The Rhythm of Numbers

Continued Fractions, Primes, Music and Self-Similarity

Teddy Rook

Teddy Rook Book
https://teddyrookbook.com/connections/rhythm-of-numbers/

November 2026

Abstract

This monograph grows out of one guess: that π is close to $3 + k/(7k+1)$ for a suitable whole number k. The guess is right at $k = 16$, where it gives $355/113$, the fraction Zu Chongzhi recorded in the fifth century AD, and it is right for a reason: $k/(7k+1)$ is a continued fraction cut off after two steps. We follow that reason in three directions.

First, *rhythm*. Playing a number α as a beat (a hit wherever $\lfloor n\alpha \rfloor$ steps up) produces a Sturmian word, and we give a self-contained proof that its blocks nest exactly as the continued fraction says: the blocks at level n are q_n beats long and hold p_n hits. For $\pi - 3 = [0; 7, 15, 1, 292, \dots]$ this gives bars of 7 beats, phrases of 106 and 113, and a single short phrase that begins at beat 32,996 and closes a block of 33,102 beats. Reading the levels as bar, phrase, section and movement turns a continued fraction into musical form; Teddy Rook used it to compose the six-movement *Continued-Fraction Suite* for piano, which we describe as a method.

Second, *primes*. The same machine (the Euclidean algorithm) generates Euclidean rhythms, tunings from $\log_2(3/2)$, mode-locking staircases and the factoring method of Lehmer, Powers, Morrison and Brillhart. Coprimality ties these to $6/\pi^2$, and we derive the exact correlation of two pulse trains, $(\gcd(a, b) - 1)/\sqrt{(a-1)(b-1)}$. We then pull two open threads with real computations. (a) *Prime convergents.* Combining Lévy's theorem with the equidistribution of convergent denominators modulo m (Szüsz; Moeckel; Jager and Liardet), we derive the heuristic that a typical real number has about $2 \ln 2 \cdot \ln N \approx 1.3863 \ln N$ prime denominators among its first N convergents. A sample of 400 random reals fits the predicted growth; π has 9 prime denominators among its first 5,000 convergents against 10.9 expected, which is typical. (b) *Spacings.* We place a crystal, the rhythm of π (a one-dimensional quasicrystal with two gaps), 5,000 zeros of the Riemann zeta function, random GUE matrices, the primes and a Poisson process on one scale of rigidity, using nearest-neighbor spacings, pair correlation and number variance. The picture organizes the Montgomery–Dyson connection; it proves nothing about the Riemann hypothesis.

Third, *self-similarity*. A continued fraction contains a copy of itself after every term, which is the idea behind Ford circles, the Stern–Brocot tree, Minkowski's question-mark function, the devil's staircase, continued-fraction Cantor sets (whose dimension we recompute as 0.531280506277), the Farey bulbs of the Mandelbrot set and the Fibonacci word fractal. A chi-square test of π's first 10,000 terms against the Gauss–Kuzmin law gives $p = 0.33$ (12 degrees of freedom).

A final chapter audits what is classical, what is known to specialists, what is a new presentation, and what is new. The mathematics is classical or known; what we did not find elsewhere is the reading of a number's continued fraction as a complete musical form, the Suite as an artwork, the prime-convergent experiment at this scale with an explicit constant, and the single spacing spectrum. Every number in the text was computed by the scripts listed in the appendices.

Keywords: continued fractions; Sturmian words; Christoffel words; Euclidean rhythms; mu-

sical form; coprimality; prime numbers; Lévy's constant; pair correlation; random matrices; quasicrystals; fractals; Gauss–Kuzmin law.

How to read this work

Each chapter is written to stand alone and to become one segment of a narrated video; a boxed note *For the video* at the end of each chapter lists what to show. Chapters 1–3 build the main tool. Chapters 4–7 follow it into music, coprimality and the primes, mostly retelling classical results with proofs or precise references. Chapters 8 and 9 are the two research threads, with new computations. Chapter 10 is about self-similarity and fractals, Chapter 11 about the music, Chapter 12 about the wider Connections diagram. Chapter 13 says what is new and what is not; Chapter 14 discusses limits and open questions.

Results that are classical carry a citation in their heading or the sentence that introduces them. Short standard proofs are included so that the text can be read by a strong undergraduate with calculus and a first course in number theory; longer ones are cited. A "Heuristic" is an argument that predicts a number without proving it, and we always say so. Interactive versions of most figures are on the companion pages [68, 69].

Every number in this monograph was computed or checked by the scripts in Appendix C. Numbers that come from data, from the terms of π to the last decimal of a p-value, are inserted into the text automatically from the scripts' output; small numbers that appear in the prose (such as $355/113$ or $292 \times 113 = 32{,}996$) are asserted by the same scripts. All figures were drawn for this work from computed data.

Contents

Abstract **i**

How to read this work **iii**

1 Introduction: a guess about pi **1**
1.1 The guess . . . 1
1.2 Why the guess works . . . 1
1.3 Three directions . . . 2
1.4 What is new and what is not . . . 3
1.5 Conventions . . . 3

2 Continued fractions: the machine **4**
2.1 The algorithm . . . 4
2.2 Convergents and the recurrence . . . 5
2.3 How close the convergents are . . . 5
2.4 Patterns in the terms . . . 6
2.5 The statistics of a typical number . . . 7

3 The rhythm of a number **8**
3.1 Playing a number as a beat . . . 8
3.2 Sturmian and Christoffel words . . . 8
3.3 The nested-block theorem . . . 9
3.4 Pi's levels, and the beat where the surprise lands . . . 10
3.5 From levels to musical form . . . 11
3.6 Three gaps . . . 12

4 Music from the same ladder **14**
4.1 Tuning: the fifth and its convergents . . . 14
4.2 Euclidean rhythms . . . 15
4.3 Polyrhythm and coprimality . . . 16
4.4 Mode locking and the circle map . . . 16

5 Coprimality and $6/\pi^2$ **18**
5.1 The orchard . . . 18
5.2 Coprime means uncorrelated . . . 18
5.3 Hydrogen's $1/n^2$: a shared formula, not shared physics . . . 20

6 Primes in sound and space **22**
6.1 Primes as notes, and the bias of neighbors . . . 22

6.2 A sunflower seeded with primes . . . 22
6.3 Diffraction: a quasicrystal and the primes . . . 24
6.4 Primes as a sum of waves . . . 24

7 Tools that hunt primes **27**
7.1 Factoring with the continued fraction of $\sqrt{N}$ . . . 27
7.2 Mersenne necklaces and Fermat's little theorem . . . 28
7.3 The prime hunt . . . 28

8 Thread (a): prime convergents **31**
8.1 The question . . . 31
8.2 A heuristic with an explicit constant . . . 31
8.3 The test on random reals . . . 33
8.4 Is pi typical? . . . 34
8.5 Do the site's tools help? . . . 35
8.6 What is known, and what is not . . . 36

9 Thread (b): one spectrum, from crystal to randomness **37**
9.1 The question . . . 37
9.2 The point sets . . . 37
9.3 Nearest-neighbor spacings . . . 38
9.4 Pair correlation . . . 39
9.5 Number variance: a scale of rigidity . . . 39
9.6 The spectrum, and Dyson's bridge . . . 40
9.7 What this does and does not show . . . 41

10 Self-similarity and fractals **42**
10.1 Teddy Rook's picture . . . 42
10.2 Ford circles . . . 43
10.3 The Stern–Brocot tree . . . 44
10.4 Minkowski's question-mark function . . . 44
10.5 The devil's staircase . . . 45
10.6 Cantor dust: numbers with small terms . . . 46
10.7 The Mandelbrot set's Farey bulbs . . . 47
10.8 The Fibonacci word fractal . . . 47
10.9 Six numbers, six pictures . . . 47
10.10 Is pi typical? . . . 49

11 Composition: the Continued-Fraction Suite as a method **51**
11.1 The work . . . 51
11.2 The method . . . 51
11.3 What the method gives a composer . . . 53

12 Connections: music, patterns, math and chemistry **55**
12.1 Music and math: ratios, clocks and rhythms . . . 55
12.2 Music and chemistry: octaves of elements . . . 55
12.3 Music and patterns: sand that draws sound . . . 55
12.4 Math and patterns, patterns and chemistry . . . 56
12.5 Math and chemistry: why 2, 8, 18, 32 . . . 56

12.6 All four: waves and symmetry . . . 56
12.7 Where the Rhythm of Numbers sits . . . 56

13 What is new here, and what is re-found **58**
13.1 The audit . . . 58
13.2 The two questions the owner asked . . . 61

14 Discussion **62**
14.1 What the work shows . . . 62
14.2 Limitations . . . 62
14.3 Open questions . . . 63
14.4 Future work . . . 63

A Algorithms **64**

B Computational methods **67**

C Reproducibility **69**

Bibliography **71**

List of Figures

1.1 The starting guess $\pi \approx 3 + k/(7k+1)$ . . . 2

3.1 The nested blocks of π's rhythm . . . 11
3.2 The three-gap theorem for $\pi - 3$ . . . 13

4.1 The fifth in equal temperaments . . . 15
4.2 Euclidean rhythms as necklaces . . . 16
4.3 Arnold tongues of the circle map . . . 17

5.1 The orchard of visible lattice points . . . 19
5.2 The pulse-train correlation matrix . . . 20
5.3 Two roads to $\pi^2/6$ . . . 21

6.1 The last digits of consecutive primes . . . 23
6.2 Pi's sunflower, seeded with primes . . . 23
6.3 Diffraction of π's rhythm and of the primes . . . 25
6.4 The prime staircase rebuilt from zeta zeros . . . 26

7.1 The prime hunt . . . 30

8.1 Convergent denominators avoid small primes . . . 32
8.2 Prime denominators among convergents . . . 34
8.3 Where π falls among random reals . . . 35

9.1 Nearest-neighbor spacings . . . 38
9.2 Pair correlation of zeta zeros . . . 39
9.3 Number variance: one scale of rigidity . . . 40
9.4 The zeros diffract at logarithms of prime powers . . . 41

10.1 The nested tail of π . . . 43
10.2 Ford circles near $\pi - 3$ . . . 44
10.3 The Stern–Brocot tree . . . 45
10.4 Minkowski's question-mark function . . . 45
10.5 The devil's staircase of the circle map . . . 46
10.6 Cantor dust of continued fractions with terms 1 and 2 . . . 47
10.7 The Farey bulbs of the Mandelbrot set . . . 48
10.8 The Fibonacci word fractal and its relatives . . . 48
10.9 The Suite's numbers as square-cutting ladders . . . 49
10.10 Pi's terms against the Gauss–Kuzmin law . . . 50
10.11 Khinchin and Lévy for π . . . 50

11.1 The forms of the six movements 52

12.1 Chladni figures 56

List of Tables

2.1 The first convergents of π 6

3.1 The levels of the rhythm of $\pi - 3$ 11

4.1 Euclidean rhythms 16

7.1 Continued-fraction factoring of $2^{32} + 1$ 28
7.2 The prime hunt in rhythm prefixes 29

8.1 Divisibility of convergent denominators by small primes 32
8.2 Prime convergents of five constants 34

11.1 The six movements 52

13.1 The prior-art audit 58

C.1 The scripts 70

CHAPTER 1

Introduction: a guess about pi

1.1 The guess

This work began with a guess written on a scrap of paper. Pi is a little more than 3. The leftover, $\pi - 3 = 0.14159\ldots$, is a little less than $1/7$. So try

$$\pi \approx 3 + \frac{k}{7k+1} \tag{1.1}$$

for some whole number k. When k is large, $k/(7k+1)$ is just under $1/7$; when k is small it is further below. Somewhere in between it might hit $\pi - 3$ well. Teddy Rook tried values of k one at a time, and Fig. 1.1 shows what happens. The error falls steadily, reaches a sharp minimum at $k = 16$, and then rises again. At $k = 16$ the guess gives

$$3 + \frac{16}{113} = \frac{355}{113} = 3.14159292\ldots,$$

which differs from π by 2.67e-7. That is six correct decimal places from a fraction with a three-digit denominator.

The fraction 355/113 is famous. The Chinese mathematician Zu Chongzhi recorded it in the fifth century AD; it is called *Milü*, "the close ratio" [54, p. 281]; see also Berggren et al. [7]. Two other values of k give classical fractions too: $k = 10$ gives 223/71, the lower bound that Archimedes proved, and the limit $k \to \infty$ gives 22/7, his upper bound. So the guess (1.1) quietly contains three classical fractions for π.

1.2 Why the guess works

The guess works because of a small piece of algebra:

$$\frac{k}{7k+1} = \frac{1}{7 + \dfrac{1}{k}}. \tag{1.2}$$

The right-hand side is a *continued fraction* stopped after two steps. Solving $1/(7+1/k) = \pi - 3$ for k gives

$$k^* = \frac{\pi - 3}{1 - 7(\pi - 3)} = 15.99659,$$

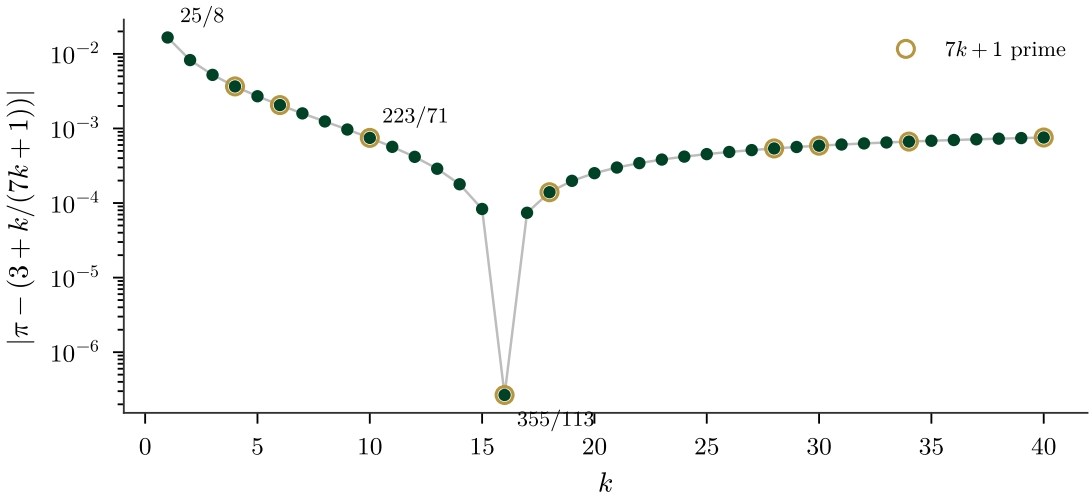


Figure 1.1. The starting guess. The error $|\pi - (3 + k/(7k + 1))|$ for $k = 1, \ldots, 40$ on a logarithmic scale. The sharp dip at $k = 16$ is 355/113; $k = 10$ gives Archimedes' lower bound 223/71 and $k = 1$ gives 25/8. Gold rings mark the k for which $7k + 1$ is prime.

not a whole number but very close to 16. Rounding k^* to the nearest integer is exactly the next step of the continued fraction of π, and the reason the rounding error is so small is that the step after that is huge. Writing out the first few steps,

$$\pi = 3 + \cfrac{1}{7 + \cfrac{1}{15 + \cfrac{1}{1 + \cfrac{1}{292 + \cdots}}}}, \tag{1.3}$$

and $k^* = 15 + 1/(1 + 1/(292 + \cdots)) \approx 16$. The terms of π's continued fraction begin

$$\pi = [3;\ 7, 15, 1, 292, 1, 1, 1, 2, 1, 3, 1, 14, 2, 1, 1, 2, 2, 2, 2, 1, 84, \ldots].$$

Chapter 2 develops this machinery from the beginning.

1.3 Three directions

The rest of this work follows the continued fraction of π, and of a few other numbers, in three directions.

Rhythm and music. Play a number α between 0 and 1 as a beat: on beat n, sound a hit if $\lfloor (n+1)\alpha \rfloor > \lfloor n\alpha \rfloor$, and stay silent otherwise. For $\alpha = \pi - 3$ the hits come every 7 beats, with an occasional gap of 8. Listen longer and the gaps of 8 come in a pattern of their own, and that pattern has a pattern, and so on. Chapter 3 proves that the levels of this rhythm are exactly the levels of the continued fraction: the terms $7, 15, 1, 292, \ldots$ are repeat counts. Read as bar, phrase, section and movement, they give a musical form; Chapter 4 connects the same machinery to tuning, Euclidean rhythms and mode locking, and Chapter 11 describes the *Continued-Fraction Suite*, six piano pieces Teddy Rook composed this way.

Primes. The algorithm that builds a continued fraction is Euclid's algorithm, and Euclid's algorithm is the oldest test for whether two numbers share a prime factor. Chapters 5–7 follow that link through coprime pairs and $6/\pi^2$, the patterns of primes in residues, spirals and diffraction, the zeros of the zeta function, factoring with continued fractions, and a hunt for primes in rhythms. Two questions from the companion site get chapters of their own. Does π's continued fraction have infinitely many prime denominators (Chapter 8)? And what is the right way to picture the famous match between zeros of the zeta function and energy levels of heavy nuclei (Chapter 9)?

Self-similarity. Teddy Rook pictured π as "a number, then a smaller number after it, then a smaller one after that, and so on forever." Chapter 10 makes this precise: the tail of a continued fraction is again a continued fraction, built by the same rule. That one fact underlies a family of classical fractals, and we draw each of them for π.

1.4 What is new and what is not

Almost all of the mathematics here is classical, and we cite it as such. The rhythm of a number is a Sturmian word [53]; the even spreading of beats is the Euclidean rhythm of Bjorklund and Toussaint [79]; Ford circles, the Stern–Brocot tree and the question-mark function date from 1858 to 1938. Chapter 13 carries out a careful audit, idea by idea. In brief, the parts we did not find elsewhere are: the systematic reading of a number's whole continued fraction as nested musical form (bar, phrase, section, movement), together with its use as a compositional method; the Suite itself, which is new as an artwork whatever the prior art on its method; the explicit constant $2\ln 2$ in the heuristic count of prime convergents and its test on hundreds of random reals; and the placement of several point sets on a single spacing spectrum. Even for these we say "we did not find prior work," not "this is new," wherever the literature is too wide to search completely.

1.5 Conventions

Throughout, $\lfloor x \rfloor$ is the floor of x and $\{x\} = x - \lfloor x \rfloor$ its fractional part. A continued fraction is written $[a_0; a_1, a_2, \dots]$. For a number α between 0 and 1 we often write $\alpha = [0; a_1, a_2, \dots]$ and work with $\pi - 3$, $\varphi - 1$, $\sqrt{2} - 1$, $e - 2$ and $\log_2(3/2)$, the five numbers of the Suite. We count beats from 1 when we state the rhythm rule, but we index positions in a beat string from 0, as the companion site does; Section 3.4 says exactly what each count means. "Prime" for numbers with more than a few hundred digits means *probable prime* by the Baillie–PSW test (Appendix B).

For the video. Open on the scrap of paper: $3 + k/(7k+1)$. Slide k from 1 to 40 while Fig. 1.1 draws itself, and stop at the dip. Show the identity (1.2) as two nested fractions. Instrument: "The starting guess" on the Rhythm of Numbers page. About 6 minutes.

CHAPTER 2

Continued fractions: the machine

This chapter collects the facts about continued fractions that the rest of the work uses. All of them are classical; good sources are Khinchin's short book [41] and Chapter X of Hardy and Wright [36]. We include the short proofs because the identities in them are the ones that turn into rhythms later.

2.1 The algorithm

Take a real number x. Write it as its whole part plus a leftover, and turn the leftover upside down:

$$x = a_0 + \frac{1}{x_1}, \qquad a_0 = \lfloor x \rfloor, \quad x_1 = \frac{1}{x - a_0} > 1. \tag{2.1}$$

Now do the same to x_1, then to x_2, and so on: $x_n = a_n + 1/x_{n+1}$ with $a_n = \lfloor x_n \rfloor$. The whole numbers $a_0, a_1, a_2, \ldots$ are the *terms* (or partial quotients) of x, and the real numbers $x_n > 1$ are its *complete quotients*. The process stops exactly when some leftover is zero, which happens exactly when x is rational. Writing the steps inside one another gives

$$x = a_0 + \cfrac{1}{a_1 + \cfrac{1}{a_2 + \cfrac{1}{\ddots + \cfrac{1}{x_n}}}} = [a_0; a_1, \ldots, a_{n-1}, x_n].$$

For π the first complete quotients are $x_1 = 7.06251$, $x_2 = 15.9966$, $x_3 = 1.00342$ and $x_4 = 292.6346$. The size of x_4 is the famous 292 at work.

When $x = a/b$ is rational, the algorithm is Euclid's algorithm in disguise: dividing a by b gives quotient a_0 and remainder r_1, then b by r_1 gives a_1, and so on. Running it on 355 and 113 gives the quotients $3, 7, 16$, so $355/113 = [3; 7, 16]$, and the last nonzero remainder is $\gcd(355, 113) = 1$.

For $0 < x < 1$ the step from x_n to x_{n+1} is one iteration of the *Gauss map*

$$T(x) = \frac{1}{x} - \left\lfloor \frac{1}{x} \right\rfloor, \qquad 0 < x < 1, \tag{2.2}$$

which reads off the term $\lfloor 1/x \rfloor$ and keeps the leftover. Most of the statistics in this work are statements about long runs of T.

2.2 Convergents and the recurrence

Stopping after the term a_n gives a fraction $p_n/q_n = [a_0; a_1, \dots, a_n]$, the nth *convergent*. Its numerator and denominator obey a two-term recurrence that is the engine of everything that follows.

Proposition 2.1 (the recurrence; see 41, Thm. 1)**.** *Set $p_{-1} = 1$, $q_{-1} = 0$, $p_0 = a_0$, $q_0 = 1$, and for $n \geq 1$*

$$p_n = a_n p_{n-1} + p_{n-2}, \qquad q_n = a_n q_{n-1} + q_{n-2}. \tag{2.3}$$

Then $[a_0; a_1, \dots, a_n] = p_n/q_n$ for every n, and more generally $[a_0; a_1, \dots, a_{n-1}, t] = (tp_{n-1} + p_{n-2})/(tq_{n-1} + q_{n-2})$ for every real $t > 0$.

Proof. In matrix form, (2.3) says

$$\begin{pmatrix} p_n & p_{n-1} \\ q_n & q_{n-1} \end{pmatrix} = \begin{pmatrix} a_0 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} a_1 & 1 \\ 1 & 0 \end{pmatrix} \cdots \begin{pmatrix} a_n & 1 \\ 1 & 0 \end{pmatrix}. \tag{2.4}$$

A matrix $\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ acts on numbers by $t \mapsto (at + b)/(ct + d)$, and products of matrices act by composition. The matrix $\left(\begin{smallmatrix} a & 1 \\ 1 & 0 \end{smallmatrix}\right)$ sends t to $a + 1/t$, which is one layer of a continued fraction. Composing the layers gives the general statement, and $t = a_n$ in the last layer gives the first. □

Taking determinants in (2.4), each factor has determinant -1, so

$$p_n q_{n-1} - p_{n-1} q_n = (-1)^{n-1}. \tag{2.5}$$

Two consequences will matter again and again. First, any common divisor of p_n and q_n divides ± 1, so every convergent is in lowest terms. Second, the same is true of neighboring denominators: $\gcd(q_n, q_{n-1}) = 1$. In the language of Chapter 3, the length of a rhythmic block and its number of hits never share a factor, and neither do the lengths of neighboring blocks.

2.3 How close the convergents are

Proposition 2.2 (the error of a convergent; 41, Thms. 9 and 13)**.** *If x is irrational, then*

$$x - \frac{p_n}{q_n} = \frac{(-1)^n}{q_n \left(x_{n+1} q_n + q_{n-1}\right)}, \tag{2.6}$$

and therefore

$$\frac{1}{q_n(q_n + q_{n+1})} < \left| x - \frac{p_n}{q_n} \right| < \frac{1}{q_n q_{n+1}} \leq \frac{1}{q_n^2}. \tag{2.7}$$

Proof. By Proposition 2.1 with $t = x_{n+1}$, $x = (x_{n+1}p_n + p_{n-1})/(x_{n+1}q_n + q_{n-1})$. Subtract p_n/q_n, put everything over a common denominator, and use (2.5). For the bounds, $a_{n+1} < x_{n+1} < a_{n+1} + 1$ gives $q_{n+1} < x_{n+1}q_n + q_{n-1} < q_{n+1} + q_n$. □

The formula (2.6) explains the 292. The error of $p_3/q_3 = 355/113$ is about $1/(113 \cdot x_4 \cdot 113)$, and $x_4 = 292.6346$ is large, so the error is tiny. A big term is a sign that the convergent just before it is unusually good. Table 2.1 lists the first convergents of π with their errors and the normalized error $q_n^2 |\pi - p_n/q_n|$, which by (2.6) is roughly $1/a_{n+1}$.

Convergents are not only close; they are the best possible for their size.

Table 2.1. The first convergents of π. The last column is $q_n^2\,|\pi - p_n/q_n|$, which is close to $1/a_{n+1}$; it is smallest just before the large term 292.

n	a_n	p_n	q_n	$\lvert\pi - p_n/q_n\rvert$	$q_n^2\,\lvert\pi - p_n/q_n\rvert$
0	3	3	1	0.142	0.1416
1	7	22	7	0.001,26	0.06196
2	15	333	106	8.32×10^{-5}	0.9351
3	1	355	113	2.67×10^{-7}	0.003406
4	292	103,993	33,102	5.78×10^{-10}	0.6332
5	1	104,348	33,215	3.32×10^{-10}	0.3659
6	1	208,341	66,317	1.22×10^{-10}	0.5381
7	1	312,689	99,532	2.91×10^{-11}	0.2887
8	2	833,719	265,381	8.72×10^{-12}	0.6138

Theorem 2.3 (best approximation; Lagrange, see 41, Thms. 16–17)**.** *If $1 \le q < q_{n+1}$, then $|qx - p| \ge |q_n x - p_n|$ for every integer p, with equality only for $p/q = p_n/q_n$. In particular no fraction with denominator smaller than q_{n+1} is closer to x than p_n/q_n.*

For π this means that no fraction with denominator below 33,102 beats 355/113; our script checks it directly by scanning every denominator. A partial converse is due to Legendre.

Theorem 2.4 (Legendre; 36, Thm. 184)**.** *If p/q is in lowest terms and $|x - p/q| < 1/(2q^2)$, then p/q is a convergent of x.*

We confirmed it for π by brute force: among all fractions with denominator up to 3,000, the only ones within $1/(2q^2)$ of π are convergents. Legendre's theorem reappears in Chapter 10 as a statement about circles.

2.4 Patterns in the terms

Some numbers have continued fractions with simple patterns:

$$\varphi = \tfrac{1+\sqrt{5}}{2} = [1;1,1,1,\dots], \qquad \sqrt{2} = [1;2,2,2,\dots],$$
$$e = [2;1,2,1,1,4,1,1,6,1,\dots], \quad \log_2 \tfrac{3}{2} = [0;1,1,2,2,3,1,5,2,23,2,2,1,1,55,\dots].$$

For φ the recurrence (2.3) with all terms equal to 1 is the Fibonacci recurrence, and the denominators are $1,1,2,3,5,8,13,21,34,55,89,144,\dots$. For $\sqrt{2}$ they are the Pell numbers $1,2,5,12,29,70,169,408,985,\dots$. The pattern of e, with $2,4,6,\dots$ in every third place, is due to Euler; our scripts verify it for the first 3,900 terms. The number $\log_2(3/2)$, which measures a musical fifth (Chapter 4), shows no pattern, and neither does π.

Theorem 2.5 (Euler and Lagrange; 36, Thms. 176–177)**.** *The continued fraction of x is eventually periodic if and only if x is a quadratic irrational, that is, an irrational root of a quadratic equation with integer coefficients.*

So φ and $\sqrt{2}$ have repeating terms, while π and e, which are not quadratic irrationals, do not. This dichotomy returns in every chapter: repeating terms give exactly self-similar rhythms, music and pictures; non-repeating terms give the same *rule* at every level but a new *shape*.

2.5 The statistics of a typical number

Pick x at random between 0 and 1. What do its terms look like? Gauss conjectured, and Kuzmin proved in 1928 [44], with a sharper version by Lévy [50], that for almost every x the terms follow one fixed law.

Theorem 2.6 (Gauss–Kuzmin law)**.** *For almost every $x \in (0,1)$, the proportion of indices $n \le N$ with $a_n = k$ tends, as $N \to \infty$, to*

$$P(k) = -\log_2\!\left(1 - \frac{1}{(k+1)^2}\right). \tag{2.8}$$

So about 41.5% of the terms are 1, 17.0% are 2 and 9.3% are 3. Behind the law is the *Gauss measure* $d\mu = dx/((1+x)\ln 2)$, which the Gauss map (2.2) leaves invariant, and the fact that T is ergodic for it; the frequency of $a_n = k$ is the Gauss measure of the interval $(1/(k+1), 1/k]$, which is (2.8). Two more laws follow from the same ergodic theory.

Theorem 2.7 (Khinchin 1935; Lévy 1936)**.** *For almost every x,*

$$(a_1 a_2 \cdots a_n)^{1/n} \to K_0 = 2.685452001\ldots \qquad \textit{and} \qquad \frac{\ln q_n}{n} \to \frac{\pi^2}{12\ln 2} = 1.18657\ldots \tag{2.9}$$

The first limit is Khinchin's constant [42]; the second is Lévy's constant [51], and it says that the denominators grow like $e^{1.18657n} \approx 3.27582^n$. Lévy's theorem is the key input to the prime-convergent heuristic of Chapter 8. No one knows whether π obeys any of these laws; nobody has even proved that its terms are unbounded. What we can do is count. From π computed to 11,600 digits, our scripts extract its first 10,000 terms (two precisions agree on every one), and find 4,206 ones (42.1%), a geometric mean of 2.6632 and $\ln q_{10000}/10000 = 1.17941$. Chapter 10 turns these counts into a formal test.

For the video. Animate the algorithm on π: peel off 3, flip 0.14159 to 7.0625, peel off 7, flip again. Show Table 2.1 filling in, with the error column shrinking and dipping at 355/113. End with the slogan "a big term means the fraction before it was unusually good." Instrument: "The ladder." About 8 minutes.

CHAPTER 3

The rhythm of a number

3.1 Playing a number as a beat

Let $0 < \alpha < 1$. Imagine a drummer who keeps a running total, adding α on every beat, and strikes the drum whenever the total passes a whole number. On beat n the total goes from $n\alpha$ to $(n+1)\alpha$, so the rule is:

$$\text{beat } n \text{ is a hit} \iff \lfloor (n+1)\alpha \rfloor > \lfloor n\alpha \rfloor , \qquad n = 1, 2, 3, \dots \tag{3.1}$$

We write a hit as `x` and a rest as `.`, or as 1 and 0. The resulting infinite word $c_\alpha = c_\alpha(1)\, c_\alpha(2)\, c_\alpha(3) \cdots$ is called the *characteristic word* of α. Since the total grows by less than 1 per beat, there is at most one hit per beat, and after L beats there have been exactly $\lfloor (L+1)\alpha \rfloor$ hits. The hits are spread as evenly as the whole numbers allow: the density of hits is α, and any two stretches of the same length differ by at most one hit. For $\alpha = \pi - 3$ the first 40 beats are

```
......x......x......x......x......x.....
```

a hit every 7 beats, because $1/\alpha = 7.06251$ lies between 7 and 8. The leftover $0.0625 \approx 1/16$ means that roughly every 16th gap is an 8 instead of a 7. That is already the start of a hierarchy, and this chapter shows that it continues forever, level by level, exactly as the continued fraction says.

3.2 Sturmian and Christoffel words

Words like c_α have a long history in combinatorics, where they are called *Sturmian words*; the standard reference is Chapter 2 of Lothaire's book, written by Berstel and Séébold [53], and the books of Allouche and Shallit [2, Ch. 9] and Berstel, Lauve, Reutenauer and Saliola [9].

Definition 3.1. A binary word is *balanced* if any two factors (blocks of consecutive letters) of the same length contain numbers of 1s that differ by at most 1. An infinite word is *Sturmian* if it is balanced and not eventually periodic. A *mechanical word* of slope α and intercept ρ has nth letter $\lfloor (n+1)\alpha + \rho \rfloor - \lfloor n\alpha + \rho \rfloor$; the characteristic word is the case $\rho = 0$. For a rational slope p/q in lowest terms, one period of the mechanical word with $\rho = 0$, read as $\lfloor (i+1)p/q \rfloor - \lfloor ip/q \rfloor$ for $i = 0, \dots, q-1$, is the *lower Christoffel word* of slope p/q.

For irrational α, the mechanical words of slope α are exactly the Sturmian words [53, Ch. 2]. A Christoffel word is the finite, rational version: it spreads p hits over q beats as evenly as possible. In music these are the maximally even rhythms of Clough and Douthett [20], and, up to rotation, the Euclidean rhythms of Chapter 4.

3.3 The nested-block theorem

The construction that the companion site uses to play a number builds the rhythm from its continued fraction rather than from the floor rule (3.1). It is the classical construction of *standard words*.

Definition 3.2 (standard words)**.** Let $\alpha = [0; a_1, a_2, \dots]$. Set $s_{-1} = \texttt{x}$, $s_0 = \texttt{.}$, and

$$s_1 = s_0^{a_1 - 1}\, s_{-1}, \qquad s_n = s_{n-1}^{a_n}\, s_{n-2} \quad (n \geq 2), \tag{3.2}$$

where w^a means a copies of w in a row.

In words: the level-1 block is $a_1 - 1$ rests and a hit; each later block is *(term) copies of the block just built, followed by the block before it.* For $\pi - 3$, s_1 is the 7-beat block `.....x`; s_2 is fifteen of those and one extra rest, 106 beats; $s_3 = s_2 s_1$ is 113 beats; $s_4 = s_3^{292} s_2$ is 33,102 beats. In pseudocode (Algorithm 3 in Appendix A):

```
this, prev = "."*(a1-1) + "x", "."; for a in a2, a3, ...: this, prev = this*a
                                    + prev, this
```

The theorem below says that this construction and the floor rule (3.1) produce the same rhythm, and that the block lengths are the convergent denominators.

Lemma 3.3 (floors near a convergent)**.** *Let $n \geq 1$ and $1 \leq m < q_n + q_{n+1}$. Then $\lfloor m\alpha \rfloor = \lfloor mp_n/q_n \rfloor - \varepsilon_n(m)$, where $\varepsilon_n(m) = 1$ if q_n divides m and n is odd, and $\varepsilon_n(m) = 0$ otherwise.*

Proof. By (2.5), p_n/q_n and p_{n+1}/q_{n+1} satisfy $|p_n q_{n+1} - p_{n+1} q_n| = 1$, and α lies strictly between them (Proposition 2.2). A fraction a/b strictly between two such fractions has $b \geq q_n + q_{n+1}$: indeed $|a/b - p_n/q_n| \geq 1/(bq_n)$ and $|p_{n+1}/q_{n+1} - a/b| \geq 1/(bq_{n+1})$, and the two distances add up to $1/(q_n q_{n+1})$, which forces $b \geq q_n + q_{n+1}$. Now take $m < q_n + q_{n+1}$. If an integer k lay strictly between $m\alpha$ and mp_n/q_n, then k/m would lie strictly between α and p_n/q_n, which is impossible. So if mp_n/q_n is not an integer, both numbers have the same floor. If it is an integer (that is, $q_n \mid m$), the floors agree when $\alpha > p_n/q_n$ and differ by one when $\alpha < p_n/q_n$; by (2.6) the convergent lies above α exactly when n is odd. □

Lemma 3.4 (near-periodicity)**.** *For $n \geq 1$ and $1 \leq j \leq q_{n+1} - 2$, $c_\alpha(j + q_n) = c_\alpha(j)$.*

Proof. Let $f(m) = \lfloor mp_n/q_n \rfloor - \varepsilon_n(m)$. Then $f(m + q_n) = f(m) + p_n$, because ε_n has period q_n. By Lemma 3.3, $c_\alpha(j) = f(j+1) - f(j)$ whenever $j + 1 < q_n + q_{n+1}$. For $j \leq q_{n+1} - 2$ both $j+1$ and $j + q_n + 1$ are below $q_n + q_{n+1}$, so $c_\alpha(j + q_n) = f(j + q_n + 1) - f(j + q_n) = f(j+1) - f(j) = c_\alpha(j)$. □

Theorem 3.5 (nested blocks; Fraenkel, Mushkin and Tassa; Brown)**.** *Let $\alpha = [0; a_1, a_2, \dots]$ be irrational, with convergents p_n/q_n and standard words s_n. Then for every $n \geq 0$:*

(i) s_n is q_n beats long and contains p_n hits;

(ii) the length and the number of hits of s_n are coprime, and so are the lengths of s_n and s_{n-1};

(iii) for $n \geq 1$, s_n is the beginning of the characteristic word c_α. Hence the first q_{n+1} beats of the rhythm consist of a_{n+1} copies of the level-n block followed by one level-$(n-1)$ block, and $c_\alpha = \lim_{n\to\infty} s_n$.

Part (iii) is due to Fraenkel, Mushkin and Tassa [29] and Brown [13]; a proof through Sturmian morphisms is in Lothaire [53, Ch. 2]. We give a direct proof from the two lemmas, since the same counting reappears in every musical application.

Proof. (i) The lengths satisfy $|s_{-1}| = |s_0| = 1$, $|s_1| = a_1 = q_1$, and $|s_n| = a_n|s_{n-1}| + |s_{n-2}|$, which is the recurrence (2.3) for q_n with $q_0 = 1$. The hit counts satisfy $0, 1, a_n \cdot (\text{previous}) + (\text{one before})$, starting from s_0 with $0 = p_0$ hits and s_1 with $1 = p_1$ hit, which is the recurrence for p_n.

(ii) follows from (i) and the determinant identity (2.5).

(iii) For $n = 1$: α lies between $1/(a_1+1)$ and $1/a_1$, so $(j+1)\alpha < 1$ for $j < a_1$ and $(a_1+1)\alpha > 1$; the first a_1 beats are $a_1 - 1$ rests and a hit, which is s_1. For $n = 2$ when $a_1 = 1$ (so $q_1 = 1$ and the lemma below does not apply): then $1-\alpha$ lies between $1/(a_2+2)$ and $1/(a_2+1)$, hence $\lfloor j\alpha \rfloor = j - 1$ for $1 \leq j \leq a_2 + 1$ and $\lfloor (a_2+2)\alpha \rfloor = a_2$; the first $a_2 + 1$ beats are a_2 hits and a rest, which is $s_2 = s_1^{a_2} s_0$.

For the inductive step, suppose $n \geq 1$, $q_n \geq 2$, and s_n and s_{n-1} are beginnings of c_α (for $n = 1$ we use $s_0 = \,.\, = c_\alpha(1)$, which holds when $a_1 \geq 2$, the only case with $q_1 \geq 2$). Since $q_n \geq 2$, every position $j \leq q_{n+1} - q_n$ satisfies $j \leq q_{n+1} - 2$, so Lemma 3.4 lets us shift by q_n repeatedly: the first $q_{n+1} = a_{n+1}q_n + q_{n-1}$ beats are a_{n+1} copies of the first q_n beats followed by the first q_{n-1} beats. By hypothesis these are s_n and s_{n-1}, so the first q_{n+1} beats are $s_n^{a_{n+1}} s_{n-1} = s_{n+1}$. Since $q_n \geq 2$ for all $n \geq 2$, the induction runs forever. □

We also confirmed the theorem by computer: for 200 random term sequences (terms drawn from $\{1, 2, 3, 4, 5, 7, 12\}$, words up to several thousand beats) and for $\pi - 3$ over the first 140,000 beats, the standard words agree letter for letter with the floor rule (3.1), and every block has q_n beats and p_n hits.

Remark 3.6. The theorem is the precise form of the statement that "every term of the continued fraction is a repeat count." The term a_{n+1} says how many times the level-n block repeats before the level-$(n-1)$ block closes the level-$(n+1)$ block. Big terms mean long runs of identical blocks; a term equal to 1 means the level adds a single block. For the golden ratio every term is 1, so every level is "one long block, then one short block," and the levels are the Fibonacci numbers.

3.4 Pi's levels, and the beat where the surprise lands

Table 3.1 lists the first levels of the rhythm of $\pi - 3$, and Fig. 3.1 draws them.

The two artifacts on which the companion site is based disagreed about where the first short phrase of level 4 "lands": one said beat 32,996, the other 33,102. Both are right about different things. Number the beats from 0, as the site does. Level 4 is $s_4 = s_3^{292} s_2$: 292 phrases of 113 beats followed by one phrase of 106. The 292 long phrases fill beats 0 through $292 \times 113 - 1 = 32{,}995$, so the short phrase *begins* at beat 32,996; it fills beats 32,996 to 33,101, so the level-4 block *ends* after beat 33,101, and the next block begins at beat $33{,}102 = q_4$. Our

Table 3.1. The levels of the rhythm of $\pi - 3 = [0; 7, 15, 1, 292, 1, 1, \dots]$. Each block is (term) copies of the block above it in the table, then the block two rows up. The fraction is hits over beats, the nth convergent of $\pi - 3$; adding 3 gives the convergent of π.

level n	term a_n	block s_n	beats q_n	hits p_n	convergent of π
1	7	$.^{6}\mathtt{x}$	7	1	22/7
2	15	$s_1^{15}\, s_0$	106	15	333/106
3	1	$s_2\, s_1$	113	16	355/113
4	292	$s_3^{292}\, s_2$	33,102	4,687	103993/33102
5	1	$s_4\, s_3$	33,215	4,703	104348/33215

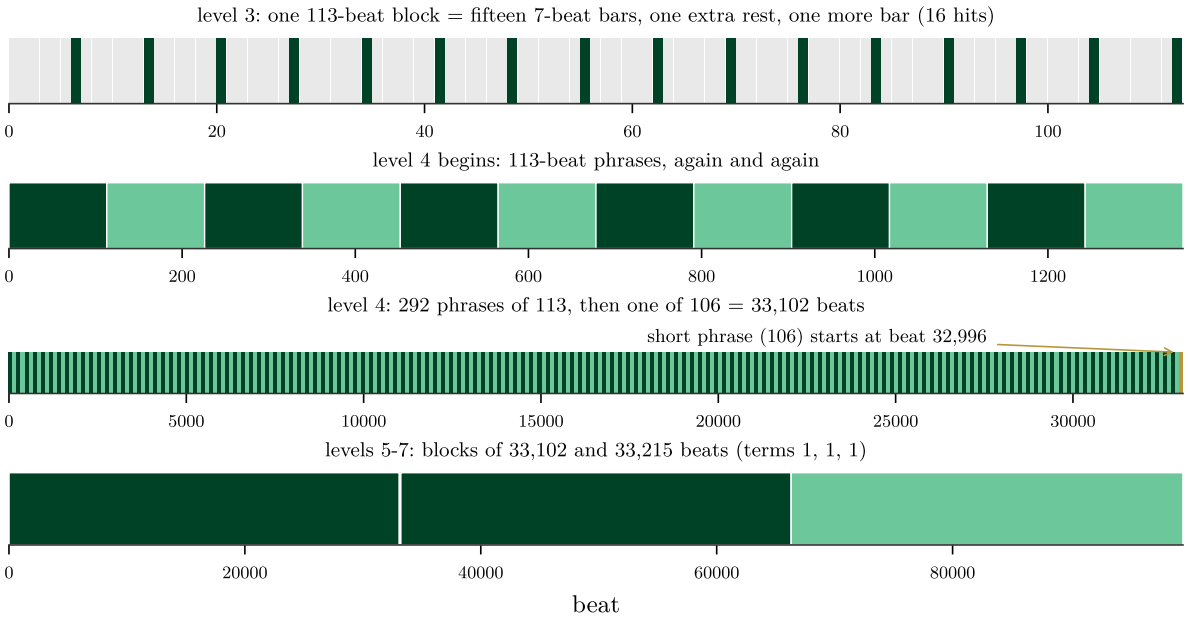


Figure 3.1. The nested blocks of π's rhythm, drawn to scale at four zoom levels. Top: one 113-beat block (hits in green). Second row: the start of level 4, a run of 113-beat phrases. Third row: all of level 4, 292 phrases of 113 beats and then the single 106-beat phrase (gold), which begins at beat 32,996 and closes the block at 33,102. Bottom: levels 5 to 7, where each term is 1.

script locates every short phrase in the first 140,000 beats: they begin at beats 32,996, 66,211, 99,313, 132,528. In the first 80,000 beats there are 706 complete phrases of 113 beats and 2 of 106.

When each beat is one bar of music, as in the fifth and sixth movements of the Suite (Chapter 11), bars are numbered from 1, and the same moment is the bar line after bar 32,996: the first bar of the short cycle is printed as bar 32,997, and the short cycle ends with bar 33,102. At a tempo of one bar every 2.5 seconds, reaching it takes 22.9 hours of continuous playing.

3.5 From levels to musical form

Theorem 3.5 turns any continued fraction into a hierarchy of blocks, and a hierarchy of blocks is what musicians call *form*: notes group into bars, bars into phrases, phrases into sections,

sections into movements. The reading proposed on the companion site, and used by Teddy Rook for the Suite, is simply to name the levels: the shortest block is a bar, the next a phrase, the next a section, and so on. Each term then says how many of one unit make the next, and the second, shorter block at each level is a contrasting unit, a "B" section that the number brings in on its own schedule.

Two small choices make the reading musical. First, a level that adds a single unit to the one below (a term equal to 1 at a low level, or the "extra rest" in $s_2 = s_1^{15} s_0$) can be absorbed: in the Composer on the site, π's bars are 7 beats, its phrases 113, and its sections 33,102. Second, the rhythm itself can be kept inside each bar, so that the accents of a piece follow the number at the smallest scale while its form follows the number at the largest. Chapter 11 shows both choices at work.

3.6 Three gaps

The rhythm of α has only two kinds of gaps between hits, $\lfloor 1/\alpha \rfloor$ and $\lfloor 1/\alpha \rfloor + 1$ (7 and 8 for $\pi - 3$), because it is balanced. A related classical fact concerns the points $\{\alpha\}, \{2\alpha\}, \ldots, \{N\alpha\}$ placed on a circle of circumference 1.

Theorem 3.7 (three-gap theorem; Sós, Świerczkowski)**.** *For any irrational α and any N, the points $0, \{\alpha\}, \ldots, \{N\alpha\}$ cut the circle into arcs of at most three different lengths, and when there are three, the largest is the sum of the other two.*

The statement was conjectured by Steinhaus. Proofs and the history are in Sós [73], Świerczkowski [75] and the survey of Alessandri and Berthé [1], which also explains the link with Sturmian words. The gap lengths are built from the distances $\|q_k \alpha\|$ of the multiples $q_k \alpha$ to the nearest integer, so the continued fraction controls them too. Figure 3.2 shows the distinct gap lengths for $\pi - 3$ and $N \le 120$; our script checks the theorem for every $N \le 1{,}200$. At $N = 100$ the three lengths are $0.0089, 0.0177, 0.0265$. In Chapter 9 this rigidity, two gaps in the rhythm and at most three on the circle, puts π's rhythm at the "quasicrystal" end of a scale whose other end is pure randomness.

For the video. Play the pi rhythm slowly and count the 7s. Then speed up and show Fig. 3.1 zooming out level by level, ending on the single gold phrase after 292 long ones. Explain "beat 32,996 vs. 33,102" with a ruler: one is where the short phrase starts, the other where the block ends. Instruments: "The levels" and the Composer's "Jump to the rarest block." About 10 minutes.

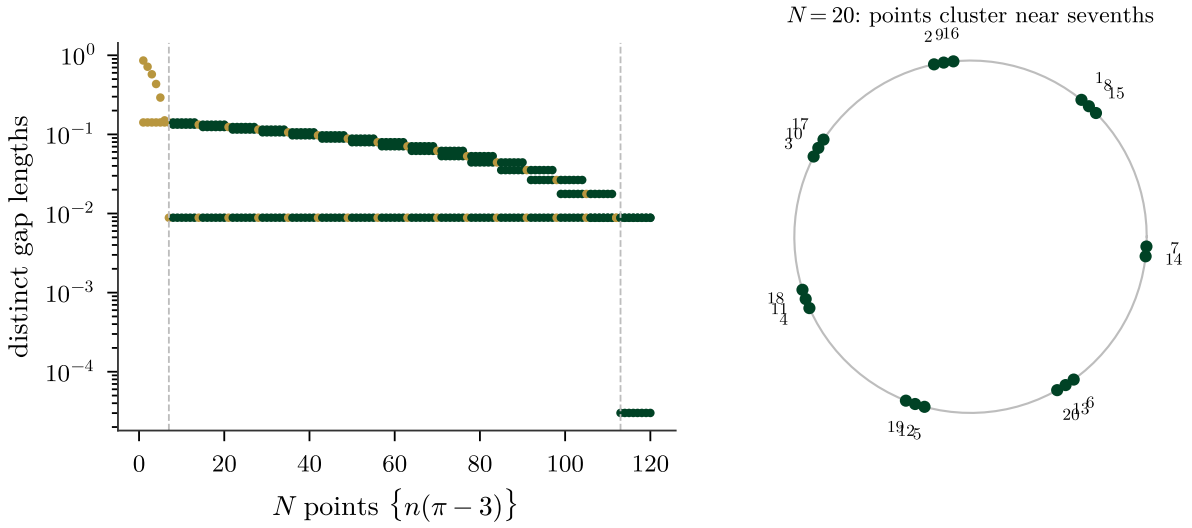


Figure 3.2. The three-gap theorem for $\alpha = \pi - 3$. Left: the distinct arc lengths among the points $0, \{\alpha\}, \ldots, \{N\alpha\}$ for $N = 1, \ldots, 120$ (green: three lengths; gold: two; dashed lines at the denominators $N = 7$ and $N = 113$). The count is two whenever the number of points, $N + 1$, is a denominator such as 7, 106 or 113. Right: the first 20 points, which cluster near multiples of $1/7$ because 7α is close to 1.

CHAPTER 4

Music from the same ladder

The continued fraction is not only a rhythm generator. The same ladder of fractions decides how musical scales are tuned, the same algorithm spreads drum hits around the world's traditional rhythms, and the same fractions are where coupled oscillators lock together. This chapter retells these classical connections in the language of Chapters 2 and 3.

4.1 Tuning: the fifth and its convergents

Two notes sound smooth together when their frequencies are in a simple ratio. The octave is $2:1$ and the fifth is $3:2$. Measured in octaves, a fifth is

$$\log_2 \tfrac{3}{2} = 0.584963\ldots = [0; 1, 1, 2, 2, 3, 1, 5, 2, 23, 2, 2, 1, 1, 55, \ldots].$$

A scale that divides the octave into n equal steps (an *n-note equal temperament*, or n-EDO) contains a good fifth when some k/n is close to $\log_2(3/2)$. The convergents of $\log_2(3/2)$ are

$$1/2, 3/5, 7/12, 24/41, 31/53, 179/306, \ldots$$

and they are the classical answers: $7/12$ is "a fifth is 7 of 12 half steps," and 41 and 53 are the next good stops, long known to theorists [6]. In cents (hundredths of an equal-tempered half step, so 1200 cents per octave), the 12-note fifth is 1.955 cents flat; the 41-note fifth misses by 0.484 cents and the 53-note fifth by 0.068.

Figure 4.1 shows the error of the best fifth in every n-note scale up to $n = 72$. The scales that set a new record for the fifth are $n = 1, 2, 3, 5, 7, 12, 29, 41, 53$. All are convergent denominators except 29, whose fifth $17/29$ is an *intermediate fraction* between the convergents $7/12$ and $24/41$. That is exactly what the theory of best approximations predicts: the record-setters for $|nx - k|$ are the convergents, and the record-setters for $|x - k/n|$ also include some intermediate fractions [41, Ch. II].

The Pythagorean comma. Twelve pure fifths overshoot seven octaves by a ratio, $3^{12}/2^{19} = 531441/524288$, called the Pythagorean comma, which is 23.46 cents. Equal temperament shares the comma out equally: each of its twelve fifths is $23.46/12 = 1.955$ cents flat, which our script confirms to nine decimals. The comma exists because $\log_2(3/2)$ is irrational: no number of pure fifths can ever equal a whole number of octaves.

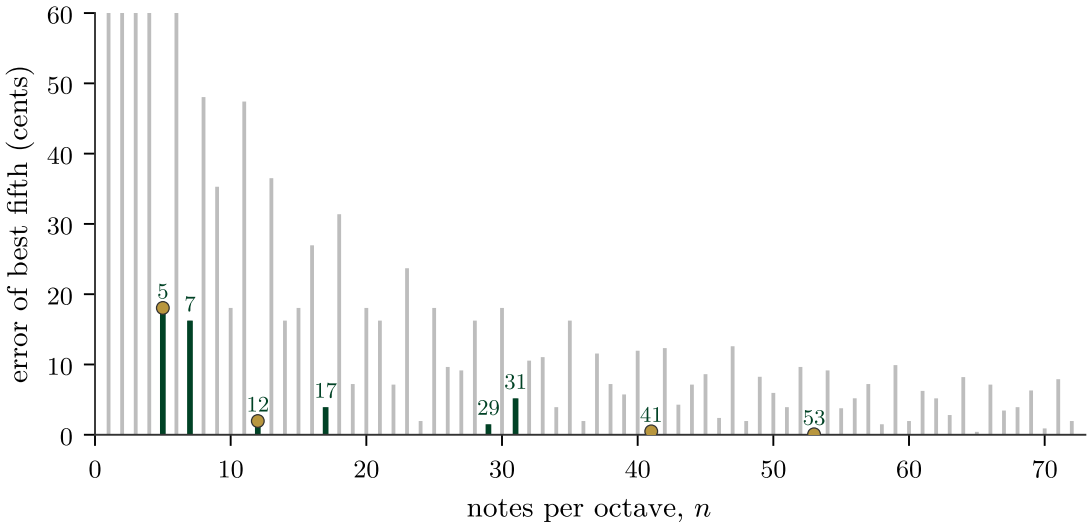


Figure 4.1. The error of the best fifth in every n-note equal temperament, $n = 1, \ldots, 72$, in cents. Green bars and labels mark well-known scales; gold dots mark the convergent denominators of $\log_2(3/2)$: 2, 5, 12, 41, 53.

Thirds, and a prime number of notes. The same comparison for the major third $5 : 4$ gives a 12-note third 13.69 cents sharp, while the 31-note scale, which Christiaan Huygens advocated [6], has a third within 0.78 cents of pure. Stepping around an n-note circle by s notes visits every note exactly when $\gcd(s, n) = 1$, so the number of step sizes that visit all notes is Euler's $\varphi(n)$. The circle of fifths visits all 12 notes because $\gcd(7, 12) = 1$. In a scale with a prime number of notes, such as 7, 31 or 53, every step size visits every note.

4.2 Euclidean rhythms

Suppose we want k drum hits in a cycle of n steps, spread as evenly as possible. In 2003 Eric Bjorklund described an algorithm for exactly this problem, for spacing pulses in the timing system of the Spallation Neutron Source accelerator; in 2005 Godfried Toussaint observed that its outputs include many traditional rhythms from around the world and named them *Euclidean rhythms*, because the algorithm has the structure of Euclid's algorithm [79]. Bjorklund's method starts with k groups `[x]` and $n - k$ groups `[.]`, and repeatedly appends one leftover group to each of the others until at most one leftover group remains. The group counts follow Euclid's algorithm on (n, k).

Proposition 4.1. *For $1 \le k < n$, the Euclidean rhythm $E(k, n)$ is a rotation of the lower Christoffel word of slope k/n, the period of the rhythm of k/n in the sense of* (3.1).

This is known; Gómez-Martín, Taslakian and Toussaint relate Euclidean rhythms to the maximally even sets of Clough and Douthett and to other even-spacing constructions [20, 31]. Our script confirms it for all 780 pairs with $n \le 40$. Two consequences connect the drum machine to π. Running Euclid's algorithm on 355 and 113 produces the quotients $3, 7, 16$, the beginning of π's continued fraction. And $E(16, 113)$ is a rotation of the level-3 block of π's rhythm (Table 3.1): the drum machine with 16 hits in 113 steps plays π at level 3. Figure 4.2 draws some of these rhythms as necklaces.

Table 4.1. Euclidean rhythms $E(k, n)$ from Bjorklund's algorithm, as computed by our script, with the gaps between hits. The names follow Toussaint [79]. The son clave (gaps $3, 3, 4, 2, 4$) is not a rotation of $E(5, 16)$, so it is not Euclidean.

name	rhythm	pattern	gaps
tresillo	$E(3, 8)$	`x..x..x.`	3, 3, 2
cinquillo	$E(5, 8)$	`x.xx.xx.`	2, 1, 2, 1, 2
West African bell	$E(7, 12)$	`x.xx.x.xx.x.`	2, 1, 2, 2, 1, 2, 2
ruchenitza	$E(3, 7)$	`x.x.x..`	2, 2, 3
khafif-e-ramal	$E(2, 5)$	`x.x..`	2, 3
Venda clapping	$E(5, 12)$	`x..x.x..x.x.`	3, 2, 3, 2, 2
samba	$E(7, 16)$	`x..x.x.x..x.x.x.`	3, 2, 2, 3, 2, 2, 2

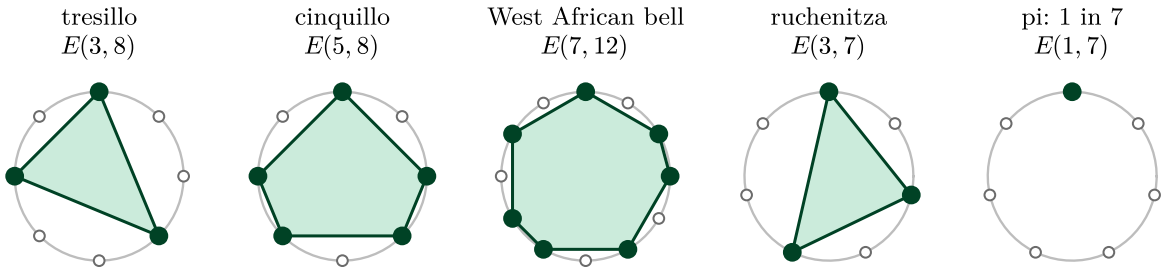


Figure 4.2. Euclidean rhythms drawn as necklaces: steps run clockwise from the top, hits are filled. The hits of each rhythm form a polygon that is as close to regular as the steps allow. The last necklace is the first level of π's rhythm, one hit in seven.

4.3 Polyrhythm and coprimality

Two pulse trains, one with a hit every a steps and one every b steps, coincide every $\operatorname{lcm}(a, b)$ steps. When a and b are coprime, they meet only once per cycle of ab steps: three against four meet only on the downbeat of a 12-step cycle. When they share a factor, they meet more often; four against six meet every 12 steps, twice per cycle of 24. The Polyrhythm layerer on the site stacks the rhythms of several convergents as concentric rings; for π, the rings $1/7$, $15/106$ and $16/113$ play three versions of one beat that drift apart and realign. Because $\gcd(7, 113) = 1$, the 7-ring and 113-ring meet only at the start of their common cycle. Chapter 5 turns "meeting less often" into an exact correlation.

4.4 Mode locking and the circle map

In 1665 Christiaan Huygens noticed that two pendulum clocks hanging from the same beam fell into step; a modern experimental study is Bennett, Schatz, Rockwood and Wiesenfeld [5]. Coupled oscillators tend to lock at simple frequency ratios, and the simplest model of this is the *sine circle map*

$$\theta_{t+1} = \theta_t + \Omega - \frac{K}{2\pi} \sin(2\pi\theta_t), \tag{4.1}$$

where Ω is the natural rotation per step and $K \geq 0$ the strength of the coupling. The *rotation number* $W = \lim(\theta_t - \theta_0)/t$ is the average rotation actually achieved. For $K = 0$, $W = \Omega$. For

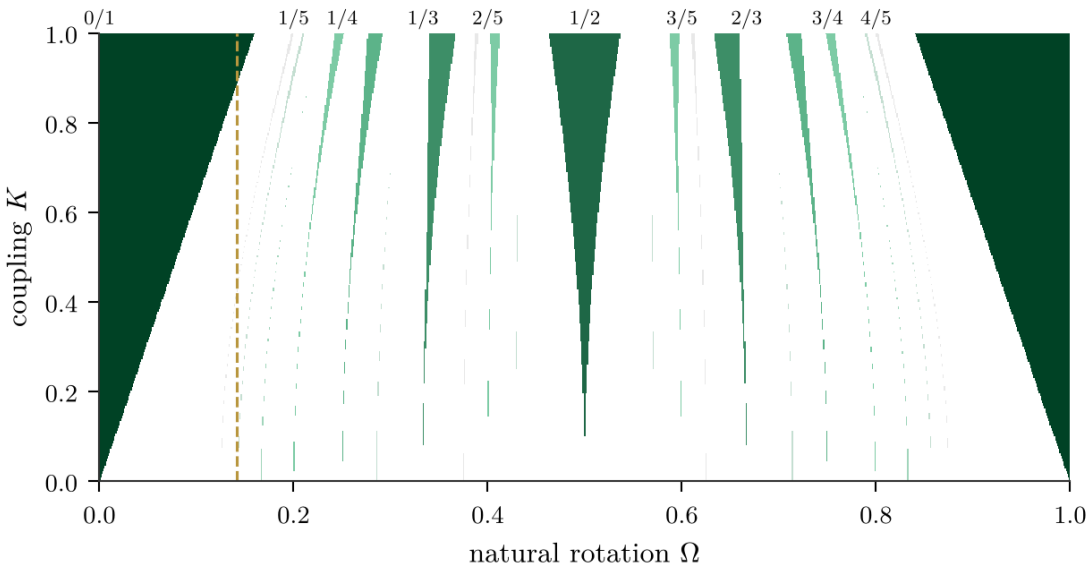


Figure 4.3. Arnold tongues of the sine circle map (4.1). Shaded points (Ω, K) have rotation number p/q with $q \le 8$ (darker: smaller q), computed from 400 iterations after a transient. The dashed gold line is $\Omega = \pi - 3$, which enters the 0/1 tongue at $K = 2\pi(\pi - 3) \approx 0.8897$.

$0 < K \le 1$ the map is a homeomorphism of the circle and W exists; it is locked to a fraction p/q on an interval of Ω around p/q, and these intervals widen with K into the *Arnold tongues* of Fig. 4.3. At $K = 1$ the locked intervals fill the whole line except a Cantor set of measure zero, a "complete devil's staircase" [40], which Chapter 10 explores.

The 0/1 tongue can be computed exactly. A fixed point of (4.1) (modulo 1) needs $\Omega = (K/2\pi)\sin(2\pi\theta)$, which has a solution if and only if $|\Omega| \le K/(2\pi)$. So at $K = 1$ the 0/1 step extends $1/(2\pi) = 0.1592$ on each side of 0; our numerical staircase gives 0.1592. The number $\pi - 3 = 0.14159\ldots$ is inside it, so a system whose natural rotation is $\pi - 3$ is captured into the 0/1 lock once $K \ge 2\pi(\pi - 3) = 0.8897$. Numerically, at $K = 1$ the 1/2 step is 0.074 wide and the 1/3 step 0.031, and the steps with denominators up to 10 already cover a fraction 0.598 of the interval $[0, 1]$.

Counting the tongues. Each tongue belongs to a fraction p/q in lowest terms, so the number of tongues with $0 \le p/q \le 1$ and $q \le N$ is $1 + \sum_{q \le N} \varphi(q)$, the length of the Farey sequence of order N. A classical theorem [36, Thm. 330] gives $\sum_{q \le N} \varphi(q) \sim 3N^2/\pi^2$. For $N = 50$, 100 and 200 our counts are 775, 3,045 and 12,233, against $3N^2/\pi^2 \approx 760$, 3040 and 12159. Here π appears through $6/\pi^2$, the subject of the next chapter.

For the video. Play a pure fifth against a 12-tone fifth and let the beating be heard. Show Fig. 4.1 with the convergent denominators lighting up. Then the drum machine: tresillo, cinquillo, and "pi: 16 in 113." Finish on the circle map with the tongues growing as K rises and the pi line being captured. Instruments: "Circle of fifths," "Euclid & Bjorklund," "Mode locking," the Polyrhythm layerer and the drum machine. About 10 minutes.

CHAPTER 5

Coprimality and $6/\pi^2$

Every level of a rhythm is a coprime pair: by Theorem 3.5(ii) a block of q_n beats holds p_n hits with $\gcd(p_n, q_n) = 1$. This chapter is about coprimality itself: how common it is, how it appears as statistical independence, and where π enters through $6/\pi^2$.

5.1 The orchard

Stand at the origin of the integer lattice and look out over an orchard with a tree at every point (a, b) with $a, b \geq 1$. The tree at (a, b) is visible exactly when no other tree stands on the segment from the origin to it, that is, when $\gcd(a, b) = 1$. Figure 5.1 shows the visible trees in a 40×40 corner, 0.6119 of the grid. For a 1000×1000 corner the share is 0.608383.

Theorem 5.1 (36, Thm. 332)**.** *The proportion of pairs* (a, b) *with* $1 \leq a, b \leq N$ *and* $\gcd(a, b) = 1$ *tends to* $6/\pi^2 = 0.607927\ldots$ *as* $N \to \infty$.

The heuristic behind it is short. A prime p divides both a and b with probability $1/p^2$, and divisibility by different primes behaves independently, so the probability that no prime divides both is

$$\prod_{p} \left(1 - \frac{1}{p^2}\right) = \frac{1}{\zeta(2)}, \qquad \zeta(2) = \sum_{n \geq 1} \frac{1}{n^2} = \frac{\pi^2}{6}. \tag{5.1}$$

The first equality is Euler's product formula [27]: expanding each factor $(1 - p^{-2})^{-1} = 1 + p^{-2} + p^{-4} + \cdots$ and multiplying out, unique factorization produces each $1/n^2$ exactly once. The second is Euler's solution of the Basel problem [26]. The theorem makes the heuristic rigorous by counting with the Möbius function. Numerically, the product over the primes below one million is 0.60792714, against $6/\pi^2 = 0.60792710$.

So π, a number about circles, measures how often two whole numbers share no prime. The convergents of $\pi - 3$, read as pairs (beats, hits) such as $(7, 1)$, $(106, 15)$ and $(113, 16)$, are all visible trees.

5.2 Coprime means uncorrelated

The orchard looks like a matrix, and it is one. Treat "a hit every ath step" as a signal $x_t^{(a)} = 1$ if $a \mid t$ and 0 otherwise, and compare two such signals over one common period.

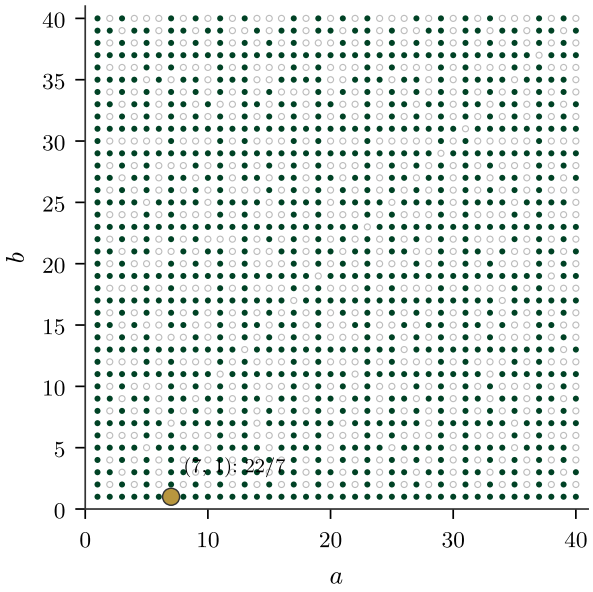


Figure 5.1. The orchard: lattice points (a, b) with $1 \le a, b \le 40$. Filled points are visible from the origin ($\gcd(a, b) = 1$); open circles are hidden. The gold point $(7, 1)$ is the first level of π's rhythm, read as (beats, hits).

Proposition 5.2 (the correlation of two pulse trains)**.** *For integers $a, b \ge 2$, the Pearson correlation of $x^{(a)}$ and $x^{(b)}$ over a full period of* $\operatorname{lcm}(a, b)$ *steps is*

$$\operatorname{corr}(a, b) = \frac{\gcd(a, b) - 1}{\sqrt{(a - 1)(b - 1)}}. \tag{5.2}$$

In particular the correlation is zero exactly when a and b are coprime, and it is 1 when $a = b$.

Proof. Over a period, $x^{(a)}$ has mean $1/a$ and, being 0 or 1, variance $\frac{1}{a}(1 - \frac{1}{a}) = (a - 1)/a^2$. The product $x_t^{(a)} x_t^{(b)}$ is 1 exactly when t is a multiple of $\operatorname{lcm}(a, b) = ab/\gcd(a, b)$, so its mean is $\gcd(a, b)/(ab)$. The covariance is therefore

$$\frac{\gcd(a, b)}{ab} - \frac{1}{a} \cdot \frac{1}{b} = \frac{\gcd(a, b) - 1}{ab},$$

and dividing by $\sqrt{(a - 1)(b - 1)}/(ab)$, the product of the standard deviations, gives (5.2). □

Our script checks (5.2) against a direct computation for every pair $2 \le a \ne b \le 30$. For example $\operatorname{corr}(4, 6) = 1/\sqrt{15} = 0.258$, which is why four against six "keeps colliding," while three against four, being coprime, is uncorrelated. In the 39×39 matrix of pulse trains $2, \ldots, 40$ (Fig. 5.2), the share of off-diagonal zeros is 0.607, close to $6/\pi^2$ for the same reason as the orchard. A prime's row is zero except at its multiples: primes are the "independent" pulse trains of arithmetic.

For periods 2 to 40 the matrix is positive definite: its eigenvalues run from 0.072 to 5.38. Random-matrix theory is used in finance to separate real correlations among stock returns from noise [45, 66]; how the spectrum of this arithmetic correlation matrix compares with such random benchmarks is a question we leave open (Chapter 14). Formula (5.2) itself is an elementary exercise; we did not find it stated in this form in the literature we searched, but we expect it is known, and we claim no novelty for it (Chapter 13).

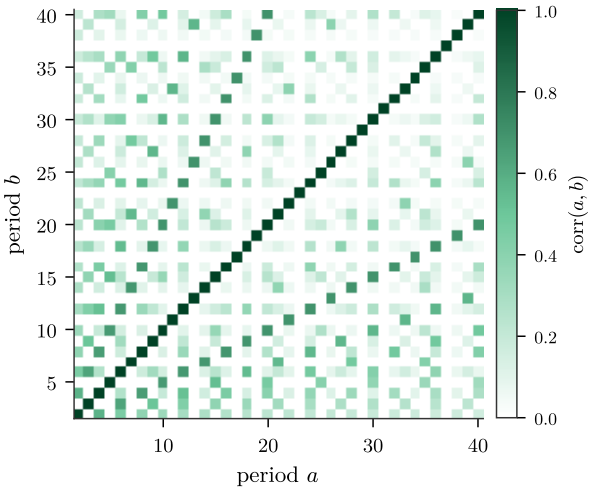


Figure 5.2. The correlation matrix (5.2) of pulse trains with periods 2 to 40. White cells are zero, exactly the coprime pairs; the pattern is the orchard of Fig. 5.1 seen as a matrix.

5.3 Hydrogen's $1/n^2$: a shared formula, not shared physics

Hydrogen's energy levels are proportional to $-1/n^2$, and its visible spectral lines, the Balmer series, come from the differences $1/4 - 1/n^2$. In wavelengths,

$$\lambda_n = B\,\frac{n^2}{n^2-4}, \qquad B = 364.56 \text{ nm}, \tag{5.3}$$

which is the formula Johann Balmer found in 1885 [4]. It gives $\lambda_3 = 656.2\,\text{nm}$ (the red H-alpha line) and $\lambda_4 = 486.1\,\text{nm}$.

Add up the reciprocal squares that set the levels, $1 + \frac{1}{4} + \frac{1}{9} + \cdots$, and the answer is $\zeta(2) = \pi^2/6$, the reciprocal of the orchard's $6/\pi^2$, and also, by (5.1), a product over the primes (Fig. 5.3). It is tempting to say that "hydrogen knows about the primes." It does not. The match is a match of formulas: the sum of $1/n^2$ is the same series wherever it appears. No physical quantity of hydrogen is the sum $\sum 1/n^2$, and nothing in the atom depends on prime factorization. We include it because it is a good illustration of how one series can sit inside several subjects, which is the theme of Chapter 12.

For the video. Walk into the orchard and count visible trees; show the share settling near 61%. Derive (5.2) on screen in four lines. Then 3 against 4 versus 4 against 6, heard and seen in the matrix. End with the Balmer lines and the careful sentence "same formula, not same physics." Instruments: "The orchard," "The correlation matrix," "Chemistry: hydrogen." About 8 minutes.

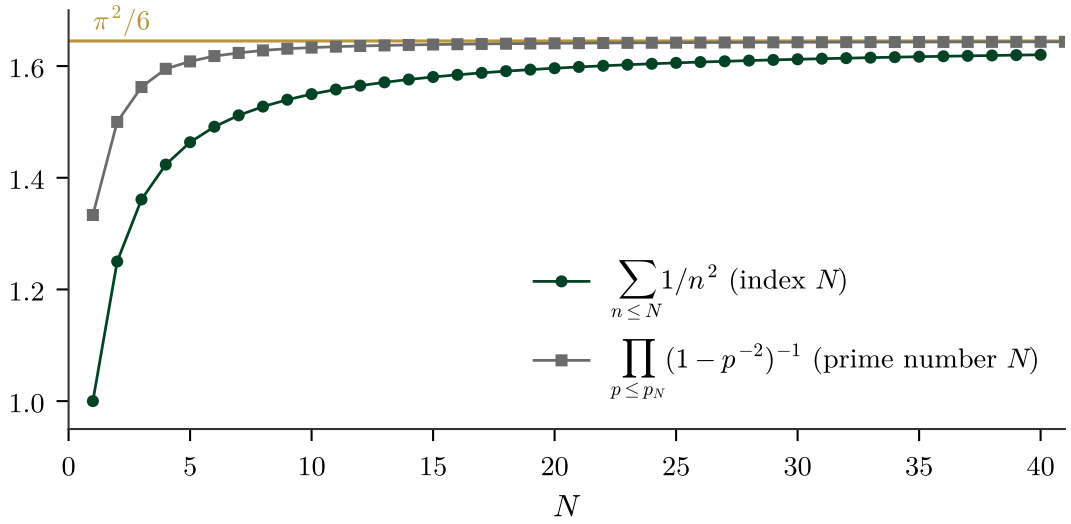


Figure 5.3. Two roads to $\pi^2/6$: the partial sums $\sum_{n \le N} 1/n^2$ (green, the reciprocal squares behind hydrogen's levels) and the partial Euler products $\prod (1 - p^{-2})^{-1}$ over the first N primes (grey). The product converges much faster.

CHAPTER 6

Primes in sound and space

This chapter collects the places on the companion site where the primes are heard or seen: as notes, as arms of a spiral, as columns of a wheel, as peaks in a diffraction pattern and as a sum of waves. Each is a classical or recent published result; our contribution is to compute each one again and to tie it to the rhythm of π.

6.1 Primes as notes, and the bias of neighbors

Number the twelve notes of the octave $0, \dots, 11$. Every prime above 3 leaves remainder 1, 5, 7 or 11 on division by 12, because the other remainders share a factor with 12; these are the $\varphi(12) = 4$ notes C♯, F, G and B. Playing the primes in order as notes, with each gap to the next prime as a duration, gives the "prime music" instrument on the site.

If the remainders of consecutive primes were independent and equally likely, a prime would repeat the remainder of the one before it one time in four. They do not. Lemke Oliver and Soundararajan [48] showed numerically that consecutive primes avoid repeating their remainder, and explained the bias by a conjecture of Hardy–Littlewood type that predicts it fades only like $1/\ln x$. Our counts agree (Fig. 6.1). Modulo 10, among the primes from 7 up to 10^5, 10^6 and 10^7, the share of repeats is 13.9%, 15.4% and 16.3%, against 25% for independence; modulo 12 it is 12.8%, 15.0% and 16.4%. The share creeps up toward 25% as the range grows, slowly, as predicted.

6.2 A sunflower seeded with primes

Place seed number k at angle $2\pi k\alpha$ and distance $\sqrt{k}$ from the center. With $\alpha = \varphi - 1$ (the golden angle) this is the familiar model of a sunflower head; with $\alpha = \pi - 3$ the seeds fall into 7 spokes that curl into 113 arms (Fig. 6.2). The reason is the rhythm again: $113\alpha = 113\pi - 339$ differs from the integer 16 by only 3.0×10^{-5}, so seeds k and $k + 113$ sit at almost the same angle, and arm j of the 113 holds the seeds $k \equiv j \pmod{113}$.

Now light only the prime seeds. A prime larger than q can lie in arm j of q only if $\gcd(j, q) = 1$, each of those $\varphi(q)$ arms holds infinitely many primes by Dirichlet's theorem of 1837 [36, Thm. 15], and by the prime number theorem for arithmetic progressions they share the primes equally in the long run. Since 113 is prime, primes avoid only the arm $j = 0$; among

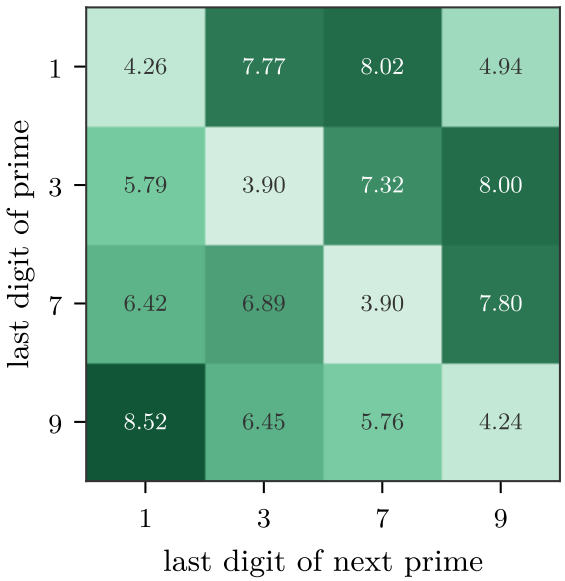


Figure 6.1. Last digits of consecutive primes up to 10^7 (664,576 primes from 7 on). Each cell is the percentage of pairs (prime, next prime) with the given last digits; independence would give 6.25% everywhere. The diagonal, where the digit repeats, is depleted.

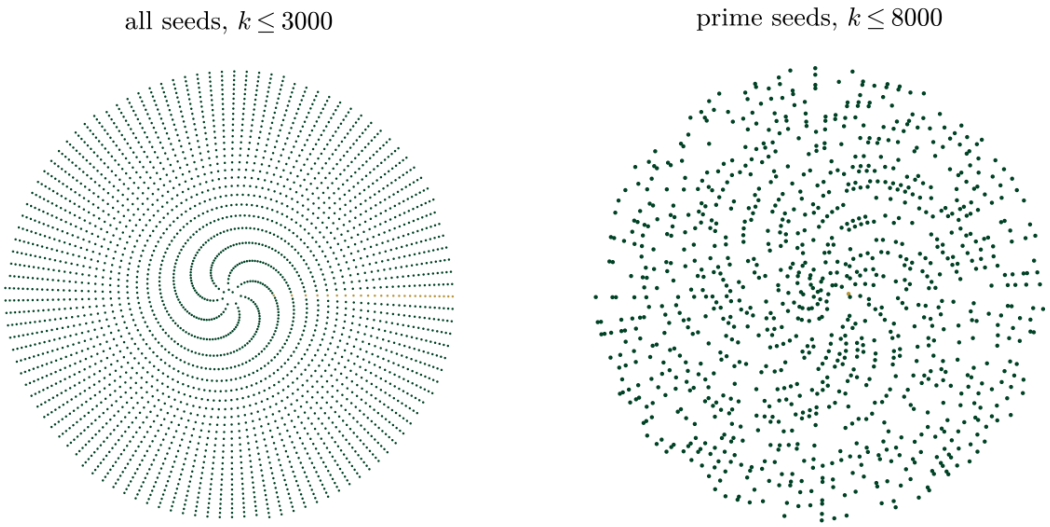


Figure 6.2. Pi's sunflower: seed k at angle $2\pi k(\pi - 3)$ and radius $\sqrt{k}$. Left, all seeds up to 3,000 form 7 spokes that curl into 113 arms (seeds divisible by 113 in gold). Right, only the prime seeds up to 8,000: they fill the arms evenly.

the primes up to 8,000 they occupy 112 of the 113 arms. With 106 arms, which is 2×53, primes occupy only $52 = \varphi(106)$ arms, and half the spiral goes dark.

Wheels. The same fact, straightened out, is *wheel factorization*: lay the integers in rows of width w, and the primes beyond w fall only in the $\varphi(w)$ columns coprime to w. With $w = 30$ a sieve can skip 73% of all numbers without testing them, and with $w = 210$, 77.1%. Pi's

rhythm, laid in rows of width 113, forms 16 nearly straight columns, one for each hit in a 113-beat block.

6.3 Diffraction: a quasicrystal and the primes

Crystallographers read the order of a material from its diffraction pattern, which is (up to constants) the squared modulus of the Fourier transform of the atom positions. For N points $x_1, \dots, x_N$ on a line, define the structure factor

$$S(k) = \frac{1}{N} \left| \sum_{j=1}^{N} e^{2\pi i k x_j} \right|^2 . \tag{6.1}$$

A random set gives $S(k) \approx 1$ for all $k \neq 0$; a periodic crystal gives sharp peaks of height about N at the multiples of its reciprocal lattice. In 1982 Dan Shechtman observed sharp diffraction from a metal alloy whose symmetry no periodic crystal can have [72]; such ordered but aperiodic materials are *quasicrystals*, and he received the 2011 Nobel Prize in Chemistry.

Pi's rhythm is a one-dimensional quasicrystal. Put an atom at each hit of π's rhythm. Sturmian point sets are standard examples of one-dimensional quasicrystals: their diffraction is pure point, with Bragg peaks at the points $m\alpha + n$ of the "Fourier module" $\mathbb{Z} + \alpha\mathbb{Z}$ [3]. In Fig. 6.3 (top) the peaks sit at $\{m(\pi - 3)\}$, folded into $[0, \frac{1}{2}]$; the peak at $k = \pi - 3$ has height 397 for the hits among the first 3,000 beats.

The primes are "effectively limit-periodic." Torquato, Zhang and de Courcy-Ireland [78] computed the structure factor of the primes in long intervals and found dense Bragg-like peaks at rational wavenumbers $k = a/q$ with *square-free* q, and none at $1/4$ or $1/9$; they called this order "effectively limit-periodic." The rational positions and the square-free rule have a short explanation. If the primes in a window are equidistributed among the $\varphi(q)$ residue classes coprime to q, then for $\gcd(a, q) = 1$

$$\sum_{p} e^{2\pi i a p/q} \approx \frac{N}{\varphi(q)} \sum_{\substack{r \bmod q \\ \gcd(r,q)=1}} e^{2\pi i a r/q} = \frac{N\,\mu(q)}{\varphi(q)},$$

because the inner sum is a Ramanujan sum, which equals the Möbius function $\mu(q)$ [36, §16.6]. So $S(a/q) \approx N\mu(q)^2/\varphi(q)^2$, which vanishes unless q is square-free. For the 1,472 primes in $[10^6, 10^6 + 2 \cdot 10^4)$ our computation gives $S(1/2) = 1472$ (all are odd), $S(1/3) = 369$ (the prediction $N/4$ is 368), but only $S(1/4) = 1.20$ and $S(1/9) = 0.57$, at the level of the random background.

Freeman Dyson suggested that the zeros of the zeta function, if the Riemann hypothesis is true, form a one-dimensional quasicrystal in a suitable sense, and that classifying such quasicrystals might lead somewhere [24]. Chapter 9 takes up that remark.

6.4 Primes as a sum of waves

Let $\psi(x) = \sum_{p^m \le x} \ln p$, a staircase that climbs by $\ln p$ at every prime power p^m. Riemann's 1859 paper [67] introduced the idea, and von Mangoldt proved [80], that for $x > 1$ not a prime

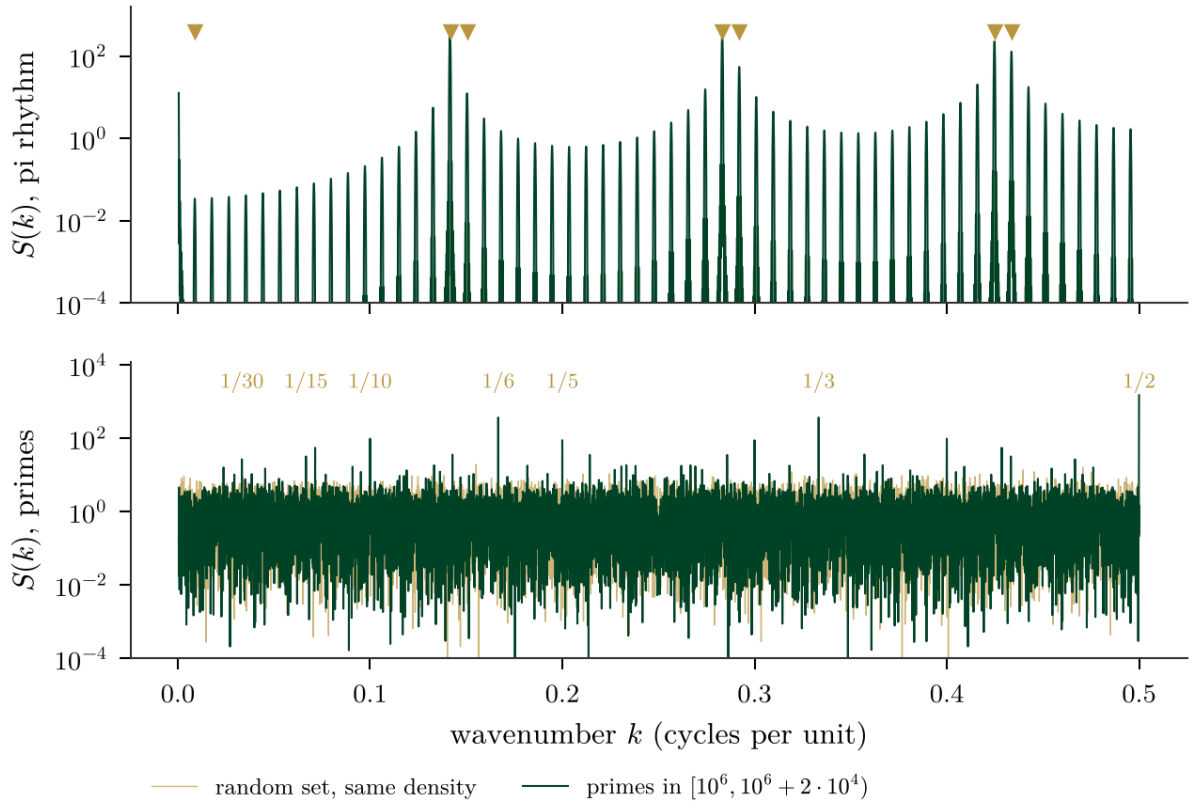


Figure 6.3. Structure factors (6.1) on a logarithmic scale (computed with a smooth taper to suppress side lobes). Top: atoms at the hits of π's rhythm in the first 3,000 beats; gold triangles mark $\{m(\pi-3)\}$ for $m = 1, \dots, 7$, folded into $[0, \frac{1}{2}]$. Bottom: atoms at the primes in $[10^6, 10^6 + 2 \cdot 10^4)$ (green) and at a random set of the same size (gold); the prime peaks sit at fractions with square-free denominators.

power

$$\psi(x) = x - \sum_{\rho} \frac{x^{\rho}}{\rho} - \ln(2\pi) - \tfrac{1}{2} \ln\bigl(1 - x^{-2}\bigr)\,, \tag{6.2}$$

where ρ runs over the nontrivial zeros of the zeta function, summed symmetrically in pairs $\rho, \bar{\rho}$. Writing $\rho = \frac{1}{2} + i\gamma$ (as all computed zeros are), each pair contributes a wave $-2\,\mathrm{Re}(x^{\rho}/\rho)$ with amplitude about $\sqrt{x}/\gamma$ and frequency γ in the variable $\ln x$. The first zeros have $\gamma = 14.13473$, 21.02204 and 25.01086. Figure 6.4 rebuilds ψ from 3, 20 and 100 pairs of zeros: the waves add up to the staircase. The $\ln 2\pi$ in (6.2) is one more place where π sits next to the primes.

The Montgomery–Dyson observation, stated carefully. In 1972 Hugh Montgomery, studying the spacings between the γ's, derived a pair-correlation function $1 - (\sin \pi u / \pi u)^2$ for suitably normalized zeros; he proved it for a restricted class of test functions assuming the Riemann hypothesis, and conjectured it in general [59]. Freeman Dyson recognized the function as the pair correlation of eigenvalues of large random Hermitian matrices from the Gaussian Unitary Ensemble (GUE), a model of the energy levels of heavy nuclei [25, 55]. Odlyzko's large-scale computations showed excellent agreement [61]. What is known is the agreement of statistics; there is no known Hermitian operator whose eigenvalues are the zeros, and no proof that the statistics must agree. Chapter 9 computes these statistics for ourselves.

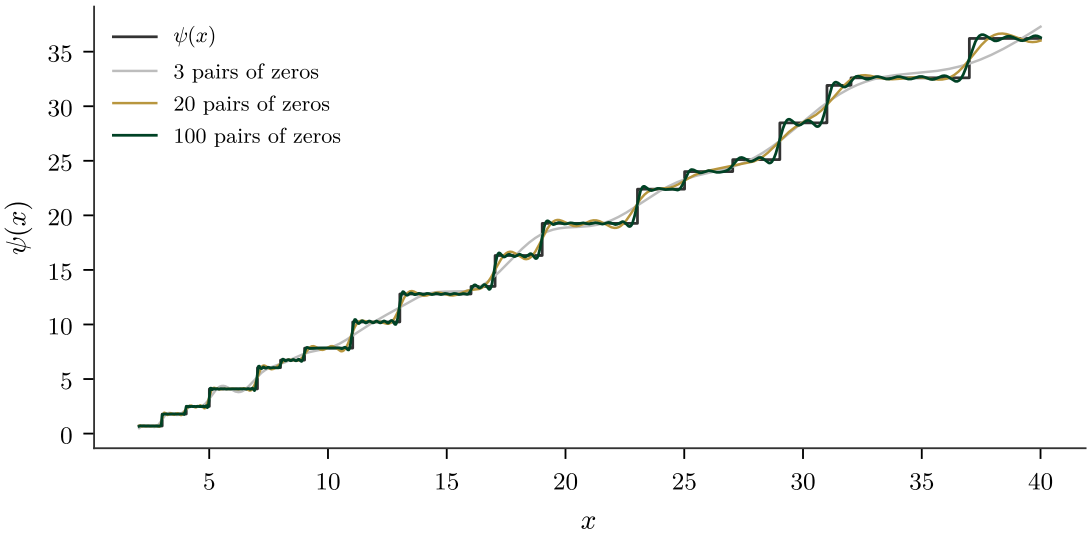


Figure 6.4. The explicit formula (6.2): the staircase $\psi(x)$ (black) and its reconstruction from the first 3, 20 and 100 pairs of zeta zeros.

For the video. Play the primes as notes and show the diagonal of Fig. 6.1 fading. Grow the sunflower seed by seed, then switch to primes only. Show both diffraction patterns, and point to the missing peaks at 1/4 and 1/9. End with Fig. 6.4 adding waves one by one, and say clearly what Montgomery and Dyson found and what is still unknown. Instruments: "Prime music," "Pi's sunflower," "Sieves & periods," "Quasicrystals," "Zeta zeros & nuclei." About 11 minutes.

CHAPTER 7

Tools that hunt primes

The continued fraction found 355/113. Can the same machine find primes? This chapter tries three tools from the companion site on the primes: a factoring method that is built from continued fractions, a picture of Fermat's little theorem made of rhythms, and an experiment that reads rhythms as binary numbers and tests them for primality. The first two are classical; the third is our experiment, and its result is a clear null.

7.1 Factoring with the continued fraction of $\sqrt{N}$

To factor an odd number N it is enough to find x and y with $x^2 \equiv y^2 \pmod N$ and $x \not\equiv \pm y$, because then $\gcd(x-y, N)$ is a proper factor. Continued fractions produce many congruences of the form $x^2 \equiv$ (small number) $\pmod N$. Expand $\sqrt{N}$ (or $\sqrt{kN}$ for a small multiplier k) with the integer recurrence

$$P_{n+1} = a_n Q_n - P_n, \qquad Q_{n+1} = \frac{kN - P_{n+1}^2}{Q_n}, \qquad a_{n+1} = \left\lfloor \frac{\lfloor \sqrt{kN} \rfloor + P_{n+1}}{Q_{n+1}} \right\rfloor, \tag{7.1}$$

starting from $P_0 = 0$, $Q_0 = 1$, $a_0 = \lfloor \sqrt{kN} \rfloor$. If A_n are the numerators of the convergents, reduced modulo N, then

$$A_{n-1}^2 \equiv (-1)^n Q_n \pmod N, \qquad 0 < Q_n < 2\sqrt{kN}. \tag{7.2}$$

These are standard facts about the continued fraction of a square root. Using them to factor, by searching the Q_n for squares, was described by D. H. Lehmer and R. E. Powers in 1931 [47]. In 1975 Michael Morrison and John Brillhart made it practical by combining many Q_n that factor over a small set of primes, using linear algebra modulo 2 to find a product that is a square, and with it factored the seventh Fermat number $F_7 = 2^{128} + 1$ [60].

Our script implements the simple Lehmer–Powers search. It is instructive to run it on $F_5 = 2^{32} + 1$, which Euler showed is divisible by 641 [36, §2.5]. Because $F_5 = 65536^2 + 1$, the continued fraction of $\sqrt{F_5}$ is $[65536; 131072, 131072, \dots]$ with every $Q_n = 1$, a square that gives only the trivial factorization. With multiplier $k = 2$ the search succeeds: at step $n = 340$ it finds a square Q_n, and the gcd gives $F_5 = 641 \times 6{,}700{,}417$. Table 7.1 shows the first steps.

Table 7.1. The first steps of the continued fraction of $\sqrt{2(2^{32}+1)}$ in the recurrence (7.1). By (7.2) each Q_n is, up to sign, a square modulo N; the search waits for one that is a perfect square at an even step.

n	a_n	Q_n	square?
1	1	166,833	
2	9	18,530	
3	289	639	
4	1	181,815	
5	62	2,930	
6	3	49,383	
7	1	109,871	
8	2	52,014	

A product of two six-digit primes, $999{,}983 \times 1{,}000{,}003$, falls after 2 steps, because two factors this close make N nearly a square.

This is the clearest case of "pi's machine doing prime work": the algorithm that, run on π, finds 22/7 and 355/113 is the algorithm that, run on $\sqrt{N}$, finds factors. It does not use any property of π, of course.

7.2 Mersenne necklaces and Fermat's little theorem

Take $m = 2^p - 1$ points on a circle and join each point x to $2x \bmod m$. Write x in binary with p digits. Since $2^p \equiv 1 \pmod{m}$, doubling modulo m simply rotates the digits by one place. So every orbit of the doubling map is a binary word of length p together with all its rotations: a *necklace*, or in our language a p-beat rhythm and all its rotations.

Proposition 7.1 (Fermat's little theorem for base 2, by necklaces)**.** *If p is prime, then p divides $2^p - 2$.*

Proof. Exclude the two constant words (all rests and all hits). Every other binary word of prime length p has exactly p distinct rotations, since a word equal to a nontrivial rotation of itself has a period dividing p, hence period 1. So the $2^p - 2$ nonconstant words split into orbits of size exactly p. □

The same argument with an alphabet of a letters gives $p \mid a^p - a$; this combinatorial proof is due to Golomb [30]. For $p = 7$ there are $(2^7 - 2)/7 = 18$ necklaces, each a 7-beat rhythm. For the composite length $p = 6$ the orbits have mixed sizes: 1 of length 2, 2 of length 3, 9 of length 6. The picture does *not* make $2^p - 1$ prime: $2^{11} - 1 = 2047 = 23 \times 89$. The test that does is the Lucas–Lehmer test, and with it the Great Internet Mersenne Prime Search found the largest known prime, $2^{136{,}279{,}841} - 1$, with 41,024,320 digits, in October 2024 [35].

7.3 The prime hunt

Here is an experiment from the companion site. Read the first n beats of a number's rhythm as a binary number, a hit as 1 and a rest as 0, and test it for primality as n grows. Does the rhythm of π attract primes?

Table 7.2. The prime hunt: probable primes among the binary prefixes of the first 1,500 beats of each rhythm. "Expected" is $\sum 2/\ln v_n$ over odd prefixes (the site's simpler $\sum 2/(n \ln 2)$ in parentheses). The last two columns are Poisson probabilities of a count at least as low, or at least as high, as the one found.

number	found	expected	$P(X \le \text{found})$	$P(X \ge \text{found})$
$\pi - 3$	1	2.4 (2.4)	0.30	0.91
$e - 2$	14	14.8 (14.3)	0.48	0.62
$\sqrt{2} - 1$	11	8.3 (8.8)	0.86	0.22
$\varphi - 1$	13	12.3 (11.8)	0.64	0.46

Method. For each of $\alpha = \pi - 3$, $e - 2$, $\sqrt{2} - 1$ and $\varphi - 1$ we built the first 1,500 beats with Theorem 3.5, formed the prefixes v_n (the binary number of the first n beats) for $n = 2, \ldots, 1500$, and tested every odd prefix with the Baillie–PSW probable-prime test (Appendix B). A prefix is odd exactly when beat n is a hit. For a random odd number of size v, the prime number theorem gives a probability of about $2/\ln v$ of being prime, so the expected count is $\sum 2/\ln v_n$ over the odd prefixes. The site uses the slightly cruder $\sum 2/(n \ln 2)$, which treats every prefix as an n-digit number; both are reported.

Results. Table 7.2 and Fig. 7.1 give the counts. The rhythm of $\pi - 3$ produced a single probable prime, at beat 49, where about 2.4 were expected; its rhythm is mostly rests, so its prefixes are odd only once every seven beats or so, and the expected count is small. The other three constants found close to their expected counts. In total the four rhythms gave 39 probable primes against 37.9 expected. No count is unusual: every Poisson tail probability in the table is above 0.2.

Interpretation. The rhythm of π does not attract primes, and neither does any other rhythm we tried. This is a null result, and we report it as one. The comparison is with Mersenne numbers $2^p - 1$, whose binary form is the all-hit rhythm: they *are* prime far more often than random numbers of their size, because every prime factor of $2^p - 1$ must be $\equiv 1 \pmod{2p}$, which removes all the small potential divisors. A construction that clearly beat its chance count in this test would deserve a closer look; none did. The same verdict, by a different route, comes from the statistics of π's terms in Chapter 10 and the prime convergents in Chapter 8.

For the video. Factor $2^{32} + 1$ live: show the Q_n scrolling past until a square appears and 641 drops out. Draw the doubling circle for $p = 7$ and let each necklace play as a rhythm. Run the prime hunt for π and watch it find one prime, then for e and watch it keep pace with chance. Say plainly: no attraction. Instruments: "Factoring," "Mersenne circles," "Rhythm → number," "Prime hunt." About 10 minutes.

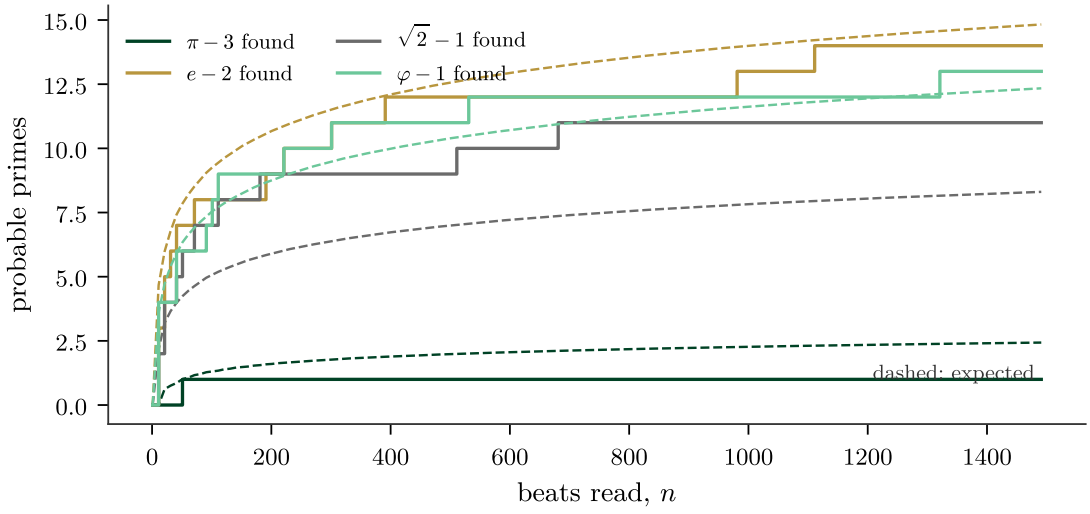


Figure 7.1. The prime hunt: cumulative probable primes among binary prefixes of each rhythm (steps) against the expected count $\sum 2/\ln v_n$ (dashed), for the first 1,500 beats.

CHAPTER 8

Thread (a): prime convergents

8.1 The question

The three classical fractions for π, 22/7, 223/71 and 355/113, all have prime denominators, and two of them are convergents. Among the convergents p_n/q_n of π, the prime denominators begin

$$q_1 = 7, \quad q_3 = 113, \quad q_8 = 265{,}381, \quad q_{32} = 842{,}468{,}587{,}426{,}513{,}207,$$

a list recorded in the OEIS [64]; the prime numerators begin $3, 103993, 833719, 4272943, \ldots$ [63]. Does the list go on forever? Nobody knows. The same question for the golden ratio asks whether there are infinitely many Fibonacci primes, and for $\sqrt{2}$ whether there are infinitely many Pell primes; both are famous open problems.

We cannot answer the question. What we can do is the next best thing: derive a heuristic prediction, with an explicit constant, for how many prime denominators a *typical* real number should have among its first N convergents; test the prediction on random numbers; and see where π, e, $\sqrt{2}$, φ and $\log_2(3/2)$ fall. This chapter does all three.

8.2 A heuristic with an explicit constant

Step 1: size. By Lévy's theorem (Theorem 2.7), for almost every x the denominators grow like

$$\ln q_n \sim Ln, \qquad L = \frac{\pi^2}{12 \ln 2} = 1.18657\ldots \tag{8.1}$$

Step 2: the naive guess. By the prime number theorem, a "random" integer near q is prime with probability about $1/\ln q$. If the q_n behaved like random integers of their size, the expected number of prime q_n with $n \le N$ would be about $\sum_{n \le N} 1/(Ln) \approx (1/L) \ln N$.

Step 3: local corrections. The q_n are not random integers. Consecutive denominators are coprime by (2.5), so, for example, two consecutive q_n are never both even. The right description of their local behavior is the following equidistribution theorem.

Theorem 8.1 (Szüsz; Moeckel; Jager and Liardet)**.** *Let $m \ge 2$. For almost every x, the pairs (q_{n-1}, q_n) reduced modulo m are equidistributed among the pairs (u, v) of residues with* $\gcd(u, v, m) = 1$.

Table 8.1. The share of convergent denominators q_n divisible by a prime ℓ: for π ($n \le 10{,}000$), for 200 random reals ($n \le 2{,}000$ each), the prediction $1/(\ell+1)$ of Corollary 8.2, and $1/\ell$ for random integers.

ℓ	π	random reals	$1/(\ell+1)$	$1/\ell$
2	0.3344	0.3331	0.3333	0.5000
3	0.2510	0.2499	0.2500	0.3333
5	0.1620	0.1666	0.1667	0.2000
7	0.1284	0.1243	0.1250	0.1429
11	0.0820	0.0833	0.0833	0.0909
13	0.0680	0.0708	0.0714	0.0769

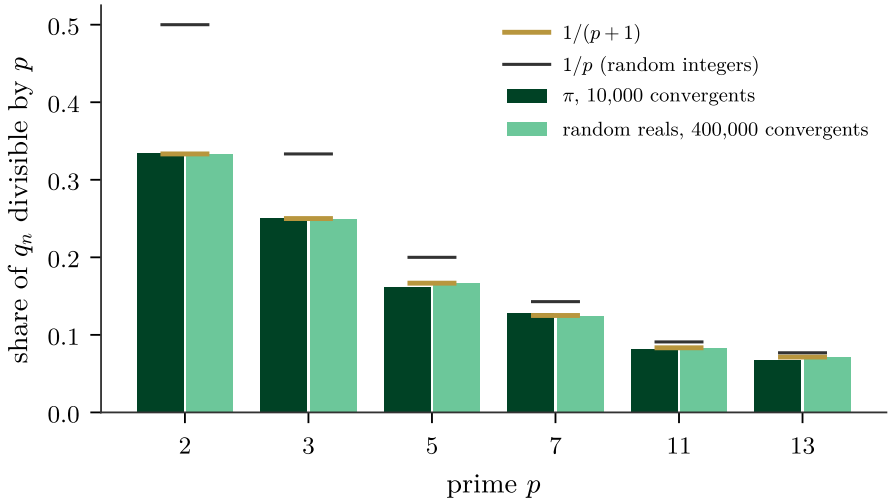


Figure 8.1. Convergent denominators avoid small prime factors: the share of q_n divisible by ℓ follows $1/(\ell+1)$ (gold), not $1/\ell$ (black), both for π and for random reals.

Szüsz proved the statement for pairs of denominators in 1962 [76, Satz 3.2]; Moeckel obtained it up to sign through the geodesic flow on modular surfaces [57]; Jager and Liardet proved equidistribution of the whole matrix $\left(\begin{smallmatrix} p_n & p_{n-1} \\ q_n & q_{n-1} \end{smallmatrix}\right)$ modulo m [37]. Borda's recent paper surveys this history and extends it [11]. The case $m = 2$ says that, for almost every x, exactly one third of the denominators are even.

Corollary 8.2. *For almost every x and every prime ℓ, the proportion of n with $\ell \mid q_n$ is $1/(\ell+1)$, not the $1/\ell$ of random integers.*

Proof. Modulo a prime ℓ, the admissible pairs are the $\ell^2 - 1$ nonzero pairs (u, v). Those with $v = 0$ are the $\ell - 1$ pairs $(u, 0)$, $u \neq 0$. By Theorem 8.1 the proportion is $(\ell-1)/(\ell^2-1) = 1/(\ell+1)$. □

Table 8.1 and Fig. 8.1 check the corollary on the first 10,000 convergents of π and on 400,000 convergents of 200 random reals. Both match $1/(\ell+1)$ to within sampling error; π is typical here too.

Step 4: the correction factor is $\zeta(2)$. In the usual way of refining the prime number theorem for structured sequences, the probability that q_n is prime is $1/\ln q_n$ multiplied, for

each prime ℓ, by the ratio of the chance that $\ell \nmid q_n$ to the chance that ℓ does not divide a random integer:

$$\mathfrak{S} = \prod_{\ell} \frac{1 - \frac{1}{\ell+1}}{1 - \frac{1}{\ell}} = \prod_{\ell} \frac{\ell}{\ell+1} \cdot \frac{\ell}{\ell-1} = \prod_{\ell} \frac{1}{1-\ell^{-2}} = \zeta(2) = \frac{\pi^2}{6}. \tag{8.2}$$

The product converges absolutely, so the correction is a single constant. It is a pleasant surprise that the constant is $\pi^2/6$, the same number as in Chapter 5, and for a related reason: the correction measures how much more often than random a denominator avoids each prime, and the reason it does is the coprimality of neighbors.

Step 5: the count. Putting the steps together, for a typical x,

$$\mathbb{E}\,\#\{n \le N : q_n \text{ prime}\} \approx \sum_{n \le N} \frac{\zeta(2)}{\ln q_n} \approx \frac{\zeta(2)}{L} \ln N = \frac{\pi^2}{6} \cdot \frac{12 \ln 2}{\pi^2} \ln N = 2 \ln 2 \cdot \ln N. \tag{8.3}$$

Heuristic 8.3. For almost every real x, the number of prime denominators among the first N convergents grows like $c \ln N$ with

$$c = 2 \ln 2 = 1.3863\ldots$$

The same holds for prime numerators, because Theorem 8.1 in the matrix form of Jager and Liardet gives the numerators the same local behavior.

Three remarks. First, the constant is independent of everything except the two laws used, which is why it is so clean: $\zeta(2)$ from the arithmetic and $\pi^2/(12 \ln 2)$ from the dynamics, and the two π^2 cancel. Second, a count growing like $\ln N$ is very slow: it predicts infinitely many prime denominators for a typical number, but only about one new one each time N is multiplied by $e^{1/c} \approx 2.06$, roughly one per doubling. Third, this is a heuristic. It treats primality of different q_n as independent events with the stated probabilities. It says nothing about any particular number such as π unless π is typical, which is unproved.

8.3 The test on random reals

To test Heuristic 8.3 we need many truly random real numbers with known continued fractions. We draw a random integer r with $D = 1{,}620$ digits and expand $r/10^D$ by exact integer arithmetic. Every real number in the interval $[r/10^D, (r+1)/10^D)$ shares the first terms, and since the error of the nth convergent shrinks like $q_n^{-2} \approx e^{-2Ln}$, about $D \ln 10/(2L) \approx 0.97D$ terms are shared, far more than the 1,000 we use. So the first 1,000 terms are those of a uniformly random real. We drew 400 such reals and counted probable primes among $q_1, \ldots, q_{1000}$ and $p_1, \ldots, p_{1000}$ (Appendix B).

Results. The mean number of prime denominators was 9.90 (standard deviation 2.87), and of prime numerators 9.68 (2.99). The growth is the test of the constant: between $n = 100$ and $n = 1000$ the mean count rose by 3.29, against the predicted $2 \ln 2 \cdot \ln 10 = 3.19$; between $n = 200$ and $n = 1000$ it rose by 2.22 against 2.23. A least-squares fit of the mean count against $\ln n$ over $50 \le n \le 1000$ gives slope 1.45 ± 0.04, slightly above $2 \ln 2$ because the earliest convergents, where q_n is small, still carry some weight. The absolute level has an additive constant that the heuristic does not fix: the pathwise sum $\sum \zeta(2)/\ln q_n$ averages 10.78 at $N = 1000$, a little above the observed mean, because $1/\ln q$ is a poor guide for the first few, very small denominators. The heuristic's *slope* is what matters, and it fits.

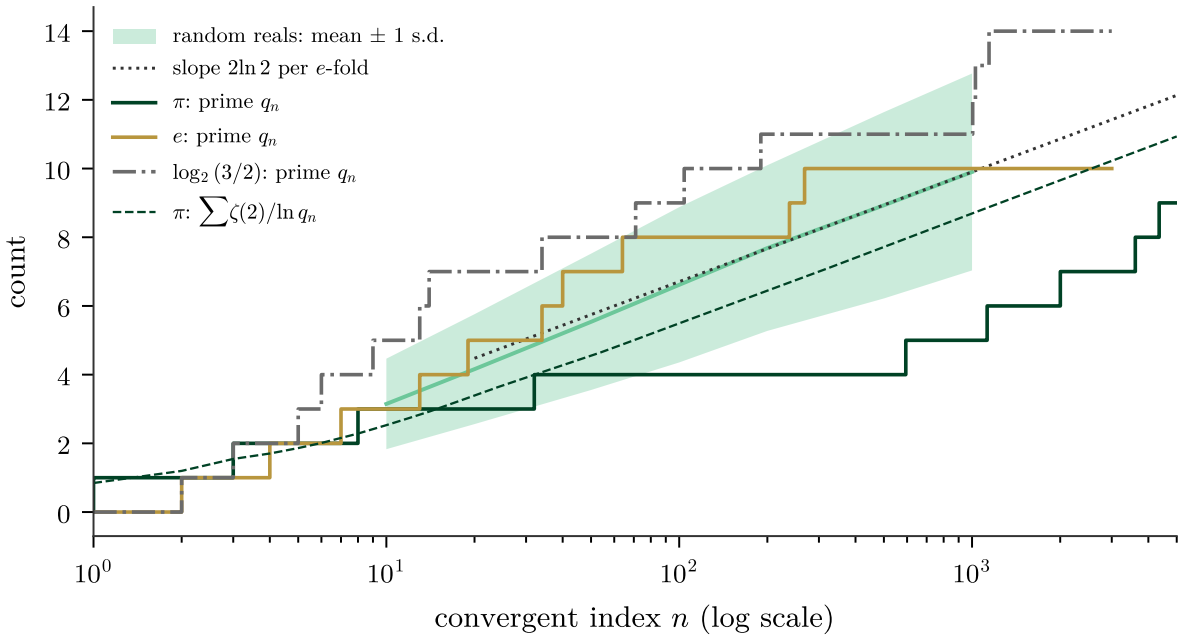


Figure 8.2. Prime convergent denominators, counted cumulatively against the index n (log scale). The shaded band is the mean ± one standard deviation over 400 random reals; the dotted line has the predicted slope $2\ln 2$ per e-fold. Steps: π (to $n = 5{,}000$), e and $\log_2(3/2)$ (to 3,000); dashed: π's own expected count $\sum\zeta(2)/\ln q_n$.

Table 8.2. Probable primes among convergents. Columns: the number of convergents N examined; prime denominators q_n, $n \le N$, and the heuristic expectation $\sum_{n\le N}\zeta(2)/\ln q_n$; prime numerators; the counts of prime q_n and p_n for $n \le 1000$ with the expectation; and the percentile of the q_n count at $n \le 1000$ among 400 random reals (share of random reals with at most as many). For $\sqrt{2}$ and φ the typical-number heuristic does not apply.

		$n \le N$			$n \le 1000$			
number	N	q_n	exp.	p_n	q_n	p_n	exp.	pct.
π	5,000	9	10.9	13	5	10	8.7	6
e	3,000	10	12.4	10	10	10	11.6	61
$\log_2(3/2)$	3,000	14	13.9	11	11	10	12.4	74
$\sqrt{2}$	3,000	23	16.8	18	17	16	14.7	–
φ	3,000	22	27.7	22	21	21	23.9	–

8.4 Is pi typical?

Table 8.2 gives the counts for the five constants of the Suite, and Figs. 8.2 and 8.3 compare them with the random sample.

Pi. Among its first 5,000 convergents, π has 9 prime denominators, at $n = 1$, 3, 8, 32, 594, 1126, 2002, 3610, 4355, against 10.9 expected; the last has 2,239 digits. It has 13 prime numerators, at $n = 4$, 8, 10, 15, 86, 229, 333, 593, 839, 852, 1148, 2725, 3787 (the largest has 1,949 digits). The first four denominators and the first numerators agree with the OEIS

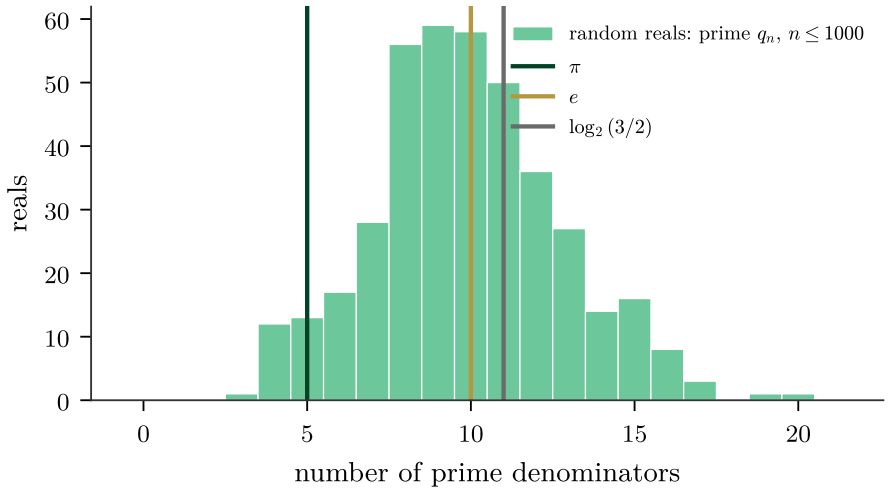


Figure 8.3. The number of prime denominators among the first 1,000 convergents, for 400 random reals (histogram), with π, e and $\log_2(3/2)$ marked.

lists [63, 64]. At $n \le 1000$, π's count of 5 prime denominators sits at the 6th percentile of the random sample: on the low side, but well inside the range of ordinary numbers, and its count of prime numerators is about average. Its even denominators make up 0.3330 of the first 10,000 (prediction 1/3) and those divisible by 3 make up 0.2484 (prediction 1/4). By every measure here, π looks like a typical real number. Heuristic 8.3 predicts that it has infinitely many prime denominators, with the next few spread out roughly one per doubling of n.

The fifth and e. The number $\log_2(3/2)$ has 14 prime denominators up to $n = 3{,}000$ against 13.9 expected; e has 10 against 12.4. Both are unremarkable, although e is not a typical number (its terms grow, and its denominators are even exactly one third of the time because of its pattern, not because of Theorem 8.1).

Quadratic irrationals. For φ, q_n is the Fibonacci number F_{n+1}; for $\sqrt{2}$ it is the Pell number P_{n+1}. Our counts recover the known lists of indices of Fibonacci and Pell primes [62, 65] up to index 3,001: the Fibonacci indices are 3, 4, 5, 7, 11, 13, 17, 23, 29, 43, 47, 83, 131, 137, 359, 431, 433, 449, 509, 569, 571, 2971, and the Pell indices are 2, 3, 5, 11, 13, 29, 41, 53, 59, 89, 97, 101, 167, 181, 191, 523, 929, 1217, 1301, 1361, 2087, 2273, 2393. These sequences break the typical-number heuristic in a structured way. They are divisibility sequences ($F_m \mid F_n$ whenever $m \mid n$), so F_n can be prime only when n is prime (or $n = 4$), and their denominators are not equidistributed modulo m: half of the Pell denominators are even, not a third. Grantham and Granville [34] develop heuristics for primes in such linear recurrences; for Fibonacci numbers they predict about $e^{\gamma} \log_{\varphi} N$ primes F_n with $n \le N$, which is 29.6 for $N = 3{,}001$, against the 22 found (they note that the Fibonacci data run below the prediction). Whether either list is infinite is open.

8.5 Do the site's tools help?

The owner of the companion site asked whether the tools assembled there help with this question. Here is a frank assessment.

- *Rhythms and Sturmian words.* The denominators are block lengths (Theorem 3.5), which makes the question easy to state musically: does π's form contain infinitely many sections whose length in bars is prime? But the rhythm adds no information about primality beyond the recurrence.
- *Coprimality and $6/\pi^2$.* This is the tool that does real work: the coprimality of neighboring blocks is exactly what produces the correction factor $\zeta(2)$ in (8.2).
- *Necklaces, diffraction, the prime hunt.* These do not bear on the question. The prime hunt is a different experiment with the same verdict: rhythms of π neither attract nor repel primes.
- *Self-similarity.* For quadratic irrationals, exact self-similarity means a linear recurrence, and the recurrence's divisibility structure is exactly what governs Fibonacci and Pell primes. For typical numbers, self-similarity in rule but not in shape is what makes the random model plausible.

In short, the tools organize the heuristic and explain its constant. They do not bring a proof closer. A proof that π has infinitely many prime denominators would require, at the least, control of π's terms that nobody has: it is not even known that they are unbounded.

8.6 What is known, and what is not

Known: the equidistribution of convergents modulo m for almost every x (Theorem 8.1); Lévy's theorem; the published lists of prime numerators and denominators for π and of Fibonacci and Pell primes; the Grantham–Granville heuristics for linear recurrences. Open: whether π, e, φ or $\sqrt{2}$ has infinitely many prime convergent denominators. We did not find in the literature we searched the explicit heuristic constant $2\ln 2$ for typical reals, or a test of it on a large random sample; both are offered here as a heuristic and an experiment, not a theorem.

For the video. State the question with π's four prime denominators on screen. Build the heuristic in five steps, ending with the two π^2's cancelling to leave $2\ln 2$. Show Fig. 8.2 and the histogram with π marked. Close with the frank assessment: the tools explain the constant but do not prove anything. Instrument: "The ladder," with the bar switched between π, φ and $\sqrt{2}$. About 12 minutes.

CHAPTER 9

Thread (b): one spectrum, from crystal to randomness

9.1 The question

Chapter 6 ended with the Montgomery–Dyson observation: the zeros of the zeta function seem to be spaced like the eigenvalues of large random Hermitian matrices, the model physicists use for energy levels of heavy nuclei. The companion site lists this as an open thread and asks whether the tools assembled here help. They cannot prove anything about the zeros. What they can do is supply a scale on which to place the zeros: at one end a perfect crystal, next the rhythm of π (a quasicrystal), then the zeros and random matrices, and at the other end pure randomness. This chapter computes the statistics that define that scale, for every point set on it, with one set of tools.

9.2 The point sets

All point sets are *unfolded*, that is, rescaled so that the mean spacing is 1, before any statistic is computed.

- *Crystal:* the integers.
- *Quasicrystal:* the hits of π's rhythm in the first 33,102 beats (Theorem 3.5), multiplied by $\pi - 3$. The spacings take two values, $7(\pi - 3) = 0.9911$ and $8(\pi - 3) = 1.1327$. The three-gap points $\{n(\pi-3)\}$ of Theorem 3.7 belong here too; for $N = 1{,}000$ their spacings, in units of $1/N$, take the values $0.030, 8.610, 8.640$.
- *Zeta zeros:* the imaginary parts $\gamma_1 < \gamma_2 < \cdots$ of the first 5,000 nontrivial zeros, up to height 5447.86, computed with mpmath and checked against the zero-counting function. They are unfolded with the smooth part of the Riemann–von Mangoldt counting formula, $\bar{N}(T) = \frac{T}{2\pi}\ln\frac{T}{2\pi e} + \frac{7}{8}$; the unfolded spacings have mean 1.0000. Odlyzko's tables [61] reach far higher; we used only zeros we computed ourselves.
- *GUE:* the eigenvalues of 16 random 1000×1000 complex Hermitian matrices with independent Gaussian entries, from the central half of each spectrum, unfolded with the semicircle law [55]; mean spacing 0.9999.

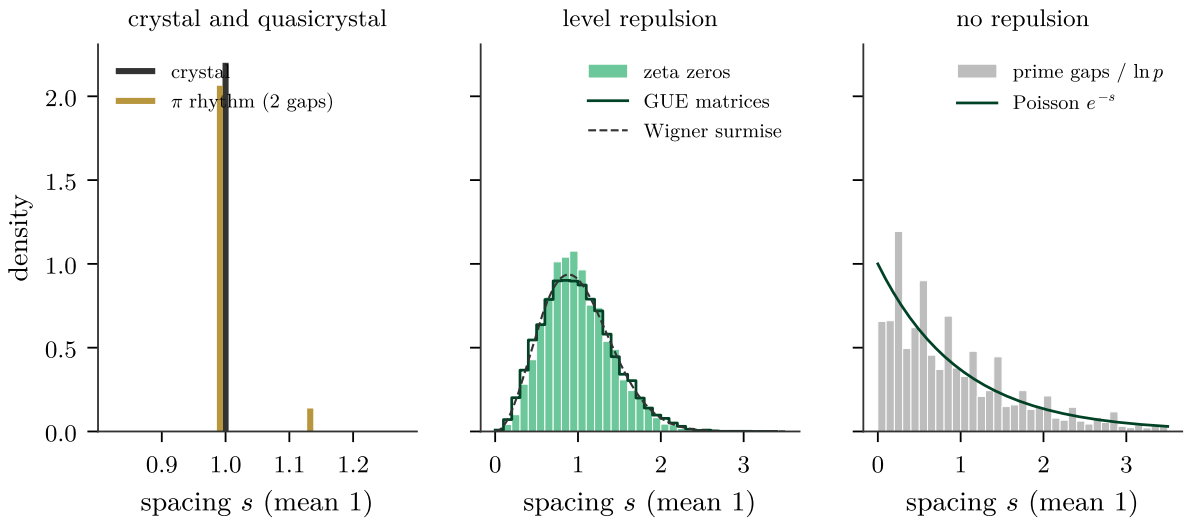


Figure 9.1. Nearest-neighbor spacing distributions, all with mean spacing 1. Left: a crystal (one spacing) and π's rhythm (two spacings, heights proportional to their frequencies; note the zoomed axis). Center: 5,000 zeta zeros (bars) and GUE eigenvalues (steps), with the Wigner surmise. Right: prime gaps divided by $\ln p$ near 10^9, with the Poisson law e^{-s}.

- *Primes:* the gaps between the 14,491 primes in $[10^9, 10^9 + 3 \cdot 10^5)$, divided by $\ln p$ and normalized to mean 1.
- *Poisson:* independent exponential spacings, the model of points with no correlation at all.

9.3 Nearest-neighbor spacings

Figure 9.1 shows the distribution of the spacings s between neighbors. The crystal has a single spacing and the quasicrystal two. The zeros and the GUE eigenvalues share a distribution that vanishes like s^2 at $s = 0$: close pairs are rare, a phenomenon called *level repulsion*. The curve drawn with them is the Wigner surmise $P(s) = \frac{32}{\pi^2} s^2 e^{-4s^2/\pi}$, a simple and very accurate approximation to the exact GUE spacing law [55]. The prime gaps, by contrast, look like the Poisson law e^{-s}, with the most likely spacing near zero.

A simple number captures the repulsion: the share of spacings below $1/4$. It is 0.009 for the zeros, 0.016 for the GUE matrices and 0.016 for the Wigner surmise, but 0.212 for the Poisson process and 0.129 for the normalized prime gaps (whose discreteness, all gaps being even, distorts the smallest bins). The Kolmogorov–Smirnov distance between the zero spacings and the GUE spacings is 0.036, against 0.305 between zeros and Poisson. The small remaining difference between zeros and GUE is expected: at the low heights we can reach, the zeros behave like random matrices of a modest effective size, a finite-height effect analyzed by Bogomolny, Bohigas, Leboeuf and Monastra [10]. Correspondingly, the variance of the zero spacings, 0.152, is a little below the GUE value 0.183.

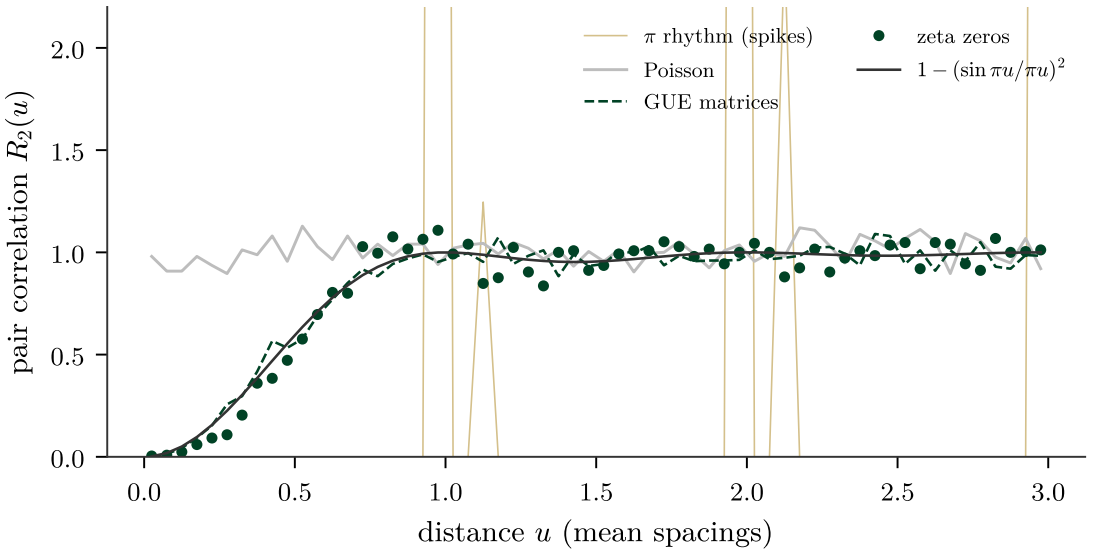


Figure 9.2. Pair correlation $R_2(u)$: 5,000 zeta zeros (dots), GUE eigenvalues (dashed), Poisson points (grey), π's rhythm (gold spikes), and Montgomery's function (9.1) (black).

9.4 Pair correlation

The pair correlation $R_2(u)$ measures how likely it is to find another point at distance u from a given point, relative to a uniform density. Montgomery's conjecture [59] is that for the unfolded zeros

$$R_2(u) = 1 - \left(\frac{\sin \pi u}{\pi u}\right)^2, \tag{9.1}$$

which is exactly the GUE pair correlation in the limit of large matrices [25, 55]. Figure 9.2 shows our estimate from the first 5,000 zeros, the GUE sample, a Poisson sample (flat at 1) and π's rhythm (sharp spikes at the sums of its two spacings). The root-mean-square difference between the zeros' histogram and (9.1) over $0 \le u \le 3$ is 0.063, comparable to the sampling noise of the histogram.

9.5 Number variance: a scale of rigidity

The most telling statistic is the *number variance* $\Sigma^2(L)$: drop a window of length L (in mean spacings) at a random place and count the points inside; $\Sigma^2(L)$ is the variance of that count. It measures how rigid the arrangement is over a range of L spacings.

- For a Poisson process, $\Sigma^2(L) = L$: the counts fluctuate as much as they can.
- For the GUE, $\Sigma^2(L) \approx \frac{1}{\pi^2}\left(\ln(2\pi L) + \gamma + 1\right)$ for large L, with γ Euler's constant [55]: the fluctuations grow only logarithmically.
- For a crystal, $\Sigma^2(L) = \{L\}(1-\{L\}) \le \frac{1}{4}$, bounded. Point sets with bounded number variance are called *hyperuniform* [77]; crystals and many quasicrystals, including Sturmian ones like π's rhythm, are of this kind.

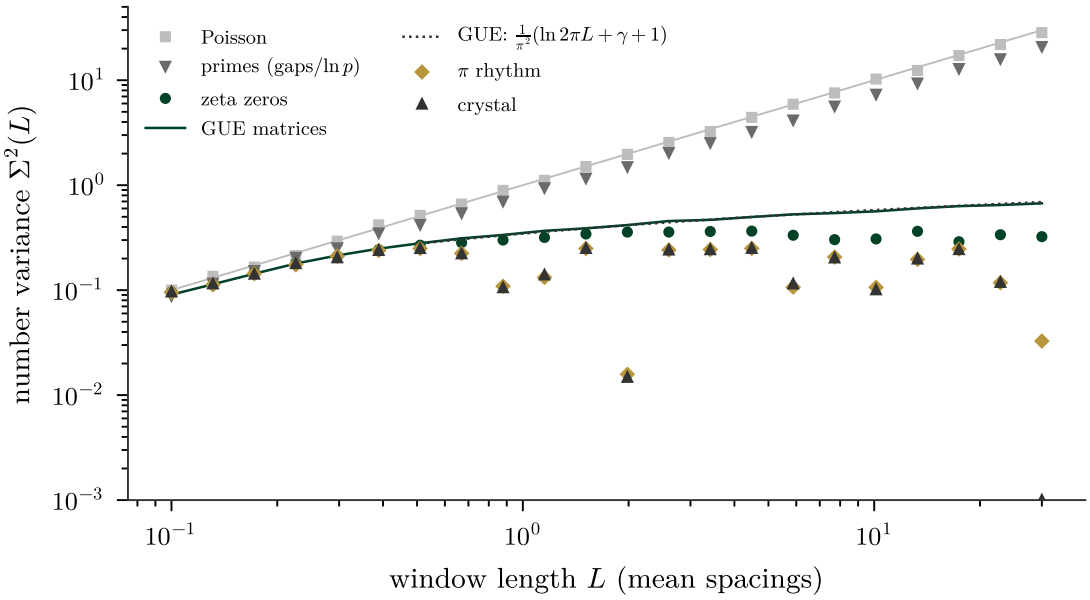


Figure 9.3. Number variance $\Sigma^2(L)$ on log–log axes for six point sets. From bottom to top: crystal and π's rhythm (bounded, with dips where L is close to a period), zeta zeros and GUE (logarithmic growth; dotted: the asymptotic GUE formula), normalized prime gaps and Poisson (linear growth).

For the zeros, Berry predicted that $\Sigma^2(L)$ follows the GUE law for small L and then saturates at a level set by the height of the zeros [8]. Figure 9.3 shows all six point sets on one logarithmic plot. At $L = 10.1$ we find $\Sigma^2 = 0.10$ for the crystal, 0.11 for π's rhythm, 0.31 for the zeros, 0.56 for the GUE sample (asymptotic formula 0.58), 7.36 for the normalized prime gaps and 10.26 for Poisson. The zeros lie below the GUE curve at this L, consistent with Berry's saturation at our low heights, while the GUE sample matches its formula.

9.6 The spectrum, and Dyson's bridge

The three statistics put the point sets in one order of decreasing rigidity:

$$\text{crystal} \;\to\; \text{quasicrystal } (\pi\text{'s rhythm}) \;\to\; \text{GUE and zeta zeros} \;\to\; \text{Poisson (and the primes at short range).}$$

The crystal has one spacing and bounded fluctuations. The rhythm of π has two spacings, at most three on the circle, and still bounded fluctuations: it is ordered without being periodic. The zeros have a continuous spacing law but repel each other and fluctuate only logarithmically. Poisson points do not interact at all, and the primes, seen through their gaps at short range, look Poisson-like, although Chapter 6 showed their long-range diffraction is far from random.

Where do the zeros belong, then? Freeman Dyson, in his lecture "Birds and Frogs," observed that if the Riemann hypothesis is true then the zeros form a one-dimensional quasicrystal in a precise sense, and suggested that classifying one-dimensional quasicrystals might lead to a proof [24]. His sense is about the Fourier transform of the zeros, not their spacings, and the explicit formula (6.2) explains it: the sum of the waves $\cos(\gamma t)$ over the zeros is concentrated at $t = \ln p^m$, the logarithms of prime powers, just as the diffraction of a quasicrystal is concentrated on a discrete set. Figure 9.4 shows this with our 5,000 zeros: the sum $\sum_\gamma \cos(\gamma t)$,

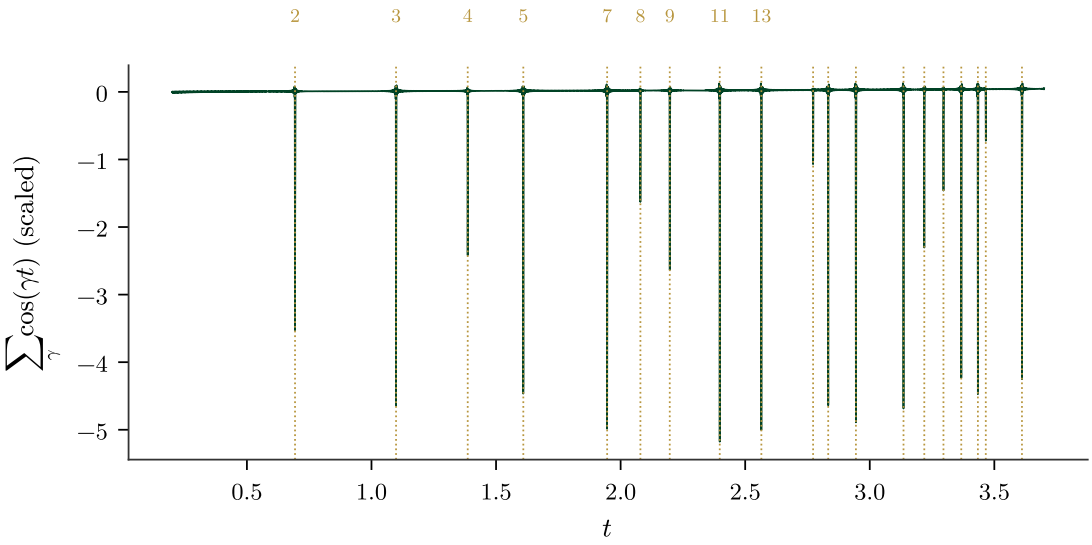


Figure 9.4. The zeros "diffract" at logarithms of prime powers: $\sum_\gamma \cos(\gamma t)$ over the first 5,000 zeros with a smooth cutoff, scaled by $1/\sqrt{\text{number of zeros}}$. Dotted gold lines mark $t = \ln n$ for the prime powers $n \le 40$; the dips sit exactly there, as the explicit formula predicts.

smoothly cut off, has sharp dips at $\ln 2$, $\ln 3$, $\ln 4$, $\ln 5$, $\ln 7$, … (depths -3.5 and -4.6 at $\ln 2$ and $\ln 3$ on the plotted scale, against a background fluctuation of about 0.01). So the zeros are GUE-like at short range and quasicrystal-like in their Fourier transform. The rhythm of π is quasicrystal-like in both. That is the bridge Dyson pointed to, and this picture organizes it; it does not cross it.

9.7 What this does and does not show

Everything in this chapter is a measurement of known phenomena with our own data, arranged so that a reader can see a crystal, π's rhythm, the zeros, random matrices and the primes on one scale. It illuminates the Montgomery–Dyson connection and makes Dyson's remark concrete. It proves nothing about the Riemann hypothesis, about the pair-correlation conjecture, or about any operator whose eigenvalues might be the zeros. Of the site's tools, the ones that help are the rhythm (as the quasicrystal end of the scale) and the diffraction picture (as the language for Dyson's remark); the others play no role.

For the video. Line up six strips of dots, from evenly spaced to random, and let the viewer guess which is which. Then show the three statistics, each with all six sets. End on Fig. 9.4: the zeros' waves spike at $\ln 2$, $\ln 3$, $\ln 5$; this is Dyson's quasicrystal remark made visible. Say clearly that this organizes the mystery and does not solve it. Instruments: "Zeta zeros & nuclei," "Quasicrystals." About 12 minutes.

CHAPTER 10

Self-similarity and fractals

10.1 Teddy Rook's picture

Teddy Rook pictured π as "a number, then a slightly smaller number after it, at least a decimal place over, then a smaller one after that, and so on." Does that idea fit with fractals? It does, and the precise version is one line long. By (2.1),

$$x = a_0 + \frac{1}{x_1}, \qquad x_1 = [a_1; a_2, a_3, \dots].$$

The "smaller number after it," the tail $1/x_1$, is built from a continued fraction made by exactly the same rule; so is its tail, and so on forever. An object that contains a copy of itself, made by the same rule, at every scale is *self-similar*, and self-similarity is the defining feature of a fractal. And each step really is smaller: by (2.7) the nth convergent misses by less than $1/q_n^2$. For π the complete quotients are $x_1, \dots, x_6 = 7.0625, 15.9966, 1.0034, 292.6346, 1.5758, 1.7367$ (Fig. 10.1).

Exact self-similarity, and self-similarity of rule. When the terms repeat, the tail is literally the number again: $\varphi = 1 + 1/\varphi$, so every complete quotient of φ is φ, and every complete quotient of $\sqrt{2} = [1; 2, 2, \dots]$ after the first is $1 + \sqrt{2}$. Every picture we draw from such a number repeats exactly under zoom. For π the terms $7, 15, 1, 292, 1, 1, 1, 2, \dots$ follow no known pattern, so the *rule* is the same at every level but the *shape* is new each time. That is not a flaw in the picture; it is one of the most interesting facts about π.

A twist: patterned fractions for π. If the numerators of a continued fraction are allowed to be numbers other than 1, π does have beautiful patterns, and then Teddy Rook's "never-ending string of continued fractions" holds with a clean rule all the way down. Brouncker found

$$\frac{4}{\pi} = 1 + \cfrac{1^2}{2 + \cfrac{3^2}{2 + \cfrac{5^2}{2 + \cdots}}} \tag{10.1}$$

in 1655, at the request of John Wallis, who published it in his *Arithmetica Infinitorum*; Khrushchev has reconstructed Brouncker's proof [43]. Lange found $\pi = 3 + 1^2/(6 + 3^2/(6 +$

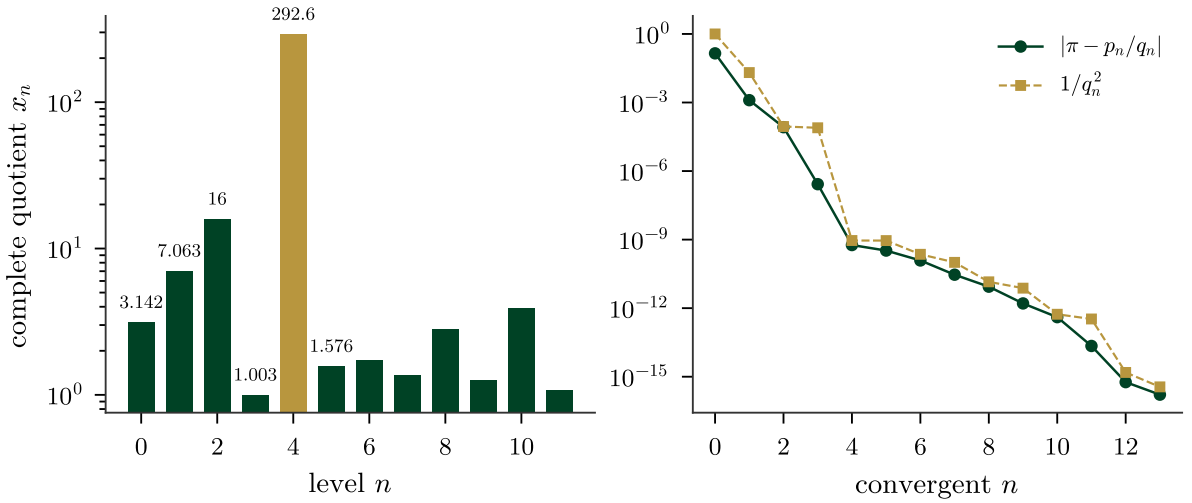


Figure 10.1. The nested tail of π. Left: the complete quotients x_n (log scale); each is the number "seen from" level n, and its whole part is the next term. The gold bar is $x_4 = 292.63$. Right: the error of each convergent (green) stays below $1/q_n^2$ (gold).

$5^2/(6+\cdots)))$ [46], and the continued fraction of the arctangent gives $\pi = 4/(1+1^2/(3+2^2/(5+3^2/(7+\cdots))))$. They converge at three very different speeds. After 10, 100 and 1000 levels Brouncker's fraction misses π by 0.0907, 0.0099 and 0.000999; its stopping points are exactly the partial sums of the Leibniz series $4(1-\frac{1}{3}+\frac{1}{5}-\cdots)$, as our script confirms with exact fractions. Lange's misses by 0.000186, $2.43e-7$ and $2.49e-10$ (its stopping points are the partial sums of Nilakantha's series). The arctangent fraction misses by 0.00013, $1.9e-8$ and $4.3e-16$ after 5, 10 and 20 levels; each level shrinks the error by a factor $0.1716 = (\sqrt{2}-1)^2$.

The rest of this chapter turns the rule "a number contains a copy of itself after every term" into pictures. Each is a classical fractal; for each we say who found it and draw it for π.

10.2 Ford circles

In 1938 Lester Ford drew, for every fraction p/q in lowest terms, a circle $C(p/q)$ tangent to the real axis at p/q with radius $1/(2q^2)$ [28].

Proposition 10.1 (Ford)**.** *Two Ford circles $C(p/q)$ and $C(r/s)$ never overlap. They are tangent if and only if $|ps-qr|=1$, that is, if p/q and r/s are neighbors in some Farey sequence; and between two tangent circles, the largest circle tangent to both is $C\big((p+r)/(q+s)\big)$.*

Proof of the first two statements. The squared distance between the centers minus the square of the sum of the radii is

$$\Big(\frac{p}{q}-\frac{r}{s}\Big)^2+\Big(\frac{1}{2q^2}-\frac{1}{2s^2}\Big)^2-\Big(\frac{1}{2q^2}+\frac{1}{2s^2}\Big)^2=\frac{(ps-qr)^2-1}{q^2s^2},$$

which is ≥ 0 for distinct fractions, with equality exactly when $|ps-qr|=1$. □

Zoom in anywhere and the same pattern of circles appears again; the picture is self-similar under the maps $x \mapsto (ax+b)/(cx+d)$ with integer entries and determinant ± 1. Ford circles turn Legendre's theorem into geometry. A vertical line through x passes through $C(p/q)$ exactly

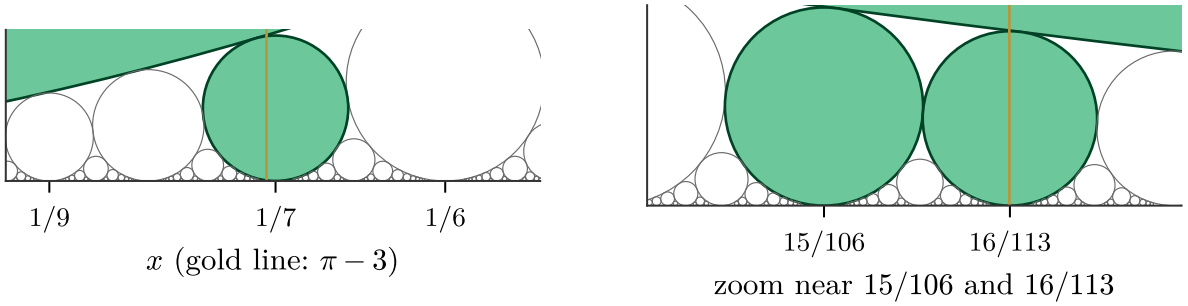


Figure 10.2. Ford circles near $\pi - 3$ (gold line), drawn to scale. Filled circles belong to convergents of $\pi - 3$. Left: the circle at $1/7$ with its neighbors. Right: a zoom by a factor of about 300, where the circles at $15/106$ and $16/113$ touch; the line passes through the second and misses the first.

when $|x - p/q| < 1/(2q^2)$, so by Theorem 2.4 every circle it pierces belongs to a convergent; and of any two consecutive convergents, at least one is pierced [41, Ch. II]. For $\pi - 3$ our script finds the pierced circles among the first convergents to be $0/1, 1/7, 16/113, 4703/33215, 14093/99532$, while $15/106$ and $4687/33102$ are convergents whose circles the line just misses (Fig. 10.2). Falling down the line toward $\pi - 3$, one drops through circle after circle, and after $16/113$ the fall is very long, because $355/113$ is so close: that is the 292 again.

10.3 The Stern–Brocot tree

Start with $0/1$ and $1/0$ and repeatedly insert the *mediant* $(p + r)/(q + s)$ between neighbors p/q and r/s. The result, described by Moritz Stern in 1858 [74] and independently by the clockmaker Achille Brocot in 1861 [12], who used it to design gear trains, is a binary tree that contains every positive fraction exactly once, in lowest terms [33, §4.5]. To locate a number, start at $1/1$ and go left (L) if the number is smaller and right (R) if larger. The path is a sequence of runs, and the run lengths are the terms of the continued fraction. For $355/113 = [3; 7, 16]$ the path is $R^3L^7R^{15}$ (the last run is one short because the path stops on the fraction itself), and for π our script finds $R^3L^7R^{15}LR^{292}LR\ldots$, with the 292 as a run of 292 Rs. The golden ratio zigzags $RLRL\ldots$ forever and $\sqrt{2}$ goes $R, LL, RR, LL, \ldots$. Every subtree is a copy of the whole tree with the fractions changed, which is self-similarity again (Fig. 10.3).

10.4 Minkowski's question-mark function

In 1904 Hermann Minkowski introduced a function that turns continued fractions into binary expansions [56]. For $x = [0; a_1, a_2, a_3, \ldots]$,

$$?(x) = 2\left(2^{-a_1} - 2^{-(a_1+a_2)} + 2^{-(a_1+a_2+a_3)} - \cdots\right), \tag{10.2}$$

which in binary is $a_1 - 1$ zeros, then a_2 ones, then a_3 zeros, and so on: each term becomes one run of binary digits. The function is continuous and strictly increasing from $?(0) = 0$ to $?(1) = 1$, maps rationals to dyadic rationals and quadratic irrationals to rationals, and yet its derivative is zero almost everywhere: it is a "singular" function, all of whose increase happens on a set of measure zero [21, 71]. It satisfies $?(x/(1 + x)) = ?(x)/2$, so the left half of

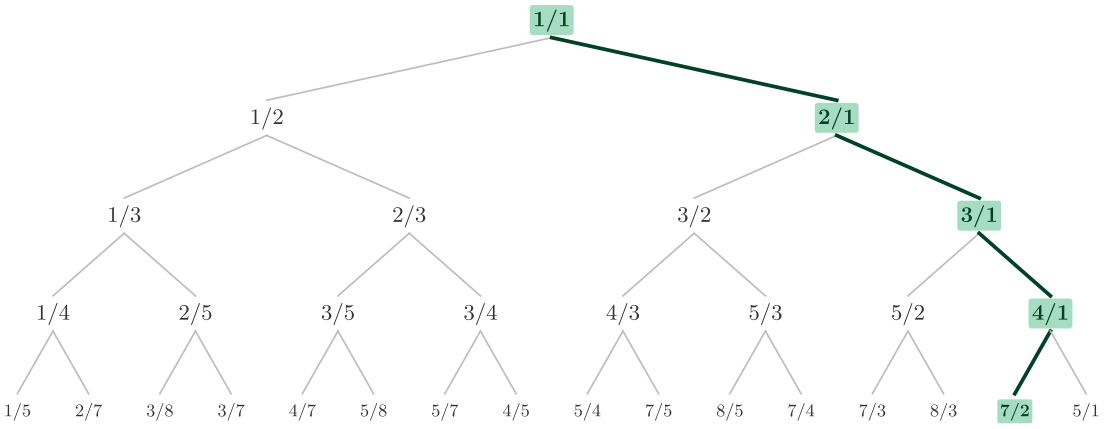


Figure 10.3. The first five levels of the Stern–Brocot tree, with the beginning $RRRL$ of the path to π highlighted: $1/1 \to 2/1 \to 3/1 \to 4/1 \to 7/2$. The path continues with six more Ls to $22/7$.

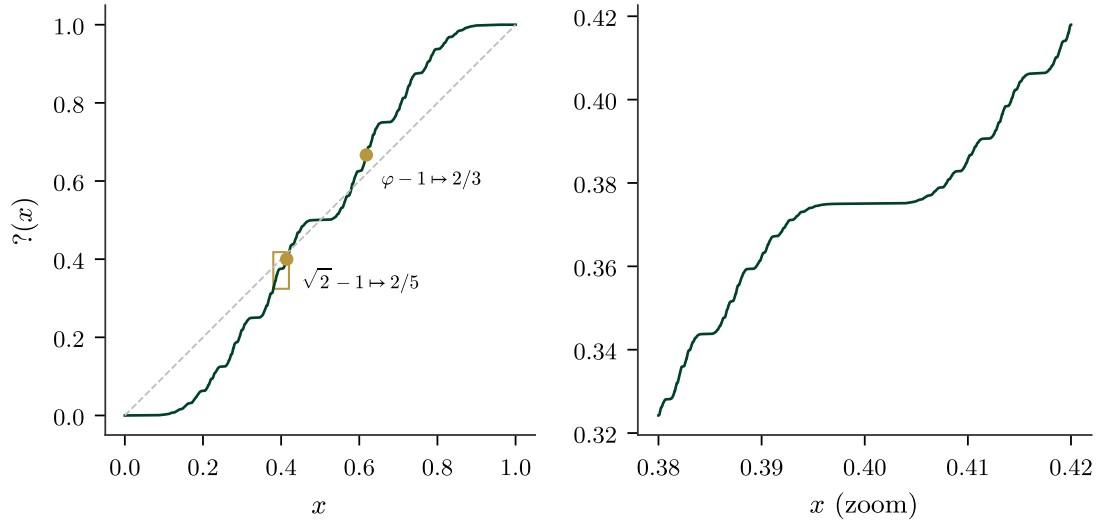


Figure 10.4. Minkowski's $?(x)$. Left: the whole graph with the diagonal dashed; quadratic irrationals land on fractions. Right: the boxed region magnified, showing flat stretches and bursts of growth at every scale.

its graph is the whole graph squashed by a factor 2. Our script confirms $?(\varphi - 1) = 2/3$ and $?(\sqrt{2} - 1) = 2/5$, and finds $?(\pi - 3) = 0.01562476\ldots$, just under $1/64$: six zeros, then fifteen ones (Fig. 10.4).

10.5 The devil's staircase

Return to the circle map (4.1) of Chapter 4, at full coupling $K = 1$. Plot the rotation number W against Ω: the result is a staircase with a flat step at every fraction (Fig. 10.5). At $K = 1$ the steps fill the whole interval except a Cantor set of measure zero, whose fractal dimension

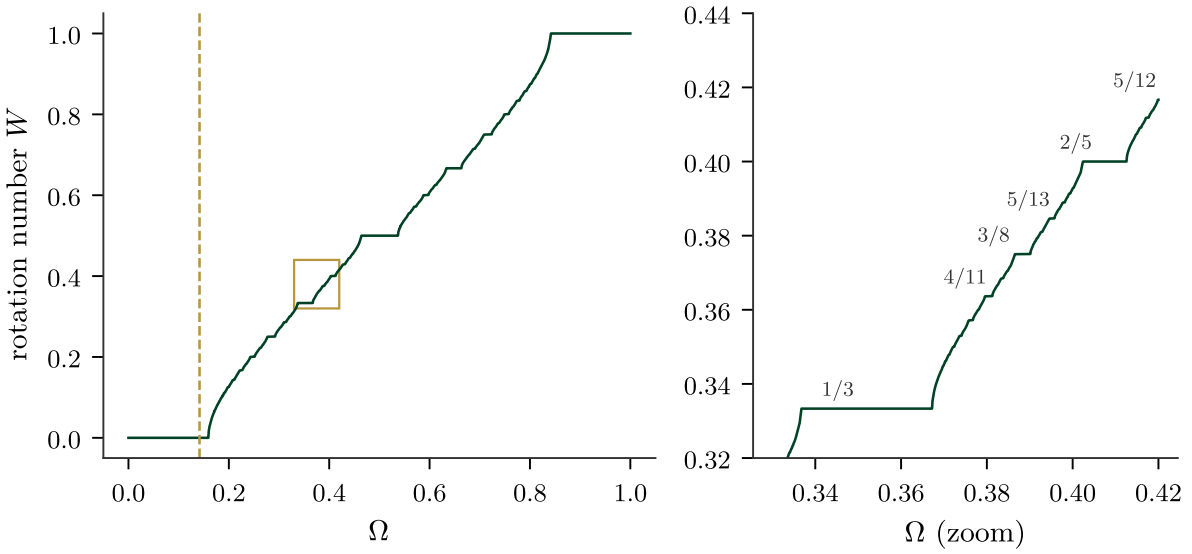


Figure 10.5. The devil's staircase: rotation number W of the sine circle map at $K = 1$ against Ω (2,520 iterations per point). The dashed line is $\Omega = \pi - 3$, inside the wide 0/1 step. Right: the boxed region, with steps labeled by their fractions; between 1/3 and 2/5 the widest step is 3/8, between 3/8 and 2/5 it is 5/13.

Jensen, Bak and Bohr measured as about 0.87 [40]. Between two adjacent steps p/q and r/s the widest step is the one at the mediant $(p+r)/(q+s)$, the Stern–Brocot rule once more, so zooming between two steps reproduces the whole staircase in miniature.

10.6 Cantor dust: numbers with small terms

Which numbers in $(0, 1)$ have continued fractions that use only the terms 1 and 2? Fixing the first term to be 1 or 2 keeps two intervals; fixing the second keeps two sub-intervals of each, and so on. What survives forever is a Cantor-like "dust" E_2 (Fig. 10.6). It has length zero, but more than a point's worth of size: its Hausdorff dimension is $\dim_H E_2 = 0.5312805062772051\ldots$. I. J. Good bounded it between 0.5306 and 0.5320 in 1941 [32]; Jenkinson and Pollicott computed it to 25 digits with a transfer-operator method [38] and later to 100 rigorously certified digits [39].

The method is elegant. The dimension of the set E_A of numbers whose terms all lie in a finite set A is the unique s for which the transfer operator

$$(\mathcal{L}_s f)(x) = \sum_{a \in A} (a + x)^{-2s} f\left(\frac{1}{a + x}\right)$$

has spectral radius 1. We discretize $\mathcal{L}_s$ on Chebyshev nodes and bisect on s (Appendix B). With 60 nodes our script gives $\dim_H E_{\{1,2\}} = 0.531280506277$, agreeing with the published value in all twelve printed digits. It also gives $\dim_H E_{\{1,2,3\}} = 0.70566$, $\dim_H E_{\{1,2,3,4\}} = 0.78895$, $\dim_H E_{\{1,\ldots,6\}} = 0.86762$ and $\dim_H E_{\{2,3\}} = 0.33744$. Since $\pi - 3$ starts with a 7 and contains a 292, it lies in none of these small-term dusts; the golden ratio's $\varphi - 1$ (all 1s) and $\sqrt{2} - 1$ (all 2s) are single points of $E_{\{1,2\}}$.

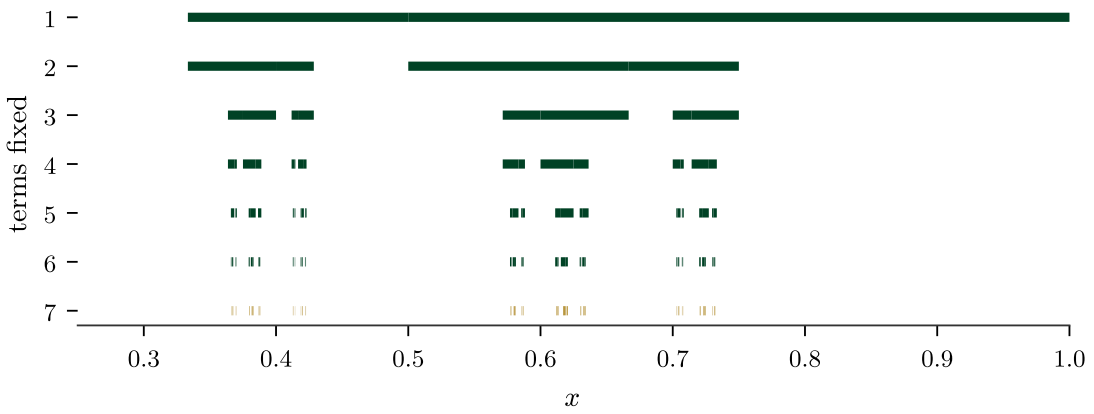


Figure 10.6. The numbers whose first k terms are all 1 or 2, for $k = 1, \ldots, 7$. Each row keeps two pieces of each piece of the row above; the limit is a Cantor set of dimension 0.5313.

10.7 The Mandelbrot set's Farey bulbs

The Mandelbrot set consists of the complex c for which $z \mapsto z^2 + c$, started at $z = 0$, stays bounded. Its heart-shaped main body is fringed with bulbs, and each bulb carries a fraction: the bulb attached at the point $c = e^{2\pi i p/q}/2 - e^{4\pi i p/q}/4$ of the main cardioid is the p/q bulb, and inside it the orbit settles on an attracting cycle of period q. Devaney explained how these fractions are arranged [22]: between the p/q and r/s bulbs, the largest bulb (measured by the share of the outside that leads into it) is the $(p + r)/(q + s)$ bulb, the mediant once more, which puts the Fibonacci fractions $1/2, 1/3, 2/5, 3/8, 5/13, \ldots$ on a path toward the golden-mean point. Our script places a test point just inside each bulb with $q \le 12$ and finds an attracting cycle of period exactly q in every case (0 exceptions). Figure 10.7 labels the larger bulbs. Walking $\pi - 3 = 0.14159$ of the way around the cardioid lands just short of the $1/7$ bulb, since $\pi - 3 < 1/7$, and π's convergents $15/106$ and $16/113$ are bulbs much smaller and much closer to that spot.

10.8 The Fibonacci word fractal

In 2009 Alexis Monnerot-Dumaine drew the Fibonacci word as a path [58]: take one unit step for each letter, and after the kth letter, if it is a 0, turn a quarter turn, left if k is even and right if k is odd. The Fibonacci word 0100101001001… is exactly the rhythm of $\varphi - 1$ from Chapter 3 with hits and rests exchanged (our script checks this over 30,000 letters), so this is the golden rhythm, walked. The path is a fractal curve made of smaller copies of itself, of dimension $3 \ln \varphi / \ln(1 + \sqrt{2}) = 1.6379$ [58]. Figure 10.8 draws the same rule for three numbers: the golden rhythm makes the fractal, $\sqrt{2} - 1$ weaves a grid, and $e - 2$ draws a zigzag of growing arms. For $\pi - 3$ the turns nearly cancel within each 7- or 8-beat gap, and over the first 33,102 beats the path is a long, thin zigzag.

10.9 Six numbers, six pictures

The Continued-Fraction Suite of Chapter 11 uses five numbers. One construction draws them all on the same footing: cut a $1 \times x$ rectangle into as many squares as fit (that count is the first

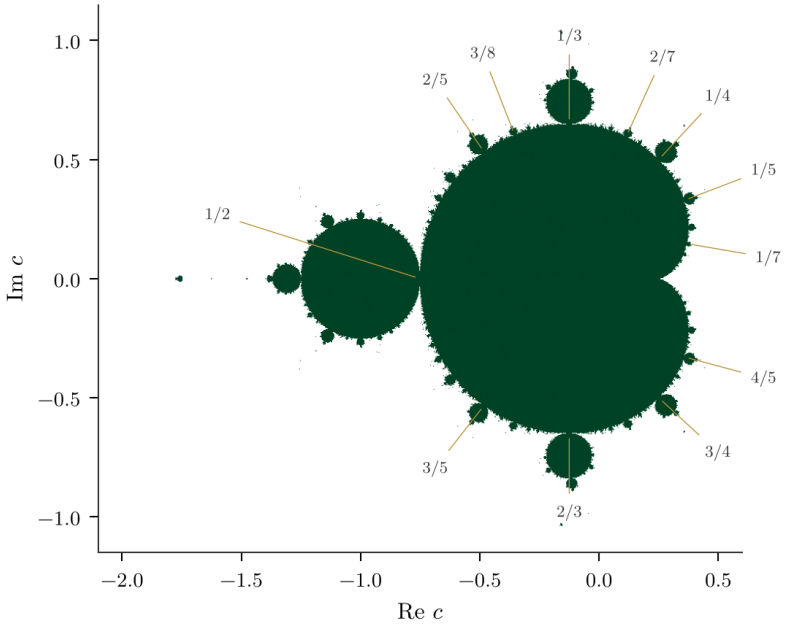


Figure 10.7. The Mandelbrot set with its main bulbs labeled p/q by the point where they attach to the cardioid. Inside the p/q bulb the orbit of 0 settles on a cycle of q points.

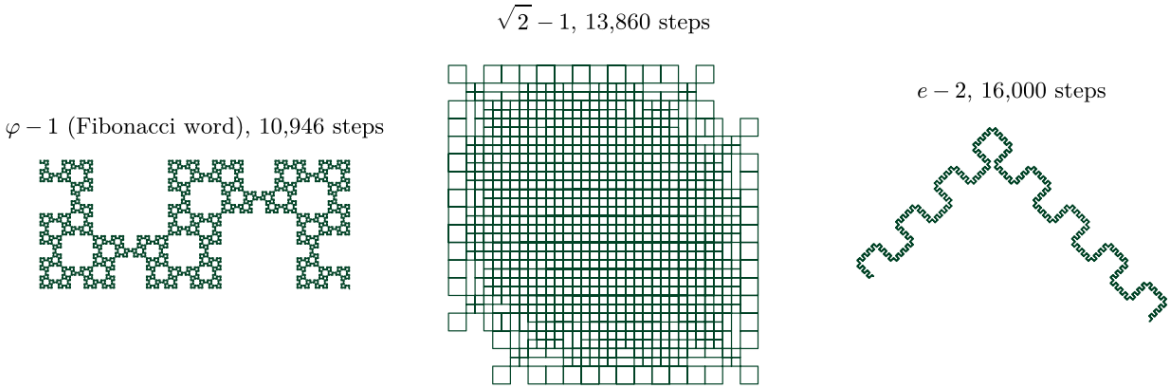


Figure 10.8. Rhythms walked as paths with Monnerot-Dumaine's rule. Left: the golden rhythm, the Fibonacci word fractal. Center: $\sqrt{2} - 1$. Right: $e - 2$.

term), turn the leftover a quarter turn, and repeat. Figure 10.9 shows the first six levels of each number, each level's rectangle stretched to the same height so that thin levels stay visible. The golden ratio cuts one square at every level and the silver ratio two: exactly self-similar. The fifth and e change their counts, e in its stretching pattern $1, 2, 1, 1, 4, 1, 1, 6$. Pi cuts 7, then 15, then 1, then 292 squares, the last a dense gold band. The sixth panel zooms to level 4 of π, the 292, which is where the sixth movement lives.

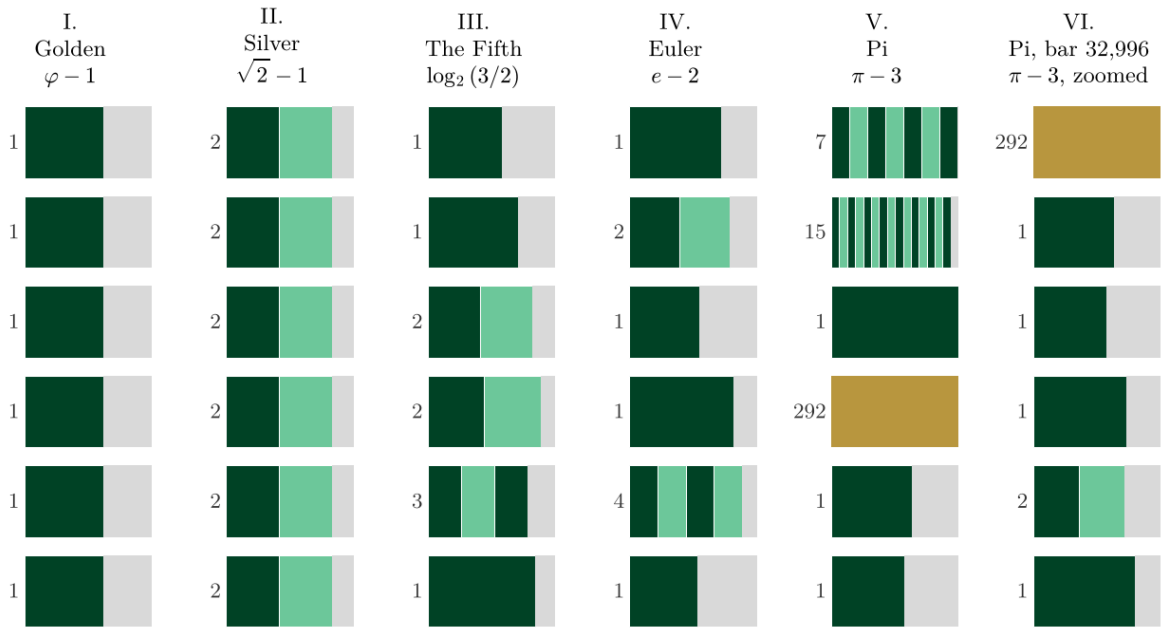


Figure 10.9. The Suite's numbers as square-cutting ladders. Each row is one level: the number of squares cut (printed at the left) is the term, and the grey sliver is what passes to the next level. Rows are stretched to equal height. A count above 50 is drawn as a gold band.

10.10 Is pi typical?

The Gauss map (2.2) reads off the terms, and the Gauss–Kuzmin law (Theorem 2.6) says what a typical number's terms look like. Here is a formal test of π's first 10,000 terms against it. Group the terms as $1, 2, \ldots, 12$ and 13+; the expected count in each group is 10,000 times its Gauss–Kuzmin probability. The chi-square statistic is $\chi^2 = 13.5$ on 12 degrees of freedom, with $p = 0.33$. With the finer grouping $1, \ldots, 20$ and 21+, $\chi^2 = 25.3$ on 20 degrees of freedom, $p = 0.19$. Either way the counts are fully consistent with the law (Fig. 10.10): 4,206 ones (42.1%, law 41.5%), 1,672 twos and 882 threes. The geometric mean of the terms is 2.6831 after 100 terms, 2.6627 after 1,000 and 2.6632 after 10,000, wandering around Khinchin's $K_0 = 2.685452001$ as a random number's would; and $\ln q_{10000}/10000 = 1.17941$ against Lévy's 1.18657 (Fig. 10.11). The largest term so far is 20,776, at position 431.

Nobody has proved that π obeys these laws, or even that its terms are unbounded, but by every count we can make it behaves like a typical number. This fits the null result of the prime hunt (Section 7.3) and the typical behavior of its prime convergents (Chapter 8). The number e, by contrast, fails the test on purpose: its terms $1, 2, 1, 1, 4, 1, 1, 6, \ldots$ have far too many ones and a geometric mean that grows without bound.

For the video. Begin with Teddy Rook's sentence and the one-line formula. Then six short visits: falling through Ford circles to π; walking the Stern–Brocot tree as L and R; the question-mark function heard as pitch; zooming into the devil's staircase; the Cantor dust building row by row; the Fibonacci word fractal drawing itself. Finish with Fig. 10.9 and the chi-square verdict: π looks typical. Instruments: all of Chapter VII on the fractals page. About 14 minutes.

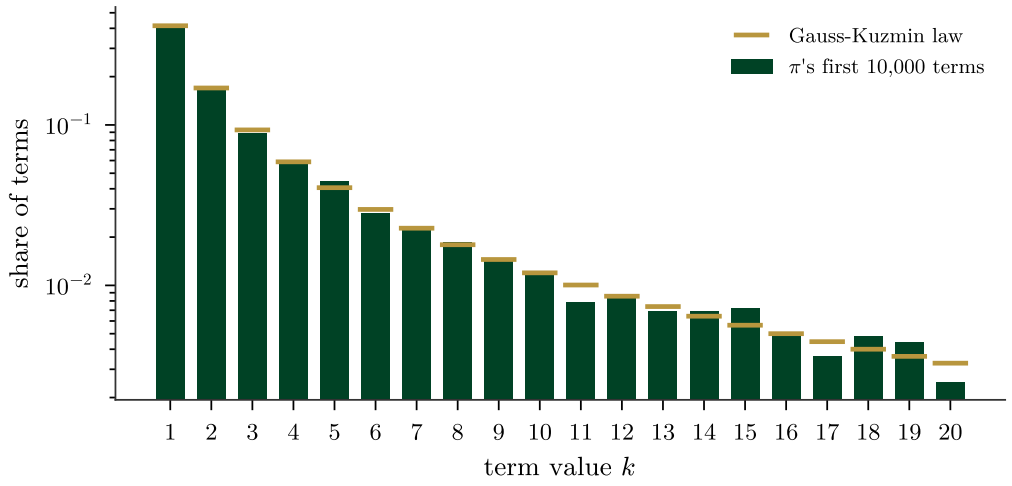


Figure 10.10. The shares of the values $1, \ldots, 20$ among π's first 10,000 terms (bars) against the Gauss–Kuzmin law (gold).

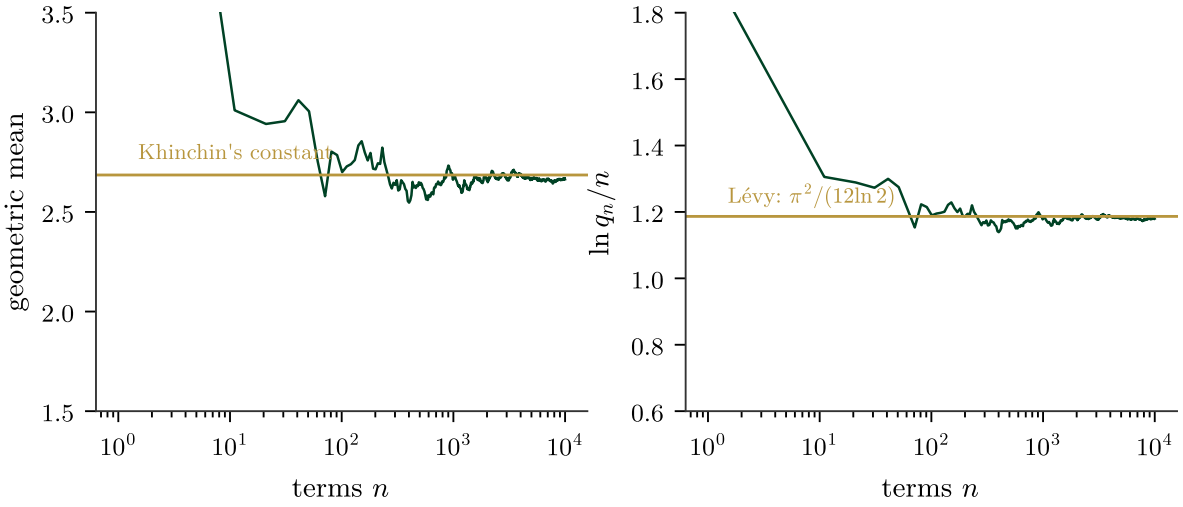


Figure 10.11. Running geometric mean of π's terms against Khinchin's constant (left) and $\ln q_n/n$ against Lévy's constant $\pi^2/(12 \ln 2)$ (right).

CHAPTER 11

Composition: the Continued-Fraction Suite as a method

11.1 The work

The *Continued-Fraction Suite* is an original composition by Teddy Rook: six short pieces for piano, one number each, engraved with LilyPond and published as a free score with a recorded performance [70]. It grew out of the Composer and "Compose a piece" instruments of the companion site, which play a number's rhythm with one voice per level. Whatever the prior art on its method (Chapter 13), the Suite is a new artwork. This chapter describes it as a method, in words and form diagrams; it does not reproduce the score. The musical facts below (keys, meters, tempos, bar counts, rehearsal marks) were read from the score by the site's builders and recorded in the site's data, and our script checks every bar count against the continued fractions.

11.2 The method

Each movement applies the same five decisions to its number.

1. *Choose the number and its levels.* Compute the terms $a_1, a_2, \ldots$ and the block lengths q_n (Theorem 3.5).

2. *Choose which level is a bar.* The block lengths then become lengths in bars of phrases, sections and larger units. Levels that are too short to be phrases are passed over: Golden starts its phrases at 3 and 5 bars, Silver at 2 and 5, Pi at 7.

3. *Lay out the form as the standard word.* At each level the long block appears a_{n+1} times and is closed by one short block. The score marks the long unit A, the short one B, and a unit that the number brings in only rarely as a "rare block."

4. *Put the rhythm inside each bar.* The left hand carries accents on the eighth notes where the number's own beat word has a hit (with α replaced by $1-\alpha$ when $\alpha > \frac{1}{2}$, so the accents stay sparse). So the same number shapes the music at the smallest scale and the largest.

Table 11.1. The six movements of the Continued-Fraction Suite. "Bars" is the length of the movement as printed; Movement VI begins at bar 32,968.

movement	number	key	meter, tempo	bars
I. Golden	$\varphi - 1 = [0; 1, 1, 1, \dots]$	G major	3/4, Andante cantabile	89
II. Silver	$\sqrt{2} - 1 = [0; 2, 2, 2, \dots]$	D major	4/4, Adagio	70
III. The Fifth	$\log_2 \frac{3}{2} = [0; 1, 1, 2, 2, 3, \dots]$	F major	4/4, Allegretto	53
IV. Euler	$e - 2 = [0; 1, 2, 1, 1, 4, \dots]$	A minor	3/4, Tempo di valzer	71
V. Pi	$\pi - 3 = [0; 7, 15, 1, 292, \dots]$	C major	4/4, Moderato	113
VI. Pi, bar 32,996	$\pi - 3$, at level 4	C major	4/4, Moderato	50

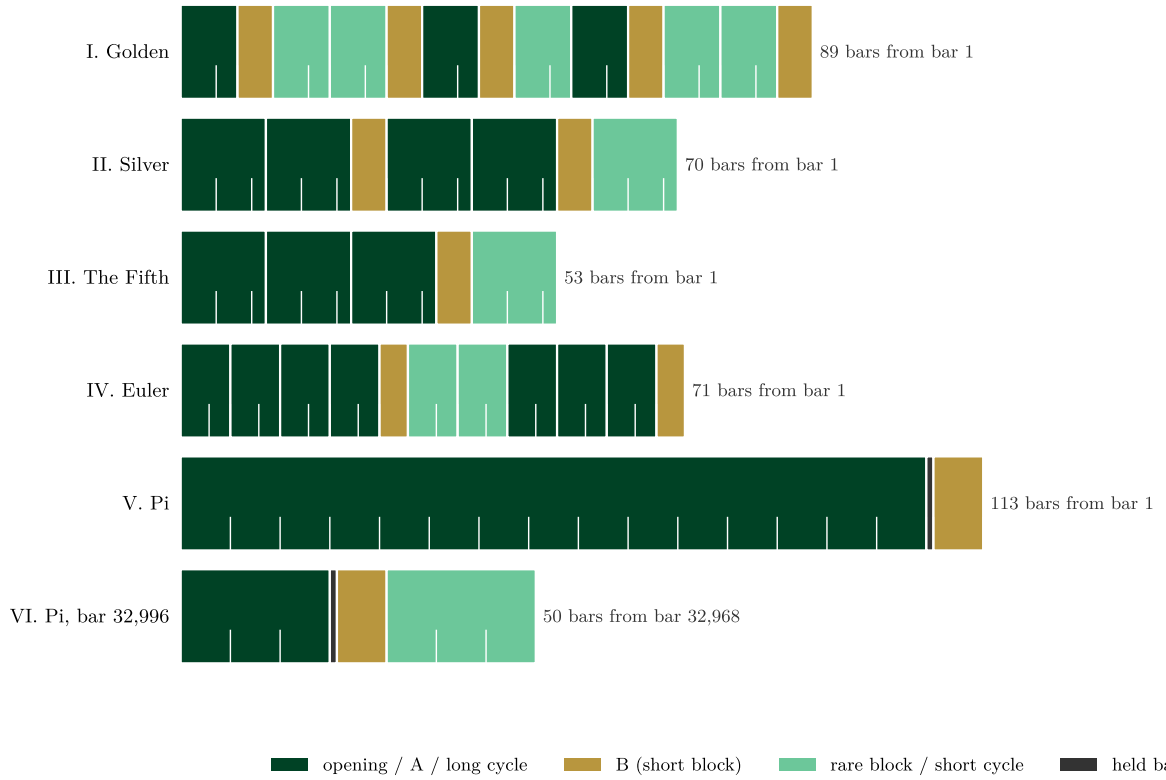


Figure 11.1. The forms of the six movements, drawn to scale (one unit per bar) from the section list read from the score. Colors follow the rehearsal marks: green for the opening and A sections (and Movement VI's long cycle), gold for B, mint for rare blocks and the short cycle, dark grey for a held bar. White ticks mark phrase boundaries.

5. *Let the melody walk the digits.* The score's preface says that the melody follows the number's decimal digits; this is a free compositional choice and carries no structural weight in what follows.

Figure 11.1 draws the six forms to scale, and the rest of this section reads each one.

I. Golden. Every term of $\varphi - 1$ is 1, so each block is the two blocks before it joined: the Fibonacci rule. Phrases are 5 and 3 bars, sections are $8 = 5 + 3$ and 5 bars, and the sections follow the Fibonacci word: $8, 5, 8, 8, 5, 8, 5, 8, 8, 5, 8, 8, 5$, which is eight 8s and five 5s, $64 + 25 =$

89 bars, itself a Fibonacci number. Our script checks that this sequence is the golden rhythm, letter for letter. Because every term is 1, no section ever comes three times in a row.

Golden also shows a subtlety of labeling that we call *two-level labels*. In the other movements the short block is truly rare. Here every level is "one long block, then one short block," and the sizes slide up the Fibonacci numbers: an 8-bar unit is the *long* block at the level of sections (8 against 5) and the *short* block at the level above ($13 = 8 + 5$ against 8). The score names each 8-bar unit by whichever level's boundary falls there, so the same 8 bars of music are marked A in some places and "rare block" in others. That is the self-similarity of the golden ratio, written into rehearsal marks.

II. Silver. Every term of $\sqrt{2} - 1$ is 2: "two of these, then one of those." Phrases are 5 and 2 bars, sections $12 = 5 + 5 + 2$ and 5 bars, and the whole is $29 + 29 + 12 = 70$ bars, where $29 = 12 + 12 + 5$. The numbers $2, 5, 12, 29, 70$ are the Pell numbers, the denominators of $\sqrt{2}$.

III. The Fifth. The terms of $\log_2(3/2)$ change: $1, 1, 2, 2, 3, 1, 5, \ldots$ Phrases are 5 and 2 bars and sections $12 = 5{+}5{+}2$; then the 3 takes over, three 12-bar sections and a 5 make 41, and one more 12 makes 53. The section sequence $12, 12, 12, 5, 12$ builds the denominators $5, 12, 41, 53$ of the best fractions for the fifth (Chapter 4) directly into the form.

IV. Euler. The terms of $e - 2$ grow in a pattern: $1, 2, 1, 1, 4, 1, 1, 6, \ldots$ Phrases are 4 and 3 bars, sections $7 = 4 + 3$ and 4; the 4 in the continued fraction makes a 32-bar block of four 7s and a 4; then $39 = 32 + 7$ and $71 = 39 + 32$. The denominators $7, 32, 39, 71$ are those of the convergents 5/7, 23/32, 28/39, 51/71 of $e - 2$.

V. Pi. The first three terms of $\pi - 3$ build a 113-bar cycle: fifteen 7-bar phrases, one held bar, and a last 7-bar phrase marked B, so $15 \times 7 + 1 + 7 = 113$. The held bar is the "one extra rest" of the standard word $s_2 = s_1^{15} s_0$, and the B phrase is the closing s_1 of $s_3 = s_2 s_1$.

VI. Pi, bar 32,996. The fourth term, 292, says that the 113-bar cycle repeats 292 times before the first short cycle of 106 bars. If Movement V never stopped, that moment would come after $292 \times 113 = 32{,}996$ bars, about 22.9 hours at the marked tempo (2.5 seconds per bar). The last movement jumps there. It plays bars 32,968 to 33,017: the end of the 292nd cycle (three phrases, the held bar and the B phrase, here in F major) and the opening three phrases of the short cycle, which begins at the printed bar 32,997. A short cycle is a long one without its closing phrase, so on paper it starts exactly like a long one; the composer marks the moment instead with a change of key and a hush. Section 3.4 explains why "beat 32,996" on the site and "bar 32,997" in the score name the same moment.

11.3 What the method gives a composer

The method gives proportions without dictating content. A number fixes how long each unit is and when contrast arrives; the composer still chooses every note. Quadratic irrationals give forms that repeat exactly at every level (Golden, Silver); numbers with changing terms give forms in which each level follows a different rule (The Fifth, Euler); and a number with one huge term gives a form with a surprise that is postponed almost indefinitely (Pi). In the language of Chapter 10, the first kind is self-similar in shape and the others only in rule.

For the video. Play eight bars of each movement while its strip from Fig. 11.1 scrolls underneath with a playhead. Pause on Golden's two-level labels. For Movement VI, show a counter running from bar 1 to 32,996 at high speed, then cut to the real music at bar 32,968. Instruments: the Composer's "Jump to the rarest block," "Compose a piece," and the six piano lessons in the Music section. About 10 minutes.

CHAPTER 12

Connections: music, patterns, math and chemistry

The Rhythm of Numbers is the deep dive of a larger page on the companion site, a four-way Venn diagram of music, patterns, mathematics and chemistry whose fifteen regions each show a place where the subjects meet [68]. This chapter places the work of the earlier chapters in that diagram, briefly and only with claims we have checked.

12.1 Music and math: ratios, clocks and rhythms

This region is the home of most of this work. Simple frequency ratios (2:1, 3:2, 4:3) sound smooth together; the twelve notes form a clock in which the step of a fifth, 7, is coprime to 12 and so visits every note; and the comma that keeps the clock from closing exactly is the irrationality of $\log_2(3/2)$ (Chapter 4). Euclidean rhythms, the drum machine and the Suite belong here too.

12.2 Music and chemistry: octaves of elements

In 1865 John Newlands arranged the known elements by atomic weight and noticed that chemically similar elements recurred every eighth place, which he compared with the octave of a musical scale and called the law of octaves [52]. He presented the arrangement to the Chemical Society of London on March 1, 1866, which declined to publish it [52]; the periodic table of Mendeleev, with gaps for undiscovered elements, followed in 1869. His "octave" was a way of counting, not a sound. Molecules, on the other hand, do have characteristic vibration frequencies, far above hearing; infrared spectra are their fingerprints.

12.3 Music and patterns: sand that draws sound

In 1787 Ernst Chladni bowed sand-covered metal plates and recorded the patterns the sand made when the plate rang [18]. The sand collects on the nodal lines of a standing wave. For a square plate a classical approximation to the modes is

$$z(x,y) = \cos(n\pi x)\cos(m\pi y) \pm \cos(m\pi x)\cos(n\pi y), \qquad 0 \le x, y \le 1,$$

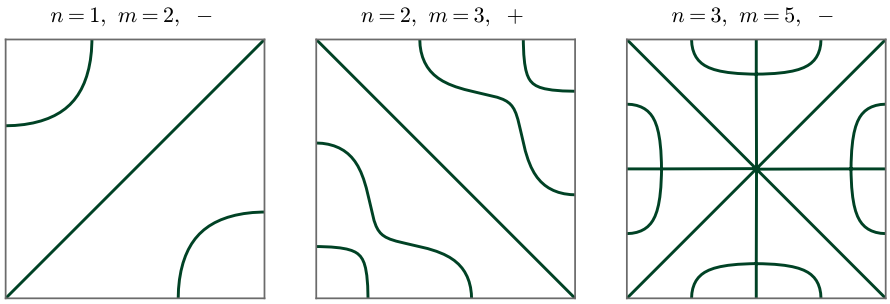


Figure 12.1. Chladni figures: the nodal curves $z = 0$ of the approximate square-plate modes for three choices of n, m and the sign.

with the sand on the curve $z = 0$ (Fig. 12.1). Real plates differ in detail near the edges. Melodies have symmetries of the same kind: transposition, inversion and retrograde are the translations and reflections of a pitch–time grid.

12.4 Math and patterns, patterns and chemistry

Periodic patterns in the plane and in space are classified by their symmetry groups, and crystals by the symmetries of their unit cells. Quasicrystals (Chapter 6) do not fit that classification: they are ordered, with sharp diffraction, but have no repeating cell [72]. On the site, the stained glass of cathedral windows shows the chemistry of color: dissolved metal ions absorb some wavelengths, and tiny metal particles scatter others.

12.5 Math and chemistry: why 2, 8, 18, 32

The electron shell n of an atom has subshells $l = 0, \dots, n-1$, each with $2l+1$ orbitals holding two electrons apiece. Since $1 + 3 + \cdots + (2n-1) = n^2$, shell n holds $2n^2$ electrons: $2, 8, 18, 32$. The squares here are the same squares whose reciprocals sum to $\pi^2/6$ in Chapter 5; as there, this is a shared formula, not shared physics.

12.6 All four: waves and symmetry

At the center of the diagram is the standing wave. A string's modes are whole-number multiples of its lowest frequency, which is why it has a clear pitch. A circular drum's modes are not: their frequencies are proportional to zeros of Bessel functions, and the lowest four are in the ratios $1 : 1.59 : 2.14 : 2.30$, which is why a drum thumps rather than sings. An electron in an atom is a standing wave in three dimensions; its orbitals sort by their number of nodal surfaces just as drum modes sort by nodal lines, and their counts, filled in order, give the shape of the periodic table.

12.7 Where the Rhythm of Numbers sits

The Rhythm of Numbers lives mainly in the music–math region and reaches into the others: tuning and rhythm (music and math), coprimality and diffraction (math and patterns, patterns

and chemistry), the $1/n^2$ of hydrogen and the n^2 of shells (math and chemistry), and self-similarity, which is the subject of patterns itself. What the diagram adds is a reminder that the same few structures (ratios, standing waves, symmetry, and the continued fraction that measures how far a ratio is from simple) appear in each subject for different reasons.

For the video. Show the four-way diagram, then light up each region for a few seconds with its one picture: the circle of fifths, Newlands' table, a Chladni plate, a quasicrystal pattern, the $2n^2$ squares, a drum and an orbital. Then zoom into the music–math region and say "this is where the rest of the film lives." Instruments: the Connections hub and its region pages. About 6 minutes.

CHAPTER 13

What is new here, and what is re-found

A work that walks through this much classical mathematics owes its readers a careful account of what is old. This chapter audits every idea in the work and in the companion site against the literature, and sorts it into four classes:

Classical: a standard result with a long history; we cite a primary source or a standard text.

Specialist: known in a research literature, though not in standard texts; we cite it.

New presentation: a new arrangement, picture, synthesis or computation of known results.

New: something we did not find anywhere: an original artifact, a computation, or a heuristic that we could not locate in the literature we searched.

The searches covered the standard texts cited throughout, the music-theory literature on maximally even sets, well-formed scales, Christoffel words and Sturmian rhythms (searching for, among others, "Sturmian words music," "continued fractions musical form," "Christoffel words rhythm"), the OEIS, and the number-theory literature on convergent denominators. No search of this kind is complete. When we say "we did not find prior work," we mean exactly that, and nothing stronger.

13.1 The audit

Table 13.1. The prior-art audit, idea by idea.

idea	class	notes and sources
The guess $\pi \approx 3+k/(7k+1)$ and $k = 16$	classical	The guess is the continued fraction of π stopped after two terms (Chapter 1); 355/113 is Zu Chongzhi's Milü [54]; 223/71 and 22/7 are Archimedes' bounds. Re-found independently by Teddy Rook.
Continued fractions, convergents, best approximation, Legendre, Lagrange	classical	Hardy and Wright [36], Khinchin [41].

idea	class	notes and sources
Gauss–Kuzmin, Khinchin, Lévy	classical	Khintchine [42], Kuzmin [44], Lévy [50, 51].
The "rhythm of a number"	classical	It is the characteristic Sturmian word of α, the cutting sequence of a line of slope α [2, 9, 53]. The rhythm of π in particular is the characteristic word of $\pi - 3$; we found no claim that it is new and make none.
Nested blocks: block lengths are the q_n	classical	Standard words; Brown [13], Fraenkel et al. [29]; our proof in Chapter 3 is a direct version for completeness.
The beats 32,996 and 33,102	new presentation	Elementary consequences of the standard word $s_4 = s_3^{292} s_2$; computing and reconciling the two counts is ours.
The continued fraction as musical *form* (bar, phrase, section, movement)	specialist / new presentation	Closely related work exists. Carey and Clampitt studied self-similar pitch structures and their rhythmic analogues [16, 17]; Clampitt and Noll use Sturmian morphisms and Christoffel words for modes [19, 23]; Callender built rhythmic canons from Sturmian words whose voices and relations are set by the continued fraction of the slope [14, 15]; Fibonacci proportions in form have a long analytic literature [49]. We did not find the specific reading of a number's *whole* continued fraction as a nested hierarchy of formal units, each term a repeat count with a closing contrasting unit, used to lay out complete pieces.
The Composer and "Compose a piece" instruments	new (artifact)	Software on the companion site.
Euclid's algorithm and Euclidean rhythms	specialist	Clough and Douthett [20], Gómez-Martín et al. [31], Toussaint [79]; Bjorklund's 2003 algorithm as reported by Toussaint. That $E(16, 113)$ is π's level-3 block is an observation that follows at once.
Tuning from $\log_2(3/2)$, the comma, 41 and 53	classical	Benson [6].
Circle map, Arnold tongues, devil's staircase	classical	Jensen et al. [40]; Huygens's clocks [5].
$6/\pi^2$, the Euler product, $\zeta(2)$	classical	Euler [26, 27], Hardy and Wright [36].
Correlation of pulse trains, $(\gcd - 1)/\sqrt{(a-1)(b-1)}$	new presentation	An elementary computation. We did not find it stated in this form, but expect it is known; the "orchard as correlation matrix" reading is a presentation, not a theorem.

idea	class	notes and sources
Hydrogen's $1/n^2$ and $\pi^2/6$	classical, with a caveat	Balmer [4]; a shared formula, not shared physics (Chapter 5).
Residue bias of consecutive primes	specialist	Lemke Oliver and Soundararajan [48]; our counts re-done.
The prime sunflower of π	new presentation	Dirichlet's theorem and wheel sieves are classical [36]; seeding the $\pi-3$ spiral with primes is a picture of them.
Diffraction of π's rhythm and of the primes	specialist	Sturmian diffraction [3]; primes [78]. The Ramanujan-sum explanation of the square-free rule is standard reasoning; our computation re-done.
Explicit formula and zeros	classical	Riemann [67], von Mangoldt [80].
Continued-fraction factoring	classical	Lehmer and Powers [47], Morrison and Brillhart [60].
Mersenne necklaces and Fermat's little theorem	classical	Golomb [30].
The prime hunt in rhythm prefixes	new (experiment)	We did not find prior work testing binary prefixes of Sturmian words for primality. The outcome is a null result (Chapter 7).
The prime-convergent heuristic with $c = 2\ln 2$	new (heuristic)	Built entirely from known theorems [37, 51, 57, 76]. We did not find the constant or the identification of the correction factor as $\zeta(2)$ in the literature we searched. Related heuristics for linear recurrences: Grantham and Granville [34].
The random-real test of the heuristic, and π's placement	new (computation)	Our computation (Chapter 8); the lists for π match the OEIS [63, 64].
The spacing spectrum (crystal $\to$ quasicrystal $\to$ GUE $\to$ Poisson)	new presentation	Every ingredient is known: Berry [8], Bogomolny et al. [10], Dyson [25], Mehta [55], Montgomery [59], Odlyzko [61]; ordering point patterns by number variance is the hyperuniformity program of Torquato and Stillinger [77]; Dyson's quasicrystal remark is in Dyson [24]. Putting π's rhythm, the zeros, GUE and the primes on one scale with one set of tools is our synthesis.
Teddy Rook's picture of the nested tail	classical	The self-similarity of continued fractions; Chapter 10.
Patterned continued fractions for π	classical	Brouncker [43]; Lange [46].

idea	class	notes and sources
Ford circles, Stern–Brocot, $?(x)$, Cantor dust, Farey bulbs, Fibonacci word fractal	classical / specialist	Brocot [12], Denjoy [21], Devaney [22], Ford [28], Good [32], Jenkinson and Pollicott [38, 39], Minkowski [56], Monnerot-Dumaine [58], Salem [71], Stern [74]. Our dimension computation re-does the published method.
The six-numbers pictures	new presentation	Square cutting is a classical picture of Euclid's algorithm; applying one construction to the Suite's numbers is ours.
Gauss–Kuzmin test of π's first 10,000 terms	new presentation	Statistical checks of π's terms are a known kind of computation; the specific statistics are ours.
The Continued-Fraction Suite	new (artwork)	An original composition by Teddy Rook [70].
Golden's two-level labels	new presentation	An observation about labeling form in a self-similar piece.

13.2 The two questions the owner asked

Is the rhythm of π new? No. The rhythm of any number, in the sense of Chapter 3, is its characteristic Sturmian word, a classical object, and the nested-block structure is a classical theorem. The rhythm of π is one instance. What is less common is to *listen* to it and to name its levels as musical form; that reading has close relatives in the music-theory literature, notably the Sturmian canons of Callender, in which the continued fraction governs the relations between voices, and the work of Carey, Clampitt and Noll on Sturmian structures in scales and rhythms. We did not find the full form reading (every term a repeat count, from bar to movement) laid out as a method for composing complete pieces.

Is the new music new? Yes, as music. The Continued-Fraction Suite is an original work by Teddy Rook, and its originality as an artwork does not depend on the novelty of its method, any more than a fugue's does on the novelty of counterpoint. Of its method, the Fibonacci proportions of Golden have a long history in analysis and composition; Sturmian rhythms and continued-fraction-governed canons exist in the specialist literature. We did not find a prior composition whose whole form is laid out level by level from a number's continued fraction, nor one that sets out to reach the 292nd return of π's 113-bar cycle.

For the video. A brisk tour of Table 13.1, colored by class. Spend most of the time on the two questions: the rhythm of π is a classical object, heard in a new way; the Suite is new music. About 6 minutes.

CHAPTER 14

Discussion

14.1 What the work shows

One guess about π led, through one algorithm, to a surprising amount of mathematics and music. The guess $3+k/(7k+1)$ is a continued fraction; a continued fraction is a rhythm whose levels are its terms (Theorem 3.5); the same algorithm builds tunings, Euclidean rhythms, mode-locking staircases and a factoring method; coprimality, which every level of the rhythm obeys, is where $6/\pi^2$ enters; and the self-similarity of the continued fraction is the common root of a family of classical fractals. Along the way we found the precise meaning of two numbers that the site's sources disagreed about (32,996 and 33,102), and we checked every numerical claim by script.

The two research threads gave concrete answers. For prime convergents, a heuristic built from Lévy's theorem and the equidistribution of convergents modulo m predicts $2\ln 2 \cdot \ln N$ prime denominators among the first N convergents of a typical number, with the correction factor turning out to be $\zeta(2)$; a sample of random reals fits the predicted growth, and π is typical. For the Montgomery–Dyson connection, one set of statistics places a crystal, π's rhythm, the zeta zeros, random matrices, the primes and Poisson points on a single scale, and the explicit formula makes Dyson's quasicrystal remark visible. Neither thread proves a theorem about π or the zeta function, and we have been careful to say so.

14.2 Limitations

- *Heuristics are not theorems.* Heuristic 8.3 treats primality of different denominators as independent events; the agreement with data supports it but does not prove it.
- *Finite samples.* The prime-convergent test uses 400 random reals and 1,000 convergents each; the slope fit has a standard error of 0.04 and is slightly above the prediction because of early terms. A larger sample, and longer expansions, would sharpen the test.
- *Low zeros.* We used the first 5,000 zeros, which we computed ourselves. At these heights finite-size corrections are visible [8, 10]; Odlyzko's tables of zeros at much greater heights [61] would show closer agreement with the GUE.
- *Probable primes.* Large "primes" are Baillie–PSW probable primes. No composite is known to pass that test, but it is not a proof of primality.

- *Score facts.* The facts about the Suite come from the site's data, which its builders read from the engraved score; we checked the bar arithmetic but not each note.
- *Literature search.* The audit of Chapter 13 is as careful as we could make it, but the music-theory literature in particular is large, and related work may exist that we did not find.

14.3 Open questions

1. Does π have infinitely many prime convergent denominators? Infinitely many prime numerators? (Open; Chapter 8.)
2. Are there infinitely many Fibonacci primes? Pell primes? (Open.)
3. Can Heuristic 8.3 be proved in a model? For example, for almost every x, is the number of $n \le N$ with q_n prime asymptotic to $2 \ln 2 \cdot \ln N$? Even an upper bound of that order for almost every x would be interesting.
4. Are π's terms unbounded? Do they obey the Gauss–Kuzmin law? (Open; every computation says "typical.")
5. How does the spectrum of the arithmetic correlation matrix of Section 5.2 compare with random-matrix benchmarks, and does $6/\pi^2$ appear in its bulk?
6. Rank all equal temperaments up to 200 notes by how well they approximate several just intervals at once. Do prime sizes stand out? (We expect not; a clear "no" would still be a result.)

14.4 Future work

On the mathematical side: extend the prime-convergent test to 10^4 convergents for hundreds of reals and to the numerators; repeat the spacing statistics with zeros at large height; and look for a rigorous version of the heuristic in a probabilistic model of continued fractions. On the musical side: the method of Chapter 11 invites pieces in tunings whose number of notes is itself a convergent denominator (a pi movement in a 113-note tuning, for instance), and pieces built from numbers with chosen terms, whose form a composer could design by picking the continued fraction. The companion site's instruments already do much of this; the score of each new piece is a continued fraction made audible.

For the video. Close the film: return to the scrap of paper with $3+k/(7k+1)$, list the open questions on screen, and play the last bars of Movement VI. About 4 minutes.

APPENDIX A

Algorithms

This appendix gives the main algorithms in pseudocode. The Python scripts of Appendix C implement each one.

Algorithm 1: Certified continued-fraction terms of a real constant

Input: a procedure that evaluates x to d decimal digits; the number of terms n
Output: the first n terms of x, each guaranteed correct

repeat
  evaluate x at precision d and at precision $d + 200$;
  expand both: $a \leftarrow \lfloor y \rfloor$, $y \leftarrow 1/(y - a)$, repeated;
  keep the common prefix of the two term lists, minus its last term;
  if *fewer than n terms are kept* **then** $d \leftarrow 1.6\,d$;
until *at least n terms are certified*;

Algorithm 2: Convergents (Proposition 2.1)

Input: terms $a_0, a_1, \ldots, a_N$

$p_{-1}, q_{-1}, p_0, q_0 \leftarrow 1, 0, a_0, 1$;
for $n \leftarrow 1$ **to** N **do**
  $p_n \leftarrow a_n p_{n-1} + p_{n-2}$; $q_n \leftarrow a_n q_{n-1} + q_{n-2}$;

Algorithm 3: The rhythm of a number (standard words, Theorem 3.5)

Input: terms $a_1, a_2, \ldots, a_k$ of $\alpha = [0; a_1, a_2, \ldots]$
Output: the level words $s_0, \ldots, s_k$; s_n has q_n beats and p_n hits

$s_{-1} \leftarrow$ `x`; $s_0 \leftarrow$ `.`;
$s_1 \leftarrow s_0^{a_1 - 1}\, s_{-1}$;
for $n \leftarrow 2$ **to** k **do**
  $s_n \leftarrow s_{n-1}^{a_n}\, s_{n-2}$; `// repeat, then close`

Algorithm 4: Bjorklund's algorithm for the Euclidean rhythm $E(k, n)$

Input: $1 \le k < n$

$A \leftarrow k$ groups `[x]`; $B \leftarrow n - k$ groups `[.]`;
while *B has more than one group* **do**
  $m \leftarrow \min(|A|, |B|)$;
  append the ith group of B to the ith group of A, for $i = 1, \dots, m$;
  the unused groups (from A if $|A| > m$, else from B) become the new B;
return *the concatenation of the groups of A and then of B*

Algorithm 5: Lehmer–Powers square search on the continued fraction of $\sqrt{kN}$

Input: odd composite N, multiplier k

$a_0 \leftarrow \left\lfloor \sqrt{kN} \right\rfloor$; $P \leftarrow 0$; $Q \leftarrow 1$; $a \leftarrow a_0$; $A_{-1} \leftarrow 1$; $A_0 \leftarrow a_0 \bmod N$;
for $n \leftarrow 1, 2, \dots$ **do**
  $P \leftarrow aQ - P$; $Q \leftarrow (kN - P^2)/Q$; $a \leftarrow \lfloor (a_0 + P)/Q \rfloor$;
  if *n is even and $Q = s^2$ for an integer s* **then**
    $g \leftarrow \gcd(A_{n-1} - s, N)$;
    if $1 < g < N$ **then return** g;
  $A_n \leftarrow (aA_{n-1} + A_{n-2}) \bmod N$;

Algorithm 6: Counting prime convergents of random reals (Chapter 8)

Input: number of reals M, number of convergents N

for $i \leftarrow 1$ **to** M **do**
  draw a uniform integer $r \in [1, 10^D)$ with $D = \lfloor 1.6N \rfloor + 20$;
  expand $r/10^D$ by exact Euclidean division; keep the first $N + 1$ terms;
  compute p_n, q_n for $n \le N$ (Algorithm 2);
  foreach $n \le N$ **do**
    $E \leftarrow E + \zeta(2)/\ln q_n$; `// heuristic expectation`
    if *q_n has no prime factor below* 3000 *and passes Baillie–PSW* **then**
      count q_n as prime

Algorithm 7: Hausdorff dimension of a continued-fraction Cantor set

Input: allowed terms A; number of Chebyshev nodes m

nodes $x_1, \dots, x_m$ on $[0, 1]$; interpolation matrix V^{-1} for polynomials of degree $< m$;
`rho(`s`)`: matrix $M = \sum_{a \in A} \operatorname{diag}\big((a + x_i)^{-2s}\big)\, V_a V^{-1}$, where V_a evaluates the basis at $1/(a + x_i)$; **return** *spectral radius of M*;
bisect on $s \in [0, 1]$ for $\rho(s) = 1$;

Algorithm 8: Number variance of an unfolded point set

Input: sorted points $u_1 < \cdots < u_N$ with mean spacing 1; window lengths L; trials T

foreach L **do**
- draw T uniform window starts t in $[u_1, u_N - L]$;
- count the points in each window $[t, t + L)$ by binary search;
- $\Sigma^2(L) \leftarrow$ sample variance of the counts;

APPENDIX B

Computational methods

Software. All computations use Python 3 with mpmath 1.3 (arbitrary-precision reals and the zeros of the zeta function), sympy (factorizations, small primality tests), gmpy2 (fast big-integer arithmetic and the Baillie–PSW test), NumPy and SciPy (linear algebra, statistics), and matplotlib (figures). The typesetting uses XeLaTeX with BibTeX.

Continued fractions. The terms of π come from π evaluated to 11,600 decimal digits, expanded by the Gauss map at two working precisions; only terms on which the two expansions agree, less one, are kept (Algorithm 1), which certifies the first 10,000 terms. The constants e, $\sqrt{2}$, φ and $\log_2(3/2)$ are expanded the same way to 4,001 terms, with the precision raised automatically when needed. The pattern of e and the constant terms of $\sqrt{2}$ and φ are asserted, which checks the method.

Rhythms. Rhythms are built with Algorithm 3 and compared, letter for letter, with the floor rule (3.1) evaluated in multiple precision: 140,000 beats for $\pi - 3$ and complete words for 200 random term sequences.

Primality. A number is called prime if it is below 3,000 and in the list of small primes, or if it has no prime factor below 3,000 (one gcd with the primorial) and passes the Baillie–PSW probable-prime test of gmpy2. For the Fibonacci and Pell indices our lists agree with the OEIS [62, 65], and for π with [63, 64].

Random reals. A random real is represented by $r/10^D$ with r a uniformly random integer of $D = 1{,}620$ digits (Algorithm 6). Its first 1,000 terms are shared by every real in $[r/10^D, (r+1)/10^D)$, as explained in Chapter 8. The generator is Python's `random.Random` with a fixed seed, so the sample is reproducible.

Zeta zeros. The first 5,000 zeros were computed with `mpmath.zetazero` at 20 digits; the script checks that `mpmath.nzeros` counts exactly 5,000 zeros up to the last one, so none is missing, and recomputes a few at higher precision. The explicit-formula figure uses the first 100 zeros.

Random matrices. GUE matrices are $H = (A + A^*)/2$ with A a 1000×1000 matrix of independent standard complex Gaussians, so that $E|H_{ij}|^2 = 1$ off the diagonal; eigenvalues are divided by $\sqrt{1000}$ (semicircle of radius 2) and unfolded with the semicircle distribution function, keeping the central half.

Statistics. Chi-square tests use SciPy's χ^2 survival function; the Kolmogorov–Smirnov distance is the two-sample statistic; Poisson tail probabilities use SciPy. The pair-correlation histogram counts all pairs within 3 mean spacings among the first 6,000 points of each set; the number variance uses 1,500 to 4,000 random windows per length (Algorithm 8).

Structure factors. $S(k)$ is computed directly from (6.1) on a grid of 40,000 wavenumbers, with positions measured from the start of the window; the figures use a Hann taper (weights $2\sin^2(\pi x/\ell)$ across the window of length ℓ) to suppress side lobes, while the peak values quoted in the text are untapered and evaluated exactly at the rational wavenumbers.

Circle map. Rotation numbers are averages over 2,520 iterations (a multiple of every $q \le 10$, so locked orbits give exact fractions) after a transient of 400; step widths are measured on a grid of 40,001 values of Ω.

Dimension. Algorithm 7 with 40 and 60 Chebyshev nodes; the two agree to 13 digits for $E_{\{1,2\}}$, and the 60-node value matches the published $0.53128050627720514\ldots$ [39] to the twelve digits we print.

APPENDIX C

Reproducibility

The full source of this work, including the LaTeX files, the bibliography and every script, is kept with the companion site's source under `docs/papers/rhythm-of-numbers/`. Running the scripts in order regenerates every figure (`paper/figures/`) and every number in the text: each script writes a file of LaTeX macros to `paper/gen/`, and the text uses only those macros for computed values. Scripts stop with an error if any asserted fact fails.

Building. With Python 3 and the packages of Appendix B installed, run `sh run_all.sh` in `compute/`, then `latexmk -xelatex main.tex` in `paper/`. Random choices use fixed seeds, so a rebuild reproduces the same numbers, except that the GUE samples (seeded too) and timings may differ across library versions in the last digit.

Provenance of non-computed facts. Dates, names and bibliographic details were checked against the sources listed in the bibliography. The musical facts about the Suite come from the companion site's data file `data/music/suite.yaml`, compiled from the engraved score.

Table C.1. The scripts, in the order `run_all.sh` runs them, with what each computes.

script	computes
`common.py`	shared helpers: certified continued fractions, convergents, standard words, the floor rule, macro output, figure style
`c1_cf.py`	terms of π, e, $\sqrt{2}$, φ, $\log_2(3/2)$; the starting guess; convergent table; Legendre and best-approximation checks; nested-block checks; beats 32,996 and 33,102; three-gap check; Gauss–Kuzmin, Khinchin and Lévy statistics
`c2_music.py`	equal temperaments, comma, thirds; Bjorklund vs. Christoffel; Euclidean table; circle map, tongues, staircase and Farey rule; Farey counts
`c3_coprime.py`	orchard, Euler product, correlation formula and matrix, Balmer, residue bias to 10^7, sunflower and wheels, structure factors, explicit formula
`c4_tools.py`	Lehmer–Powers factoring, necklaces, the prime hunt and its Poisson tails
`zeros.py`	the first 5,000 zeta zeros (about one hour)
`c5_spacing.py`	spacing distributions, pair correlation, number variance, the Fourier picture of the zeros
`c6_fractals.py`, `c6b_gcf.py`	complete quotients, Ford circles, Stern–Brocot, $?(x)$, dimensions, Mandelbrot bulbs, Fibonacci word, six pictures; patterned fractions for π
`c7_primeconv.py`, `c7b_primeconv_figs.py`	prime convergents of five constants and 400 random reals (several minutes); divisibility shares; figures and tables
`c8_suite.py`	the Suite's forms checked against the continued fractions; form figure
`c9_misc_figs.py`, `c10_connections.py`	remaining figures; drum-mode ratios

Bibliography

[1] Pascal Alessandri and Valérie Berthé. Three distance theorems and combinatorics on words. *L'Enseignement Mathématique*, 44:103–132, 1998.

[2] Jean-Paul Allouche and Jeffrey Shallit. *Automatic Sequences: Theory, Applications, Generalizations*. Cambridge University Press, 2003.

[3] Michael Baake and Uwe Grimm. *Aperiodic Order. Volume 1: A Mathematical Invitation*, volume 149 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, 2013.

[4] J. J. Balmer. Notiz über die Spectrallinien des Wasserstoffs. *Annalen der Physik und Chemie*, 25:80–87, 1885. Third series.

[5] Matthew Bennett, Michael F. Schatz, Heidi Rockwood, and Kurt Wiesenfeld. Huygens's clocks. *Proceedings of the Royal Society A*, 458:563–579, 2002.

[6] Dave Benson. *Music: A Mathematical Offering*. Cambridge University Press, 2006.

[7] Lennart Berggren, Jonathan Borwein, and Peter Borwein. *Pi: A Source Book*. Springer, New York, 1997.

[8] M. V. Berry. Semiclassical formula for the number variance of the Riemann zeros. *Nonlinearity*, 1(3):399–407, 1988.

[9] Jean Berstel, Aaron Lauve, Christophe Reutenauer, and Franco V. Saliola. *Combinatorics on Words: Christoffel Words and Repetitions in Words*, volume 27 of *CRM Monograph Series*. American Mathematical Society, 2008.

[10] E. Bogomolny, O. Bohigas, P. Leboeuf, and A. G. Monastra. On the spacing distribution of the Riemann zeros: corrections to the asymptotic result. *Journal of Physics A: Mathematical and General*, 39(34), 2006. doi:10.1088/0305-4470/39/34/010.

[11] Bence Borda. Equidistribution of continued fraction convergents in $\mathrm{sl}(2,\mathbb{Z}_m)$ with an application to local discrepancy. *Journal of Modern Dynamics*, 2025. doi:10.3934/jmd.2025005; arXiv:2303.08504.

[12] Achille Brocot. Calcul des rouages par approximation, nouvelle méthode. *Revue Chronométrique*, 1861. Volume and pages not confirmed.

[13] T. C. Brown. Descriptions of the characteristic sequence of an irrational. *Canadian Mathematical Bulletin*, 36(1):15–21, 1993.

[14] Clifton Callender. Sturmian canons. In J. Yust, J. Wild, and J. A. Burgoyne, editors, *Mathematics and Computation in Music (MCM 2013)*, volume 7937 of *Lecture Notes in Computer Science*. Springer, 2013. doi:10.1007/978-3-642-39357-0_5.

[15] Clifton Callender. Infinite rhythmic tiling canons. In *Proceedings of Bridges 2015: Mathematics, Music, Art, Architecture, Culture*, pages 399–402, 2015.

[16] Norman Carey and David Clampitt. Aspects of well-formed scales. *Music Theory Spectrum*, 11(2):187–206, 1989.

[17] Norman Carey and David Clampitt. Self-similar pitch structures, their duals, and rhythmic analogues. *Perspectives of New Music*, 34(2), 1996.

[18] Ernst Florens Friedrich Chladni. Entdeckungen über die Theorie des Klanges. Leipzig, 1787.

[19] David Clampitt and Thomas Noll. Modes, the height-width duality, and Handschin's tone character. *Music Theory Online*, 17(1), 2011.

[20] John Clough and Jack Douthett. Maximally even sets. *Journal of Music Theory*, 35: 93–173, 1991.

[21] Arnaud Denjoy. Sur une fonction réelle de Minkowski. *Journal de Mathématiques Pures et Appliquées*, 17:105–151, 1938.

[22] Robert L. Devaney. The Mandelbrot set, the Farey tree, and the Fibonacci sequence. *American Mathematical Monthly*, 106(4):289–302, 1999.

[23] Manuel Domínguez, David Clampitt, and Thomas Noll. WF scales, ME sets, and Christoffel words. In *Mathematics and Computation in Music*. Springer, Berlin, Heidelberg, 2009. doi:10.1007/978-3-642-04579-0_48.

[24] Freeman Dyson. Birds and frogs. *Notices of the American Mathematical Society*, 56(2): 212–223, 2009.

[25] Freeman J. Dyson. Statistical theory of the energy levels of complex systems. I. *Journal of Mathematical Physics*, 3(1):140–156, 1962.

[26] Leonhard Euler. De summis serierum reciprocarum. *Commentarii academiae scientiarum Petropolitanae*, 7:123–134, 1740. Enestrom index E41.

[27] Leonhard Euler. Variae observationes circa series infinitas. *Commentarii academiae scientiarum Petropolitanae*, 9:160–188, 1744. Enestrom index E72.

[28] L. R. Ford. Fractions. *American Mathematical Monthly*, 45(9):586–601, 1938.

[29] A. S. Fraenkel, M. Mushkin, and U. Tassa. Determination of $[n\theta]$ by its sequence of differences. *Canadian Mathematical Bulletin*, 21(4):441–446, 1978.

[30] Solomon W. Golomb. Combinatorial proof of Fermat's "little" theorem. *American Mathematical Monthly*, 63(10):718, 1956.

[31] Francisco Gómez-Martín, Perouz Taslakian, and Godfried T. Toussaint. Structural properties of Euclidean rhythms. *Journal of Mathematics and Music*, 2009. doi:10.1080/17459730902819566.

[32] I. J. Good. The fractional dimensional theory of continued fractions. *Proceedings of the Cambridge Philosophical Society*, 37:199–228, 1941.

[33] Ronald L. Graham, Donald E. Knuth, and Oren Patashnik. *Concrete Mathematics: A Foundation for Computer Science*. Addison-Wesley, second edition, 1994.

[34] Jon Grantham and Andrew Granville. Fibonacci primes, primes of the form $2^n - k$ and beyond. *Journal of Number Theory*, 2024. doi:10.1016/j.jnt.2024.02.002; arXiv:2307.07894.

[35] Great Internet Mersenne Prime Search. GIMPS discovers largest known prime number: $2^{136,279,841} - 1$. Press release, `https://mersenne.org/primes/press/M136279841.html`, 2024.

[36] G. H. Hardy and E. M. Wright. *An Introduction to the Theory of Numbers*. Oxford University Press, sixth edition, 2008. Revised by D. R. Heath-Brown and J. H. Silverman.

[37] H. Jager and P. Liardet. Distributions arithmétiques des dénominateurs de convergents de fractions continues. *Indagationes Mathematicae (Proceedings)*, 50:181–197, 1988.

[38] Oliver Jenkinson and Mark Pollicott. Computing the dimension of dynamically defined sets: E_2 and bounded continued fractions. *Ergodic Theory and Dynamical Systems*, 21(5): 1429–1445, 2001.

[39] Oliver Jenkinson and Mark Pollicott. Rigorous effective bounds on the Hausdorff dimension of continued fraction Cantor sets: a hundred decimal digits for the dimension of E_2. *Advances in Mathematics*, 325:87–115, 2018.

[40] M. H. Jensen, P. Bak, and T. Bohr. Complete devil's staircase, fractal dimension, and universality of mode-locking structure in the circle map. *Physical Review Letters*, 50:1637, 1983.

[41] A. Ya. Khinchin. *Continued Fractions*. Dover Publications, Mineola, NY, 1997. Translation of the third Russian edition; first English edition University of Chicago Press, 1964.

[42] A. Khintchine. Metrische Kettenbruchprobleme. *Compositio Mathematica*, 1:361–382, 1935.

[43] Sergey Khrushchev. A recovery of Brouncker's proof for the quadrature continued fraction. *Publicacions Matemàtiques*, 50(1):3–42, 2006.

[44] R. O. Kuzmin. On a problem of Gauss. *Doklady Akademii Nauk SSSR*, pages 375–380, 1928.

[45] Laurent Laloux, Pierre Cizeau, Jean-Philippe Bouchaud, and Marc Potters. Noise dressing of financial correlation matrices. *Physical Review Letters*, 83:1467, 1999.

[46] L. J. Lange. An elegant continued fraction for π. *American Mathematical Monthly*, 106 (5):456–458, 1999.

[47] D. H. Lehmer and R. E. Powers. On factoring large numbers. *Bulletin of the American Mathematical Society*, 37(10):770–776, 1931.

[48] Robert J. Lemke Oliver and Kannan Soundararajan. Unexpected biases in the distribution of consecutive primes. *Proceedings of the National Academy of Sciences*, 113(31):E4446–E4454, 2016.

[49] Ernő Lendvai. *Béla Bartók: An Analysis of his Music*. Kahn & Averill, London, 2000.

[50] Paul Lévy. Sur les lois de probabilité dont dépendent les quotients complets et incomplets d'une fraction continue. *Bulletin de la Société Mathématique de France*, 57:178–194, 1929.

[51] Paul Lévy. Sur le développement en fraction continue d'un nombre choisi au hasard. *Compositio Mathematica*, 3:286–303, 1936.

[52] Linda Hall Library. Scientist of the day: John R. Newlands. https://www.lindahall.org/about/news/scientist-of-the-day/john-r-newlands/, 2026. Accessed November 2026.

[53] M. Lothaire. *Algebraic Combinatorics on Words*, volume 90 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, 2002.

[54] Jean-Claude Martzloff. *A History of Chinese Mathematics*. Springer, 2006.

[55] Madan Lal Mehta. *Random Matrices*, volume 142 of *Pure and Applied Mathematics*. Elsevier/Academic Press, third edition, 2004.

[56] Hermann Minkowski. Zur Geometrie der Zahlen. In *Verhandlungen des III. Internationalen Mathematiker-Kongresses in Heidelberg*, pages 164–173, 1904.

[57] Richard Moeckel. Geodesics on modular surfaces and continued fractions. *Ergodic Theory and Dynamical Systems*, 2(1):69–83, 1982.

[58] Alexis Monnerot-Dumaine. The Fibonacci word fractal. HAL preprint hal-00367972, 2009.

[59] H. L. Montgomery. The pair correlation of zeros of the zeta function. In *Analytic Number Theory*, volume 24 of *Proceedings of Symposia in Pure Mathematics*, pages 181–193. American Mathematical Society, 1973.

[60] Michael A. Morrison and John Brillhart. A method of factoring and the factorization of F_7. *Mathematics of Computation*, 29(129):183–205, 1975.

[61] A. M. Odlyzko. On the distribution of spacings between zeros of the zeta function. *Mathematics of Computation*, 48(177):273–308, 1987.

[62] OEIS Foundation. Sequence A001605: Indices of prime Fibonacci numbers. The On-Line Encyclopedia of Integer Sequences, https://oeis.org/A001605, 2026. Accessed November 2026.

[63] OEIS Foundation. Sequence A086785: Primes found among the numerators of the continued fraction rational approximations to Pi. The On-Line Encyclopedia of Integer Sequences, https://oeis.org/A086785, 2026. Accessed November 2026.

[64] OEIS Foundation. Sequence A086788: Primes found among the denominators of the continued fraction rational approximations to Pi. The On-Line Encyclopedia of Integer Sequences, https://oeis.org/A086788, 2026. Accessed November 2026.

[65] OEIS Foundation. Sequence A096650: Indices of prime Pell numbers. The On-Line Encyclopedia of Integer Sequences, https://oeis.org/A096650, 2026. Accessed November 2026.

[66] Vasiliki Plerou, Parameswaran Gopikrishnan, Bernd Rosenow, Luís A. Nunes Amaral, and H. Eugene Stanley. Universal and nonuniversal properties of cross correlations in financial time series. *Physical Review Letters*, 83:1471, 1999.

[67] Bernhard Riemann. Ueber die Anzahl der Primzahlen unter einer gegebenen Grösse. *Monatsberichte der Königlich Preußischen Akademie der Wissenschaften zu Berlin*, November 1859.

[68] Teddy Rook. Where music, patterns, math and chemistry meet. Teddy Rook Book, https://teddyrookbook.com/connections/, 2026.

[69] Teddy Rook. The rhythm of numbers. Teddy Rook Book, https://teddyrookbook.com/connections/rhythm-of-numbers/, 2026.

[70] Teddy Rook. The Continued-Fraction Suite: six pieces for piano, one number each. Score, Teddy Rook Book, https://teddyrookbook.com/music/continued-fraction-suite/, 2026.

[71] Raphaël Salem. On some singular monotonic functions which are strictly increasing. *Transactions of the American Mathematical Society*, 53(3):427–439, 1943.

[72] D. Shechtman, I. Blech, D. Gratias, and J. W. Cahn. Metallic phase with long-range orientational order and no translational symmetry. *Physical Review Letters*, 53:1951, 1984.

[73] V. T. Sós. On the distribution mod 1 of the sequence $n\alpha$. *Annales Universitatis Scientiarum Budapestinensis de Rolando Eötvös Nominatae, Sectio Mathematica*, 1:127–134, 1958.

[74] M. A. Stern. Ueber eine zahlentheoretische Funktion. *Journal für die reine und angewandte Mathematik*, 55:193–220, 1858.

[75] S. Świerczkowski. On successive settings of an arc on the circumference of a circle. *Fundamenta Mathematicae*, 46:187–189, 1958.

[76] P. Szüsz. Verallgemeinerung und Anwendungen eines Kusminschen Satzes. *Acta Arithmetica*, 7:149–160, 1962.

[77] Salvatore Torquato and Frank H. Stillinger. Local density fluctuations, hyperuniformity, and order metrics. *Physical Review E*, 68:041113, 2003.

[78] Salvatore Torquato, Ge Zhang, and Matthew de Courcy-Ireland. Uncovering multiscale order in the prime numbers via scattering. *Journal of Statistical Mechanics: Theory and Experiment*, page 093401, 2018. doi:10.1088/1742-5468/aad6be.

[79] Godfried T. Toussaint. The Euclidean algorithm generates traditional musical rhythms. In *Renaissance Banff: Mathematics, Music, Art, Culture (Proceedings of Bridges 2005)*, pages 47–56, 2005.

[80] H. von Mangoldt. Zu Riemanns Abhandlung "Ueber die Anzahl der Primzahlen unter einer gegebenen Grösse". *Journal für die reine und angewandte Mathematik*, 114:255–305, 1895.