in the next chapter.

In spite of the importance of the upper Lyapunov exponent γ , γ is known exactly for very few examples. Kingman, one of the pioneers of subadditive ergodic theory of which the theory of random matrix products is a special case, wrote [70, 1973]:

Pride of place among the unsolved problems of subadditive ergodic theory must go to the calculation of the constant γ (. . .). In none of the applications described here is there an obvious mechanism for obtaining an exact numerical value, and indeed this usually seems to be a problem of some depth.

One of the applications Kingman refers to is the general problem of finding γ for random matrix products. For this and other applications, Kingman’s problem is still unsolved. Bougerol [16, p. 33] and Lima and Rahibe [77] calculate γ for some examples. The work of Chassaing, Letac and Mora [22] is closer to our determination of γ_f . But in all their examples, the matrices, unlike B in (2.1), have only non-negative entries. In our opinion, the random Fibonacci recurrence is more natural than these examples. In fact, the random Fibonacci recurrence in a more general form appears as a motivating example in the very first paragraph of Furstenberg’s famous paper [46].

In Section 2.2, we present a formula for γ_f due to Furstenberg that forms the basis for this chapter. The matrices A and B map a direction in the real plane of slope m to directions of slope $1+1/m$ and $-1+1/m$, respectively. Since μ_f picks A or

B with probability $1/2$, it induces the random walk which sends a direction of slope m to a direction of slope $1 + 1/m$ or $-1 + 1/m$ with probability $1/2$. The invariant probability measure for this random walk is central to Furstenberg's formula.

In Section 2.3, we find that invariant probability measure, denoted by ν_f , using the Stern-Brocot division of the real line. See Figures 2.3, 2.4. The measure ν_f gives a probability measure over the real line R because the slope m can be any real number. Since the backward maps for the random walk are $m \rightarrow 1/(\pm 1 + m)$, the invariance condition requires

$$\nu_f([a, b]) = \frac{1}{2}\nu_f\left(\frac{1}{-1 + [a, b]}\right) + \frac{1}{2}\nu_f\left(\frac{1}{1 + [a, b]}\right),$$

for any interval $[a, b]$. The notations $1/(-1 + [a, b])$ and $1/(1 + [a, b])$ are not standard; they denote the images of the interval $[a, b]$ under the maps $m \rightarrow 1/(-1 + m)$ and $m \rightarrow 1/(1 + m)$, respectively. Since the slopes of the backward maps vary in magnitude from 0 to ∞ , not only is ν_f self-similar [101], the self-similarity equation has multiple scales. Self-similar functions, especially ones with multiple scales, usually turn out to be fractals. For example, Weierstrass's nowhere-differentiable but continuous functions, which are commonly used examples of fractal graphs, satisfy $f(t) = \lambda^{s-2} \sin(\lambda t) + \lambda f(\lambda t)$ with $1 < s < 2$, $\lambda > 1$, and λ large enough [43]. Another remarkable example is Daubechies' wavelets; the solution to

$$\begin{aligned} f(t) = & \frac{1 + \sqrt{3}}{4}f(2t) + \frac{3 + \sqrt{3}}{4}f(2t - 1) + \frac{3 - \sqrt{3}}{4}f(2t - 2) \\ & + \frac{1 - \sqrt{3}}{4}f(2t - 3) \end{aligned}$$

has an irregular graph which is a type of fractal, yet it can be used to construct wavelets that approximate smooth functions very well [100, p. 437]. Repetition of the same structure at finer scales and an irregular appearance in Figure 2.4 suggest that ν_f too may be a type of fractal.

In Section 2.4, we use Furstenberg's formula and the invariant measure ν_f given in Section 2.3 and arrive at Theorem 2.6 ($e^{\gamma_f} = 1.13198824\dots$). The proof of Theorem 2.6 depends on a computer calculation. Thus its correctness depends not only upon mathematical arguments that can be checked line by line, but also upon a program that can also be checked line by line and the correct implementation of various software and hardware components of the computer system. The most famous of theorems whose proofs depend on computer calculations is the four color theorem. The first proof of the four color theorem (all planar graphs can be colored using only four colors so that no two adjacent vertices have the same color) by Appel, Haken and Koch caused controversy and aroused great interest because it relied on producing and checking 1834 graphs using 1200 hours of 1976 computer time [6] [72]. In spite of improvements (for example, the number 1834 was brought down to 1482 soon afterwards by Appel and Haken themselves), all proofs of the four color theorem still rely on the computer.

Computer assisted proofs are more common now. Our computation uses floating point arithmetic, which is inexact owing to rounding errors. Thus it becomes necessary to bound the effect of the rounding errors, which we do in the appendix. An early example of rigorous use of floating point arithmetic is due to Brent [19]. Lanford's proof of Feigenbaum's conjecture about the period doubling route to chaos used interval arithmetic [73]. The computer assisted proof of chaos in the Lorenz equations announced by Mischaikow and Mrozek [87] is another notable example. We will discuss the use of floating point arithmetic and other issues related to our Theorem 2.6 in Section 2.4.

Besides random matrix products, random Fibonacci sequences are connected to many areas of mathematics. For example, the invariant measure ν_f is also the distribution of the continued fractions

$$\pm 1 + \frac{1}{\pm 1 + \frac{1}{\pm 1 + \dots}}$$

with each ± 1 independent and either $+1$ or -1 with probability $1/2$. The matrices A and B in (2.1) can both be thought of as Möbius transformations of the complex plane; then the random matrix product and the exponential growth of $|t_n|$ in (2.2) and (2.3) correspond to the dynamics of complex numbers acted upon by a composition of the Möbius transformations A and B [16, p. 38]. Also, the random walk on slopes $m \rightarrow \pm 1 + 1/m$ can be thought of as a random dynamical system [7]. These different interpretations amount merely to a change of vocabulary as far as the computation of γ_f is concerned; but each interpretation offers a different point of view.

The study of random matrix products, initiated by Bellman [10, 1954], has led to many deep results and applications. Applications have been made to areas as diverse as Schrödinger operators, image generation, and demography [28][33][111]. Our own interest in random recurrences was aroused by their connection to random triangular matrices described in Chapter 4. Furstenberg and Kesten [47, 1960], Furstenberg [46, 1963], Osseledac [89, 1968], Kingman [70, 1973], Guivarc'h and Raugi [59, 1985], and Le Page [75, 1982] [76, 1989] are some of the profound contributions to this area. We enthusiastically recommend the lucid, elegant and well-organized account by Bougerol [16]. For a more modern treatment, see [13]. For the basics of probability, our favorite is Breiman [18].

2.2 Furstenberg's Formula

To determine γ_f , we use a formula from the theory of random matrix products that complements (2.2). Three things that will be defined below — the notation \bar{x} for directions in the real plane R^2 , $\text{amp}(\bar{x})$, which is a smooth function of \bar{x} (Figure on p. 27), and $\nu_f(\bar{x})$, which is a probability measure over directions \bar{x} (Figure 2.4) —

μ_f -invariance in terms of slopes is

$$\nu_f([a, b]) = \frac{1}{2}\nu_f\left(\frac{1}{-1 + [a, b]}\right) + \frac{1}{2}\nu_f\left(\frac{1}{1 + [a, b]}\right), \quad (2.5)$$

where $[a, b]$ is any interval in the real line. See Figure 2.2.


The function $\text{amp}(\bar{x})$ defined by

$$\text{amp}(\bar{x}) = \frac{1}{2} \log \frac{\|Ax\|}{\|x\|} + \frac{1}{2} \log \frac{\|Bx\|}{\|x\|}$$

gives the average amplification in the direction \bar{x} when x is multiplied by A or B with probability $1/2$. Recall that $\|\cdot\|$ was taken to be the 2-norm. In terms of slopes,

$$\begin{aligned} \text{amp}(m) &= \frac{1}{4} \log \left(\frac{m^2 + (-1 + m)^2}{1 + m^2} \right) + \frac{1}{4} \log \left(\frac{m^2 + (1 + m)^2}{1 + m^2} \right) \\ &= \frac{1}{4} \log \left(\frac{1 + 4m^4}{(1 + m^2)^2} \right). \end{aligned}$$

The figure below plots $\text{amp}(m)$ vs. m .



Furstenberg's formula (2.4) can now be put in a concrete form using slopes to parameterize directions \bar{x} :

$$\gamma_f = \int_{-\infty}^{\infty} \text{amp}(m) d\nu_f(m) = \frac{1}{4} \int_{-\infty}^{\infty} \log \left(\frac{1 + 4m^4}{(1 + m^2)^2} \right) d\nu_f(m). \quad (2.6)$$

If we were to use a norm other than the 2-norm for vectors in the real plane, $\text{amp}(m)$ and $\text{amp}(\bar{x})$ would be different functions. But Furstenberg's formula (2.4) holds for any norm, even though the measure ν_f is independent of the norm. Our choice of the 2-norm is one of many equally suitable alternatives. For the weighted 2-norm

$$\left\| \begin{pmatrix} a \\ b \end{pmatrix} \right\| = \sqrt{a^2 + \frac{1 + \sqrt{5}}{2} b^2},$$

$\text{amp}(m) > 0$ for all m except $m = \pm\sqrt{(\sqrt{5}-1)/2}$, and $\text{amp}(m) = 0$ at those two points.

To illustrate how (2.6) is used, we verify quickly that $\gamma_f > 0$. The invariance condition (2.5) applied to the set $[-\infty, -1] \cup [1, \infty]$ implies $\nu_f(|m| \geq 1) \geq 1/2$ because $1/(-1 + [1, \infty]) = [0, \infty]$ and $1/(1 + [-\infty, -1]) = [-\infty, 0]$. Now,

$$\begin{aligned} \gamma_f &= \int_{-\infty}^{\infty} \text{amp}(m) d\nu_f(m) \\ &> \min_{|m| < 1} \text{amp}(m) \nu_f(|m| < 1) + \min_{|m| \geq 1} \text{amp}(m) \nu_f(|m| \geq 1) \\ &= -\frac{1}{4} \log\left(\frac{5}{4}\right) \nu_f(|m| < 1) + \frac{1}{4} \log\left(\frac{5}{4}\right) \nu_f(|m| \geq 1) \\ &\geq 0. \end{aligned}$$

The first inequality above is strict because ν_f must be continuous and $\text{amp}(m)$ is not a constant function. Minimizing $\text{amp}(m)$ over $|m| < 1$ and $|m| \geq 1$ is basic calculus: the minima occur at the points $m = \pm 1/2$ and $m = \pm 1$. The second inequality follows from $\nu_f(|m| \geq 1) \geq 1/2$. Actually, it will be shown in Section 2.3 that $\nu_f(|m| \geq 1) = (\sqrt{5}-1)/2$.

2.3 The Stern-Brocot Tree and Construction of the Invariant Measure ν_f

Assuming $\pm 1 \notin (a, b)$, we write the invariance condition using slopes (2.5) in a more explicit form:

$$\nu_f([a, b]) = \frac{1}{2} \nu_f\left(\left[\frac{1}{-1+b}, \frac{1}{-1+a}\right]\right) + \frac{1}{2} \nu_f\left(\left[\frac{1}{1+b}, \frac{1}{1+a}\right]\right). \quad (2.7)$$

Our goal in this section is to find ν_f , the unique probability measure on the real line R satisfying (2.7) for all intervals $[a, b]$ not containing ± 1 . Since ν_f must be continuous, it does not matter whether we take the intervals in (2.7) to be open or closed or half-closed.

The construction of ν_f is based on the Stern-Brocot tree shown in Figure 2.3. The Stern-Brocot tree is an infinite binary tree that divides R recursively. Represent ∞ as $\frac{1}{0}$ and 0 as $\frac{0}{1}$, and write negative fractions with the numerator negative. Then the root of the Stern-Brocot tree is the real line $[\frac{-1}{0}, \frac{1}{0}]$. Its left and right children are $[\frac{-1}{0}, \frac{0}{1}]$ and $[\frac{0}{1}, \frac{1}{0}]$, the positive and negative halves of R . The rest of the tree is defined by dividing any node $[\frac{a}{b}, \frac{c}{d}]$ other than the root into a left child $[\frac{a}{b}, \frac{a+c}{b+d}]$ and a right child $[\frac{a+c}{b+d}, \frac{c}{d}]$. For example, the root's left child $[\frac{-1}{0}, \frac{0}{1}]$ divides into $[\frac{-1}{0}, \frac{-1}{1}]$ and $[\frac{-1}{1}, \frac{0}{1}]$.

The Stern-Brocot tree was discovered and reported independently by the mathematician Moriz Stern in 1858 and by the watchmaker Achille Brocot in 1860 [99] [20]. Unaware of its existence, we found it again while trying to construct ν_f . We

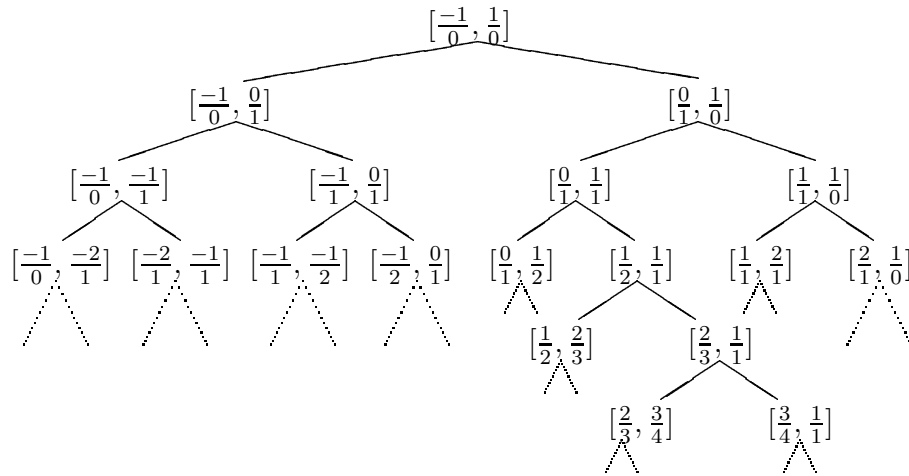


Figure 2.3: The Stern-Brocot tree; its nodes are intervals of the real line R . The division of any interval $[\frac{a}{b}, \frac{c}{d}]$, except the root, into two children is done by inserting the point $\frac{a+c}{b+d}$.

summarize some basic facts about it in Lemma 2.1. The Stern-Brocot tree and its connections with continued fractions are discussed in detail by Graham, Knuth, and Patashnik [53]. Their definition of the Stern-Brocot tree is slightly different from ours. We adopt their notation $a \perp b$ to say that integers a and b are relatively prime.

Lemma 2.1. (a) *The Stern-Brocot tree is symmetric about 0, with its right half positive and its left half negative.*

(b) *If $[\frac{a}{b}, \frac{c}{d}]$ is a node in the positive half of the Stern-Brocot tree, then $bc - ad = 1$, $a \perp b$, and $c \perp d$.*

(c) *Conversely, if a/b and c/d are non-negative rational numbers with zero and infinity, represented as $\frac{0}{1}$ and $\frac{1}{0}$, respectively, and $bc - ad = 1$, then $[\frac{a}{b}, \frac{c}{d}]$ occurs as a node in the Stern-Brocot tree. Consequently, every rational number a/b , $a \perp b$, appears as an endpoint of a Stern-Brocot interval of finite depth.*

Proof. (a) is obvious; see Figure 2.3. The proof of (b) is an easy induction on the depth of the tree. (c) is a little bit less easy. Its proof is related to Euclid's algorithm for computing the greatest common divisor of two integers. See [53]. \square

We adopt a labelling scheme for Stern-Brocot intervals (nodes of the Stern-Brocot tree) that differs only a bit from that in [53]. The root $[-\frac{1}{0}, \frac{1}{0}]$ has the empty label. Its left and right children $[-\frac{1}{0}, \frac{0}{1}]$ and $[\frac{0}{1}, \frac{1}{0}]$ are labelled l and r , respectively. The left child of l , $[-\frac{1}{0}, \frac{-1}{1}]$, is labelled lL . The right child of lL , $[-\frac{2}{1}, \frac{-1}{1}]$, is labelled lLR , and so on. Only the first letter of a label is in small case because the division of the root is special.

We use $l\alpha$ or $r\alpha$ to denote the labels of Stern-Brocot intervals other than the root, with α being a possibly empty sequence of L s and R s. The sequence obtained by changing α 's L s to R s and R s to L s is denoted $\bar{\alpha}$. For example, the reflection of the positive interval $r\alpha$ about 0 is the negative interval $l\bar{\alpha}$. The length of α is denoted by $|\alpha|$. We take the depth of $l\alpha$ or $r\alpha$ to be $1 + |\alpha|$.

Lemmas 2.2 and 2.3 express the maps $m \rightarrow 1/m$ and $m \rightarrow \pm 1 + m$ succinctly for Stern-Brocot intervals. They allow us to reduce the invariance requirement (2.7) for Stern-Brocot intervals to an infinite system of linear equations (see (2.8)). That reduction is the first step in constructing ν_f .

Lemma 2.2. *The image of the interval $[a/b, c/d]$ under the map $m \rightarrow 1/m$ — which is $[d/c, b/a]$ if 0 is not an interior point — is given by the following rules for Stern-Brocot intervals:*

$$l\alpha \rightarrow l\bar{\alpha}, \quad r\alpha \rightarrow r\bar{\alpha}.$$

Proof. We give the proof for intervals of type $r\alpha$ using induction on the depth of $r\alpha$ in the Stern-Brocot tree. The proof for intervals $l\alpha$ is similar.

The base case $r \rightarrow r$ holds because $m \in [0, \infty]$ if and only if $1/m \in [0, \infty]$.

For the inductive case, note that $[a/b, c/d]$, its left child $[a/b, \frac{a+c}{b+d}]$, and its right child $[\frac{a+c}{b+d}, c/d]$ are mapped by $m \rightarrow 1/m$ to $[d/c, b/a]$, its right child $[\frac{b+d}{a+c}, b/a]$, and its left child $[d/c, \frac{b+d}{a+c}]$, respectively. Therefore, if $r\alpha \rightarrow r\bar{\alpha}$ then $r\alpha L \rightarrow r\bar{\alpha}R$ and $r\alpha R \rightarrow r\bar{\alpha}L$. \square

Unlike the inversion operation $m \rightarrow 1/m$ in the previous lemma, both the operations $m \rightarrow \pm 1 + m$ in the following lemma change the depth of Stern-Brocot intervals.

Lemma 2.3. *The image of Stern-Brocot intervals under the map $m \rightarrow -1 + m$ is given by the following rules:*

$$l\alpha \rightarrow lL\alpha, \quad rL\alpha \rightarrow lR\alpha, \quad rR\alpha \rightarrow r\alpha.$$

Similarly, the image of Stern-Brocot intervals under the map $m \rightarrow 1 + m$ is given by the following rules:

$$lL\alpha \rightarrow l\alpha, \quad lR\alpha \rightarrow rL\alpha, \quad r\alpha \rightarrow rR\alpha.$$

Proof. Similar to the previous proof. We will outline the proof for $m \rightarrow 1 + m$ only.

The base cases, adding 1 to the intervals lL , lR and r , are easy to check.

For the induction, we note that $[a/b, c/d]$ is divided in the Stern-Brocot tree at the point $\frac{a+c}{b+d}$, and its map under $m \rightarrow 1 + m$, $[1 + \frac{a}{b}, 1 + \frac{c}{d}]$, is divided in the Stern-Brocot tree at the point $1 + \frac{a+c}{b+d}$. Thus $[a/b, c/d]$, its left child, and its right child map to $[1 + \frac{a}{b}, 1 + \frac{c}{d}]$, its left child, and its right child, respectively. \square

By Lemma 2.3, subtraction and addition of 1 to intervals in the Stern-Brocot tree correspond to left and right rotation of the tree. Tree rotations are used to implement balanced trees in computer science [27].

Thanks to Lemmas 2.2 and 2.3, the backward maps $m \rightarrow 1/(\pm 1 + m)$ can be performed on Stern-Brocot intervals easily. For example, $1/(1 + lLRL) = 1/lRL = lLR$. The invariance requirement (2.7) for Stern-Brocot intervals becomes an infinite set of linear equations for $\nu_f(I)$, I being any Stern-Brocot interval:

$$\begin{aligned}
\nu_f(l) &= \frac{1}{2}\nu_f(lR) + \frac{1}{2}(\nu_f(l) + \nu_f(rR)) \\
\nu_f(r) &= \frac{1}{2}(\nu_f(r) + \nu_f(lL)) + \frac{1}{2}\nu_f(rL) \\
\nu_f(lL\alpha) &= \frac{1}{2}\nu_f(l\overline{LL\alpha}) + \frac{1}{2}\nu_f(l\overline{\alpha}) \\
\nu_f(lR\alpha) &= \frac{1}{2}\nu_f(l\overline{LR\alpha}) + \frac{1}{2}\nu_f(r\overline{L\alpha}) \\
\nu_f(rL\alpha) &= \frac{1}{2}\nu_f(l\overline{R\alpha}) + \frac{1}{2}\nu_f(r\overline{RL\alpha}) \\
\nu_f(rR\alpha) &= \frac{1}{2}\nu_f(r\overline{\alpha}) + \frac{1}{2}\nu_f(r\overline{RR\alpha}).
\end{aligned} \tag{2.8}$$

We guessed the solution of (2.8). Even though the linear system (2.8) has only rational coefficients, its solution involves $\sqrt{5}$, an irrational number! Let $g = (1 + \sqrt{5})/2$. Since ν_f is a probability measure, we require that $\nu_f([-\infty, \infty]) = 1$. The solution is:

$$\begin{aligned}
\nu_f(r) &= 1/2 \\
\nu_f(r\alpha L) &= \frac{1}{1+g}\nu_f(r\alpha) \text{ if } |\alpha| \text{ is even} \\
&= \frac{g}{1+g}\nu_f(r\alpha) \text{ if } |\alpha| \text{ is odd} \\
\nu_f(r\alpha R) &= \frac{g}{1+g}\nu_f(r\alpha) \text{ if } |\alpha| \text{ is even} \\
&= \frac{1}{1+g}\nu_f(r\alpha) \text{ if } |\alpha| \text{ is odd} \\
\nu_f(l\alpha) &= \nu_f(r\overline{\alpha}).
\end{aligned} \tag{2.9}$$

For example, $\nu_f(r) = 1/2$, $\nu_f(rL) = (1+g)^{-1}/2$, $\nu_f(rLL) = g(1+g)^{-2}/2$. Since $\nu_f(l\alpha) = \nu_f(r\overline{\alpha})$ by (2.9), the measure ν_f is symmetric about 0. The same features of ν_f repeat at finer and finer scales. See Figure 2.4.

Theorem 2.4. *The measure ν_f defined by (2.9) satisfies the invariance requirement (2.7) for every Stern-Brocot interval. Further, with directions parameterized by slopes, ν_f defined by (2.9) gives the unique μ_f -invariant probability measure over directions in the real plane R^2 .*

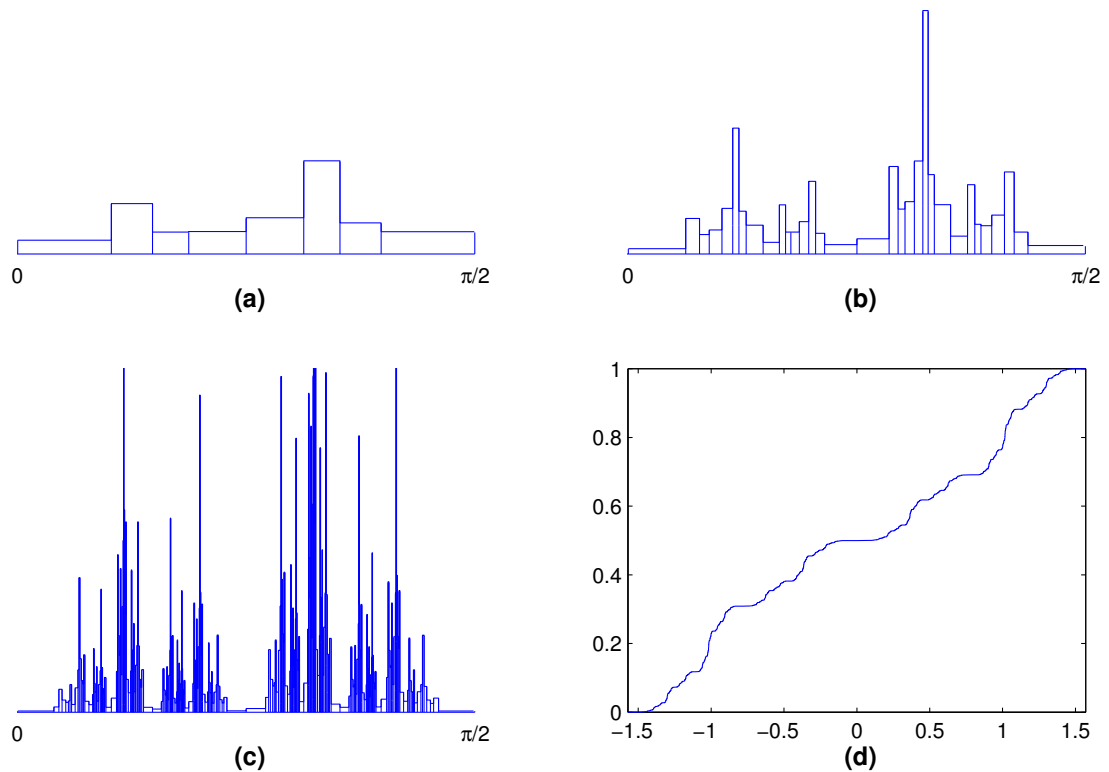


Figure 2.4: (a), (b), (c) show the measure ν_f over directions in R^2 . In these figures, the interval $[0, \infty]$ is divided into 2^3 , 2^5 , and 2^8 Stern-Brocot intervals of the same depth, and then slopes are converted to angles in the interval $[0, \pi/2]$. The area above an interval gives its measure under ν_f . Because of symmetry, ν_f in the directions $[-\pi/2, 0]$ can be obtained by reflecting (a), (b) and (c). Some of the spikes in (c) were cut off because they were too tall. (d) is the distribution function for ν_f with directions parameterized using angles.

Proof. To show that ν_f is μ_f -invariant, it is enough to show that ν_f satisfies the invariance conditions (2.8) for Stern-Brocot intervals. The reason is — ν_f is obviously a continuous measure, every rational appears in the Stern-Brocot tree at a finite depth by Lemma 2.1c, and the rationals are dense in R . For the uniqueness of ν_f , see [16, p. 31].

It is enough to prove the invariance condition for positive intervals $r\alpha$. The validity of the invariance condition for negative Stern-Brocot intervals follows from symmetry. Assume the invariance condition for the interval $rL\alpha$:

$$\nu_f(rL\alpha) = \frac{1}{2}\nu_f(l\overline{R\alpha}) + \frac{1}{2}\nu_f(r\overline{RL\alpha}).$$

Then the invariance condition for $rL\alpha L$,

$$\nu_f(rL\alpha L) = \frac{1}{2}\nu_f(l\overline{R\alpha R}) + \frac{1}{2}\nu_f(r\overline{RL\alpha R}),$$

is also true, because the three fractions

$$\frac{\nu_f(rL\alpha L)}{\nu_f(rL\alpha)}, \quad \frac{\nu_f(l\overline{R\alpha R})}{\nu_f(l\overline{R\alpha})}, \quad \frac{\nu_f(r\overline{RL\alpha R})}{\nu_f(r\overline{RL\alpha})},$$

are all either $g/(1+g)$ or $1/(1+g)$ according as $|\alpha|$ is even or odd. By a similar argument, if the invariance condition (2.8) holds for all positive Stern-Brocot intervals at depth $d \geq 2$, then the invariance condition holds for all positive Stern-Brocot intervals at depth $d+1$.

Therefore, it suffices to verify (2.8) for r , rL , and rR . For r , (2.8) requires

$$\frac{1}{2} = \frac{1}{2}\left(\frac{1}{2} + \frac{1}{2(1+g)}\right) + \frac{1}{2}\left(\frac{g}{2(1+g)}\right),$$

which is obviously true. For rL , (2.8) requires,

$$\frac{1}{2(1+g)} = \frac{g}{4(1+g)} + \frac{1}{4(1+g)^2},$$

which is true because $g = (1+\sqrt{5})/2$. The invariance condition for rR can be verified similarly. Thus the invariance condition (2.8) holds for all Stern-Brocot intervals, and we can say that ν_f is the unique μ_f -invariant probability measure. \square

Because of symmetry, the measure ν_f over slopes given by (2.9) is invariant even for the distribution that picks one of $\begin{pmatrix} 0 & 1 \\ \pm 1 & \pm 1 \end{pmatrix}$ with probability 1/4. Moreover, Furstenberg's integral for the Lyapunov exponent γ of this distribution is also given by (2.6).

For some distributions supported on 2-dimensional matrices with non-negative entries, the infinite linear system analogous to (2.8) is triangular, or in other words, the invariance requirement for a Stern-Brocot interval involves only intervals at a lesser depth. For a typical example, choose $\begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$ with probability p , $0 < p < 1$, and

$\begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}$ with probability $1 - p$. In this example, the invariant measure over directions parameterized by slopes is supported on $[0, \infty]$, the slope m is mapped to $1/(1 + m)$ and $1 + 1/m$ respectively, and the ranges of those two maps ($[0, 1]$ and $[1, \infty]$) are disjoint. Chassaing, Letac and Mora [22] have found the invariant measure for several 2-dimensional random matrix products that fit into this framework. All their matrices have non-negative entries. Moreover, since the linear systems for finding the invariant measure are triangular for all the examples in [22], the solution can contain irrational numbers only if the original problem does.

According to historical remarks in [22], measures similar to ν_f have been studied by Denjoy, Minkowski, and de Rham. But is ν_f a fractal? To make this question precise, we need the definition

$$\dim(\nu_f) = \inf\{\dim(S) \mid \nu_f \text{ is supported on } S\},$$

where $\dim(S)$ is the Hausdorff dimension of $S \subset \mathbb{R}$. To show that ν_f is a fractal, it is necessary to prove that $0 < \dim(\nu_f) < 1$. It is known that $0 < \dim(\nu_f)$ [16, p. 162]. David Allwright of Oxford University has shown us a short proof that ν_f is singular with respect to the Lebesgue measure. We note that Allwright's proof relies on Theorems 30 and 31 of Khintchine [69]. The Hausdorff dimensions of very similar measures have been determined by Kinney and Pitcher [71]. There is also a conjecture by Ledrappier about $\dim(\nu_f)$ [74] [16, p. 162].

2.4 $e^{\gamma_f} = 1.13198824\dots$

Furstenberg's integral for γ_f (2.6) can be written as

$$\gamma_f = 2 \int_0^\infty \frac{1}{4} \log\left(\frac{1 + 4m^4}{(1 + m^2)^2}\right) d\nu_f(m)$$

because both the integrand and ν_f are symmetric about 0. In this section, we use this formula to compute γ_f with the help of a computer. Thus the determination of e^{γ_f} to be $1.13198824\dots$ is computer assisted. We will explain later why we report this result as a theorem (Theorem 2.6), even though it is computer assisted.

Let I_j^d , $1 \leq j \leq 2^d$, be the 2^d positive Stern-Brocot intervals at depth $d + 1$. Then,

$$p_d = 2 \sum_{j=1}^{2^d} \min_{m \in I_j^d} \text{amp}(m) \nu_f(I_j^d) < \gamma_f < q_d = 2 \sum_{j=1}^{2^d} \max_{m \in I_j^d} \text{amp}(m) \nu_f(I_j^d). \quad (2.10)$$

The inequalities above are strict because $\text{amp}(m)$ is not constant, and ν_f is continuous. Also, (2.10) defines p_d and q_d . Since γ_f is trapped in the intervals (p_d, q_d) , and the interval length $|q_d - p_d|$ shrinks to 0 as d increases, we can find γ_f to any desired accuracy by computing p_d and q_d for large enough d .

We computed p_d and q_d with $d = 28$ on a computer using IEEE double precision arithmetic (the C program used is described in the appendix). Computations in

floating point arithmetic are not exact, but when done carefully, they give an answer that is close to the exact answer. If $\text{fl}(e)$ denotes the number obtained by evaluating the expression e in floating point arithmetic, $\text{fl}(e)$ depends both on the type of floating point arithmetic used and the algorithm used to evaluate e . Our computations using IEEE double precision arithmetic [66] and an algorithm described in the appendix to this chapter gave

$$\text{fl}(p_{28}) = 0.1239755981508, \quad \text{fl}(q_{28}) = 0.1239755994406. \quad (2.11)$$

Below we give the hexadecimal codes for the 64 bits of $\text{fl}(p_{28})$ and $\text{fl}(q_{28})$ in IEEE double precision format.

$$\begin{aligned} \text{Hex code for } \text{fl}(p_{28}) &= 3fbfbcdd638f4d87 \\ \text{Hex code for } \text{fl}(q_{28}) &= 3fbfbcdd6919756d. \end{aligned} \quad (2.12)$$

The appendix will explain how to reproduce our computation to get exactly these two numbers. We will now upper bound the errors $|\text{fl}(p_{28}) - p_{28}|$ and $|\text{fl}(q_{28}) - q_{28}|$ to realize our aim of obtaining bounds for γ_f from (2.11).

IEEE double precision arithmetic (defined by the standard IEEE-754 [66]) can represent all real numbers of binary form $(-1)^s b_0 . b_1 \dots b_{52} 2^{e-1023}$ exactly. Here, $b_0 = 1$, the bits b_1 to b_{52} can be 1 or 0, the sign bit s can be 1 or 0, and the biased exponent e can be any integer in the range $0 < e < 2047$. The number 0 can also be represented exactly. In fact, the values $e = 0$ and $e = 2047$ are used to implement special features that we do not describe. From here on, floating point arithmetic always refers to IEEE double precision arithmetic, and a floating point number refers to a number in that arithmetic. Thus if a is a real number in the range $[2^{-1022}, (1 + 2^{-1} + \dots + 2^{-52})2^{1023}]$, a can be represented in such a way that $\text{fl}(a) = a(1 + E)$ with the relative error E satisfying $|E| < 2^{-52}$ [64, p. 42].

The IEEE standard treats $+$, $-$, \times , \div , $\sqrt{}$ as basic operations. The basic operations cannot always be performed exactly. For example, the sum of two floating point numbers may not have an exact floating point representation. However, all these basic operations are performed as if an intermediate result correct to infinite precision is coerced into a representable number by rounding. We assume the ‘‘round to nearest’’ mode which is the default type of rounding. Thus if a and b are floating point numbers,

$$\begin{aligned} \text{fl}(a + b) &= (a + b)(1 + E) \\ \text{fl}(a - b) &= (a - b)(1 + E) \\ \text{fl}(a/b) &= (a/b)(1 + E) \\ \text{fl}(a \times b) &= (a \times b)(1 + E) \\ \text{fl}(\sqrt{a}) &= (\sqrt{a})(1 + E), \end{aligned} \quad (2.13)$$

where the relative error E may depend upon a , b , and the operation performed, but $|E| < 2^{-52}$. For convenience, we denote 2^{-52} by u ². For (2.13) to be valid, however,

²The bounds on $|E|$ can be taken as 2^{-53} [64, p. 42], but with the current choice the relative error of Tang’s log function (see (2.14)) has the same bound as that of the basic operations.

the operation should not overflow and produce a number that is too big to be represented, or underflow and produce a number that is too small to be represented.

The C program we give in the appendix uses a function $\mathbf{tlog}(x)$ to compute $\log x$. This becomes necessary because \log is not a basic operation in the IEEE standard. What is special about $\mathbf{tlog}()$ is that it is implemented so that

$$\mathbf{fl}(\log a) = \log a(1 + E) \quad (2.14)$$

with $|E| < u$ whenever a is a positive floating point number. For the clever ideas that go into $\mathbf{tlog}()$ and the error analysis, see the original paper by Tang [105].

The proof of the following lemma is given in the appendix.

Lemma 2.5. *Assume that (2.13) and (2.14) hold with $0 < u < 1/10$ for the floating point arithmetic used. Then for the algorithm to compute the sums p_d and q_d described in the appendix,*

$$\begin{aligned} |\mathbf{fl}(p_d) - p_d| &< \frac{\log 4}{4}(e^{u(d+1)} - 1) + \frac{33}{4}ue^{u(d+1)}, \\ |\mathbf{fl}(q_d) - q_d| &< \frac{\log 4}{4}(e^{u(d+1)} - 1) + \frac{33}{4}ue^{u(d+1)}. \end{aligned}$$

In the theorem below, by 1.13198824... we mean a number in the interval $[1.13198824, 1.13198825)$.

Theorem 2.6. (a) *The constant γ_f lies in the interval*

$$(0.1239755980, 0.1239755995).$$

(b) $e^{\gamma_f} = 1.13198824\dots$

(c) *As $n \rightarrow \infty$,*

$$\sqrt[n]{|t_n|} \rightarrow 1.13198824\dots$$

with probability 1.

Proof. In the computation leading to $\mathbf{fl}(p_{28})$ and $\mathbf{fl}(q_{28})$, there are no overflows or underflows, and hence, (2.13) and (2.14) always hold. Therefore, we can use $u = 2^{-52}$ and $d = 28$ in Lemma 2.5 to get

$$|\mathbf{fl}(p_{28}) - p_{28}| < 10^{-14}, \quad |\mathbf{fl}(q_{28}) - q_{28}| < 10^{-14}.$$

Now the values of $\mathbf{fl}(p_{28})$ and $\mathbf{fl}(q_{28})$ in (2.11) imply (a). (b) is implied by (a). In fact, we can also say that the digit of e^{γ_f} after the last 4 in (b) must be an 8 or a 9. (c) follows from earlier remarks. \square




Figure 2.5: The Lyapunov exponent $\gamma_f(p)$ vs. p . $\gamma_f(p)$ is determined by numerically approximating the correct invariant distribution for the given p . For a description of the numerical method, sometimes called Ulam’s method, see [42] or [65].

Theorem 2.6 above is the main result of this chapter. We arrived at Theorem 2.6 using Lemma 2.5 and rounding error analysis. An alternative is to use interval arithmetic to validate the computation [5]. Instead of rounding the computations to the nearest floating point number, interval arithmetic carefully rounds the various stages of the computation either upwards or downwards to compute a lower bound for p_d and an upper bound for q_d . As a result, were we to use interval arithmetic there would be no need for rounding error analysis. A disadvantage would be that the manipulation of rounding modes necessary for implementing interval arithmetic would make it significantly more expensive on most computers. Our approach exposes the ideas behind floating point arithmetic and shows that floating point arithmetic can be rigorous too. Besides, the rounding error analysis as summarized by Lemma 2.5 gives a clear idea of the error due to rounding. This tells us, for example, that the rounding errors $|\text{fl}(p_{28}) - p_{28}|$ and $|\text{fl}(q_{28}) - q_{28}|$, which are both less than 10^{-14} , are much smaller than the discretization error $|p_{28} - q_{28}|$, which is about 10^{-8} .

Since the proof of Theorem 2.6 relies on a computer calculation, the validity of the proof requires some comment. The construction of ν_f in Section 2.2, the program and the rounding error analysis given in the appendix can all be checked line by line. However, Theorem 2.6 still assumes the correct implementation of various software and hardware components including the standard IEEE–754. We did the computation on two entirely different systems — SUN’s Sparc server 670 MP, and Intel’s i686 with the Linux operating system. In both cases, the results were exactly the same as given in (2.11); the hex codes for $\text{fl}(p_d)$ and $\text{fl}(q_d)$ matched the hex codes given in (2.12). As it is very unlikely that two systems with such different architectures may have the same bug, we feel that the correctness of Theorem 2.6 should, at worst, be doubted no more than that of tedious and intricate proofs that can be checked line by line. Though the use of floating point arithmetic to prove

a theorem may be unusual, the proof of Theorem 2.6 is only as dependent on the correctness of the computer system as, say, the proof of the four-color theorem; in other words, assuming the implementation of IEEE arithmetic to be correct is just like assuming the implementation of a memory-to-register copy instruction to be correct.

Besides, all components of a computer system, like mathematical proofs, can be checked in careful line by line detail, and this is done many times during and after their implementation. However, experience has shown that some bugs can defy even the most careful scrutiny. A great deal of research has gone into developing systems to verify that hardware and software implementations meet their specification [23].

To conclude, we ask: Is there a short analytic description of γ_f ? The fractal quality of γ_f suggests no. But let $\gamma_f(p)$ be the Lyapunov exponent of the obvious generalization $t_1 = t_2 = 1$, and for $n \geq 2$, $t_n = \pm t_{n-1} \pm t_{n-2}$ with each \pm sign independent and either $+$ with probability p or $-$ with probability $1 - p$. Unfortunately, the techniques described in this chapter for $\gamma_f(1/2)$ do not seem to generalize easily to $\gamma_f(p)$, $0 < p < 1$. A beautiful result of Peres [92] implies that $\gamma_f(p)$ is a real analytic function of p . See Figure 2.5. The analyticity of $\gamma_f(p)$ vs. p seems to increase the possibility that there might be a short analytic description of γ_f .

2.5 Appendix : Rounding Error Analysis

The main steps in the computation of p_d and q_d are the computation of $\nu_f(I_j^d)$, where I_j^d , $1 \leq j \leq 2^d$, are the 2^d positive Stern-Brocot intervals of depth $d + 1$; the minimization and maximization of $\text{amp}(m)$ over I_j^d ; and the summation over $1 \leq j \leq 2^d$ as in the defining equation (2.10). We describe some aspects of the computation and then give a rounding error analysis to prove Lemma 2.5. A C program for computing p_d and q_d for $d = 28$ is given at the end of this section so that our computation can be reproduced; its perusal is not necessary for reading this section.

Lemma 2.2 implies that the denominators of the 2^d positive Stern-Brocot intervals of depth $d + 1$ occur in an order that is the reverse of the order of the numerators. For example, the positive Stern-Brocot intervals of depth 4 are defined by divisions at the points $\frac{0}{1}, \frac{1}{3}, \frac{1}{2}, \frac{2}{3}, \frac{1}{1}, \frac{3}{2}, \frac{2}{1}, \frac{3}{1}, \frac{1}{0}$, the numerators for that depth occur in the order 0, 1, 1, 2, 1, 3, 2, 3, 1, and the denominators occur in the reverse order 1, 3, 2, 3, 1, 2, 1, 1, 0. We use this fact to avoid storing the denominators of the Stern-Brocot divisions. The numerators are stored in the array `num[]` by the C program.

To compute p_d and q_d , we use (2.10) in the following form:

$$\begin{aligned} p_d &= \sum_{j=1}^{2^d} \min_{m \in I_j^d} \left(\log \frac{1 + 4m^4}{(1 + m^2)^2} \right) \frac{\nu_f(I_j^d)}{2}, \\ q_d &= \sum_{j=1}^{2^d} \max_{m \in I_j^d} \left(\log \frac{1 + 4m^4}{(1 + m^2)^2} \right) \frac{\nu_f(I_j^d)}{2}. \end{aligned} \quad (2.15)$$

By (2.9), $\nu_f(I_j^d)/2$ is one of the $d + 1$ numbers $g^{d-i}(1 + g)^{-d}/4$, $0 \leq i \leq d$, where $g = (1 + \sqrt{5})/2$. The array `table[]` in the C program is initialized after precomputing these $d + 1$ numbers to very high accuracy in the symbolic algebra system *Mathematica* so that `table[i] = (g^{d-i}(1 + g)^{-d}/4)(1 + E)` with the relative error E satisfying $|E| < u$. The index i into `table[]` for getting $\nu_f(I_j^d)/2$ is obtained by taking the binary representation of j , flipping all the odd bits if d is even and all the even bits if d is odd with the least significant bit taken as an even bit, and then counting the number of 1s; correctness of this procedure can be proved easily using induction.

The minimization or the maximization of $4 \operatorname{amp}(m)$ over I_j^d in (2.15) are easy to do. Since $\operatorname{amp}(m)$ has its only local minimum for $m \geq 0$ at $m = 1/2$ (see Figure on page 6), both the minimum and the maximum are at the endpoints of I_j^d .

The summations in (2.15) are performed pairwise, not left to right. The pairwise summation of 2^d numbers is done by dividing the 2^d numbers into 2^{d-1} pairs of adjacent numbers, adding each pair to get 2^{d-1} numbers, and then reducing the 2^{d-1} numbers to 2^{d-2} numbers similarly, and so on until a single number is obtained. Rounding error analysis leads to smaller upper bounds on $|\operatorname{fl}(p_d) - p_d|$ and $|\operatorname{fl}(q_d) - q_d|$ for pairwise summation than for term-by-term left to right summation [64, p. 92]. The bounds for left to right summation are not small enough to give e^{7f} correctly to the 8 decimal digits shown in Theorem 2.6.

Lemmas 2.7 and 2.8 help simplify the proof of Lemma 2.5.

Lemma 2.7. *Assume $0 < f_1(u) < 1 + e_1 < g_1(u)$ and $0 < f_2(u) < 1 + e_2 < g_2(u)$.*

(a) *If $a > 0$, $b > 0$, and $a(1 + e_1) + b(1 + e_2) = (a + b)(1 + E)$, then $\min(f_1(u), f_2(u)) < 1 + E < \max(g_1(u), g_2(u))$.*

(b) *If $1 + E = (1 + e_1)(1 + e_2)$, then $f_1(u)f_2(u) < 1 + E < g_1(u)g_2(u)$.*

(c) *If $1 + E = (1 + e_1)/(1 + e_2)$, then $f_1(u)/g_2(u) < 1 + E < g_1(u)/f_2(u)$.*

Proof. To prove (a), note that $1 + E$ is the weighted mean of $1 + e_1$ and $1 + e_2$. (b) and (c) are trivial. \square

Consider the computation $\operatorname{fl}(m^2)$:

$$\operatorname{fl}(m^2) = \operatorname{fl}(m) \operatorname{fl}(m)(1 + e') = m^2(1 + e')(1 + e'')^2,$$

where e'' is the relative error in representing m , and e' is the relative error caused by rounding the multiplication. By (2.13) and remarks in the paragraph preceding it, $1 - u < 1 + e'$, $1 + e'' < 1 + u$. Lemma 2.7b allows us to gather the factors $1 + e'$ and $(1 + e'')^2$ together and write

$$\text{fl}(m^2) = m^2(1 + e_0), \quad (2.16)$$

with $(1 - u)^3 < 1 + e_0 < (1 + u)^3$.

Consider the computation $\text{fl}(1 + m^2)$:

$$\text{fl}(1 + m^2) = (1 + \text{fl}(m^2))(1 + e''') = (1 + m^2(1 + e'))(1 + e'')^2(1 + e'''),$$

where e''' is the relative error in the addition $1 + m^2$, and e'' , e' are, as before, the relative errors in representing m and the multiplication $m \times m$, respectively. As it was with $1 + e'$ and $1 + e''$, $1 - u < 1 + e''' < 1 + u$ by (2.13), and we can use Lemma 2.7a to pull $(1 + e')(1 + e'')^2$ out of the sum $1 + m^2$, and Lemma 2.7b to multiply $(1 + e')(1 + e'')^2(1 + e''')$ to get

$$\text{fl}(1 + m^2) = (1 + m^2)(1 + e'_0), \quad (2.17)$$

with $(1 - u)^4 < 1 + e'_0 < (1 + u)^4$.

Thus Lemma 2.7 allows us to pull factors like $(1 + e_i)$ out of sums (Lemma 2.7a), or to multiply them together (Lemma 2.7b), or to divide between them (Lemma 2.7c). Rounding error analyses of simple computations, like the analyses of $\text{fl}(m^2)$ and $\text{fl}(1 + m^2)$ given above, feature three steps. First, relative errors e_i caused by rounding are assigned to all the basic operations. Second, $1 + e_i$ are bounded using (2.13) or (2.14). Third, factors like $(1 + e_i)$ are gathered together using Lemma 2.7. In the proof of Lemma 2.5, we always spell out the first step in detail, but sometimes omit details for the second and third steps.

The inequalities in Lemma 2.8 below are used in the proof of Lemma 2.5.

Lemma 2.8. (a) *If $0 < u < 1/4$, $\log \frac{1+u}{1-u} < 3u$.*

(b) *$(1 + \alpha)^d < e^{\alpha d}$ for $\alpha > 0$ and d a positive integer.*

Proof. It is easy to prove (a) by expanding $\log((1 + u)/(1 - u))$ in a series. (b) can be proved by comparing the binomial expansion of $(1 + \alpha)^d$ with the series expansion of $e^{\alpha d}$. \square

The summations in the proof below are all over $1 \leq j \leq 2^d$.

Proof of Lemma 2.5. We will prove the upper bound only for $|\text{fl}(p_d) - p_d|$. The proof for $|\text{fl}(q_d) - q_d|$ is similar.

Firstly, consider the computation of $4 \text{amp}(m) = \log \frac{1+4m^4}{(1+m^2)^2}$:

$$\text{fl}\left(\log \frac{1 + 4m^4}{(1 + m^2)^2}\right) = \log\left(\frac{(1 + 4m^4(1 + e_0)^2(1 + e_1)(1 + e_2))(1 + e_3)}{(1 + m^2)^2(1 + e'_0)^2(1 + e_4)}(1 + e_5)\right) \\ (1 + e_6),$$

where e_0 and e'_0 are the relative errors in $\text{fl}(m^2)$ and $\text{fl}(1+m^2)$ as in (2.16) and (2.17), respectively, e_1, e_2 are the relative errors of the two multiplications $(4 \times m^2) \times m^2$, e_3 of the addition $1 + 4m^4$, e_4 of the multiplication $(1 + m^2) \times (1 + m^2)$, e_5 of the division $(1 + 4m^4)/(1 + m^2)^2$, and e_6 of taking the log. By assumptions (2.13) and (2.14), $1 - u < 1 + e_i < 1 + u$ for $1 \leq i \leq 6$. Lemma 2.7 gives

$$\text{fl}\left(\log \frac{1 + 4m^4}{(1 + m^2)^2}\right) = \left(\log \frac{1 + 4m^4}{(1 + m^2)^2}\right)(1 + E_1) + E_2, \quad (2.18)$$

with $1 - u < 1 + E_1 < 1 + u$ and $|E_2| < (1 + u) \log((1 + u)^{10}(1 - u)^{-9})$. A weaker, but simpler, bound is $|E_2| < 10(1 + u) \log((1 + u)/(1 - u))$. Now, the assumption $u < 1/10$ implies $10(1 + u) < 11$, which together with Lemma 2.8b, gives the simple bound $|E_2| < 33u$.

Secondly, recall that $\nu_f(I_j^d)/2$ is obtained by precomputing $g^{d-i}(1 + g)^{-d}/4$ to high precision. Therefore,

$$\text{fl}(\nu_f(I_j^d)/2) = \frac{\nu_f(I_j^d)}{2}(1 + E_3), \quad (2.19)$$

with $|E_3| < u$.

Finally, consider the pairwise summation to compute p_d . Let m_j be the endpoint of I_j^d where $\text{amp}(m)$ is minimum. Then,

$$\text{fl}(p_d) = \sum \left(\log \frac{1 + 4m_j^4}{(1 + m_j^2)^2} (1 + E_1^j) + E_2^j \right) \left(\frac{\nu_f(I_j^d)}{2} (1 + E_3^j) \right) (1 + E_4^j)$$

where E_1^j and E_2^j are the relative errors in computing $\log((1 + 4m_j^4)(1 + m_j^2)^{-2})$, and therefore, are bounded like E_1 and E_2 in (2.18); E_3^j is the relative error in computing $\nu_f(I_j^d)/2$ and is bounded like E_3 in (2.19); and the factors $1 + E_4^j$ take up the errors in the pairwise summation. By Higham [64, p. 91], E_4^j can be chosen so that $(1 - u)^d < 1 + E_4^j < (1 + u)^d$. Lemma 2.7 gives

$$\text{fl}(p_d) = \frac{1}{2} \sum \log \frac{1 + 4m_j^4}{(1 + m_j^2)^2} \nu_f(I_j^d) (1 + E_a^j) + \frac{1}{2} \sum \nu_f(I_j^d) E_b^j \quad (2.20)$$

with $(1 - u)^{d+2} < 1 + E_a^j < (1 + u)^{d+2}$ and $|E_b^j| < 33u(1 + u)^{d+1}$.

Bounding $|\text{fl}(p_d) - p_d|$ is now a simple matter:

$$\begin{aligned} |\text{fl}(p_d) - p_d| &< \frac{1}{2} \sum \left| \log \frac{1 + 4m_j^4}{(1 + m_j^2)^2} \nu_f(I_j^d) \right| |E_a^j - 1| + \frac{1}{2} \sum \nu_f(I_j^d) |E_b^j| \\ &< \frac{\log 4}{4} ((1 + u)^{d+2} - 1) + \frac{33}{4} u (1 + u)^{d+1} \\ &< \frac{\log 4}{4} (e^{u(d+2)} - 1) + \frac{33}{4} u e^{u(d+1)}. \end{aligned}$$

The second inequality above uses $\sum \nu_f(I_j^d) = 1/2$, $|\log \frac{1+4m^4}{(1+m^2)^2}| < \log 4$, $|E_a^j - 1| < (1+u)^{d+2} - 1$, and $|E_b^j| < 33u(1+u)^{d+1}$. The bound on $|E_a^j - 1|$ can be derived easily from $(1-u)^{d+2} < 1 + E_a^j < (1+u)^{d+2}$. The final inequality follows from Lemma 2.8b. \square

Upper bounding $|\text{fl}(q_d) - q_d|$ involves a small, additional detail. For the rightmost positive Stern-Brocot interval I_j^d , $\text{amp}(m)$ is maximum at $m = \infty$. This causes no difficulty, however, because $\log((1+4m^4)/(1+m^2)^2)$ is taken as $\log 4$ at $m = \infty$ by the computation, and as a result, the bounds in (2.18) still hold.

A program to compute p_d and q_d is given on page 50 so that the computation leading to (2.11) can be easily reproduced. The program uses up 1.1 gigabytes of memory. It can be written using only a small amount of memory, but then it would be harder to read. For finding logs, we used the version of Tang's algorithm [105] that does not precompute and store $1/F$ for $F = 1 + j2^{-7}$, $0 \leq j \leq 128$. Though we do not give the code here because it is machine dependent, the guidelines given in [105] are enough to reproduce that log function (called `tlog()` in the program) exactly.

```
#include <stdlib.h>
#include <stdio.h>
#define D 28
#define N 268435456
#define NRT 16384
unsigned int filter = 0xAFFFFFFF;

#define bitcount(x,b) \
{b = 0; \
 for( ; x!=0; x&=(x-1)) \
 b++; \
}
double tlog(double);
double sum(double *, int);

static double table[D+1] = {
  3.51792099313013395856e-7,
  2.17419474349120812252e-7,
  1.34372624963892583604e-7,
  8.30468493852282286483e-8,
  5.13257755786643549553e-8,
  3.17210738065638736930e-8,
  1.96047017721004812623e-8,
  1.21163720344633924307e-8,
  7.48832973763708883155e-9,
```

```

4.62804229682630359918e-9,
2.86028744081078523237e-9,
1.76775485601551836682e-9,
1.09253258479526686555e-9,
6.75222271220251501272e-10,
4.17310313575015364275e-10,
2.57911957645236136997e-10,
1.59398355929779227278e-10,
9.85136017154569097184e-11,
6.08847542143223175599e-11,
3.76288475011345921584e-11,
2.32559067131877254014e-11,
1.43729407879468667570e-11,
8.88296592524085864439e-12,
5.48997486270600811265e-12,
3.39299106253485053174e-12,
2.09698380017115758091e-12,
1.29600726236369295083e-12,
8.00976537807464630088e-13,
4.95030724556228320737e-13};

```

```

void main()
{
    int i,j,*num;
    double lower,upper,larray1[NRT],larray2[NRT],
           uarray1[NRT],uarray2[NRT];
    unsigned int *lptr, *uptr;

    num = (int *)malloc(sizeof(int)*(N+1));
    num[0] = 1; num[1]=1;
    for(i=2;i<N;i=i+2){
        num[i] = num[i/2];
        num[i+1] = num[i/2]+num[i/2+1];}
    num[N] = 1;

    for(i=0; i<NRT; i++){
        unsigned int k,b,x; double m, m2, m2p1,
            left, right, measure;

        k = i*NRT; m =(double)num[k]/(double)num[N-k];
        m2 = m*m; m2p1 = m2+ 1.0;

```

```

left = tlog((1+4*m2*m2)/(m2p1*m2p1));

if (i < NRT/4)
  for(j=0; j<NRT; j++){
    k = i*NRT+j;
    m = (double)num[k+1]/(double)num[N-k-1];
    m2 = m*m;
    m2p1 = 1 + m2;
    right = tlog((1+4*m2*m2)/(m2p1*m2p1));
    x = k^filter;
    bitcount(x,b);
    measure = table[b];
    larray1[j] = measure*right; uarray1[j] = measure*left;
    left = right;}
else if(i < NRT-1)
  for(j=0; j<NRT; j++){
    k = i*NRT+j;
    m = (double)num[k+1]/(double)num[N-k-1];
    m2 = m*m;
    m2p1 = 1 + m2;
    right = tlog((1+4*m2*m2)/(m2p1*m2p1));
    x = k^filter;
    bitcount(x,b);
    measure = table[b];
    larray1[j] = measure*left; uarray1[j] = measure*right;
    left = right;}
else /* i == NRT-1 */
  for(j=0; j<NRT; j++){
    k = i*NRT+j;
    if(j==NRT-1)
      right = tlog(4.0);
    else{
      m = (double)num[k+1]/(double)num[N-k-1];
      m2 = m*m;
      m2p1 = 1 + m2;
      right = tlog((1+4*m2*m2)/(m2p1*m2p1));}
    x = k^filter;
    bitcount(x,b);
    measure = table[b];
    larray1[j] = measure*left; uarray1[j] = measure*right;
    left = right;}
larray2[i] = sum(larray1,NRT); uarray2[i] = sum(uarray1,NRT);}

lower = sum(larray2,NRT);

```

```
upper = sum(uarray2,NRT);

lptr = (unsigned int *)&lower;
uptr = (unsigned int *)&upper;
printf("(l,r)= (%.17E, %.17E)\n",lower, upper);
printf("(l,u) in hex = (%x %x, %x %x)\n",*lptr,*(lptr+1),
*uptr, *(uptr+1));
}

/* sums a list, length being a power of 2 */
double sum(double *list, int length)
{
    int i,step;

    for(step = 1; step < length; step = 2*step)
        for(i=0; i < length; i += 2*step)
            list[i]+= list[i+step];

    return list[0];
}
```

Chapter 3

A Brief Overview of Random Matrix Products

Our analysis of random Fibonacci sequences led to a product of random 2×2 matrices, which is one example of a random matrix product. The product $P_n = M_n M_{n-1} \dots M_1$ can be formed by picking M_i from a probability distribution μ over $d \times d$ real matrices. The extension of the theory to complex matrices, which is straightforward, will not be discussed. When $d = 1$, $\log|P_n|$, being a sum of independent and identically distributed scalar variables, is governed by the strong law of large numbers and the central limit theorem in the limit $n \rightarrow \infty$. Analogous results about $\log\|P_n\|$ in the limit $n \rightarrow \infty$ for $d > 1$ form the foundation of the theory of random matrix products.

The account in this chapter lacks the vivid description of detail and comprehensiveness necessary for satisfactory understanding. Our aim is merely to highlight some important elements and provide support for comments made elsewhere in this thesis. None of the results in this chapter is original, but since their importance is so well established, we omit references to the original papers. The original papers [46], [47], [59], [75] are all most remarkable. Bougerol's monograph [16] is detailed, comprehensive, and historically accurate.

3.1 Existence of Lyapunov Exponents and the Theorem of Furstenberg and Kesten

The following theorem is basic. For a real-valued function f , $f^+(x) = \max(f(x), 0)$.

Theorem 3.1. *Assume that $E \log^+ \|M_1\| < \infty$ for a distribution μ on $R^{d,d}$. Then*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \|P_n\| = \gamma$$

with probability 1 for a constant $\gamma < \infty$.

A key step in the proof of Theorem 3.1 is an application of the ergodic theorem. We will only exhibit the stationary sequence that the ergodic theorem is applied to. Consider the sequence

$$((M_n)_{n \geq 1}, Y_1), ((M_n)_{n \geq 2}, Y_2), ((M_n)_{n \geq 3}, Y_3), \dots,$$

where $\|Y_i\| = 1$, $Y_k = \frac{M_{k-1}Y_{k-1}}{\|M_{k-1}Y_{k-1}\|}$ for $k \geq 2$, and the distribution of Y_1 is invariant under the action $Y_1 \rightarrow \frac{M_1 Y_1}{\|M_1 Y_1\|}$, the distribution of M_1 being μ . There is always such an invariant distribution for Y_1 . It is easy to show this sequence is stationary. Consequently, the entrywise function of this sequence $\log \|M_1 Y_1\|, \log \|M_2 Y_2\|, \log \|M_3 Y_3\|, \dots$ is also stationary. Since $\sum_{i=1}^n \log \|M_i Y_i\| = \log \|P_n Y_1\|$, Birkhoff's ergodic theorem implies that $\frac{1}{n} \log \|P_n Y_1\|$ converges both almost surely and in the L^1 norm to an integrable random variable. The proof relies on both kinds of convergence. The proof simplifies significantly if one makes the more restrictive assumption $E|\log \|M_1\|| < \infty$ about μ .

There is a result of basic probability whose proof is similar to that of Theorem 3.1. Let x_i , $i \geq 1$, be independent, identically distributed variables which take only integer values. Let $S_k = x_1 + \dots + x_k$, and let N_k be the cardinality of the set $\{S_1, \dots, S_k\}$. Assume further that $E|x_1| < \infty$. Then,

$$\lim_{k \rightarrow \infty} \frac{N_k}{k} = P(S_1 \neq 0, S_2 \neq 0, S_3 \neq 0, \dots)$$

with probability 1. For the proof, see Breiman [18]. We note that it is also possible to define a subadditive process corresponding to N_k .

The existence of the full set of Lyapunov exponents can be easily derived from Theorem 3.1. One has only to apply Theorem 3.1 to compound matrices formed from the M_i .

Theorem 3.2. *Assume $E \log^+ \|M_1\| < \infty$ as before. Let $\sigma_i(M)$ be the i th largest singular value of the matrix M . Then,*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \sigma_i(P_n) = \gamma_i$$

with probability 1 for $1 \leq i \leq d$ and constants $\gamma_1 \geq \dots \geq \gamma_d$.

The constants γ_i , $1 \leq i \leq d$, are the Lyapunov exponents of the random matrix product formed using μ . For convenience, we will call them the Lyapunov exponents of μ .

3.2 Furstenberg's Theorem

Let us denote the group of invertible matrices in $R^{d,d}$ by $\text{Gl}(d, R)$. In the rest of this chapter, μ is a distribution over $\text{Gl}(d, R)$ unless stated otherwise. Every matrix M in $\text{Gl}(d, R)$ induces a bijection from the projective space $P(R^d)$ to itself; an element of $P(R^d)$ in the direction of vector $x \in R^d$ is mapped to the element of $P(R^d)$ that gives the direction of Mx . Let us denote the element of $P(R^d)$ in the direction x by \bar{x} as in Chapter 2. As in Chapter 2, x can be any nonzero vector in the direction \bar{x} . Let ν be a probability distribution over $P(R^d)$. Then the probability distribution $\mu * \nu$ on $P(R^d)$ is defined by

$$\mu * \nu(A) = \int \chi_A(\overline{Mx}) d\mu(M) d\nu(\bar{x}),$$

where A is any Borel measurable subset of $P(R^d)$ and χ_A is its indicator function. The distribution ν is said to be μ -invariant if $\mu * \nu = \nu$.

An easy compactness argument guarantees the existence of an invariant distribution ν for every μ . But to set up Furstenberg's theory, one requires a μ -invariant measure that is also *proper*; a probability distribution ν is said to be proper if $\nu(\overline{V}) = 0$ — where V is a subspace of R^d and \overline{V} is the corresponding set of directions in $P(R^d)$ — whenever V is a proper subspace of R^d . To be sure of finding a proper μ -invariant measure, the support of μ is assumed to be *strongly irreducible*.

A subset S of $\text{Gl}(R, d)$ is said to be *strongly irreducible* if there does not exist a finite family of proper linear subspaces V_i , $1 \leq i \leq k$, of R^d of the same dimension and with $V_i \cap V_j = \{0\}$ for $i \neq j$ such that MV_i is one of the subspaces V_j , $1 \leq j \leq k$, for every $M \in S$ and $1 \leq i \leq k$. The smallest closed subgroup of $\text{Gl}(d, R)$ containing the support of μ is denoted by G_μ . The group G_μ is strongly irreducible if and only if the support of μ is strongly irreducible.

In the formula for γ_1 in the theorem below, $\|\cdot\|$ can be any vector norm over R^d . The formula is sometimes called Furstenberg's formula.

Theorem 3.3. *Let μ be a probability distribution over $\text{Gl}(d, R)$, with G_μ being strongly irreducible. Assume $E \log^+ \|M_1\| < \infty$. Then,*

(i) *For any vector $x \neq 0$,*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \|P_n x\| = \gamma_1 \quad \text{almost surely,}$$

where γ_1 is the top Lyapunov exponent of μ .

(ii) *If ν is any μ -invariant probability distribution on $P(R^d)$, then*

$$\gamma_1 = \int \int \log \frac{\|Mx\|}{\|x\|} d\mu(M) d\nu(\bar{x}).$$

The μ -invariant probability measure is not necessarily unique.

The proof of the first part of this theorem relies on a brilliant use of martingales which is essentially due to Furstenberg. Without going into details, we will show how random matrix products lead to martingales. Let ν be a μ -invariant probability distribution of $P(R^d)$. Let $f(\bar{x})$ be a bounded, measurable, real-valued function on $P(R^d)$. Let ω be a sample point in a probability space such that $M_1(\omega), M_2(\omega), \dots$ are independent matrix-valued random variables with the distribution μ . Consider the random variables defined by

$$\begin{aligned} r_1(\omega) &= \int f(\overline{M_1(\omega)x}) d\nu(\bar{x}), \\ r_2(\omega) &= \int f(\overline{M_1(\omega)M_2(\omega)x}) d\nu(\bar{x}), \\ &\vdots \\ r_n(\omega) &= \int f(\overline{M_1(\omega) \dots M_n(\omega)}) d\nu(\bar{x}). \end{aligned}$$

Note that the matrix products in the equations above are formed by multiplying on the right. It is easy to show that $r_1(\omega), r_2(\omega), \dots$ is a bounded martingale for any bounded measurable f . Some simple martingale calculations and an application of the martingale convergence theorem yield a great deal of information about the product $P_n(\omega) = M_n(\omega) \dots M_1(\omega)$ in the limit $n \rightarrow \infty$.

The following theorem is usually known as Furstenberg's theorem.

Theorem 3.4. *Let μ be a probability distribution over $\{M \in \text{Gl}(R, d) \mid |\det M| = 1\}$, with G_μ strongly irreducible. If G_μ is not compact, then $\gamma_1 > 0$.*

If G_μ is compact, the Haar measure on G_μ can be used to deduce that all matrices in G_μ are actually unitary for an appropriate inner product.

What happens when G_μ is not strongly irreducible? Then G_μ may be *reducible* — i.e., there may exist a proper subspace V of R^d which is invariant under every member of the group G_μ — or it may be *irreducible*. When G_μ is reducible, natural restrictions of G_μ to V and to the quotient space R^d/V induce two groups of matrices of dimension k and $d - k$, respectively, where k is the dimension of V . The Lyapunov exponents of the distributions induced by μ on these two groups give all the Lyapunov exponents of μ . Thus it suffices to understand Lyapunov exponents assuming G_μ to be irreducible.

Theorem 3.5. *Assume that G_μ is irreducible and that $E \log^+ \|M_1\| < \infty$.*

- (i) *Both assertions of Theorem 3.1 still hold.*
- (ii) *If there does not exist a probability distribution m on $P(R^d)$ such that $Mm = m$ for every $M \in G_\mu$, then not all Lyapunov exponents of μ are equal.*

But Theorem 3.3 is not true when G_μ is merely irreducible. Let μ pick one of the two matrices $\begin{pmatrix} 2 & 0 \\ 0 & 1/2 \end{pmatrix}$ or $\begin{pmatrix} 0 & 1/3 \\ 3 & 0 \end{pmatrix}$, the first with probability p and the second with

probability $1 - p$. If $0 < p < 1$, both Lyapunov exponents are zero even though G_μ is not compact.

Let us mention a result about the uniqueness of the μ -invariant measure over $P(R^d)$ due to Guivarc'h and Raugi. Let G_μ be strongly irreducible, and let there be a sequence of matrices in G_μ whose limit is a matrix of rank one. Then the μ -invariant measure must be unique. This result guarantees uniqueness of the μ -invariant measure for the μ we considered in Chapter 2 and for all the examples in [42].

The following result of Guivarc'h and Raugi was alluded to in Chapter 2. With $y = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and $x = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$, the theorem below implies that the growth rate of $|t_n|$ in Chapter 2 is given by the top Lyapunov exponent of the corresponding random matrix product.

Theorem 3.6. *Assume that the group G_μ is strongly irreducible, that there is a sequence of matrices in G_μ whose limit is a matrix of rank one, and that the expectation $E(\exp(\tau \log^+ \|M_1\| + \log^+ \|M_1^{-1}\|)) < \infty$ for some $\tau > 0$. Then*

$$\lim_{n \rightarrow \infty} \frac{1}{n} |y^* P_n x| = \gamma_1 \quad \text{almost surely,}$$

for any nonzero vectors $x, y \in R^d$. Here γ_1 is the top Lyapunov exponent of μ .

3.3 Computability of Lyapunov Exponents

Tsitsiklis and Blondel [110] claim that the largest Lyapunov exponent is not ‘‘algorithmically approximable.’’ The following theorem taken out their paper, with some minor modifications, clarifies their claim.

Theorem 3.7. *Consider the following algorithmic problem. The input is a pair of $d \times d$ matrices with integer entries. Let μ be a probability distribution which picks one of these two matrices, each with probability $1/2$. The desired output is an approximation $\tilde{\gamma}$ for the top Lyapunov exponent γ of μ such that*

$$|\tilde{\gamma} - \gamma| \leq \epsilon_a + \epsilon_r \gamma,$$

where ϵ_a and ϵ_r are the absolute and relative errors that are tolerated. If $d \geq 48$, there is no Turing machine that can solve this algorithmic problem for any fixed $\epsilon_a \geq 0$ and $0 \leq \epsilon_r < 1$; for example, ϵ_a and ϵ_r can both be fixed at $1/2$.

In a sense, the theorem of Tsitsiklis and Blondel is limited. Their proof works only when the class of problems includes μ that pick singular matrices. In fact, as the Theorem below suggests, approximating the top Lyapunov exponent is probably decidable when the support of μ is a finite subset of $\text{Gl}(d, R)$; but a precise proof of decidability should include many details we have not worked out. Approximating the top Lyapunov exponent can be reduced to finding the global minimum and the global maximum of a smooth function over the compact set $P(R^d)$.

Theorem 3.8. *Let G_μ be irreducible. Assume $E \log^+ \|M_1\| < \infty$ and $E \log^+ \|M_1^{-1}\| < \infty$. If γ_1 is the top Lyapunov exponent of μ , then*

$$\inf_{\bar{x} \in P(R^d)} \frac{1}{n} E \left(\log \frac{\|M_n \dots M_1 x\|}{\|x\|} \right) \leq \gamma_1 \leq \sup_{\bar{x} \in P(R^d)} \frac{1}{n} E \left(\log \frac{\|M_n \dots M_1 x\|}{\|x\|} \right),$$

and the upper and lower bounds both converge to γ_1 in the limit $n \rightarrow \infty$.

Proof. Proposition 7.2 and the proof of item (iii) of Corollary 3.4 in Chapter 3 of [16] can be used in combination to conclude that

$$\frac{1}{n} E \left(\log \frac{\|M_n \dots M_1 x\|}{\|x\|} \right) \rightarrow \gamma_1$$

as $n \rightarrow \infty$ uniformly in $\bar{x} \in P(R^d)$. Thus it suffices to prove the upper and lower bounds.

An easy calculation using Furstenberg's formula gives

$$\gamma_1 = \int \frac{1}{n} E \left(\log \frac{\|M_n \dots M_1 x\|}{\|x\|} \right) d\nu(\bar{x}),$$

where ν is μ -invariant; see [16]. The validity of the bounds is now obvious. \square

Theorem 3.8 is the only original result in this chapter. We hasten to add that it is a trivial consequence of known results, though its relevance to the computation of γ_1 was unnoticed. For a possible application of Theorem 3.8, consider a distribution μ_b which picks one of $\begin{pmatrix} 0 & 1 \\ \pm b & 1 \end{pmatrix}$ with probability 1/2. Asymptotic expressions for the top Lyapunov exponent were found using numerical experiments by Embree and Trefethen [42]; $\gamma_1 \sim -b^2/2$ as $b \rightarrow 0$ and $\gamma_1 \sim 2^{-1} \log b$ as $b \rightarrow \infty$. It appears that these expressions can be proven using Theorem 3.8 with $n = 2$ or $n = 3$.

The central limit theorem below is a product of the work of Bougerol, Le Page, and many others.

Theorem 3.9. *Let γ_1 be the top Lyapunov exponent of μ . Assume*

(i) *For some $\tau > 0$,*

$$E(e^{\tau(\log^+ \|M_1\| + \log^+ \|M_1^{-1}\|)}) < \infty$$

(ii) *G_μ is strongly irreducible*

(iii) *$\{|\det M|^{-1/d} M \mid M \in G_\mu\}$ is not compact.*

Then there exists $a > 0$ such that $\frac{1}{\sqrt{n}}(\log \|P_n x\| - n\gamma_1)$ for any $x \neq 0$ in R^d , and $\frac{1}{\sqrt{n}}(\log \|P_n\| - n\gamma_1)$ converge in distribution to $N(0, a^2)$ as $n \rightarrow \infty$.

None of the assumptions in the theorem above is restrictive, and the convergence to the normal law holds irrespective of the dimension d . Using compound matrices and further irreducibility assumptions, it is easy to show that

$$\frac{\log \sigma_1(P_n) + \cdots + \log \sigma_k(P_n) - n(\gamma_1 + \cdots + \gamma_k)}{n^{1/2}},$$

where $k \leq d$ and $\sigma_i(P_n)$ are the singular values of P_n with σ_1 being the largest, also converges in distribution to the normal law as $n \rightarrow \infty$.

Thus approximating γ_i as $\log \sigma_i(P_n)/n$ will have an error of about $n^{1/2}$. But this procedure is not algorithmically feasible without suitable normalization to avoid numerical overflows and underflows. The right way to normalize, which is similar to simultaneous iteration for computing eigenvalues, is briefly described in Section 5.2.

Chapter 4

Random Recurrences and Condition Numbers of Random Triangular Matrices ¹

4.1 Introduction

Random Fibonacci sequences have finite memory — every new term is a random sum of just the previous two terms. A generalization with unbounded memory is the following:

$$\begin{aligned}t_1 &= 1 \\t_2 &= \alpha_{21}t_1 \\t_3 &= \alpha_{31}t_1 + \alpha_{32}t_2 \\&\vdots \\t_k &= \alpha_{k1}t_1 + \cdots + \alpha_{k,k-1}t_{k-1},\end{aligned}$$

where α_{ij} , $i < j$ are independent and identically distributed. Thus, every t_k is a random sum of all the previous terms. Even if the signs of α_{ij} are random one might expect that the cancellation is rarely close to exact, and that $|t_k|$ will increase exponentially with k . We show in this chapter that

$$\sqrt[n]{|t_n|} \rightarrow C \quad \text{almost surely,}$$

for some $C > 1$, when α_{ij} are picked from a variety of distributions including the normal distribution and the Cauchy distribution. The formulas we derive for C are formally analogous to Furstenberg's formula for random matrix products.

Random recurrences with unbounded memory like the one above are related to random triangular matrices. Consider a unit lower triangular matrix L_n of dimension n whose subdiagonal entries are independent and identically distributed. If the ij th

¹This chapter is adapted from [114].




Figure 4.1: Empirical distribution functions of $\sqrt{\kappa_n}$, for triangular and unit triangular matrices with $n = 25, 50, 100$ obtained from 1000 random matrices for each n . The random entries are $N(0, 1)$ variables. The higher values of n correspond to the steeper curves. In the limit $n \rightarrow \infty$, the distribution functions converge to Heaviside step functions with jumps at the dashed lines.

entry for $i < j$ is denoted by $-\alpha_{ij}$, the entries of the first column of L_n^{-1} satisfy exactly the same recurrence as the t_i . When the diagonal entries of L_n are not fixed at 1 but are drawn from the same distribution, the recurrence satisfied by the entries of the first column of L_n^{-1} is only slightly different. This relationship between triangular inversion and random recurrences is stated in more detail in Section 4.2. If the terms t_k of the corresponding random recurrence increase exponentially with k , the norm $\|L_n^{-1}\|_2$ and the condition number $\|L_n\|_2\|L_n^{-1}\|_2$ also increase exponentially with n and at exactly the same rate. Our results for rate of growth constants (C in the previous paragraph) are derived for random recurrences and then stated for condition numbers.

For a concrete result, take L_n to be a random lower triangular matrix of dimension n with independent $N(0, 1)$ entries. If κ_n is the 2-norm condition number of L_n (defined as $\|L_n\|_2\|L_n^{-1}\|_2$), we prove that

$$\sqrt[n]{\kappa_n} \rightarrow 2 \text{ almost surely}$$

as $n \rightarrow \infty$ (Theorem 4.9). Figure 4.1a illustrates this result. The matrices that arise in the experiments reported in Figure 4.1 are so ill-conditioned that the normwise stable method of finding the condition number using the SVD [51] fails owing to rounding errors. The method used to generate the figures finds the inverse of the triangular matrix explicitly using the standard algorithm for triangular inversion, and then computes the norms of the matrix and its inverse independently. This works because the computation of each column of the inverse by the standard triangular inversion algorithm is componentwise backward stable [64].

Let us mention that random *dense* matrices are very well-conditioned. Edelman

has shown that if each of the n^2 entries of a matrix of dimension n is an independent $N(0, 1)$ variable, the probability density function (PDF) of κ_n/n , where κ_n is the 2-norm condition number of such a matrix, converges pointwise to the function

$$\frac{2x + 4}{x^3} \exp(-2x^{-1} - 2x^{-2})$$

as $n \rightarrow \infty$ [40]. Since the distribution of κ_n/n is independent of n in the limit $n \rightarrow \infty$, we can say that the condition numbers of random dense matrices increase only linearly with n . Using this PDF, it can be shown, for example, that $E(\log(\kappa_n)) = \log(n) + 1.537\dots + o(1)$ [40].

For a concrete result about unit triangular matrices, take L_n to be a unit lower triangular matrix of dimension n whose subdiagonal entries are independent $N(0, 1)$ variables. If κ_n is the condition number of L_n , then

$$\sqrt[n]{\kappa_n} \rightarrow 1.305683410\dots \text{ almost surely}$$

as $n \rightarrow \infty$ (Theorem 4.12). Obviously, in this case the ill-conditioning has nothing to do with the diagonal entries (i.e., the eigenvalues) since they are all equal to 1. The relationship of the exponential ill-conditioning of random unit triangular matrices to the stability of Gaussian elimination with partial pivoting is discussed in Section 4.7.

We will use L_n to refer to triangular matrices of various kinds — real or complex, with or without a unit diagonal. But L_n always denotes a lower triangular matrix of dimension n . If the entries of L_n are random variables, they are assumed to be independent. Thus, if we merely say that L_n has entries from a certain distribution, those entries are not only identically distributed but also independent. Of course, only the nonzero entries of L_n are chosen according to that distribution. The condition number always refers to the 2-norm condition number. However, all our results concerning the limits $\lim_{n \rightarrow \infty} \sqrt[n]{\kappa_n}$ apply to all the L_p norms, $1 \leq p \leq \infty$, since $n^{1/n} \rightarrow 1$ as $n \rightarrow \infty$ and the L_p norms differ by at most a factor of n . The 2-norm condition number of L_n , defined as $\|L_n\|_2 \|L_n^{-1}\|_2$, is denoted by κ_n . The context will make clear the distribution of the entries of L_n .

The analyses and discussions in this chapter are phrased for lower, not upper, triangular matrices. However, all the theorems are true for upper triangular matrices as well, as is obvious from the fact that a matrix and its transpose have the same condition number. The random recurrence that corresponds to the first column of the inverse of a random lower triangular matrix corresponds to the first row of the inverse of its transpose.

We obtain results about κ_n for triangular matrices with entries chosen from the complex normal distribution $\tilde{N}(0, \sigma^2)$. By $\tilde{N}(0, \sigma^2)$, we denote the complex normal distribution of mean 0 and variance σ^2 obtained by taking the real and imaginary parts as independent $N(0, \sigma^2/2)$ variables. Let L_n denote a triangular matrix with $\tilde{N}(0, \sigma^2)$ entries. Then,

$$\sqrt[n]{\kappa_n} \rightarrow e^{1/2} \text{ almost surely}$$

as $n \rightarrow \infty$ (Theorem 4.15). Since $e^{1/2} < 2$, triangular matrices with complex normal entries tend to have smaller condition numbers than triangular matrices with real normal distributed entries.

Similarly, let L_n denote a unit lower triangular matrix with $\tilde{N}(0, 1)$ subdiagonal entries. Then,

$$\sqrt[n]{\kappa_n} \rightarrow 1.347395784\dots \text{ almost surely}$$

as $n \rightarrow \infty$ (Theorem 4.16). Thus, unit triangular matrices with complex normal entries tend to have slightly bigger condition numbers than unit triangular matrices with real normal entries.

Our results are similar in spirit to results obtained by Silverstein for random dense matrices [96]. Consider a matrix of dimension $n \times (yn)$, where $y \in [0, 1]$, each of whose n^2y entries is an independent $N(0, 1)$ variable. Denote its largest and smallest singular values by σ_{\max} and σ_{\min} respectively. It is shown in [96] that

$$\frac{\sigma_{\max}}{\sqrt{n}} \rightarrow 1 + \sqrt{y}, \quad \frac{\sigma_{\min}}{\sqrt{n}} \rightarrow 1 - \sqrt{y} \text{ almost surely}$$

as $n \rightarrow \infty$. The complex analogues of these results can be found in [39]. The technique used in [96] is a beautiful combination of what is now known as the Golub-Kahan bidiagonalization step in computing the singular value decomposition with the Gerschgorin circle theorem and the Marčenko-Pastur semicircle law. The techniques we use are more direct.

Let us outline the approach for determining the rate of exponential growth of κ_n by assuming L_n triangular with $N(0, 1)$ entries. In Section 4.2, we derive the joint probability density function for the entries in any column of L_n^{-1} (Proposition 4.1). If T_k is the 2-norm of column $n - k + 1$ of L_n^{-1} , i.e., the column with k nonzero entries, both positive and negative moments of T_k are explicitly derived in Section 4.3 (Lemma 4.4). These moments allow us to deduce that $\sqrt[n]{\kappa_n}$ converges to 2 almost surely (Theorem 4.9). A similar approach is used to determine the limit of $\sqrt[n]{\kappa_n}$ for L_n unit triangular with $N(0, \sigma^2)$ entries, triangular with $\tilde{N}(0, \sigma^2)$ entries, and unit triangular with $\tilde{N}(0, \sigma^2)$ entries (Theorems 4.12, 4.15, and 4.16 respectively).

The same approach is used more generally to determine the limit of $\sqrt[n]{\kappa_n}$ as $n \rightarrow \infty$ for L_n with entries drawn from any symmetric, strictly stable distribution (Theorems 4.20 and 4.22). These theorems are specialized to the Cauchy distribution, which is symmetric and strictly stable, in Theorems 4.21 and 4.23.

In Section 4.9 we state a conjecture about random recurrences.




Figure 4.2: Entries of L_n^{-1} on the same solid line in (a) have the same PDF. Sets of entries of L_n^{-1} in the boxes in (b) have the same JPDF.

4.2 Inverse of a Random Triangular Matrix

Consider the matrix

$$L_n = \begin{pmatrix} \alpha_{11} & & & \\ -\alpha_{21} & \alpha_{22} & & \\ \vdots & \vdots & \ddots & \\ -\alpha_{n1} & -\alpha_{n2} & \dots & \alpha_{nn} \end{pmatrix},$$

where each α_{ij} is an independent $N(0, 1)$ variable. Then L_n^{-1} is also lower triangular. Denote the first k entries in the first column of L_n^{-1} by t_1, \dots, t_k . The t_i satisfy the following relations:

$$\begin{aligned} t_1 &= 1/\alpha_{11} \\ t_2 &= (\alpha_{21}t_1)/\alpha_{22} \\ t_3 &= (\alpha_{31}t_1 + \alpha_{32}t_2)/\alpha_{33} \\ &\vdots \\ t_k &= (\alpha_{k1}t_1 + \dots + \alpha_{k,k-1}t_{k-1})/\alpha_{kk}. \end{aligned} \tag{4.1}$$

This system of equations is interpreted as a system of random recurrence relations. The first entry t_1 is the reciprocal of an $N(0, 1)$ variable. The k th entry t_k is obtained by summing the previous entries t_1, \dots, t_{k-1} with independent $N(0, 1)$ variables as coefficients, and dividing that sum by an independent $N(0, 1)$ variable.

Next, consider an arbitrary column of L_n^{-1} and denote the first k entries of that column from the diagonal downwards by t_1, \dots, t_k . The entries t_i satisfy random recurrence relations similar in form to (4.1), but the α_{ij} are a different block of entries in L_n for different columns. For example, any diagonal entry of L_n^{-1} is the reciprocal of an $N(0, 1)$ variable; in particular, the k th diagonal entry is $1/\alpha_{kk}$.

These observations about triangular inversion can be represented pictorially. Every entry of L_n^{-1} at a fixed distance from the diagonal has the same probability density function (PDF). We may say that the matrix L_n^{-1} , like L_n , is “statistically Toeplitz.” See Figure 4.2a. Moreover, if we consider the first k entries of a column of L_n^{-1} from the diagonal downwards, those k entries will have the same joint probability density function (JPDF) irrespective of the column. See Figure 4.2b. The different columns of L_n^{-1} , however, are by no means independent.

The description of triangular inversion above and later arguments are stated in terms of the columns of L_n^{-1} . However, rows and columns are indistinguishable in this problem; we could equally well have framed the analysis in terms of rows.

Denote the JPDF of t_i , $1 \leq i \leq k$, by $f_k = f_k(t_1, \dots, t_k)$. In the next proposition, a recursive formula for f_k is derived. For simplicity, we introduce the further notation $T_k = \sqrt{t_1^2 + \dots + t_k^2}$. Throughout this section, L_n is the random triangular matrix of dimension n with $N(0, 1)$ entries.

Proposition 4.1. *The JPDFs $f_k = f_k(t_1, \dots, t_k)$ satisfy the following recurrence:*

$$f_1 = \frac{\exp(-1/2t_1^2)}{\sqrt{2\pi}t_1^2}, \quad (4.2)$$

$$f_k = \frac{1}{\pi} \frac{T_{k-1}}{T_k^2} f_{k-1} \quad \text{for } k > 1. \quad (4.3)$$

Proof. The t_k are defined by the random recurrence in (4.1).

The expression for f_1 is easy to get. If x is an $N(0, 1)$ variable, its PDF is

$$\frac{1}{\sqrt{2\pi}} \exp(-x^2/2).$$

The change of variable $x = 1/t_1$ gives (4.2).

To obtain the recursive expression (4.3) for f_k , consider the variable τ_k obtained by summing the variables t_1, \dots, t_{k-1} as $\sum_{i=1}^{k-1} \alpha_{ki} t_i$, where α_{ki} are independent $N(0, 1)$ variables. For fixed values of t_i , $1 \leq i \leq k-1$, the variable τ_k , being a sum of random normal variables, is itself a random normal variable of mean 0 and variance T_{k-1}^2 . Therefore, the JPDF of τ_k and t_1, \dots, t_{k-1} is given by

$$\frac{1}{\sqrt{2\pi}} \frac{\exp(-\tau_k^2/2T_{k-1}^2)}{T_{k-1}} f_{k-1}.$$

By (4.1), the variable t_k can be obtained as τ_k/α , where α is an independent $N(0, 1)$ variable. The JPDF of α , τ_k and t_1, \dots, t_{k-1} is given by

$$\frac{1}{\sqrt{2\pi}} \exp(-\alpha^2/2) \frac{1}{\sqrt{2\pi}} \frac{\exp(-\tau_k^2/2T_{k-1}^2)}{T_{k-1}} f_{k-1}.$$

Changing the variable τ_k to $t_k = \tau_k/\alpha$ and integrating out α , we obtain

$$f_k = \frac{1}{\pi} \frac{T_{k-1}}{T_{k-1}^2 + t_k^2} f_{k-1} = \frac{1}{\pi} \frac{T_{k-1}}{T_k^2} f_{k-1};$$

i.e., f_k is given by (4.3). □

Note that the form of the recurrence for f_k in Proposition 4.1 mirrors the random recurrence (4.1) for obtaining t_k from the previous entries t_1, \dots, t_{k-1} . In the following corollary, an explicit expression for f_k in terms of the t_i is stated.

Corollary 4.2. *For $k > 1$, the JPDF $f_k = f_k(t_1, \dots, t_k)$ is given by*

$$f_k = \frac{1}{\pi^{k-1} \sqrt{2\pi}} \frac{1}{(t_1^2 + \dots + t_k^2)} \frac{1}{\sqrt{t_1^2 + \dots + t_{k-1}^2}} \cdots \frac{1}{\sqrt{t_1^2 + t_2^2}} \frac{\exp(-1/2t_1^2)}{|t_1|}.$$

4.3 Moments of T_k

In this section and the next, L_n continues to represent a triangular matrix of dimension n with $N(0, 1)$ entries. As we remarked earlier, the exponential growth of $\kappa_n = \|L_n\|_2 \|L_n^{-1}\|_2$ is due to the second factor $\|L_n^{-1}\|_2$. Since the 2-norm of column $i + 1$ of L_n^{-1} has the same distribution as T_{n-i} , we derive formulas for various moments of T_k with the intention of understanding the exponential growth of $\|L_n^{-1}\|_2$ with n .

In the lemma below, we consider the expected value $E(T_k^\xi)$ for both positive and negative values of ξ . By our notation, $T_1 = |t_1|$. The notation $d\Omega_k = dt_k \dots dt_1$ is used to reduce clutter in the proof. As usual, R^k denotes the real Euclidean space of dimension k .

The next lemma is stated as a recurrence to reflect the structure of its proof. Lemma 4.4 contains the same information in a simpler form.

Lemma 4.3. *For any real $\xi < 1$, $E(T_k^\xi)$ is given by the following recurrence:*

$$E(T_1^\xi) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \frac{\exp(-1/2x^2)}{|x|^{2-\xi}} dx, \quad (4.4)$$

$$E(T_k^\xi) = \frac{E(T_{k-1}^\xi)}{\pi} \int_{-\infty}^{\infty} \frac{dx}{(1+x^2)^{1-\xi/2}} \quad \text{for } k > 1. \quad (4.5)$$

For $\xi \geq 1$ and $k \geq 1$, $E(T_k^\xi)$ is infinite.

Proof. To obtain (4.4), use $T_1 = |t_1|$ and the PDF of t_1 given by Equation (4.2). It is easily seen that the integral is convergent if and only if $\xi < 1$.

Next, assume $k > 1$. By definition,

$$E(T_k^\xi) = \int_{R^k} T_k^\xi f_k d\Omega_k.$$

Using the recursive equation (4.3) for f_k , and writing T_k in terms of t_k and T_{k-1} , we get

$$\begin{aligned} E(T_k^\xi) &= \frac{1}{\pi} \int_{R^k} \frac{T_{k-1}}{T_k^{2-\xi}} f_{k-1} d\Omega_k \\ &= \frac{1}{\pi} \int_{R^{k-1}} \int_{-\infty}^{\infty} \frac{dt_k}{(t_k^2 + T_{k-1}^2)^{1-\xi/2}} T_{k-1} f_{k-1} d\Omega_{k-1}. \end{aligned} \quad (4.6)$$

By the substitution $t_k = xT_{k-1}$, the inner integral with respect to dt_k can be reduced to

$$T_{k-1}^{\xi-1} \int_{-\infty}^{\infty} \frac{dx}{(1+x^2)^{1-\xi/2}}.$$

Inserting this in the multiple integral (4.6) gives the recursive equation (4.5) for $E(T_k^\xi)$. It is easily seen that the integral in (4.5) is convergent if and only if $\xi < 1$. \square

Define γ_ξ by

$$\gamma_\xi = \frac{1}{\pi} \int_{-\infty}^{+\infty} \frac{dx}{(1+x^2)^{1-\xi/2}}. \quad (4.7)$$

Beginning with the substitution $x = \tan \theta$ in (4.7), it can be shown that $\gamma_\xi = \pi^{-1} B((1-\xi)/2, 1/2)$, where B is the beta function. The relevant expression for the beta function $B(x, y)$ is Equation (6.2.1) in [4]. Also, if x is chosen from the standard Cauchy distribution, then $\gamma_\xi = E((1+x^2)^{\xi/2})$. We do not need γ_ξ in terms of the beta function, however; the integral expression (4.7) is more suitable for our purposes. Lemma 4.3 can be restated in a more convenient form using γ_ξ as follows:

Lemma 4.4. *For $\xi < 1$, $E(T_k^\xi) = C_\xi \gamma_\xi^k$ for a finite positive constant C_ξ . Also, $\gamma_0 = 1$, $\gamma_\xi < 1$ for $\xi < 0$, and $\gamma_\xi > 1$ for $\xi > 0$.*

Proof. The expression for $E(T_k^\xi)$ is a restatement of Lemma 4.3. By elementary integration, $\gamma_0 = 1$, and by the form of the integral in (4.7), $\gamma_\xi < 1$ for $\xi < 0$ and $\gamma_\xi > 1$ for $\xi > 0$. \square

Lemma 4.4 implies that the positive moments of T_k increase exponentially with k while the negative moments decrease exponentially with k .

Obtaining bounds for $P(T_k > M^k)$ and $P(T_k < m^k)$ is now a simple matter.

Lemma 4.5. *For $k \geq 1$, $\xi > 0$, and $m > 0$,*

$$P(T_k < m^k) < C_{-\xi} (m/\gamma_{-\xi}^{-1/\xi})^{\xi k}.$$

Proof. Since $\xi > 0$, $P(T_k < m^k) = P(T_k^{-\xi} > m^{-\xi k})$. Use Lemma 4.4 with $\xi = -\xi$ to obtain an expression for $E(T_k^{-\xi})$ and apply Markov's inequality [15]. \square

Lemma 4.6. *For $k \geq 1$, $0 < \xi < 1$, and $M > 0$,*

$$P(T_k > M^k) < C_\xi (\gamma_\xi^{1/\xi}/M)^{\xi k}.$$

Proof. As in Lemma 4.5, $\xi > 0$ implies that $P(T_k > M^k) = P(T_k^\xi > M^{\xi k})$. Again, the proof can be completed by obtaining an expression for $E(T_k^\xi)$ using Lemma 4.4 followed by an application of Markov's inequality. \square

4.4 Exponential Growth of κ_n

We are now prepared to derive the first main result of this chapter, namely, $\sqrt[n]{\kappa_n} \rightarrow 2$ almost surely as $n \rightarrow \infty$ for triangular matrices L_n with $N(0, 1)$ entries. In the sequel, *a.s.* means almost surely as $n \rightarrow \infty$. The definition of almost sure convergence for a sequence of random variables can be found in textbooks on probability, for example [15] [18]. Roughly, it means that the convergence holds for a set of sequences of measure 1.

Lemma 4.7. $\|L_n\|_2^{1/n} \rightarrow 1$ almost surely as $n \rightarrow \infty$.

Proof. The proof is easy. We provide only an outline. The Frobenius norm of L_n , $\|L_n\|_F^2$, is a sum of $n(n+1)/2$ independent χ^2 variables of mean 1. By an argument exactly analogous to the proof of the strong law of large numbers with finite fourth moment assumption [15, p. 80],

$$\frac{\|L_n\|_F^2}{n(n+1)/2} \rightarrow 1 \text{ a.s.}$$

The proof can be completed using the inequalities $n^{-1/2}\|L_n\|_F \leq \|L_n\|_2 \leq \|L_n\|_F$. Note that the suggested proof relies on the existence of the fourth moment of the χ^2 variables. \square

The proof of Lemma 4.8 uses the first Borel-Cantelli lemma in a way that is typical of several proofs in probability. We use $\liminf_{n \rightarrow \infty} x_n$ and $\limsup_{n \rightarrow \infty} x_n$ for $\lim_{n \rightarrow \infty} \inf_{k \geq n} x_k$ and $\lim_{n \rightarrow \infty} \sup_{k \geq n} x_k$ in the following lemma and later.

Lemma 4.8. As $n \rightarrow \infty$, for any $0 < \xi < 1$,

$$\gamma_{-\xi}^{-1/\xi} \leq \liminf_{n \rightarrow \infty} \sqrt[n]{\kappa_n} \leq \limsup_{n \rightarrow \infty} \sqrt[n]{\kappa_n} \leq \gamma_{\xi}^{1/\xi} \text{ almost surely.}$$

Proof. By Lemma 4.7, it suffices to show that

$$\gamma_{-\xi}^{-1/\xi} \leq \liminf_{n \rightarrow \infty} \sqrt[n]{\|L_n^{-1}\|_2} \leq \limsup_{n \rightarrow \infty} \sqrt[n]{\|L_n^{-1}\|_2} \leq \gamma_{\xi}^{1/\xi} \text{ a.s.}$$

We consider the lower bound first. The 2-norm of the first column of L_n^{-1} , which has the same distribution as T_n , is less than or equal to $\|L_n^{-1}\|_2$. Therefore, for $0 < \epsilon < 1$,

$$P(\sqrt[n]{\|L_n^{-1}\|_2} < \gamma_{-\xi}^{-1/\xi} - \epsilon) \leq P(T_n < (\gamma_{-\xi}^{-1/\xi} - \epsilon)^n).$$

Using Lemma 4.5 with $k = n$ and $m = \gamma_{-\xi}^{-1/\xi} - \epsilon$, we get

$$P(\sqrt[n]{\|L_n^{-1}\|_2} < \gamma_{-\xi}^{-1/\xi} - \epsilon) < C_{-\xi} \left(\frac{\gamma_{-\xi}^{-1/\xi} - \epsilon}{\gamma_{-\xi}^{-1/\xi}} \right)^{\xi n} = C_{-\xi} p_{\epsilon}^{\xi n},$$

where $p_{\epsilon} = \gamma_{-\xi}^{1/\xi}(\gamma_{-\xi}^{-1/\xi} - \epsilon) < 1$. Since $|p_{\epsilon}| < 1$, $\sum_{n=1}^{\infty} p_{\epsilon}^{\xi n}$ is finite. The first Borel-Cantelli lemma [15] can be applied to obtain

$$P(\sqrt[n]{\|L_n^{-1}\|_2} < \gamma_{-\xi}^{-1/\xi} - \epsilon \text{ infinitely often as } n \rightarrow \infty) = 0.$$

Taking the union of the sets in the above equation over all rational ϵ in $(0, 1)$ and considering the complement of that union, we obtain

$$P(\liminf_{n \rightarrow \infty} \sqrt[n]{\|L_n^{-1}\|_2} \geq \gamma_{-\xi}^{-1/\xi} \text{ as } n \rightarrow \infty) = 1.$$

In other words, $\gamma_{-\xi}^{-1/\xi} \leq \liminf_{n \rightarrow \infty} \sqrt[n]{\|L_n^{-1}\|_2}$ *a.s.*

The upper bound can be established similarly. At least one of the columns of L_n^{-1} must have 2-norm greater than or equal to $n^{-1/2}\|L_n^{-1}\|_2$. Since the 2-norm of column $k+1$ has the same distribution as T_{n-k} ,

$$P(\sqrt[n]{\|L_n^{-1}\|_2} > \gamma_\xi^{1/\xi} + \epsilon) \leq \sum_{k=1}^n P(T_k > n^{-1/2}(\gamma_\xi^{1/\xi} + \epsilon)^n).$$

Bounding each term in the summation using Lemma 4.6 gives

$$P(\sqrt[n]{\|L_n^{-1}\|_2} > \gamma_\xi^{1/\xi} + \epsilon) < C_\xi n^{\xi/2} \sum_{k=1}^n \left(\frac{\gamma_\xi^k}{(\gamma_\xi^{1/\xi} + \epsilon)^{\xi n}} \right).$$

Since $\gamma_\xi > 1$ by Lemma 4.4, the largest term in the summand occurs when $k = n$. Therefore,

$$P(\sqrt[n]{\|L_n^{-1}\|_2} > \gamma_\xi^{1/\xi} + \epsilon) < C_\xi n^{1+\xi/2} \left(\frac{\gamma_\xi^{1/\xi}}{\gamma_\xi^{1/\xi} + \epsilon} \right)^{\xi n}.$$

From this point, the proof can be completed in the same manner as the proof of the lower bound. \square

Theorem 4.9. *For random triangular matrices with $N(0, 1)$ entries, as $n \rightarrow \infty$,*

$$\sqrt[n]{\kappa_n} \rightarrow 2 \text{ almost surely.}$$

Proof. By an inequality sometimes called Lyapunov's [61, p. 144] [15],

$$\gamma_\beta^{1/\beta} < \gamma_\alpha^{1/\alpha}$$

for any real $\beta < \alpha$. Thus the bounding intervals $[\gamma_{-\xi}^{-1/\xi}, \gamma_\xi^{1/\xi}]$ in Lemma 4.8 shrink as ξ decreases from 1 to 0. A classical theorem [61, p. 139] says that these intervals actually shrink to the following point:

$$\begin{aligned} \lim_{\xi \rightarrow 0} \gamma_\xi^{1/\xi} &= \lim_{\xi \rightarrow 0} \left(\frac{1}{\pi} \int_{-\infty}^{\infty} \frac{1}{(1+x^2)^{1-\xi/2}} dx \right)^{1/\xi} \\ &= \exp \left(\frac{1}{2\pi} \int_{-\infty}^{\infty} \frac{\log(1+x^2)}{1+x^2} dx \right). \end{aligned}$$

The exact value of the limit can be evaluated to 2 using the substitution $x = \tan \theta$ followed by complex integration [26, p. 121]. Thus $\sqrt[n]{\kappa_n} \rightarrow 2$ *a.s.* \square

Theorem 4.9 holds in exactly the same form if the nonzero entries of L_n are independent $N(0, \sigma^2)$ variables rather than $N(0, 1)$ variables, since the condition number is invariant under scaling.

Our approach to Theorem 4.9 began by showing that $E(T_k^\xi) = C_\xi \gamma_\xi^k$ for both positive and negative ξ . Once these expressions for the moments of T_k were obtained, our arguments did not depend on how the recurrence was computed. Let us summarize the asymptotic information about a recurrence that can be obtained from a knowledge of its moments.

Let t_1, t_2, \dots be a sequence of random variables. If $E(|t_n|^\xi)$ grows exponentially with n at the rate ν_ξ^n for $\xi > 0$, then $\limsup_{n \rightarrow \infty} \sqrt[n]{|t_n|} \leq \nu_\xi^{1/\xi}$ almost surely. Similarly, if $E(|t_n|^\xi)$ decreases exponentially with n at the rate ν_ξ^n as $n \rightarrow \infty$ for $\xi < 0$, then $\nu_\xi^{1/\xi} \leq \liminf_{n \rightarrow \infty} \sqrt[n]{|t_n|}$ almost surely. Thus, knowledge of any positive moment of $|t_n|$ yields an upper bound on $\sqrt[n]{|t_n|}$ as $n \rightarrow \infty$, while knowledge of any negative moment yields a lower bound.

4.5 Unit Triangular Matrices

So far, we have considered triangular matrices whose nonzero entries are independent, real $N(0, 1)$ variables. In this section and in Section 4.7, we establish the exponential growth of the condition number for other kinds of random triangular matrices with normally distributed entries. The key steps in the sequence of lemmas leading to the analogues of Theorem 4.9 are stated but not proved. The same techniques used in Sections 4.2, 4.3, and 4.4 work here too.

Let L_n be a unit lower triangular matrix of dimension n with $N(0, \sigma^2)$ subdiagonal entries. Let s_1, \dots, s_k be the first k entries from the diagonal downwards of any column of L_n^{-1} . The entries s_i satisfy the recurrence

$$\begin{aligned} s_1 &= 1 \\ s_2 &= \alpha_{21}s_1 \\ s_3 &= \alpha_{31}s_1 + \alpha_{32}s_2 \\ &\vdots \\ s_k &= \alpha_{k1}s_1 + \dots + \alpha_{k,k-1}s_{k-1}, \end{aligned} \tag{4.8}$$

where α_{ij} , $i > j$, are $N(0, \sigma^2)$ variables. The notation $S_k = \sqrt{s_1^2 + s_2^2 + \dots + s_k^2}$ is used below.

Proposition 4.10. *The JPDP of s_1, \dots, s_k , $g_k(s_1, \dots, s_k)$, is given by the recurrence*

$$\begin{aligned} g_2 &= \frac{1}{\sqrt{2\pi}\sigma} \exp(-s_2^2/2\sigma^2), \\ g_k &= \frac{1}{\sqrt{2\pi}\sigma} \frac{\exp(-s_k^2/2\sigma^2 S_{k-1}^2)}{S_{k-1}} g_{k-1} \quad \text{for } k > 2, \end{aligned}$$

and the identity $s_1 = 1$.

Lemma 4.11. For any real ξ , $E(S_k^\xi) = \lambda_\xi^{k-1}$, where

$$\lambda_\xi = \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{\infty} (1+x^2)^{\xi/2} \exp(-x^2/2\sigma^2) dx.$$

The last paragraph of Section 4 provides part of the link from Lemma 4.11 to the following theorem about κ_n .

Theorem 4.12. For random unit triangular matrices with $N(0, \sigma^2)$ entries, as $n \rightarrow \infty$,

$$\sqrt[n]{\kappa_n} \rightarrow \exp\left(\frac{1}{2\sqrt{2\pi}\sigma} \int_{-\infty}^{\infty} \log(1+x^2) e^{-x^2/2\sigma^2} dx\right) \text{ almost surely.}$$

If this limit is denoted by $C(\sigma)$, then

$$\begin{aligned} C(\sigma) - 1 &\sim \sigma^2/2 \text{ as } \sigma \rightarrow 0, \\ C(\sigma) &\sim K\sigma \text{ as } \sigma \rightarrow \infty, \end{aligned}$$

where $K = \sqrt{\exp(-\gamma)/2} = 0.5298\dots$, with γ being the Euler constant.

Proof. The constant K is given by

$$K = \exp\left(\sqrt{\frac{2}{\pi}} \int_0^{\infty} \log x \exp(-x^2/2) dx\right).$$

To evaluate K , we used integral 4.333 of [52]. □

In contrast to the situation in Theorem 4.9, the constant that $\sqrt[n]{\kappa_n}$ converges to in Theorem 4.12 depends on σ . This is because changing σ scales only the subdiagonal entries of the unit triangular matrix L_n while leaving the diagonal entries fixed at one. For $\sigma = 1$, the case discussed in the introduction to this chapter, numerical integration shows the constant to be 1.305683410\dots

It is possible to compare the asymptotic expressions in Theorem 4.12 with results of Embree and Trefethen [42] about the recurrence $f_n = f_{n-1} \pm \beta f_{n-2}$, where each \pm sign is independent and equally likely to be $+$ or $-$. If the constant that $\sqrt[n]{|t_n|}$ converges to is denoted by $c(\beta)$, $c(\beta) - 1 \sim -\beta^2/2$ as $\beta \rightarrow 0$ and $c(\beta) \sim \beta^{1/2}$ as $\beta \rightarrow \infty$. Numerical experiments in [42] show convincingly that $c(\beta)$ vs. β , unlike $C(\sigma)$ vs. σ , can be a fractal, especially near $\beta = 0.3674$.

4.6 A Comment on the Stability of Gaussian Elimination

Numerical stability of Gaussian elimination has not been completely explained for nearly 50 years. Gaussian elimination with pivoting solves the linear system $Ax = b$

by factoring PA , where P is a permutation matrix, into LU , where L is lower triangular and U is upper triangular. The permutation matrix P is determined by the choice of pivots, and because of pivoting L will be unit triangular with all the subdiagonal entries smaller than 1 in magnitude. As shown mainly by Wilkinson, Gaussian elimination is numerically stable and useful in the presence of rounding errors if and only if $g(A) = \frac{\|U\|_2}{\|A\|_2}$ is small. The growth factor can be defined in other ways; but since the notion of smallness is that $g(A)$ is a polynomial in the dimension d of A , any definition of growth factor may be used. Examples are known where $g(A)$ is exponential in d . But for practically every example that has arisen in computing for 50 years, $g(A)$ has been small and Gaussian elimination has been successful. Attempts to understand the numerical properties of Gaussian elimination go back to von Neumann, Turing, and Wilkinson. Some examples for which Gaussian elimination is unstable are given in [45] [117].

Partial pivoting ensures that the subdiagonal entries of L are smaller than 1 in magnitude. But smallness of subdiagonal entries is an insufficient explanation by itself. We have shown that for random unit triangular matrices the condition number of L can be exponential in d even though the off-diagonal entries are very small. When L has a large condition number, $\|U\|_2$ also tends to be large since $U = L^{-1}PA$, and Gaussian elimination tends to be unstable.

But there are arguments which make the numerical stability of Gaussian elimination seem plausible. One of these is due to Day and Peterson [31]. Denote the matrix $(PA)^{-T}$ by B . Let B_k be the $k \times k$ matrix obtained using the last k rows and the last k columns of B . Then it is easily shown that of all the $k + 1$ $k \times k$ submatrices of B_{k+1} which omit its first column, B_k has the largest volume. Further, Gaussian elimination will be stable if the smallest singular value of each B_k is not much smaller than the smallest singular value of B . Geometrically, this is likely to happen because at each step B_k is picked to have a large volume.

Another plausibility argument is due to Trefethen and Schreiber [109]. They point out that when the growth factor tends to be large, Gaussian elimination with partial pivoting tends to perform large rank one updates. Large rank one updates in turn cause more cancellation. Thus there seems to be a self-stabilizing mechanism.

Trefethen and Bau have proposed yet another argument based on the column spaces of A which goes much farther. For an account of this, see [108]. A precise theorem about the stability of Gaussian elimination is stated in [108].

4.7 Complex Matrices

We now consider matrices with complex entries. Let L_n be a lower triangular matrix with $\tilde{N}(0, 1)$ entries. The complex distribution $\tilde{N}(0, 1)$ was defined in Section 4.1. Let t_1, \dots, t_k denote the first k entries from the diagonal downwards of any column of L_n^{-1} . The quantities t_k satisfy (4.1), but the α_{ij} are now independent $\tilde{N}(0, 1)$ variables. Let $r_k = |t_k|^2$, and denote $r_1 + \dots + r_k$ by R_k .

Proposition 4.13. *The JPDP of r_1, \dots, r_k , $h_k(r_1, \dots, r_k)$, is given by the recur-*

rence

$$h_1 = \frac{\exp(-1/r_1)}{r_1^2}, \quad (4.9)$$

$$h_k = \frac{R_{k-1}}{R_k^2} h_{k-1} \quad \text{for } k > 1, \quad (4.10)$$

for $r_i \geq 0$, $1 \leq i \leq k$.

Proof. We sketch only the details that do not arise in the proof of Proposition 4.1. If x and y are independent $N(0, \sigma^2)$ variables, $x = \sqrt{r} \cos(\theta)$ and $y = \sqrt{r} \sin(\theta)$, then r and θ are independent. Moreover, the distribution of r is exponential with the PDF

$$(1/2\sigma^2) \exp(-r/2\sigma^2) \quad (4.11)$$

for $r > 0$.

Consider the sum $\tau_k = \alpha_{k1}t_1 + \dots + \alpha_{k,k-1}t_{k-1}$ with α_{ki} taken as independent $\tilde{N}(0, 1)$ variables. For fixed t_1, \dots, t_{k-1} , $\text{Re}(\tau_k)$ and $\text{Im}(\tau_k)$ are independent. To see their independence, we write out the equations for $\text{Re}(\tau_k)$ and $\text{Im}(\tau_k)$ as follows:

$$\begin{aligned} \text{Re}(\tau_k) &= \sum_{i=1}^{k-1} \text{Re}(\alpha_{ki})\text{Re}(t_i) - \text{Im}(\alpha_{ki})\text{Im}(t_i), \\ \text{Im}(\tau_k) &= \sum_{i=1}^{k-1} \text{Re}(\alpha_{ki})\text{Im}(t_i) + \text{Im}(\alpha_{ki})\text{Re}(t_i). \end{aligned}$$

The linear combinations of $\text{Re}(\alpha_{ki})$ and $\text{Im}(\alpha_{ki})$ in these two equations can be realized by taking inner products with the two vectors

$$\begin{aligned} v &= [\text{Re}(t_1), \dots, \text{Re}(t_{k-1}), -\text{Im}(t_1), \dots, -\text{Im}(t_{k-1})], \\ w &= [\text{Im}(t_1), \dots, \text{Im}(t_{k-1}), +\text{Re}(t_1), \dots, +\text{Re}(t_{k-1})]. \end{aligned}$$

The independence of $\text{Re}(\tau_k)$ and $\text{Im}(\tau_k)$ is a consequence of the orthogonality of v and w , i.e., $(v, w) = vw' = 0$, and the invariance of the JPDF of independent, identically distributed normal variables under orthogonal transformation [85].

Thus for fixed t_1, \dots, t_{k-1} , the real and imaginary parts of τ_k are independent normal variables of mean 0 and variance $R_{k-1}/2$. By Equation (4.11), the PDFs of $x = |\tau_k|^2$ and $y = |\alpha_{kk}|^2$ are given by

$$\frac{1}{R_{k-1}} \exp(-x/R_{k-1}), \quad \exp(-y)$$

for positive x, y . The expression (4.10) for h_k can now be obtained using $r_k = |\tau_k|^2/|\alpha_{kk}|^2$. \square

Lemma 4.14. For any $\xi < 1$, $E(R_k^\xi) = C\mu_\xi^{k-1}$, where

$$C = \int_0^\infty \frac{\exp(-1/r_1)}{r_1^{2-\xi}} dr_1, \quad \mu_\xi = \int_0^\infty \frac{dx}{(1+x)^{2-\xi}}.$$

The constant μ_ξ in Lemma 4.14 can be reduced to $(1-\xi)^{-1}$ for $\xi < 1$. However, as with γ_ξ in Section 4.3, the integral expression for μ_ξ is more suitable for our purposes. As before, the last paragraph of Section 4.4 is an essential part of the link from the previous lemma to the following theorem about κ_n .

Theorem 4.15. For random triangular matrices with complex $\tilde{N}(0,1)$ entries, as $n \rightarrow \infty$,

$$\sqrt[n]{\kappa_n} \rightarrow \exp\left(\frac{1}{2} \int_0^\infty \frac{\log(1+x)}{(1+x)^2} dx\right) = e^{1/2} \text{ almost surely.}$$

Theorem 4.15 holds unchanged if the entries are $\tilde{N}(0, \sigma^2)$ variables. As with Theorem 4.9, this is because the condition number is invariant under scaling.

Now, let L_n be a unit lower triangular matrix of dimension n with $\tilde{N}(0, \sigma^2)$ subdiagonal entries. We state only the final theorem about κ_n .

Theorem 4.16. For random unit triangular matrices with complex $\tilde{N}(0, \sigma^2)$ entries, as $n \rightarrow \infty$,

$$\begin{aligned} \sqrt[n]{\kappa_n} &\rightarrow \exp\left(\frac{1}{4} \int_0^\infty \log(1 + \sigma^2 x/2) e^{-x/2} dx\right) \\ &= \exp(-\exp(\sigma^{-2}) \text{Ei}(-\sigma^{-2})/2) \text{ almost surely,} \end{aligned}$$

where Ei is the exponential integral. If this limit is denoted by $C(\sigma)$, then

$$\begin{aligned} C(\sigma) - 1 &\sim \sigma^2/2 \text{ as } \sigma \rightarrow 0, \\ C(\sigma) &\sim K\sigma \text{ as } \sigma \rightarrow \infty, \end{aligned}$$

where $K = \exp(-\gamma/2) = 0.7493\dots$, with γ being the Euler constant.

Proof. To obtain K , we evaluated

$$K = \exp\left(\frac{1}{4} \int_0^\infty \log(x/2) \exp(-x/2) dx\right)$$

using the Laplace transform of $\log(x)$ given by integral 4.331.1 of [52]. The explicit formula involving $\text{Ei}(\sigma^{-2})$ was obtained using integral 4.337.2 of [52]. \square

For $\sigma^2 = 1$, $\sqrt[n]{\kappa_n}$ converges to 1.347395784\dots

4.8 Matrices with Entries from Stable Distributions

The techniques used to deduce Theorem 4.9 require that we first derive the joint density function of the t_k , defined by recurrence (4.1), as was done in Proposition 4.1. That proposition made use of the fact that when the α_{ki} are independent and normally distributed, and the t_i are fixed, the sum

$$\sum_{i=1}^{k-1} \alpha_{ki} t_i$$

is also normally distributed. This property of the normal distribution holds for any stable distribution.

A distribution is said to be stable, if for X_i chosen independently from that distribution,

$$\sum_{i=1}^n X_i$$

has the same distribution as $c_n X + d_n$, where X has the same distribution as X_i and $c_n > 0$ and d_n are constants [44, p. 170]. If $d_n = 0$, the distribution is said to be strictly stable. As usual, the distribution is symmetric if X has the same distribution as $-X$. A symmetric, strictly stable distribution has exponent a if $c_n = n^{1/a}$. A standard result of probability theory says that any stable distribution has an exponent $0 < a \leq 2$. The normal distribution is stable with exponent $a = 2$ [44].

The techniques used for triangular matrices with normal entries work more generally when the entries are drawn from a symmetric, strictly stable distribution. Let L_n be a unit lower triangular matrix with entries chosen from a symmetric, strictly stable distribution. Denote the PDF of that stable distribution by $\phi(x)$. The recurrence for the entries s_i of the inverse L_n^{-1} is again given by (4.8), but α_{ki} , $k > i$, are now independent random variables with the density function $\phi(x)$.

Our program for deriving the constants that $\sqrt[n]{\kappa_n}$ converge to as $n \rightarrow \infty$ began with Lemma 4.7 in all the previous examples. A referee for [114] pointed out to us that a new proof is needed for that lemma in the present context since a stable distribution of index $a < 2$ does not have the a th or higher moments.

Lemma 4.17. *For $a < 2$, $\|L_n\|_2^{1/n} \rightarrow 1$ almost surely as $n \rightarrow \infty$.*

Proof. Define $\|L_n\|_\alpha = (\sum_{i,j} |l_{ij}|^\alpha)^{1/\alpha}$ for some $0 < \alpha < a/4$. Then the inequality

$$n^{(1/2-2/\alpha)} \|L_n\|_\alpha \leq \|L_n\|_2 \leq n \|L_n\|_\alpha,$$

and the existence of the fourth moment of $|l_{ij}|^\alpha$ make possible a proof analogous to what was outlined for Lemma 4.7. \square

The proposition, the lemma and the theorem below are analogues of Proposition 4.10, Lemma 4.11, and Theorem 4.12 respectively. If the exponent of the stable distribution is a , denote $(|s_1|^a + \dots + |s_k|^a)^{1/a}$ by S_k .

Proposition 4.18. *If $\phi(x)$ is the density function of a symmetric, strictly stable distribution with exponent a , the JPDF of s_1, \dots, s_k , $g_k(s_1, \dots, s_k)$, is given by the recurrence*

$$g_2 = \phi(s_2),$$

$$g_k = \frac{\phi(s_k/S_{k-1})}{S_{k-1}} g_{k-1} \quad \text{for } k > 2,$$

and the identity $s_1 = 1$.

Proof. The proof is very similar to the proof of Proposition 4.1. We note that if α_{ki} , $k > i$, are independent random variables with the PDF $\phi(x)$, and the s_i are fixed, then the sum

$$\alpha_{k1}s_1 + \dots + \alpha_{k,k-1}s_{k-1}$$

has the PDF $\phi(x/S_{k-1})/S_{k-1}$ [44, p. 171]. □

Lemma 4.19. *For any real ξ , $E(S_k^\xi) = \lambda_\xi^{k-1}$, where*

$$\lambda_\xi = \int_{-\infty}^{+\infty} (1 + |x|^a)^{\xi/a} \phi(x) dx,$$

with $\lambda_\xi = \infty$ for $\xi \geq a$.

Theorem 4.20. *For random unit triangular matrices with entries from a symmetric, strictly stable distribution with density function $\phi(x)$ and exponent a , as $n \rightarrow \infty$,*

$$\sqrt[n]{\kappa_n} \rightarrow \exp\left(\frac{1}{a} \int_{-\infty}^{+\infty} \log(1 + |x|^a) \phi(x) dx\right) \quad \text{almost surely.}$$

Theorem 4.12 is a special case of Theorem 4.20 when $\phi(x)$ is the density function for the symmetric, strictly stable distribution $N(0, \sigma^2)$. Another notable symmetric, strictly stable distribution is the Cauchy distribution with the density function

$$\phi(x) = \frac{1}{\pi} \frac{1}{1 + x^2}.$$

The exponent a for the Cauchy distribution is 1 [44]. Using Theorem 4.20 we obtain,

Theorem 4.21. *For random unit triangular matrices with entries from the standard Cauchy distribution, as $n \rightarrow \infty$,*

$$\sqrt[n]{\kappa_n} \rightarrow \exp\left(\frac{1}{\pi} \int_{-\infty}^{+\infty} \frac{\log(1 + |x|)}{1 + x^2} dx\right) \quad \text{almost surely.}$$

Numerical integration shows the constant to be 2.533737279....

A similar generalization can be made for triangular matrices without a unit diagonal. However, the analogue of Theorem 4.20 for such matrices involves not $\phi(x)$, but the density function $\psi(x)$ of the quotient $x = y/z$ obtained by taking y, z as independent variables with the PDF ϕ . The distribution ψ can be difficult to compute and work with. We state only the final theorem about κ_n for triangular matrices with entries drawn from a symmetric strictly stable distribution.

Theorem 4.22. *For random triangular matrices with entries from a symmetric, strictly stable distribution with density function $\phi(x)$ and exponent a , as $n \rightarrow \infty$,*

$$\sqrt[n]{\kappa_n} \rightarrow \exp\left(\frac{1}{a} \int_{-\infty}^{\infty} \log(1 + |x|^a) \psi(x) dx\right) \text{ almost surely,}$$

where $\psi(x)$ is the density function of the quotient of two independent variables with the density function $\phi(x)$.

Theorem 4.9 is a special case of Theorem 4.22 when $\phi(x)$ is the density function of the distribution $N(0, \sigma^2)$. The $\psi(x)$ corresponding to $N(0, \sigma^2)$ is the standard Cauchy distribution. To apply Theorem 4.22 for the Cauchy distribution, we note that

$$\psi(x) = \frac{2 \log |x|}{\pi^2 x^2 - 1}$$

is the density function of the quotient if the numerator and the denominator are independent Cauchy variables [60]. Therefore, Theorem 4.22 implies

Theorem 4.23. *For random triangular matrices with entries from the standard Cauchy distribution, as $n \rightarrow \infty$,*

$$\sqrt[n]{\kappa_n} \rightarrow \exp\left(\frac{2}{\pi^2} \int_{-\infty}^{\infty} \log(1 + |x|) \frac{\log |x|}{x^2 - 1} dx\right) \text{ almost surely.}$$

The constant of convergence in Theorem 4.23 is 3.063094192....

4.9 A Conjecture about Random Recurrences

Do terms of the random recurrence

$$\begin{aligned} t_1 &= 1 \\ t_2 &= \alpha_{21} t_1 \\ t_3 &= \alpha_{31} t_1 + \alpha_{32} t_2 \\ &\vdots \\ t_k &= \alpha_{k1} t_1 + \cdots + \alpha_{k,k-1} t_{k-1}, \end{aligned}$$

increase exponentially with k , when the α_{ij} are independent and identically distributed, but not stably distributed? We make the following conjecture.

Conjecture 4.24. *If the α_{ij} are independent, identically distributed, real or complex valued random variables, but the distribution is not supported on a constant α with $|1 + \alpha| \leq 1$, then*

$$\sqrt[n]{|t_n|} \rightarrow C \quad \text{almost surely}$$

as $n \rightarrow \infty$ for some $C > 1$.

We have tested this conjecture using at least 50 different distributions on a computer. Its proof may be nontrivial, but does not seem unattainable. A careful understanding of Fourier analysis is perhaps necessary to show it to be correct. There may be a central limit theorem that governs convergence to C , and perhaps even a short formula for C .

In all the formulas we have derived, $\log C$ is obtained by integrating a function with respect to a probability measure. Thus there is a formal similarity to Furstenberg's formula for the top Lyapunov exponent, which integrates an amplification function with respect to an invariant measure.

4.10 Summary




Figure 4.3: Illustration of Theorem 4.9. After 100,000 steps of the random recurrence (4.1), $\sqrt[n]{|t_n|}$ has settled to within 1% of its limiting value 2. The implementation is explained in the text.

Below is a summary of the exponential growth factors $\lim_{n \rightarrow \infty} \sqrt[n]{\kappa_n}$ that we have established for triangular matrices with normal entries:




Figure 4.4: Another illustration of Theorem 4.9. Each cross is obtained by computing the condition number κ_n for one random triangular matrix of dimension n with $N(0, 1)$ entries. The solid line represents 2^n .

Real triangular	2	Theorem 4.9
Real unit triangular, $\sigma^2 = 1$	1.305683410...	Theorem 4.12
Complex triangular	$e^{1/2} = 1.647\dots$	Theorem 4.15
Complex unit triangular, $\sigma^2 = 1$	1.347395784...	Theorem 4.16

The theorems about unit triangular matrices with normally distributed, real or complex entries apply for any variance σ^2 , not just $\sigma^2 = 1$. Constants of convergence for any symmetric, strictly stable distribution were derived in Theorems 4.20 and 4.22. Those two theorems were specialized to the Cauchy distribution in Theorems 4.21 and 4.23.

We close with two figures that illustrate the first main result of this chapter, namely, for random triangular matrices with $N(0, 1)$ entries, $\sqrt[n]{\kappa_n} \rightarrow 2$ almost surely as $n \rightarrow \infty$ (Theorem 4.9). Figure 4.3 plots the results of a single run of the random recurrence (4.1) to 100,000 steps, confirming the constant 2 to about two digits. The expense involved in implementing the full recurrence (4.1) for so many steps would be prohibitive. However, since t_k grows at the rate 2^k , we need include only a fixed number of terms in (4.1) to compute t_k to machine precision. For the figure, we used 200 terms, although half as many would have been sufficient. Careful scaling

was necessary to avoid overflow while computing this figure.

Figure 4.4 plots the condition number of a single random triangular matrix for each dimension from 1 to 200. The exponential trend at the rate 2^n is clear, but as in Figure 4.1, the convergence as $n \rightarrow \infty$ is slow.

Chapter 5

Lyapunov Exponents of the Lorenz Equations ¹

5.1 Introduction

Chapter 1 explained that Lyapunov exponents can be defined for random matrix products, linear or nonlinear stochastic systems, and deterministic dynamical systems. So far, we have concerned ourselves mainly with random recurrences and random matrix products. This chapter examines Lyapunov exponents used to quantify chaos in deterministic dynamical systems. The Lorenz equations are the main example, and the focus is on accurate computation and continuous dynamical systems of the form $\frac{dx}{dt} = f(x)$.

Let us recall how Lyapunov exponents quantify sensitive dependence on initial conditions. For a given time $t > 0$, think of $x(t) \in R^d$ as a function of the initial conditions $x(0) = x_0 \in R^d$. Then the Jacobian $\frac{\partial x(t)}{\partial x_0}$ is a $d \times d$ matrix that depends on t and x_0 . The Lyapunov exponents, $\gamma_1 \geq \gamma_2 \geq \dots \geq \gamma_d$, are defined by

$$\gamma_i = \lim_{t \rightarrow \infty} \frac{1}{t} \log \sigma_i \left(\frac{\partial x(t)}{\partial x_0} \right), \quad (5.1)$$

where σ_i is the i th singular value with σ_1 being the 2-norm [89] [25]. Obviously, $\gamma_1 > 0$ implies sensitive dependence on initial conditions. The right singular vectors of $\frac{\partial x(t)}{\partial x_0}$ also converge as $t \rightarrow \infty$, and the directions corresponding to positive Lyapunov exponents are called the stretching directions at x_0 . The first Lyapunov exponent γ_1 is also called the top Lyapunov exponent. The direction along the trajectory always corresponds to a zero Lyapunov exponent. Numerical computations show that the γ_i exist and are independent of x_0 for a given attractor. But as Chapter 1 pointed out, existing theory does not answer questions about the existence of γ_i convincingly.

Figure 5.1 illustrates sensitive dependence on initial conditions for the Lorenz equations. Similar calculations are commonly used as numerical evidence for chaos [84] [104].

¹This chapter is adapted from [113].

Table 5.1: Lyapunov exponents γ_1 , γ_2 and γ_3 of the Lorenz equations for three different choices of the parameters σ , ρ and β . *Error* is what we believe is the maximum possible error in the numbers reported. D_{KY} is the Kaplan-Yorke dimension $2 + |\gamma_1/\gamma_3|$ which is conjectured to be equal to the Hausdorff dimension of the Sinai-Ruelle-Bowen measure [37].

σ	ρ	β	γ_1	γ_2	γ_3	<i>Error</i>	D_{KY}
16	45.92	4	1.50255	0.00000	-22.50255	.00008	2.0668
16	40	4	1.37446	0.00000	-22.37446	.00006	2.0614
10	28	8/3	0.90566	0.00000	-14.57233	.00007	2.0621

The literature on computing Lyapunov exponents is extensive [12] [34] [35] [37] [49] [50] [56] [94]. Methods for computing Lyapunov exponents, which are all based on (5.1), are compared and contrasted in [49]. The literature on the related topic of estimating Lyapunov exponents from an experimental time series is even more extensive; a sampling is [1] [30] [36] [37] [38] [41] [54] [55] [90] [97] [116]. Lyapunov exponents are often used to provide numerical evidence of chaos; some examples are [48] [84] [86] [88] [104] [118]. There has also been an effort to define and work with finite-time, local Lyapunov exponents [1] [2] [3].

Accurate computation of Lyapunov exponents has received attention in the literature, but the effect of numerical discretization of continuous dynamical systems on their Lyapunov exponents has not been studied. It is the study of discretization errors that enables us to compute Lyapunov exponents more accurately than in previously reported calculations. Table 5.1 shows the Lyapunov exponents we computed for the Lorenz equations

$$\begin{aligned}\frac{dx_1(t)}{dt} &= \sigma(x_2(t) - x_1(t)), \\ \frac{dx_2(t)}{dt} &= x_1(t)(\rho - x_3(t)) - x_2(t), \\ \frac{dx_3(t)}{dt} &= x_1(t)x_2(t) - \beta x_3(t),\end{aligned}$$

with three standard choices of the parameters σ , ρ , and β . The first choice $\sigma = 16$, $\rho = 45.92$, and $\beta = 4$ occurs widely in papers about computing or estimating Lyapunov exponents [2] [34] [35] [90] [94] [116]. The second choice with the same σ and β but $\rho = 40$ is less common [30] [34] [50]. The third choice $\sigma = 10$, $\rho = 28$, and $\beta = 8/3$ is perhaps the most widely recognized, thanks to its appearance not only in Lorenz's original paper [79] but also in popular textbooks [58] [102]. For the first choice of parameters, $\gamma_1/\log(2)$ is given as 2.16 in [116] though the correct rounded value to two digits is 2.17, and as 2.164 in [90] though the correct rounded value to three digits is 2.168; γ_1 is given as 1.48804 in [34] which is off by about .01, and as 1.51 in [2] which is also off by about .01. The only one of these papers to give an estimate of the error is [2], where it is correctly implied that the value 1.51 can be off by about .01. We note that the error in γ_1 as reported in Table 5.1 is only




Figure 5.1: (a) $\|\delta x(t)\|$ is the 2-norm distance between two solutions of the Lorenz equations with the same choice of parameters as the first row of Table 5.1. $\|\delta x(t)\|$ is upper bounded by the diameter of the Lorenz attractor, which is about 100. The dashed line $e^{1.5025t}$ approximates the rate of divergence of the two solutions. (b) Direct use of (5.1) allows a more accurate estimation of the top Lyapunov exponent γ_1 . The Jacobian $\frac{\partial x(t)}{\partial x_0}$ in (5.1) was approximated using automatic differentiation and appropriate scaling.

.00008.

The longest length of integration that we could find reported for the Lorenz equations was $T = 1000$ [34] [35]. In fact, this length, which is more than a thousand times $1/\gamma_1$ (the reciprocal of the top Lyapunov exponent), is quite long compared to most other calculations. For example, the numerical evidence for chaos in the solar system given by Sussman and Wisdom [104], uses an integration length of only 100 million years, which is just 5 to 20 times $1/\gamma_1$ for the solar system². Besides, the rate of convergence of (5.1) as $t \rightarrow \infty$ is very slow; it is about $t^{-1/2}$ for the Lorenz equations as we show in Section 5.3 or as reported in [2]. Thus, unsurprisingly, the errors in most published calculations of the Lyapunov exponents are mainly due to the finite length of integration.

The length of integration we use for the Lorenz equations, $T = 10^7$, is far longer than in any previous calculation we are aware of. Consequently, the discretization errors in solving $\frac{dx(t)}{dt} = f(x)$ become as significant as errors due to the finite length of integration in computing Lyapunov exponents. Every numerical method approximates $\frac{dx(t)}{dt} = f(x)$ by a discrete dynamical system that is locally close to it [103]. The discretization error, which is the difference between the Lyapunov exponents of the continuous dynamical system and its discrete approximation, depends on the time step h of the numerical method. We show in Section 5.3 that the discretization error seems to be $O(h^r)$ if the order of accuracy of the numerical method is r .

²The age of the solar system, which is 5 billion years, is 250 to 1000 times the estimated $1/\gamma_1$.

Thus an accurate computation of Lyapunov exponents requires a combination of a long integration, a short time step, and a higher order method. For dynamical systems of even moderately high dimensions (say ≥ 50), meeting all these requirements to obtain 3 or 4 digits of the top Lyapunov exponent accurately involves too much computation to be practical. But it is still good practice to experiment with long integrations and short time steps. Using a system of coupled Ginsburg-Landau attractors studied in [88] as an example, we show that long integrations and short time steps help avoid erroneous conclusions of chaos.

The next section briefly explains the numerical method we used for computing Lyapunov exponents and is a prelude to our study of computing Lyapunov exponents accurately.

5.2 Computing Lyapunov Exponents

Computing Lyapunov exponents using the definition in (5.1) is possible only if we can approximate the $d \times d$ Jacobian matrix $\frac{\partial x(t)}{\partial x_0}$. The usual way of doing this is to numerically solve the matrix differential equation $\frac{dM(t)}{dt} = \frac{\partial f(x)}{\partial x} M(t)$, $M(t) \in R^{d \times d}$, along with $\frac{dx(t)}{dt} = f(x)$, $x(0) = x_0$. If the initial conditions $M(0)$ is taken as the $d \times d$ identity matrix, $M(t)$ will be equal to $\frac{\partial x(t)}{\partial x_0}$. This approach is clearly explained in [37]. We used *automatic differentiation* [57] [14] of the numerical ODE solver that approximately computed $x(t)$ from x_0 to approximate $\frac{\partial x(t)}{\partial x_0}$ and thus avoided the extra set of matrix differential equations. The two approaches are actually equivalent up to rounding errors for explicit methods, but automatic differentiation may be advantageous if the algebraic complexity of $f(x)$ makes the hand coding of $\frac{\partial f(x)}{\partial x}$, which is needed by any numerical method used for solving the matrix differential equation, laborious.

Even though $\frac{\partial x(t)}{\partial x_0}$ can be numerically approximated, making direct use of (5.1) and a routine for computing singular values is numerically unsafe. Such a scheme is plagued by the twin hazards of numerical overflow and numerical rank deficiency. IEEE double precision arithmetic can represent numbers as large as 2^{1023} with a relative precision of 2^{-53} [66] [64]. Therefore, for the Lorenz equations with $\sigma = 16$, $\rho = 45.92$ and $\beta = 16$ as in Table 5.1, the entries of the 3×3 matrix $\frac{\partial x(t)}{\partial x_0}$ will overflow when $e^{\gamma_1 t} \approx 2^{1023}$ or $t \approx 470$; and the 3×3 matrix itself will become numerically rank deficient when $e^{(\gamma_1 - \gamma_3)t} \approx 2^{53}$ or $t \approx 1.53$. It is numerically safe to form $\frac{\partial x(t)}{\partial x_0}$ only if $t \leq 1/4$ or so. The use of repeated QR-factorizations to avoid these numerical problems was proposed by Shimada and Nagashima [94] as well as Benettin, et al. [12] and is clearly explained in [37]. Their proposal, which is exactly analogous to simultaneous iteration for computing eigenvalues of a matrix [108], relies on the fact that if $0 = t_0 \leq t_1 \leq \dots \leq t_k$ then $\frac{\partial x(t_k)}{\partial x(t_0)}$ equals the product $\frac{\partial x(t_k)}{\partial x(t_{k-1})} \dots \frac{\partial x(t_1)}{\partial x(t_0)}$.

There are algorithms for computing Lyapunov exponents that perform QR factorization or the SVD continuously in time [50] [56]. But Geist et al. [49] conclude that these algorithms are problematic. Their conclusions were borne out by our nu-




Figure 5.2: Exponential increase of higher derivatives of a component of $x(t)$ with respect to a component of $x(0)$ with t for the Lorenz equations. The parameters were chosen as in the first row of Table 5.1. The dashed lines are $e^{\gamma_1 t}$, $e^{2\gamma_1 t}$, $e^{3\gamma_1 t}$, and $e^{4\gamma_1 t}$ with the top Lyapunov exponent γ_1 taken as 1.50255.

merical experiments, even though we implemented the suggestions in [34]. Therefore, we will not consider these continuous algorithms at all.

It is well known that it is impossible to approximate trajectories of chaotic dynamical systems accurately as $t \rightarrow \infty$ [103]. Since the top Lyapunov exponent is positive, even the inevitable local discretization errors are amplified exponentially in time. However, the numerical attractor usually approximates the actual strange attractor quite well [103].

The situation for $\frac{\partial x(t)}{\partial x_0}$ is similar. Figure 5.2 shows that for the Lorenz equations, the second derivative of $x_2(t)$ with respect to $x_1(0)$ increases exponentially with t at the rate $2\gamma_1$, which is also the rate of increase of all other second derivatives of components of $x(t)$ with respect to components of $x(0)$. Therefore, the task of approximating $\frac{\partial x(t)}{\partial x_0}$ as $t \rightarrow \infty$ is hopeless. However, the rates of growth of the singular values of $\frac{\partial x(t)}{\partial x_0}$, when properly computed, are very close to the actual Lyapunov exponents. In other words, the Lyapunov exponents of the discrete dynamical systems used by numerical methods to approximate continuous dynamical systems are good approximations to the actual Lyapunov exponents. There is no explanation why this is so, in part because there is no perturbation theory for Lyapunov exponents in this situation.

Figure 5.2 illustrates the exponential increase of higher derivatives of $x(t)$ with respect to $x(0)$ as $t \rightarrow \infty$. The higher derivatives were computed using the automatic differentiation package ADOL-C [14]. The exponential increase shown in Figure 5.2 is not surprising. Since the trajectory $x(t)$ itself is exponentially sensitive to small changes in $x(0)$, all derivatives of $x(t)$ with respect to $x(0)$ also have to be exponentially sensitive to changes in $x(0)$. It is also easy to show using tensors that the rate of increase of a k th derivative of $x(t)$ with respect to $x(0)$ is typically $k\gamma_1$.

5.3 Errors in Computed Lyapunov Exponents

Finite length of integration, discretization and rounding are the three sources of errors in computed Lyapunov exponents. The discussion of these three sources of error in this section is followed by an accurate computation of the Lyapunov exponents of the Lorenz equations in the next section.

For the Lorenz equations studied in Figures 5.3 and 5.4, γ_2 corresponds to perturbations along the trajectory and is 0. The exact value of $\gamma_1 + \gamma_2 + \gamma_3$ is $-\sigma - \beta - 1$, which is the constant rate of volume contraction for the Lorenz equations [102]. For γ_1 , we used the accurate but inexact value in Table 5.1 for computing errors. Figure 5.3 used the standard 4th order Runge-Kutta method [63, p. 138]. From Figure 5.3, it is clear that the convergence to γ_1 as $t \rightarrow \infty$ is not very clean. The convergence is like t^{-1} for $0 \leq t \leq 100$, but the error, which is about 10^{-2} at $t = 100$, increases first and then decreases again to about 10^{-3} at $t = 10000$. There is evidence in the next section to conclude that the convergence is like $t^{-1/2}$ as t increases beyond 10000. Abarbanel, et al. [2] came to the same conclusion using numerical experiments performed from a different perspective. The convergence to γ_2 as $t \rightarrow \infty$ is like t^{-1} . The discretized dynamical system also seems to contract volume at a constant rate like the Lorenz equations. Thus the error in $\gamma_1 + \gamma_2 + \gamma_3$ does not vary much with t . Figure 5.3 also implies that the error in γ_1 and γ_2 because of stopping the integration at $t = 10000$ are about 10^{-3} and 10^{-4} , respectively.

Figure 5.4 studies the discretization errors in computing γ_1 , γ_2 and γ_3 using an integration till $t = 10000$. As we explained before, discretization error is the difference between the Lyapunov exponents of the discrete numerical method and the Lyapunov exponents of the continuous dynamical system. Since our integrations are only till $t = 10000$, part of the error will be because of the finite length of integration. But when the time step is not too small the total error in approximating the Lyapunov exponents will be dominated by the discretization error; thus we can up to a point ignore the error due to the finiteness of the integration. The 2nd and 4th order Runge-Kutta methods used for the experiments in Figure 5.4 are described in [63, p. 135, p. 138]. The errors in γ_2 quickly go down to about 10^{-4} , which is the minimum error possible since all integrations are stopped at $t = 10000$. The errors in both γ_1 and $\gamma_1 + \gamma_2 + \gamma_3$ are roughly $O(h^{-2})$ for the 2nd order Runge-Kutta method with a constant step size of h . For the 4th order Runge-Kutta method, the errors are roughly $O(h^{-4})$. But the error in γ_1 levels off, as h is decreased, at 10^{-3} which is the minimum possible with an integration till $t = 10000$.

Rounding errors have not been a factor in any computation of Lyapunov exponents of chaotic systems done so far, and we expect, will not be a factor for some time to come. The errors due to rounding in a computation that takes as many as 10^{12} time steps might cause a loss of about 6 decimal digits of precision. But since IEEE double precision arithmetic allows for 16 digits of precision and it is typically hard to get even 3 or 4 digits of the Lyapunov exponents accurately, rounding errors are harmless. In their computations with the solar system, Sussman and Wisdom [104] implemented quadruple precision arithmetic only to realize that it was unnecessary.




Figure 5.3: The Lyapunov exponents γ_1 , γ_2 and γ_3 of the Lorenz equations were computed using a 4th order Runge-Kutta method. The parameters σ , ρ and β were the same as in the first row of Table 5.1. The time step $h = 1/128$ was small enough that all the errors in both γ_1 and γ_2 are mainly due to the finite time of integration.




Figure 5.4: The Lyapunov exponents γ_1 , γ_2 , γ_3 of the Lorenz equations were computed using an integration till $t = 10000$. The parameters σ , ρ and β were the same as in the first row of Table 5.1. The left and right hand side plots used a 2nd and a 4th order, constant time step Runge-Kutta discretization respectively. The markers \circ , $*$, \square denote the errors in $\gamma_1 + \gamma_2 + \gamma_3$, γ_1 , and γ_2 , respectively

It is clear from the experiments in Figures 5.3 and 5.4 that computing the Lyapunov exponents accurately for the Lorenz equations requires a long integration, a higher order discretization and a short time step. For example, computation of γ_1 for the Lorenz equations with an error less than 10^{-4} using the standard 4th order Runge-Kutta method will require an integration at least till $t = 10^6$ and a time step of about 10^{-2} .

The slow convergence as $t \rightarrow \infty$ is because of variation in the local rate of divergence of trajectories on the surface of the attractor. The variation of the local rates of divergence on the surface of a strange attractor is studied with several examples by Abarbanel, et al. [2]. Eckmann and Ruelle [38] note that the number of points needed for filling up a d dimensional volume is exponential in d , and argue that the length of the experimental time series needed to correctly infer the dimension of the attractor and the Lyapunov exponents increases exponentially with d . If the attractor is d -dimensional, a numerical method will also have to sample points from that volume to compute Lyapunov exponents, particularly because the local rates of divergence and convergence of trajectories can vary a lot on the surface of an attractor. Thus we might expect the convergence to Lyapunov exponents to be of the order $t^{-1/d}$. If it were possible to sample points on the attractor randomly, the typical Monte Carlo order of convergence of $t^{-1/2}$ may or may not be attainable depending upon whether local rates of divergence of trajectories can be combined properly. The trajectories themselves do not sample the attractor in a random, memoryless fashion.

It is possible to explain why numerical methods appear to approximate Lyapunov exponents with the same order of accuracy as their global order, if we assume

the Lyapunov exponents depend continuously on the dynamical system. If we consider one step to be taken by integrating for a fixed amount of time, say $t = 1$, then a numerical method of order r with a sufficiently small time step h approximates one step of the dynamics with an error of $O(h^r)$ [103]. Thus if the Lyapunov exponents were to depend continuously on the dynamical system, they would change upon discretization by only about $O(h^r)$. This argument cannot be made rigorous because examples with discontinuous dependence of the Lyapunov exponents on the parameters are known [37]. However, it is unclear if these examples are typical. In the related theory of random matrix products, Le Page [76] has proven continuous dependence of the Lyapunov exponents on the distribution of the matrices.

5.4 The Lorenz Equations

Since $\gamma_2 = 0$ and $\gamma_1 + \gamma_2 + \gamma_3 = -\sigma - \beta - 1$ for the Lorenz equations, determining γ_1 accurately leads to an accurate value for both γ_2 and γ_3 . The Lyapunov exponents of the Lorenz equations reported in Table 5.1 are based on computations summarized in Table 5.2. We note that the final estimates for γ_1 are nearly contained inside most of the earlier estimates. Table 5.2 also suggests that the convergence to γ_1 as $t \rightarrow \infty$ is roughly like $t^{-1/2}$. For a detailed discussion of the Lorenz equations, we refer to [98]. The Lorenz equations were recently proved to be chaotic for some parameter values [87].

All the numbers in Tables 5.1 and 5.2 would be fully reproducible (in IEEE double precision arithmetic) if we reported the initial conditions $x(0)$ for the numerical experiments of Table 5.2. But any initial condition on the attractor will lead to nearly the same numbers as in Table 5.2. To get a point on the attractor, we integrated from the initial conditions $x_1(0) = x_2(0) = x_3(0) = 1$ till $t = 1000$.

Using numerical methods other than the standard 4th order Runge-Kutta method leads to estimates for γ_1 that are compatible with Table 5.1. The initial conditions for the numerical experiments reported in Table 5.3 were all different. The Fehlberg 7th order Runge-Kutta method used is described in [63, p. 180]. We believe that the estimate of γ_1 in the first row, and maybe even the estimate in the second row, are corrupted by discretization errors. All the other estimates properly contain the estimate of γ_1 with $\sigma = 16$, $\rho = 45.92$ and $\beta = 4$ in Table 5.1.

5.5 Reliability of Numerical Evidence for Chaos

Long integrations and short time steps are a necessity to ensure the reliability of numerical computations of continuous dynamical systems. We will illustrate this using a system of coupled Ginsburg-Landau oscillators studied by Nakagawa and Kuramoto [88].

Table 5.2: Estimates of the top Lyapunov exponent γ_1 of the Lorenz equations. The method of integration was standard 4th order Runge-Kutta with a constant time step of $1/128$. min and max are the minimum and maximum estimates for γ_1 in the interval $[T/2, T]$. The last column gives $(\min + \max)/2 \pm (\max - \min)$ with appropriate rounding as the estimate of γ_1 obtained by an integration till time T . The choice of the initial conditions $x(0)$ is described in the text.

T	min	max	γ_1
10^2	1.4507464	1.5262463	$1.49 \pm .08$
10^3	1.4928845	1.5110816	$1.50 \pm .02$
10^4	1.4989339	1.5012788	$1.500 \pm .002$
10^5	1.5023502	1.5029479	$1.5026 \pm .0006$
10^6	1.5025711	1.5027434	$1.5026 \pm .0002$
10^7	1.5025132	1.5025913	$1.50255 \pm .00008$

(a) $\sigma = 16$, $\rho = 45.92$, and $\beta = 4$.

T	min	max	γ_1
10^2	1.3605200	1.4273030	$1.39 \pm .07$
10^3	1.3778507	1.3917447	$1.38 \pm .01$
10^4	1.3738285	1.3766673	$1.375 \pm .003$
10^5	1.3736034	1.3740567	$1.3738 \pm .0005$
10^6	1.3743501	1.3746253	$1.3745 \pm .0003$
10^7	1.3744344	1.3744935	$1.37446 \pm .00006$

(b) $\sigma = 16$, $\rho = 40$, and $\beta = 4$.

T	min	max	γ_1
10^2	.86057046	.99438121	$.9 \pm .1$
10^3	.90309866	.91436721	$.91 \pm .01$
10^4	.90595195	.90899048	$.907 \pm .003$
10^5	.90555825	.90698752	$.906 \pm .001$
10^6	.90557440	.90583146	$.9057 \pm .0003$
10^7	.90562532	.90568993	$.90566 \pm .00007$

(c) $\sigma = 10$, $\rho = 28$, and $\beta = 8/3$.

Table 5.3: Estimates of the top Lyapunov exponent of the Lorenz equations with $\sigma = 16$, $\rho = 45.92$ and $\beta = 4$. The method of integration was either a 7th order or the standard 4th order Runge-Kutta method with a constant time step of h till time T . The estimate for γ_1 in the last column was arrived at as in Table 5.2

Method	h	T	min	max	γ_1
RK7	1/16	10^6	1.5017319	1.5020835	$1.5019 \pm .0004$
RK7	1/32	10^6	1.5026282	1.5028236	$1.5027 \pm .0002$
RK7	1/64	10^6	1.5023734	1.5027018	$1.5025 \pm .0003$
RK4	1/128	10^6	1.5025711	1.5027434	$1.5026 \pm .0002$
RK4	1/256	10^6	1.5023756	1.5026027	$1.5025 \pm .0002$




Figure 5.5: Estimates for the top Lyapunov exponent γ_1 of 64 coupled Ginsburg-Landau oscillators with $c_1 = -2.5$, $c_2 = 3.0$ and $K = 0.445$. It appears as though γ_1 is converging to a positive value near .06, especially when the time step is 1/32, before it begins to fall off to 0 at the rate t^{-1} .

The equations

$$\frac{dw_j}{dt} = w_j - (1 + ic_2)|w_j|^2w_j + K(1 + ic_1)(\bar{w} - w_j), \quad 1 \leq j \leq N,$$

with $\bar{w} = \sum_{i=1}^N w_i/N$ define a system of globally coupled Ginsburg-Landau oscillators. The w_j are complex valued and denote the position of the j th oscillator in the complex plane. This system is extensively studied for several values of N , especially with $c_1 = -2.5$, $c_2 = 3.0$, and $K = 0.445$, in [88].

Figure 5.5 shows the computation of the top Lyapunov exponent γ_1 using the standard 4th order Runge-Kutta method. In Figure 5.5, $N = 64$. When the time step is $h = 1/32$, γ_1 appears to converge to a value near .06 till $t = 28000$ before it begins to fall off to 0. The length of the chaotic transient is roughly halved and then halved again when h is decreased to $1/64$ and then to $1/128$. Thus both long integrations and short time steps are useful. The 64 oscillators eventually settle into periodic motion, but it is not clear if the transient that appears chaotic is a numerical artifact or not. The chaotic transient (from a random starting state) persisted with a variety of ODE solvers that were implicit or adaptive or both. It seems as though the chaotic transient arises when a random starting state goes into a periodic state, but the decrease in the length of the transient with the time step suggests that the transient may be a numerical artifact.

There are numerous examples of chaotic transients. Crutchfield and Kaneko [29] show an example where the length of the transient seems to increase exponentially with the system size. Manneville [83] refers to long, seemingly chaotic transients in the Kuramoto-Sivashinsky equation with periodic boundary conditions. The Kuramoto-Sivashinsky equation is often used as a model partial differential equation which can become turbulent.

What is a reliable simulation of chaos? The issue of when to trust a numerical computation comes up in every area of numerical analysis and computational science (indeed, to the chagrin of careful numerical analysts, it frequently does not get the attention it deserves). Numerical computations routinely go beyond what is accessible using mathematical analysis, but judgements of reliability are based on mathematical concepts. In the area of differential equations, for example, energy principles, properties of waves like dispersion and dissipation, and speed of propagation of shocks are used to assess reliability of numerical methods. The roots of the numerical convergence theory of ordinary and partial differential equations are firmly planted in the mathematical study of stability and existence of solutions. But the mathematical understanding of chaos and chaotic transients is very incomplete, especially from a numerical analyst's point of view. Consequently, it is unclear how to distinguish asymptotic chaos from a chaotic transient. Perhaps there is some truth in the assertion that anything comparable to the classical numerical analysis of differential equations is not possible for chaotic phenomena without significant mathematical advances.

We propose that an accurate computation of Lyapunov exponents may be useful for identifying chaotic transients. Since this proposal has been shown to work for

just one example, the coupled Ginsburg-Landau oscillators, its effectiveness is admittedly open to question. But let us summarize the requirements for an accurate computation of Lyapunov exponents. The length of the integration must be at least a few hundred times, and if possible a few thousand times, the reciprocal of the estimated top Lyapunov exponent. Indeed, even this requirement may be too permissive for high dimensional systems. In many cases, it is good practice to repeat the calculations with smaller time steps.

References

- [1] H.D.I. Abarbanel, *Analysis of Observed Chaotic Data*, Springer-Verlag, New York, 1996.
- [2] H.D.I. Abarbanel, R. Brown and M.B. Kennel, Variation of Lyapunov exponents on a strange attractor, *J. of Nonlinear Science 1* (1991), 175-199.
- [3] H.D.I. Abarbanel, R. Brown and M.B. Kennel, Local Lyapunov exponents from observed data, *J. of Nonlinear Science 2* (1992), 343-365.
- [4] M. Abramowitz and I.A. Stegun, eds., *Handbook of Mathematical Functions*, Dover Publications, New York, 1970.
- [5] G. Alefeld and J. Herzberger, *Introduction to Interval Computations*, Academic Press, New York, 1983.
- [6] L. Appel and W. Haken, The solution of the four-color map problem, *Scientific American 237*(4) (Oct. 1977), 108-121.
- [7] L. Arnold and H. Crauel, Random dynamical systems, in *Lyapunov Exponents*, eds. L. Arnold, H. Crauel, J.-P. Eckmann, Lecture Notes in Math. 1486, Springer-Verlag, Berlin, 1991, 1-22.
- [8] L. Arnold, W. Kliemann and E. Oeljeklaus, Lyapunov exponents of linear stochastic systems, in *Lyapunov Exponents*, eds. L. Arnold and V. Wihstutz, Lecture Notes in Mathematics 1186, Springer-Verlag, Berlin, 1986, 85-128.
- [9] L. Arnold and V. Wihstutz, Lyapunov exponents: a survey, in *Lyapunov Exponents*, eds. L. Arnold and V. Wihstutz, Lecture Notes in Mathematics 1186, Springer-Verlag, Berlin, 1986, 1-26.
- [10] R. Bellman, Limit theorem for non-commutative operations, *I. Duke Math. J. 21* (1954), 491-500.
- [11] M. Benedicks, New developments in the ergodic theory of nonlinear dynamical systems, *Philos. Trans. Roy. Soc. London Ser. A 346 no. 1679* (1994), 145-157.
- [12] G. Benettin, L. Galgani, A. Giorgilli and J.-M. Strelcyn, Lyapunov characteristic exponents for smooth dynamical systems: a method for computing all of them, *Meccanica 15* (1980), 9-31.

- [13] M.A. Berger, *An Introduction to Probability and Stochastic Processes*, Springer-Verlag, Berlin, 1993.
- [14] C. Bischof and A. Griewank, Tools for the automatic differentiation of computer programs, *ICIAM/GAMM 95: issue I: Numerical Analysis, Scientific Computing, Computer Science*, eds. G. Alefeld, O. Mahrenholtz and R. Mennicken, 267-288.
- [15] P. Billingsley, *Probability and Measure*, 2nd ed., John Wiley & Sons, New York, 1986.
- [16] P. Bougerol and J. Lacroix, *Random Products of Matrices with Applications to Schrödinger Operators*, Birkhäuser, Basel, 1984.
- [17] R. Bowen and D. Ruelle, The ergodic theory of Axiom-A flows, *Invent. Math.* 29 (1975), 181-202.
- [18] L. Breiman, *Probability*, SIAM, Philadelphia, 1992.
- [19] R.P. Brent, On the zeros of the Riemann zeta function in the critical strip, *Math. of Comp.* 33 (1979), 1361-1372.
- [20] A. Brocot, Calcul des rouages par approximation nouvelle méthode, *Revue Chronometrique* 6 (1860), 186-194.
- [21] A. Carverhill, Furstenberg's theorem for nonlinear stochastic systems, *Probab. theory related fields* 74(4) (1987), 529-534.
- [22] P. Chassaing, G. Letac and M. Mora, Brocot sequences and random walks on $SL_2(R)$, in *Probability measures on groups* 7, ed. H. Heyer, Lecture Notes in Math. 1064, Springer-Verlag, Berlin, 1984, 36-48.
- [23] E. Clarke and J. Wing, Formal methods: state of the art and future directions, *ACM Computing Surveys* 28(4) (1996), 626-643.
- [24] E.A. Coddington and N. Levinson, *Theory of Ordinary Differential Equations*, McGraw-Hill, New York, 1955.
- [25] J.E. Cohen, H. Kesten and C.M. Newman, eds., *Random Matrices and their Applications*, American Mathematical Society, Providence, 1986.
- [26] J.B. Conway, *Functions of One Complex Variable*, Springer, New York, 1995.
- [27] T.H. Cormen, C.E. Leiserson and R.L. Rivest, *Introduction to Algorithms*, MIT press, Cambridge, Massachusetts, (1990).
- [28] A. Crisanti, G. Paladin and A. Vulpiani, *Products of Random Matrices in Statistical Physics*, Springer-Verlag, Berlin, 1992.

- [29] J.P. Crutchfield and K. Kaneko, Are attractors relevant to turbulence?, *Phys. Rev. Lett.* *60* (1988), 2715-2718.
- [30] A.G. Darbyshire and D.S. Broomhead, Robust estimation of tangent maps and Liapunov spectra, *Physica D* *89* (1996), 287-305.
- [31] J. Day and B. Peterson, Growth in Gaussian elimination, *Amer. Math. Monthly* *95* (1988), 489-513.
- [32] R.L. Devaney, *An Introduction to Chaotic Dynamical Systems*, Benjamin/Cummings, California, 1986.
- [33] P. Diaconis and M. Shahshahani, Products of random matrices and computer image generation, in *Random Matrices and their Applications*, eds. J.E. Cohen, H. Kesten and C.M. Newman, American Mathematical Society, Providence, 1986, 173-182.
- [34] L. Dieci, R.D. Russell and E.S. Van Vleck, On the computation of Lyapunov exponents for continuous dynamical systems, *SIAM J. Numer. Anal.* *34*(1) (1997), 402-423.
- [35] L. Dieci and E.S. Van Vleck, Computation of a few Lyapunov exponents for continuous and discrete dynamical systems, *Appl. Num. Math.* *17* (1995), 275-291.
- [36] J.-P. Eckmann, S.O. Kamphorst, D. Ruelle and S. Ciliberto, Liapunov exponents from time series, *Phys. Rev. A* *34* (1986), 4971-4979.
- [37] J.-P. Eckmann and D. Ruelle, Ergodic theory of chaos and strange attractors, *Rev. Mod. Phys.* *57*(3) (1985), 617-656.
- [38] J.-P. Eckmann and D. Ruelle, Fundamental limitations for estimating dimensions and Lyapunov exponents in dynamical systems, *Physica D* *56* (1992), 185-187.
- [39] A. Edelman, *Eigenvalues and Condition Numbers of Random Matrices*, Ph.D. dissertation and Numerical Analysis Report 89-7, Massachusetts Institute of Technology, 1989.
- [40] A. Edelman, Eigenvalues and condition numbers of random matrices, *SIAM J. Matrix Anal. Appl.* *9* (1988), 543-560.
- [41] S. Ellner, R. Gallant, D. McGaffrey and D. Nychka, Convergence rates and data requirements for the Jacobian based estimates of Lyapunov spectra, *Phys. Lett. A* *153* (1991), 357-363.
- [42] M. Embree and L.N. Trefethen, Surprising behavior of random Fibonacci sequences, *NAGp Report No. 97/22*, Oxford University Computing Laboratory.

- [43] K. Falconer, *Fractal Geometry, Mathematical Foundations and Applications*, John Wiley & Sons, New York, 1990.
- [44] W. Feller, *An Introduction to Probability Theory and Its Applications*, Vol. 2, 2nd ed., John Wiley & Sons, New York, 1971.
- [45] L.V. Foster, Gaussian elimination with partial pivoting can fail in practice, *SIAM J. Matrix Anal. Appl.* 15 (1994), 1354-1362.
- [46] H. Furstenberg, Non-commuting random products, *Trans. Amer. Math. Soc.* 108 (1963), 377-428.
- [47] H. Furstenberg and H. Kesten, Products of random matrices, *Ann. Math. Stat.* 31 (1960), 457-469.
- [48] J.A.C. Gallas, Structure of the parameter space of the Henon Map, *Phys. Rev. Lett.* 70(18) (1993), 2714-2717.
- [49] K. Geist, U. Parlitz and W. Lauterborn, Comparison of different methods for computing Lyapunov exponents, *Prog. Theor. Phys.* 83(5) (1990), 875-893.
- [50] I. Goldhirsh, P.-L. Sulem and S.A. Orszag, Stability and Lyapunov stability of dynamical systems: a differential approach and a numerical method, *Physica* 27D (1987), 311-337.
- [51] G.H. Golub and C.F. Van Loan, *Matrix Computations*, 3rd ed., Johns Hopkins Univeristy Press, Baltimore, 1996.
- [52] I.S. Gradshteyn and I.W. Ryzhik, *Table of Integrals, Series, and Products*, 4th ed., Academic Press, New York, 1965.
- [53] R. Graham, D. Knuth and O. Patashnik, *Concrete Mathematics*, Addison-Wesley, Reading, Massachusetts, 1994.
- [54] P. Grassberger and I. Procaccia, Estimating the Kolmogorov entropy from a chaotic signal, *Phys. Rev. A* 28 (1983), 2591-2593.
- [55] P. Grassberger and I. Procaccia, Measuring the strangeness of strange attractors, *Physica* 9D (1983), 189-208.
- [56] J.M. Greene and J.-S. Kim, The calculation of Lyapunov spectra, *Physica* 24D (1987), 213-225.
- [57] A. Griewank, D. Juedes and J. Utke, **ADOL-C**: a package for the automatic differentiation of algorithms written in C/C++, *ACM TOMS* 22(2) (1996), Algorithm 755, 131-167.
- [58] J. Guckenheimer and P. Holmes, *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*, Springer-Verlag, New York, 1990.

- [59] Y. Guivarc'h and A. Raugi, Frontière de Furstenberg, propriétés de contraction et théorèmes de convergence, *Zeit. für Wahrsch. und Verw. Gebiete.* 69 (1985), 187-242.
- [60] A.K. Gupta and Z. Govindarajulu, Distribution of the quotient of two independent Hotelling T2-variates, *Comm. Statist.* 4 (1975), 449-453.
- [61] G. Hardy, J.E. Littlewood and G. Polya, *Inequalities*, 2nd ed., Cambridge University Press, Cambridge, 1988.
- [62] R.Z. Hašminskii, *Stochastic Stability of Differential Equations*, Sijthoff and Noordhoff, Alphen, 1980 (translated from 1969 Russian edition).
- [63] E. Hairer, S.P. Norsett and G. Wanner, *Solving Ordinary Differential Equations I*, 2nd edition, Springer-Verlag, Berlin, 1993.
- [64] N.J. Higham, *Accuracy and Stability of Numerical Algorithms*, SIAM, Philadelphia, 1996.
- [65] F.Y. Hunt and W.M. Miller, On the approximation of invariant measures, *J. Stat. Phys.* 66(1/2) (1992), 535-548.
- [66] *IEEE Standard for Binary Floating-Point Arithmetic, ANSI/IEEE Standard 754-1985*, Institute of Electrical and Electronics Engineers, New York, 1985. Reprinted in *SIGPLAN Notices* 22(2) (1987), 9-25.
- [67] Yu.S. Il'yashenko, On the dimension of attractors of k -contracting systems in an infinite-dimensional space, *Vestn. Mosk. Univ. Ser. 1 Mat. Mekh.* 3 (1983), 52-58.
- [68] J.L. Kaplan and J.A. Yorke, Preturbulence: a regime observed in a fluid flow model of Lorenz, *Comm. Math. Phys.* 67 (1979), 93-108.
- [69] A. Ya. Khintchine, *Continued Fractions*, P. Noordhoff, Groningen, 1963.
- [70] J.F.C. Kingman, Subadditive ergodic theory, *Ann. Prob.* 1 (1973), 883-909.
- [71] J.R. Kinney and T.S. Pitcher, The dimension of some sets defined in terms of f -expansions, *Zeit. für Wahr. und Verw. Gebiete* 4 (1966), 293-315.
- [72] V. Krishnamurthy, The four color theorem, Appendix IV in F. Harary, *Graph Theory*, Indian student edition, Narosa/Addison-Wesley, New Delhi, 1988.
- [73] O.E. Lanford, A computer assisted proof of the Feigenbaum conjecture, *Bull. Amer. Math. Soc.* 6 (1982), 427-434.
- [74] F. Ledrappier, Quelques propriétés des exposants caractéristiques, in *Ecole d'été de Saint-Flour 12-1982*, ed P.L. Hennequin, Lecture Notes in Math. 1097, Springer-Verlag, Berlin, 1984.

- [75] E. Le Page, Théorèmes limites pour les produits de matrices aléatoires, in *Probability Measures on Groups*, ed. H. Heyer, Lecture Notes in Mathematics 928, Springer-Verlag, New York, 1982.
- [76] E. Le Page, Régularité du plus grand exposant caractéristique des produits de matrices aléatoires indépendantes et applications, *Ann. Inst. Henri Poincaré* 25 (1989), 109-142.
- [77] R. Lima and M. Rahibe, Exact Lyapunov exponent for infinite products of random matrices, *J Phys. A: Math. Gen.* 27 (1994), 3427-3437.
- [78] K.A. Loparo and G.L. Blankenship, Almost sure instability of a class of linear stochastic systems with jump process coefficients, in *Lyapunov Exponents*, eds. L. Arnold and V. Wihstutz, Lecture Notes in Mathematics 1186, Springer-Verlag, Berlin, 1986, 160-190.
- [79] E.N. Lorenz, Deterministic non-periodic flow, *J. Atmos. Sci.* 20 (1963), 130-141.
- [80] A.M. Lyapunov, *Problème générale de la stabilité du mouvement*, Comm. Soc. Math. Kharkov 2 (1892), 3 (1893), 265-272; Ann. Fac. Sci. Toulouse 9 (1907), 204-474; Ann. of Math. Studies 17, Princeton University Press, Princeton, 1949.
- [81] *Lyapunov Exponents, Proceedings, Bremen 1984*, edited by L. Arnold and V. Wihstutz, Lecture Notes in Mathematics 1186, Springer-Verlag, Berlin, 1986.
- [82] *Lyapunov Exponents, Proceedings, Oberwolfach 1990*, edited by L. Arnold, H. Crauel and J.-P. Eckmann, Lecture Notes in Mathematics 1486, Springer-Verlag, Berlin, 1991.
- [83] P. Manneville, Liapunov exponents of the Kuramoto-Sivashinsky model, in *Macroscopic Modelling of Turbulent Flows*, Lecture Notes in Physics 230, Springer-Verlag, Berlin, 1985, 319-326.
- [84] P.C. Matthews, R.E. Mirollo and S.H. Strogatz, Dynamics of a large system of coupled nonlinear oscillators, *Physica* 52D (1991), 293-331.
- [85] M.L. Mehta, *Random Matrices and the Statistical Theory of Energy Levels*, Academic Press, New York, 1967.
- [86] A. Milani and A.M. Nobili, An example of stable chaos in the solar system, *Nature* 357 (1992), 569-571.
- [87] K. Mischaikow and M. Mrozek, Chaos in the Lorenz equations: a computer assisted proof, *Bull. Amer. Math. Soc. (N.S.)* 33 (1995), 66-72.
- [88] N. Nakagawa and Y. Kuramoto, Anomalous Lyapunov spectrum in globally coupled oscillators, *Physica* 80D (1995), 307-316.

- [89] V.I. Osseledac, A multiplicative ergodic theorem, Ljapunov characteristic numbers for dynamical systems, *Trans. Moscow Math. Soc.* 19 (1968), 197-231.
- [90] U. Parlitz, Identification of true and spurious Lyapunov exponents from time series, *Intl. J. Bifurcation and Chaos* 2 (1992), 155-165.
- [91] L. Pastur and E.P. Feldman, Wave transmittance for a thick layer of a randomly inhomogeneous medium, *Soviet Physics JETP* 40 (1977), 241-243.
- [92] Y. Peres, Analytic dependence of Lyapunov exponents on transition probabilities, in *Lyapunov Exponents*, ed. L. Arnold, Lecture Notes in Math. 1486, Springer-Verlag, Berlin, 1986, 64-80.
- [93] O. Perron, Über Stabilität und asymptotisches Verhalten der Integrale von Differentialgleichungssystemen, *Math. Zeit.* 29 (1929), 129-160.
- [94] I. Shimada and T. Nagashima, A numerical approach to ergodic problems of dissipative dynamical systems, *Prog. Theor. Phys.* 61 (1979), 1605.
- [95] G. Sansone and R. Conti, *Non-linear Differential Equations*, Macmillan, New York, 1964.
- [96] J.W. Silverstein, The smallest eigenvalue of a large-dimensional Wishart matrix, *Ann. Prob.* 13 (1985), 1364-1368.
- [97] L.A. Smith, Intrinsic limits on dimension calculations, *Phys. Lett. A* 133(6) (1988), 283-288.
- [98] C. Sparrow, *The Lorenz Equations: Bifurcations, Chaos, and Strange Attractors*, Springer-Verlag, New York, 1982.
- [99] M.A. Stern, Ueber eine zahlentheoretische Funktion, *J. für die reine und angewandte Mathematik* 55 (1858), 193-220.
- [100] G. Strang and T. Nguyen, *Wavelets and Filter Banks*, Wellesley-Cambridge Press, Cambridge, Massachusetts, 1996.
- [101] R.S. Strichartz, A. Taylor and T. Zhang, Densities of self-similar measures on the line, *Experimental Mathematics* 4(2) (1995), 101-128
- [102] S.H. Strogatz, *Nonlinear Dynamics and Chaos*, Addison-Wesley, Reading, MA, 1994.
- [103] A.M. Stuart and A.R. Humphries, *Dynamical Systems and Numerical Analysis*, Cambridge University Press, UK, 1996.
- [104] G.J. Sussman and J. Wisdom, Chaotic evolution of the solar system, *Science* 257 (1992), 56-62.

- [105] P.T.P. Tang, Table-driven implementation of the logarithm function for IEEE floating-point arithmetic, *ACM Trans. Math. Soft.* 16(4) (1990), 378-400.
- [106] L.N. Trefethen, Psuedospectra of matrices, in *Numerical Analysis 1991, Proceedings of the 14th Dundee Conference*, eds. D.F. Griffiths and G.A. Watson, Longman, Essex, U.K., 1992, 234-266.
- [107] L.N. Trefethen, Psuedospectra of linear operators, *SIAM Review* 39 (1997), 383-406.
- [108] L.N. Trefethen and D. Bau III, *Numerical Linear Algebra*, SIAM, Philadelphia, 1997.
- [109] L.N. Trefethen and R.S. Schreiber, Average-case stability of Gaussian elimination, *SIAM J. Matrix Anal. Appl.* 11 (1990), 335-360.
- [110] J.N. Tsitsiklis and V.D. Blondel, The Lyapunov exponent and joint spectral radius of pairs of matrices are hard — when not impossible — to compute and to approximate, *Mathematics of Control, Signals and Systems* 10 (1997), 31-40.
- [111] S. Tuljapurkar, Demographic applications of random matrix products, in *Random Matrices and their Applications*, eds. J.E. Cohen, H. Kesten and C.M. Newman, American Mathematical Society, Providence, 1986, 319-326.
- [112] D. Viswanath, Random Fibonacci sequences and the number 1.13198824..., *Math. Comp.*, to appear.
- [113] D. Viswanath, Accurate computation of Lyapunov exponents, submitted.
- [114] D. Viswanath and L.N. Trefethen, Condition numbers of random triangular matrices, *SIAM J. Matrix Anal. Appl.* 19 (1998), 564-581.
- [115] J.H. Wilkinson, Error Analysis of direct methods of matrix inversion, *J. Assoc. Comput. Mach.* 8 (1961), 281-330.
- [116] A. Wolf, J.B. Swift, H.L. Swinney and J.A. Vostano, Determining Lyapunov exponents from a time series, *Physica 16D* (1985), 285-317.
- [117] S.J. Wright, A collection of problems for which Gaussian elimination with partial pivoting is unstable, *SIAM J. Sci. Comput.* 14 (1993), 231-238.
- [118] A.R. Zeni and J.A.C. Gallas, Lyapunov exponents for a Duffing Oscillator, *Physica D* 89 (1995), 71-92.