

“Random” Gaps. The statistics of nearest-neighbor spacings range from random to uniform (<’s indicate spacings too close for the figure to resolve). The second column shows the primes from 7,791,097 to 7,791,877. The third column shows energy levels for an excited heavy (Erbium) nucleus. The fourth column is a “length spectrum” of periodic trajectories for Sinai billiards. The fifth column is a spectrum of zeroes of the Riemann zeta function. (Figure courtesy of Springer-Verlag New York, Inc., “Chaotic motion and random matrix theories” by O. Bohigas and M. J. Giannoni in *Mathematical and Computational Methods in Nuclear Physics*, J. M. Gomez et al., eds., *Lecture Notes in Physics*, volume 209 (1984), pp. 1–99.)

A Prime Case of Chaos

Is one of the deepest problems in number theory tied to one of the most difficult subjects in modern physics? More and more researchers from both disciplines think—indeed, *hope*—so. The possibility has led to a surprising collaboration of physicists pursuing the implications of quantum chaos and mathematicians hunting for a proof of a famous conjecture in number theory known as the Riemann Hypothesis.

Quantum chaos is a relative newcomer in theoretical physics. “It’s a problem of extreme mathematical subtlety and complexity,” says Michael Berry, a “quantum chaologist” at the University of Bristol. “There’s a lot that’s understood, and a lot that’s not understood. One particular thing that’s lacking is a good model which has the essence of the general behavior.”

As its name implies, quantum chaos combines two revolutionary developments in modern physics: quantum mechanics and chaos theory. Quantum mechanics is an esoteric theory of particles and waves, developed by physicists in the 1920’s to account for a body of mysterious phenomena that had turned up when experimentalists started exploring the tiny world of atoms. It is currently viewed as the “true” description of nature, usurping the position held for 250 years by Isaac Newton’s “classical” mechanics. Its rules are very precise and its results have so far proved highly accurate, but much of quantum mechanics still seems counterintuitive, even to physicists steeped in the subject.

Chaos, on the other hand, is a familiar phenomenon. Its rise to prominence in the 1980’s stemmed from the realization that disorderly systems tend to be disorderly in an *orderly* fashion—and that the underlying order within chaotic phenomena can be revealed by a combination of mathematical analysis and computer simulations. Mathematically, chaos is caused by “nonlinearities” in the equations that describe a dynamical system. “Nonlinearity” means that changes in output are not necessarily proportional to changes in input (see Figure 1). Not every nonlinear system is chaotic, of course; if that were the case, science would have essentially no predictive power. But where chaos does occur (or seems to occur)—in long-range weather patterns, for example—it puts a limit on what can be predicted, and consequently calls into question the usefulness of the equations.

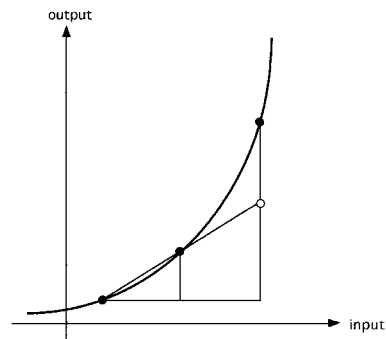


Figure 1. *In a nonlinear system, change in output is not in proportion to change in input.*

Quantum chaos is seemingly a contradiction in terms.

Classical mechanics abounds in nonlinear systems that exhibit chaos. The “double pendulum” is one such system. The classical picture of electrons orbiting an atomic nucleus is another. In computer simulations, the equations of these systems produce solutions that seem to wander aimlessly about, even though everything about the problem is completely deterministic.

You might expect that the weirdness of quantum mechanics would only compound the chaotic aspects of classically chaotic systems. But that’s not what happens—what happens is even more bizarre.

In a nutshell, quantum mechanics *forbids* chaos. Quantum chaos is seemingly a contradiction in terms.

That’s because quantum mechanics is inherently linear. At the heart of quantum mechanics is the notion that everything in nature is a sum, or superposition, of certain “base states.” The base states (also called *eigenstates*) behave in very simple ways, and the behavior of the system is just the sum of these simple components. The only things that are subject to change are the relative proportions of the various eigenstates in the superposition. These changes are made by a mathematical object called an operator—so called because it acts, or operates, on the states of the system. (When there are only finitely many base states, operators are also called matrices.) The whole theory is set up so that changes in output are proportional to changes in input. There’s no nonlinear crack through which chaos can slip in.

Or so it would seem.

But in fact quantum chaos *does* exist. It’s just a little different from what it first sounds like.

An essential feature of any quantum system is a set of numbers called eigenvalues, which characterize the simple behavior of the eigenstates. For example, in the fundamental equation of quantum mechanics known as the Schrödinger equation, the eigenvalues are positive real numbers that correspond to energy levels. The eigenstate with smallest eigenvalue is called the ground state, the one with next smallest eigenvalue is called the first excited state, and so forth. Mathematically, there are infinitely many such states, each with its own eigenvalue. The collection of eigenvalues is called the *spectrum* of the system.

For some systems—say an electron trapped inside a rectangular box—the spectrum is easy to calculate. But these are systems whose classical analogues are non-chaotic. The electron in a box

Trace Elements

Physicist Martin Gutzwiller at IBM's T. J. Watson Research Center didn't know he was doing number theory in the 1960's. But neither did Atle Selberg know he was setting the stage for quantum chaos in the 1950's.

Selberg, a native of Norway now at the Institute for Advanced Study in Princeton, New Jersey, was studying number-theoretic implications of the analytic structure of certain curved spaces. (Curiously, Selberg's most famous result is an "elementary" proof of the Prime Number Theorem—a proof that does *not* use any analytic properties of the Riemann zeta function. He and the Hungarian mathematician Paul Erdős found the first elementary proofs in 1948.) Selberg derived an equation that had the eigenvalues of a differential operator on one side and the lengths of closed curves in the space on the other side. This equation, he showed, "encodes" the number-theoretic properties that underlie the structure of the curved space.

Selberg's equation has come to be known as the Selberg Trace Formula. Number theorists have written hundreds of papers and devoted dozens of conferences to understanding and generalizing it. For Gutzwiller, it represents a Rosetta stone connecting quantum chaos with rigorous results in mathematics.

Dennis Hejhal, an expert on the trace formula at the University of Minnesota and the University of Uppsala in Sweden, has conducted extensive numerical experiments, computing eigenvalues and closed curves for various curved spaces. In one of his early calculations, Hejhal thought he might be on to a proof of the Riemann Hypothesis: Among the eigenvalues for one of Selberg's trace formulae, he saw numbers that he recognized as zeroes of the Riemann zeta function. Closer inspection revealed that they really *were* zeta zeroes. Moreover, the calculation was producing *all* of the zeroes of the zeta function. Alas, it was all a mirage. The zeta zeroes had crept into the computation by accident: The curved space had a certain "singular" point which required a minor adjustment in the differential operator; once the adjustment was made, the zeta zeroes disappeared from the spectrum.

Selberg, who recently turned 80, says he has never seriously tried to prove the Riemann Hypothesis. He's worked on the problem of the zeroes, he says, "but always with a somewhat more limited goal in mind." In fact, he adds, "there have probably been very few attempts at proving the Riemann Hypothesis, because, simply, no one has ever had any really good idea for how to go about it!"



Atle Selberg. (Photo courtesy of the Institute for Advanced Study.)

For some systems the spectrum is easy to calculate.

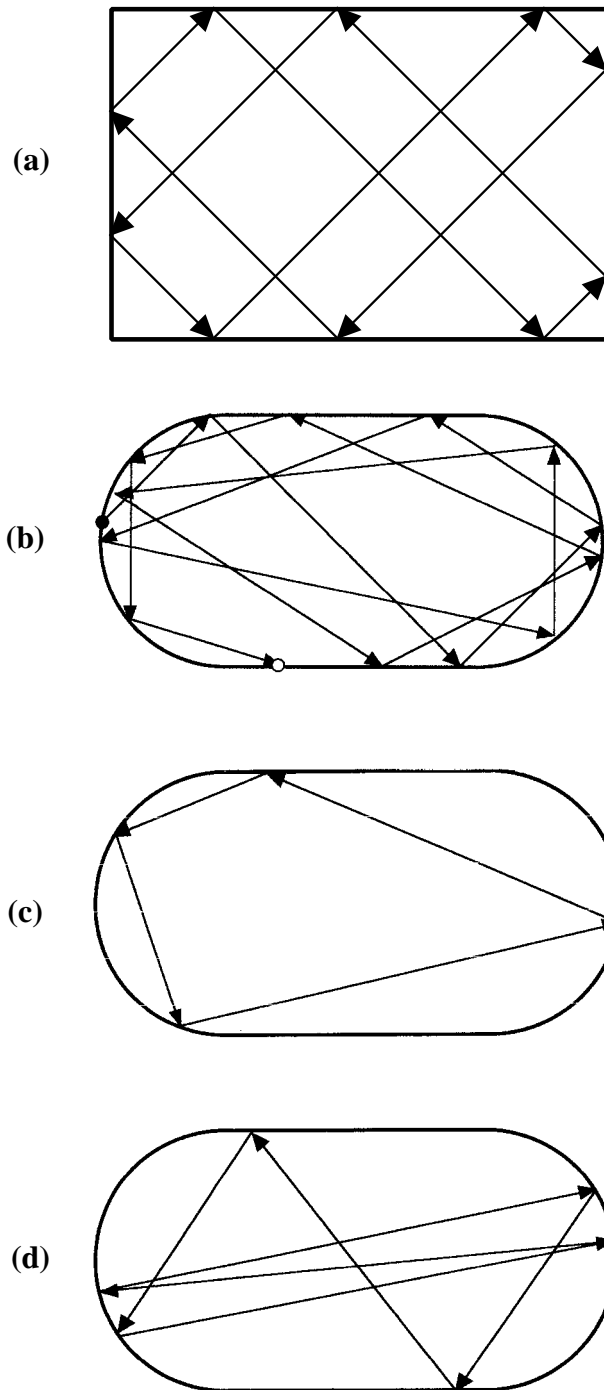


Figure 2. Trajectories on a standard, rectangular billiard table (a) are easy to analyze. The case of “stadium billiards” (b) – (d) is provably chaotic. (Figures (c) and (d) courtesy of Steve Tomsovic.)

behaves much like an ideal billiard ball on a perfect pool table: The eigenstates of the electron correspond to periodic trajectories of the billiard ball, and for a rectangular table, these are easy to compute (see Figure 2a). Put the same electron inside a “stadium”—a rectangular box capped at each end by semicircles (see Figure 2b)—and all hell breaks loose. That’s because the classical dynamics of a billiard ball rattling around inside such a shape is inherently chaotic.

Roughly speaking, quantum chaos is concerned with the spectrum of a quantum system when the classical version of the system is chaotic. Physicists such as Berry and John Keating, also at the University of Bristol, Achim Richter at the Technical University of Darmstadt, and Oriol Bohigas at the University of Paris believe—

All That Jazz

Prime numbers are music to Michael Berry’s ears.

Berry, a theoretical physicist at the University of Bristol, is one of the leading theorists in the study of quantum chaos. And that’s brought him to a keen appreciation of the Riemann zeta function.

Prime numbers are a lot like musical chords, Berry explains. A chord is a combination of notes played simultaneously. Each note is a particular frequency of sound created by a process of resonance in a physical system, say a saxophone. Put together, notes can make a wide variety of music—everything from Chopin to Spice Girls. In number theory, zeroes of the zeta function are the notes, prime numbers are the chords, and theorems are the symphonies.

Of course chords need not be concordant; a lot of vibrations are nothing more than noise. The Riemann Hypothesis, however, imposes a pleasing harmony on the number-theoretic, zeta-zero notes. “Loosely speaking, the Riemann Hypothesis states that the primes have music in them,” Berry says.

But Berry is looking for more than a musical analogy; he hopes to find the actual instrument behind the zeta function—a mathematical drum whose natural frequencies line up with the zeroes of the zeta function. The answer, he thinks, lies in quantum mechanics. “There are vibrations in classical physics, too,” he notes, “but quantum mechanics is a richer, more varied source of vibrating systems than any classical oscillators that we know of.”

What if someone finds a counterexample to the Riemann Hypothesis? “It would destroy this idea of mine,” Berry readily admits—one reason he’s a firm believer in Riemann’s remark. A counterexample would effectively end physicists’ interest in the zeta function. But one question would linger, he says: “How could it be that the Riemann zeta function so convincingly mimics a quantum system without being one?”

and have begun to marshal evidence—that there is a subtle interplay between the mathematical description of chaotic classical systems and the spacing of energy levels in their quantum-mechanical counterparts.

In the 1960's, Martin Gutzwiller, a physicist at the IBM T.J. Watson Research Center in Yorktown Heights, New York, proposed that classical chaos and quantum-mechanical eigenvalues are related by an equation called a *trace formula*. On one side of the equation is a combination of numbers representing the lengths of closed orbits in the classical system. On the other side is a combination of the eigenvalues of the quantum-mechanical version of the system. Gutzwiller's theory is that the two sides of the equation, which use such different data, really are equal.

This equality has profound consequences. By studying closed orbits in a purely classical (albeit chaotic) setting, physicists can deduce properties of a quantum-mechanical system without having to solve the system's Schrödinger equation—which is, in general, beyond any computer's ability anyway. In particular, the trace formula has implications for the spacing between energy levels: For many systems, it “discourages” eigenvalues from being too close together. When physicists draw histograms of the distribution of energy spacings, they find a curve quite different from the familiar, exponentially decreasing “Poisson” distribution (see Figure 3).

Gutzwiller's theory of the trace formula didn't come entirely out of the blue. Physicists have used various types of trace formulas over the years; Gutzwiller's formula itself was derived from an

Euclid's Finest: A Prime-Time Proof

Euclid is best known for geometry, but he also devoted a book to the rudiments of number theory. His most memorable result there is a proof that there are infinitely many prime numbers.

To say there are infinitely many prime numbers is to say that every finite list of primes omits at least one prime number. That's the tack that Euclid took.

Suppose you have a finite list of primes, say p_1, p_2, \dots, p_n . Multiply them together and add 1, Euclid tells us. The result, $p_1 p_2 \cdots p_n + 1$, is a number not divisible by any of the primes in the list, since it leaves the remainder 1 under each such division. But *every* number is divisible by *some* prime number. Thus your list omits at least one prime—and so the supply of prime numbers is infinite.

approach to solving the Schrödinger equation known as Feynman's path integral (introduced by the physicist Richard Feynman). What Gutzwiller didn't know—at the time—was that mathematicians had also developed an extensive theory of trace formulas for a completely different purpose: number theory.

Mathematicians have long been fascinated by the existence of prime numbers—numbers such as 2, 3, 5, etc., which have no factors except themselves and 1. One of the oldest facts about primes dates back to Euclid, around 300 BC: There are infinitely many of them. Euclid's theorem is easily proved (see “Euclid's Finest: A Prime Time Proof”), but it leaves another question unanswered: If you look only at the numbers up to, say N , how many primes do you find?

Number theory reached the first of many peaks in 1896, with a proof of the celebrated Prime Number Theorem, which states that the number of primes up to any number N is approximately $N/\ln N$, where $\ln N$ is the natural logarithm of N . The Prime

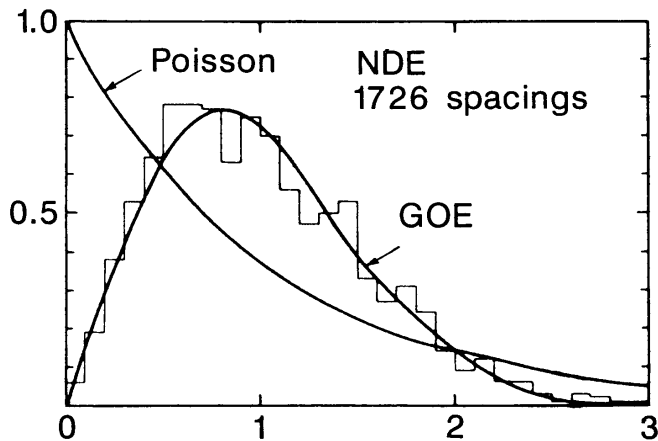


Figure 3. The nearest-neighbor spacings for a nuclear data ensemble (NDE) are similar to the statistics of a set of random matrices called the Gaussian orthogonal ensemble (GOE). The Poisson distribution is shown for comparison. (Figure courtesy of Oriol Bohigas, Institut de Physique Nucleaire, Orsay, France, and Kluwer Academic Publishers, O. Bohigas, R. U. Haq, and A. Pandey in “Nuclear Data from Science and Technology, K. H. Bockhoff, ed., Reidel, Dordrecht 1983, Figure 1, p. 809. With kind permission from Kluwer Academic Publishers.)

Number Theorem was proved independently by the French mathematician Jacques Hadamard and the Belgian mathematician Charles-Jean de la Vallée Poussin. The formula $N/\ln N$ is good to within about 5%, for example, in estimating the number of primes up to a billion. (The formula predicts about 48 million; the exact count is 50,847,534.)

The proof of the Prime Number Theorem actually implies other, more complicated formulas which are even more accurate. The main one says that the number of primes up to N is approximately equal to what's known as the logarithmic integral of N , which is defined as the area beneath the curve $y = 1/\ln x$ between $x = 2$ and $x = N$ (see Figure 4). A succinct way to state the Prime Number Theorem is $\pi(N) \approx \text{Li}(N)$, where π means "the number of primes" (up to N) and "Li" is short for "logarithmic integral." (The wavy equal sign denotes an approximation. In this case, it means that the ratio $\pi(N)/\text{Li}(N)$ gets closer and closer to 1 as N gets larger and larger.) If N is a billion, this approximation is off by only three and a third *thousandths* of a percent.

Computational evidence suggests that $\text{Li}(N)$ is an extremely good approximation to $\pi(N)$. Mathematicians have theoretical reasons to believe that the relative error incurred by the logarithmic integral is bounded by some constant multiple of $(\ln N)^2/\sqrt{N}$. But despite a century of intense study, they don't have a proof that the error actually stays that small.

What stands in the way is a "minor technicality" known as the Riemann Hypothesis. In mathematics, such seemingly insignificant details are often the acorns from which mighty theories grow. That's certainly been the case for the Riemann Hypothesis.

The Riemann Hypothesis concerns a mathematical object called the Riemann zeta function. The zeta function is defined by summing

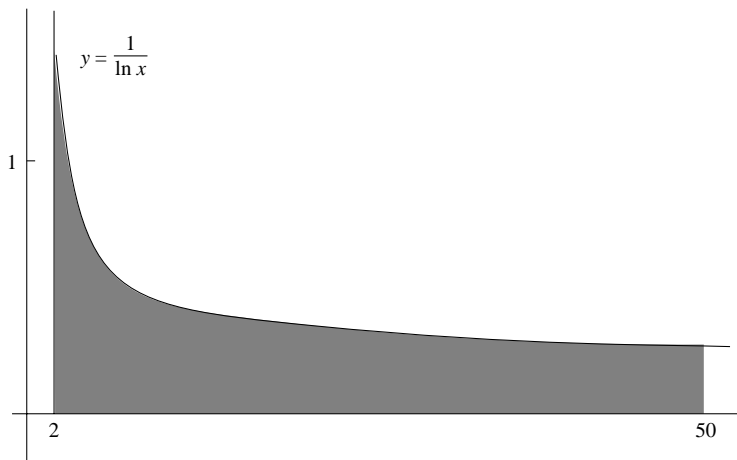


Figure 4. The number of primes up to N can be estimated by measuring the area beneath the curve $y = 1/\ln x$.

inverse powers of the positive integers:

$$\zeta(s) = 1^{-s} + 2^{-s} + 3^{-s} + 4^{-s} + \dots$$

For example, $\zeta(2) = 1 + \frac{1}{4} + \frac{1}{9} + \frac{1}{16} + \dots$. This sum, a bit surprisingly, turns out to equal $\pi^2/6$. The zeta function itself was introduced in the eighteenth century by the Swiss mathematician Leonhard Euler (who found the value $\pi^2/6$ for $\zeta(2)$), but the name Riemann has stuck because the nineteenth-century German mathematician Bernhard Riemann first showed that the properties of the zeta function are intimately intertwined with the distribution of prime numbers.

Riemann, who was primarily a mathematical physicist and geometer (he created much of the mathematics that later went into Einstein's theory of general relativity) wrote only one paper on number theory, an eight-page memoir published in 1859, titled *Über die Anzahl der Primzahlen unter einer gegebenen Grösse* ("On the Number of Prime Numbers Less Than a Given Magnitude"). In his paper, Riemann took Euler's infinite series and applied to it the then-new tools of complex analysis. He showed that the Prime Number Theorem could be proved by deriving certain complex properties of the zeta function. In short, Riemann set the stage for Hadamard and de la Vallée Poussin—but also for a century of mathematics to follow.

Riemann's key contribution was an equation linking two ways of looking at the zeta function. Euler had discovered one side of Riemann's equation. He had shown that the zeta function—initially defined as a *sum* over the positive integers—can be rewritten as a *product* over the prime numbers:

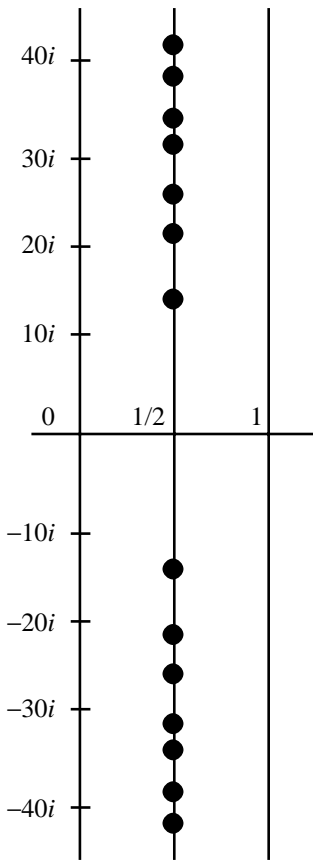
$$\zeta(s) = \frac{1}{(1 - 2^{-s})(1 - 3^{-s})(1 - 5^{-s})(1 - 7^{-s}) \dots}$$

What Riemann observed (and Hadamard later proved) was that the zeta function can also be written as a product over its *zeroes* in the complex plane:

$$\zeta(s) = f(s)(1 - s/\rho_1)(1 - s/\rho_2)(1 - s/\rho_3) \dots,$$

where ρ_1, ρ_2 , etc. are the complex numbers for which $\zeta(\rho) = 0$, and $f(s)$ is a fairly simple "fudge factor." Riemann showed that by

Riemann's key contribution was an equation linking two ways of looking at the zeta function.



equating the logarithms of these two expressions for the zeta function, it's possible to derive not just the approximation $\text{Li}(N)$ for the number of primes up to N , but a whole sequence of increasingly accurate approximations—in effect, an *exact* formula for $\pi(N)$.

Riemann's equation implies that the distribution of prime numbers can be “heard” in the zeroes of the zeta function, much as a musician can “hear” the sound of a symphony by looking at the notes on a printed score. The zeta function, it turns out, has one batch of zeroes at the negative even integers, -2 , -4 , -6 , etc. Number theorists refer to these as the “trivial” zeroes of the zeta function.

What's interesting are the “nontrivial” zeroes.

Riemann showed that all the nontrivial zeroes lie in a particular part of the complex plane: the infinite strip lying above and below the unit interval from $s = 0$ to $s = 1$ (see Figure 5). Furthermore, the zeroes are *symmetrically* arranged, so that each zero above the unit interval has a mirror image below, and each zero (if any!) in the right half of the strip has a mirror image in the left half. Riemann's paper includes a formula that estimates the number of zeroes that lie above the real axis and below any given height.

It turns out that the Prime Number Theorem is equivalent to the assertion that no (nontrivial) zero of the zeta function lies on the boundary of the “critical strip.” Furthermore, the relative error of

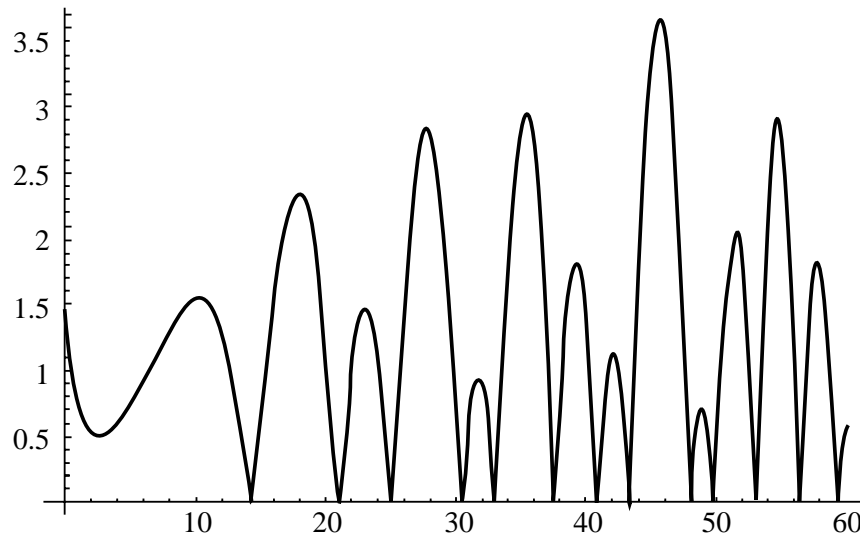


Figure 5. Zeroes of the Riemann zeta function (dark dots, top) lie in a strip. The Riemann Hypothesis says they all lie on the vertical line through $s=1/2$. The graph $|\zeta(\frac{1}{2} + it)|$ is plotted (bottom) for t between 1 and 60.

the approximation $\pi(N) \approx \text{Li}(N)$, depends only on how close to the boundary the nontrivial zeroes ever get. And here is where things get dicey.

As far as anyone knows, there may be zeroes of the zeta function extremely close to the boundary. If that's the case, then the relative error term in the Prime Number Theorem is rather large. (It doesn't mean the theorem is wrong, of course; it just means that the logarithmic integral only guarantees a good approximation to $\pi(N)$ when N is extremely large.) If, on the other hand, all the zeroes are as far from the boundary as they can be—that is, if they all lie on the “critical line” above and below $s = 1/2$ —then the relative error term is as small as it can possibly be, which turns out to be the constant multiple of $(\ln N)^2/\sqrt{N}$ which we spoke of earlier.

In his paper, Riemann called it “very likely” that all the zeroes of the zeta function lie on the critical line. This is the celebrated Riemann Hypothesis.


If true, the Riemann Hypothesis has profound implications not just for the error term in the Prime Number Theorem, but in many other parts of mathematics. The mathematical literature is rife with theorems that say “If the Riemann Hypothesis is true, then....” Many other published papers go to great lengths to prove theorems that would be simple corollaries of the Riemann Hypothesis.

Harold Stark, a number theorist at the University of California at San Diego, jokes that a proof of the Riemann Hypothesis “would make lots of my life's work irrelevant.” Nevertheless, he says, “the idea that you can presumably correctly conjecture that infinitely many numbers are on a particular line, and you can't prove it, is frustrating beyond any description. It's just unacceptable!”

But what's any of this got to do with quantum chaos? The answer lies in the two ways of looking at the zeta function: Riemann's equation is a trace formula. Physicists think it's the mother of all trace formulas for quantum chaos. They believe the zeroes of the zeta function can be interpreted as energy levels in the quantum version of some classically chaotic system. If they're right, the Riemann Hypothesis *must* be true.

There's a lot more than wishful thinking and soft analogies behind this line of reasoning. Mathematicians have collected a vast amount of data on the zeroes of the zeta function. So far, the data show every indication of perfect agreement with the statistics of energy levels in quantum chaos. That's part of the appeal for

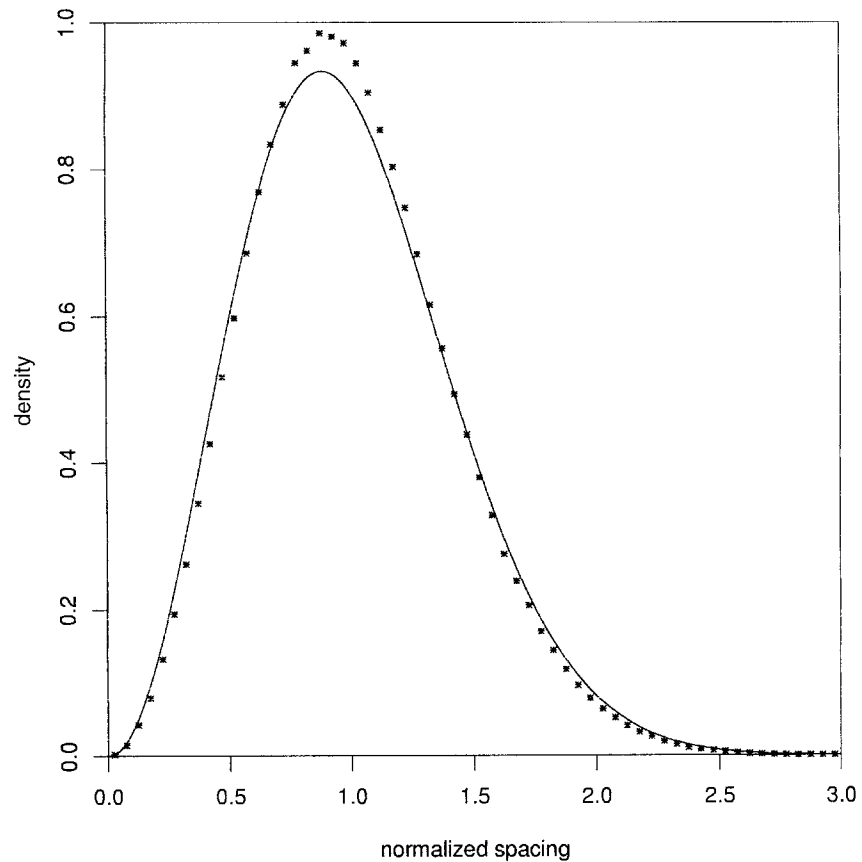
In his paper, Riemann called it “very likely” that all the zeroes of the zeta function lie on the critical line.



Andrew M. Odlyzko. (Photo courtesy of Andrew Odlyzko.)

physicists: The eigenvalues of the Schrödinger equation for actual quantum systems are hard to compute, and extremely hard to measure experimentally, so many of the predictions of quantum chaos are difficult to test. But zeroes of the zeta function are relatively easy to come by; number theorists can churn out tens of millions of them on demand.

The chief supplier of zeta zeroes is Andrew Odlyzko at AT&T Labs in Florham Park, New Jersey. Odlyzko has developed extremely efficient algorithms that make it possible to compute hundreds of millions of zeroes at various “heights” in the critical strip. Needless to say, they’ve all, so far, been found to satisfy the Riemann Hypothesis. (A single counterexample would be enough to ruin the Riemann Hypothesis. It would also end much of the interest physicists have in the zeta function.)



Actually, Odlyzko's computer has reported a number of counterexamples, but every one (so far) has been traced to a hard- or software glitch. "The Riemann Hypothesis is very sensitive to any kind of error," Odlyzko explains. "The slightest mistake, especially with the methods I use, produces counterexamples."

In purely mathematical terms, Odlyzko's computations indicate that the spacings between consecutive zeroes of the zeta function behave, statistically, like the spacings between consecutive eigenvalues of large, random matrices belonging to a class known as the Gaussian Unitary Ensemble (see Figure 6). It was in precisely this context that the zeta function first caught the eye of physicists, in the early 1970's.

Random matrices had been proposed as a way of studying the energy levels of large nuclei. The idea originated with Eugene Wigner, a theoretical physicist at Princeton University. (Wigner is best known to mathematicians for his 1960 essay, "The Unreasonable Effectiveness of Mathematics in the Natural

A single counterexample would be enough to ruin the Riemann Hypothesis.

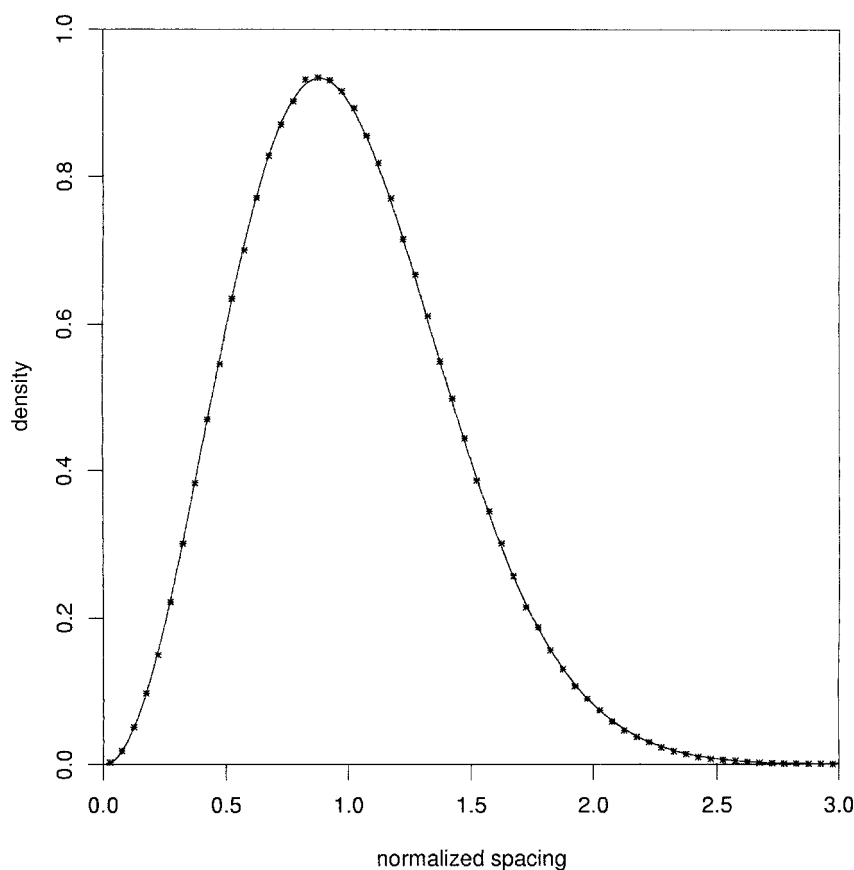


Figure 6. A comparison of the statistics of spacings between consecutive eigenvalues of random matrices (solid curve) and nearest-neighbor spacings of zeroes of the Riemann zeta function (data points). The fit is not bad for the first million zeroes (left), but nearly perfect for the spacings among 1,041,600 zeroes near the 2×10^{20} th zero. (Figure courtesy of AT&T.)

Sciences.” The zeta–chaos connection would seem to be a case in point.) More recently Oriol Bohigas and colleagues at the *Institut de Physique Nucléaire* in Orsay, France, have brought the random-matrix interpretation under the umbrella of quantum chaos. But during the 1960’s, the mathematics of random matrices was intensively studied in its own right by theoretical physicists, including Freeman Dyson at the Institute for Advanced Study in Princeton, New Jersey.

In 1972, Hugh Montgomery, a number theorist at the University of Michigan, was visiting the Institute for Advanced Study. Montgomery had been studying the distribution of zeroes of the zeta function, in hopes of gaining insight into the Riemann Hypothesis. He was able to prove that the Riemann Hypothesis had implications for the spacing of zeroes along the critical line, but his key discovery was an additional property that the zeroes seemed to have, one which implied a particularly nice formula for the average spacing between zeroes.

During tea one day at the Institute, Montgomery was introduced to Dyson and described his conjecture. Dyson immediately recognized it as the same result as had been obtained for random matrices.

“It just so happened that he was one of the two or three physicists in the world who had worked all of these things out, so I was actually talking to the greatest expert in exactly this!” Montgomery recalls.

Odlyzko’s computations agree amazingly well with Montgomery’s conjecture, notes Peter Sarnak, a number theorist at Princeton University. But they do much more than that, Sarnak insists: They are “the first phenomenological insight that the zeroes are absolutely, undoubtedly ‘spectral’ in nature.” Riemann himself would be impressed, Sarnak says.

Sarnak and Nicholas Katz, also at Princeton, have found additional evidence for a spectral interpretation of the zeta zeroes. The Riemann zeta function is just one of many number-theoretic zeta functions, each of which has an analogue of the Riemann Hypothesis. Sarnak and Katz have analyzed in detail the relation between random matrices and a class of zeta functions for which the Riemann Hypothesis has actually been proved. (Unfortunately, the proofs for these zeta functions seem to offer no clue as to a

Odlyzko’s computations agree amazingly well with Montgomery’s conjecture.

proof of “the” Riemann Hypothesis.) They have shown that the key statistics for the spacing of eigenvalues and zeroes of their zeta functions are *exactly* the same.

Of possibly greater importance, the key statistics for eigenvalues are relatively robust: Sarnak and Katz have shown that the statistics are the same for several different classes of random matrices. That’s good news for physicists. It suggests that the assumptions they make in choosing one class of random matrix over another don’t affect the properties they’re interested in.

Despite the stunning advances linking Riemann’s zeta function to twentieth-century physics, no one is predicting an imminent proof of the Riemann Hypothesis. Odlyzko’s numerical experiments and the evidence amassed by physicists have convinced everyone that a spectral interpretation of the zeta zeroes is the way to go, but number theorists say they are at least one “big idea” away from even the beginnings of a proof. Mathematicians aren’t yet sure what to aim at, says Sarnak: “What we really have to do with the Riemann Hypothesis is put it in the ballpark.”

Of course that was how things stood not long ago for Fermat’s Last Theorem (see *What’s Happening in the Mathematical Sciences*, volumes 2 and 3). In the early 1980’s no one had any idea how to prove Fermat’s famous marginal remark. Then, in the mid 80’s, number theorists discovered a deep connection between Fermat’s Last Theorem and a branch of modern mathematics known as the theory of elliptic curves. Within ten years, this discovery had culminated in Andrew Wiles’s solution of the 350-year-old problem.

Interestingly, the denouement of Fermat’s Last Theorem was preceded by a conference, held in 1980, on recent developments in the mathematics of Fermat’s Last Theorem. Number theorists are hoping that history repeats itself: In 1996, the recently founded American Institute of Mathematics (AIM) sponsored a conference devoted to the zeta function and its connections with quantum chaos. If everything goes according to “plan,” the first breakthrough idea will come in 2001, and a proof of the Riemann Hypothesis will be published in 2011. Stay tuned.

Despite the stunning advances linking Riemann’s zeta function to twentieth-century physics, no one is predicting an imminent proof of the Riemann Hypothesis.
