

THE IN A HURRY SERIES

Prime Numbers

in a Hurry

The atoms of arithmetic

Max Brocklesby

BOOKSINAHURRY.COM

Contents

The Whole Thing in One Page 4

Why You Should Care 6

The Core Ideas 8

How It Actually Works 20

What People Get Wrong 30

Use It 35

Terms 40

Go Deeper 44

Notes and Sources 46

Bibliography 53

The Whole Thing in One Page

Prime numbers are usually introduced as the awkward integers that refuse to divide cleanly. That makes them sound like leftovers from school arithmetic: useful for factor trees, divisibility tests and little else. The opposite is closer. Primes are the pieces from which every integer above one is built under multiplication, and the assembly instructions are unique. Twelve is $2 \times 2 \times 3$. No different collection of primes will make twelve. Change the operation or the number system and the metaphor needs care, but inside the positive integers these are arithmetic's atoms.

That rigid foundation produces a puzzle. Multiplication is perfectly organised by primes, yet their positions along the number line look disorderly. Euclid proved that the supply never ends. The sieve associated with Eratosthenes finds them by crossing out multiples. The prime number theorem says how quickly they thin: near a large number x , their average density is about $1/\log x$. It predicts the crowd with remarkable accuracy while remaining silent about the next individual arrival.

Remainders expose a second role. Arithmetic modulo a prime has no broken division among its non-zero values. Every non-zero residue has a multiplicative inverse, so equations behave cleanly and Fermat's little theorem appears. Replace the prime modulus with a composite one and zero divisors enter: two non-zero values can multiply to zero. Primes are therefore building blocks in one view and complete little arithmetical worlds in another.

Euler connected the two views by rewriting the zeta function as a product over every prime. Riemann then found that the deviations of prime counts from their smooth average are governed by the zeros of that function. His hypothesis that all non-trivial zeros lie on one vertical line remains unproved. It would control the fluctuations far more tightly. It would not print the next prime on demand.

The irregularity still contains patterns. Every arithmetic progression whose start and step are coprime contains infinitely many primes. Primes contain

progressions of any finite length. There are arbitrarily long deserts with none, yet published work proves that infinitely many consecutive prime pairs differ by at most 246. A 31 August 2026 preprint claims 240; no peer-reviewed publication was found when checked. Whether the gap two occurs infinitely often is the twin-prime conjecture. Whether every even integer above two is a sum of two primes is the strong Goldbach conjecture. Both statements are easy to understand and still beyond proof.

Computation adds the final distinction. Testing whether a huge number is prime can be efficient; finding the factors of a huge composite can be much harder. The largest known prime has more than forty-one million decimal digits and belongs to a Mersenne family with a decisive specialised test. RSA exploits a related asymmetry: multiplying two selected primes is easy, while reversing the product is not known to be efficient on classical computers. That difficulty is an assumption, not a theorem, and a sufficiently capable fault-tolerant quantum computer would change it.

So the atom metaphor contains both the order and the mystery. Primes give multiplication its unique grammar. Their locations resist local prediction. Their remainders create algebraic structure. Their collective distribution reaches into complex analysis. Their easy assembly and difficult recovery connect pure proof to computation. The pieces are elementary. The pattern made from them is not.

That is the book.

Why You Should Care

On 12 October 2024, an NVIDIA H100 in San Antonio completed a Lucas-Lehmer calculation showing that $2^{136,279,841} - 1$ is prime. The number has 41,024,320 decimal digits. Reading them aloud at one per second would take about 475 days without sleep. At roughly 3,000 digits per page, its expansion would exceed 13,600 pages. Yet the claim is exact, and the Great Internet Mersenne Prime Search (GIMPS) confirmed it with independent implementations on different hardware rather than trying every possible divisor.

The record is spectacle, but it reveals the subject's first surprise. A definition understood by a child reaches numbers too large to write down in any practical setting, while proof still reaches them. Prime-number research repeatedly does this. It begins with counting, division and remainders, then arrives at distributed computing, complex functions and open problems that have survived centuries.

Primes also organise ordinary arithmetic more deeply than their rarity suggests. Fractions reduce because numerator and denominator can be broken into prime powers. Greatest common divisors compare those powers. A statement about every integer often becomes a statement about each prime separately, then gets reassembled. Algebra modulo a prime produces a finite field, one of the cleanest structures in mathematics and a foundation for error-correcting codes, computational number theory and several cryptographic systems.

Then there is the question of pattern. Deterministic arithmetic generates every prime, but no comparably simple law gives the next location from the current scale. From far away, their population follows a smooth law. Close up, gaps jump, clusters appear and simple-looking conjectures resist attack. Primes are therefore a laboratory for a distinction that matters across mathematics: knowing the average behaviour of a system is not the same as knowing its next event.

The open problems are unusually democratic at the entrance. You can

understand twin primes by looking at 11 and 13. You can understand Goldbach by writing 100 as $47 + 53$. You can understand a long prime gap by noticing that factorials manufacture consecutive composites. The proofs, where they exist, may demand decades of theory. The question itself often fits on a napkin. Few fields let a newcomer stand so close to the research frontier.

Cryptography supplies the public reason people now recognise. It should not swallow the subject. Prime numbers do not make all encryption work, a large prime is not a security guarantee, and modern systems depend on protocols, randomness and implementation as much as arithmetic. Still, the gap between easy multiplication and hard factorisation is a striking conversion: a property of whole numbers becomes a way to separate public information from private capability.

Prime numbers also show what mathematical progress looks like when there is no single finish line. Euclid settled infinitude completely. Gauss guessed the correct scale of the count, and later mathematicians proved it. Riemann sharpened the question into one that remains open. Computers now certify examples larger than any person could inspect digit by digit. Each achievement answers a different question, and none is made smaller because another survives. That hierarchy is useful to anyone who wants to read mathematics intelligently: distinguish a theorem from a conjecture, an asymptotic law from an exact value, a computational verification from an infinite proof, and an efficient test from an efficient search.

By the end, you should be able to see primes in three ways at once: as the irreducible pieces of multiplication, as moduli that create clean finite arithmetic, and as a sparse deterministic sequence whose global order coexists with local surprise. The reward is larger than a catalogue of famous conjectures. It is a model of how exact structure can produce uncertainty without any randomness entering the rules.

The Core Ideas

Multiplication Has Atoms

A prime number is a positive integer greater than one with exactly two positive divisors: one and itself. Two is prime. Three is prime. Four is not, because 2 divides it. The distinction sounds like a sorting rule, but it becomes structural once multiplication enters.

Take 756. Repeated division gives

$$756 = 2 \times 2 \times 3 \times 3 \times 3 \times 7.$$

The order of those factors can change, but the collection cannot. You may begin by dividing by 7 or by 3; you will end with the same prime powers. This is the fundamental theorem of arithmetic: every integer greater than one can be expressed as a product of primes, and that expression is unique apart from order.

Existence is the easier half. If a number is composite, split it into smaller factors. If any factor is composite, split again. The pieces keep shrinking, so the process must stop at primes. Uniqueness needs a stronger fact, often called Euclid's lemma: if a prime p divides a product ab , then p divides a or p divides b . Suppose one number had two different prime factorisations. A prime on the first side divides the product on the second, so it must match one of the primes there. Cancel the match and repeat. No unmatched prime can survive.

That is why one is excluded. If one counted as prime, uniqueness would collapse into pointless multiplicity. Six could be written as 2×3 , or $1 \times 2 \times 3$, or with any number of extra ones. Calling one a unit rather than a prime keeps the atoms irreducible and the assembly record unique. Zero is neither prime nor composite: every non-zero integer divides it, so it does not fit the definition or the factorisation theorem.

Unique factorisation lets arithmetic questions be moved into exponents. If

$$360 = 2^3 \times 3^2 \times 5,$$

then every divisor of 360 is obtained by choosing an exponent from 0 to 3 for 2, from 0 to 2 for 3, and from 0 to 1 for 5. That gives $(3 + 1)(2 + 1)(1 + 1) = 24$ divisors. The greatest common divisor of two numbers takes the smaller exponent of each shared prime. The least common multiple takes the larger. A fraction is in lowest terms when numerator and denominator share no prime.

The atom metaphor therefore earns its place, with a boundary. It describes positive integers under multiplication. Under addition, one is enough to build every positive integer, so primes have no comparable privilege. Change the number system and factorisation can change. In the Gaussian integers, which include numbers of the form $a + bi$, the ordinary prime 5 splits as $(2 + i)(2 - i)$. Other algebraic systems can even lose unique factorisation. Primes are not indivisible matter in every possible mathematics. They are the irreducible grammar of ordinary integer multiplication.

This first idea creates the book's central tension. The factorisation of each integer is rigid, yet the positions of the primes themselves are not given by an equally transparent pattern. Every composite number carries a unique receipt listing its prime ingredients. The shop shelves on which those ingredients appear remain irregular.

A Prime Modulus Makes Division Work

Clock arithmetic teaches the first step. On a twelve-hour clock, 17 and 5 indicate the same position because their difference is divisible by 12. Mathematicians write $17 \equiv 5 \pmod{12}$. A modulus groups all integers into residue classes according to their remainder after division.

Addition and multiplication survive this compression. Modulo 7, the numbers 10 and 3 are interchangeable, so $10 + 12$ has the same residue as $3 + 5$, and 10×12 has the same residue as 3×5 . An infinite line of integers has become a finite system with seven positions.

Division is where the modulus reveals whether it is prime. Modulo 7, every non-zero residue has a multiplicative inverse. Two times four is eight, which is congruent to one, so four is the inverse of two. Three times five is fifteen,

also congruent to one, so five is the inverse of three. This means an equation such as $3x \equiv 4 \pmod{7}$ has one solution: multiply both sides by five to get $x \equiv 6$.

Modulo 8, the behaviour breaks. Two has no inverse because every multiple of two is even and can never leave remainder one. Worse, $2 \times 4 \equiv 0 \pmod{8}$ even though neither factor is congruent to zero. These are zero divisors. They make cancellation unsafe: from $2 \times 1 \equiv 2 \times 5 \pmod{8}$, you cannot conclude $1 \equiv 5$. Both products are congruent to two, but the original residues differ.

Arithmetic modulo a prime avoids that failure. If p is prime and $ab \equiv 0 \pmod{p}$, then p divides ab . Euclid's lemma says p divides a or b , so one factor was already zero modulo p . No non-zero zero divisors exist. Multiplication by any non-zero class cannot merge two residues; on a finite set it must therefore reach every non-zero class, including one. Every non-zero residue has an inverse, and the residue classes form a finite field.

This clean structure produces Fermat's little theorem. Choose a number a not divisible by a prime p . Multiplying the non-zero residues $1, 2, \dots, p-1$ by a merely rearranges them modulo p , because two products cannot collide without cancellation forcing the original residues to match. Their products are therefore congruent:

$$a^{(p-1)}(p-1)! \equiv (p-1)! \pmod{p}.$$

Since none of the factors in $(p-1)!$ is divisible by p , it can be cancelled, leaving

$$a^{(p-1)} \equiv 1 \pmod{p}.$$

An equivalent form is $a^p \equiv a \pmod{p}$ for every integer a . For $p = 7$ and $a = 3$, the theorem says 3^6 leaves remainder one on division by seven. It does, without requiring the full power to be written out.

The theorem supplies a necessary test for primality, but not a sufficient one. Some composite numbers imitate the congruence for particular bases, and Carmichael numbers imitate it for every base coprime to them. A number

can pass a simple Fermat test and still be composite. That failure matters later, because it separates an elegant theorem from a dependable algorithm.

Prime moduli matter because they turn remainders into algebra rather than a mere bookkeeping trick. Equations can be solved, division by non-zero values is legitimate, polynomial behaviour becomes tractable, and large integers can be studied through many small finite worlds. A prime is therefore an atom when numbers are multiplied and an architectural condition when arithmetic is compressed.

The Supply Never Ends

No table of primes is complete. Euclid proved this more than two thousand years ago with an argument that still fits in a paragraph.

Assume there are only finitely many primes, and list them as $p_1, p_2, \dots, p_n$. Multiply them together and add one:

$$N = p_1 p_2 \dots p_n + 1.$$

Divide N by any prime on the list. The product is divisible by that prime, so N leaves remainder one. None of the listed primes divides N . Yet every integer greater than one is either prime or has a prime divisor. N must therefore be prime itself or possess a prime factor missing from the list. Either result contradicts the claim that the list contained every prime.

The proof does not say that the product plus one is always prime. Begin with 2, 3, 5, 7, 11 and 13. Their product plus one is 30,031, which factors as 59×509 . The argument needs only a new prime factor, not a new prime value of the special form. This is a useful lesson in reading proofs: identify the exact conclusion that creates the contradiction rather than strengthening it into a false slogan.

Infinitude alone says little about abundance. The primes might, in principle, become so sparse that their reciprocal sum converged. Euler showed that they do not. The sum

$$1/2 + 1/3 + 1/5 + 1/7 + 1/11 + \dots$$

diverges, although it does so with painful slowness. Euler's route joined primes to the harmonic series through finite products. Fix a cutoff y and multiply $1 + 1/p + 1/p^2 + \dots$ over the primes p no larger than y . Unique factorisation makes the expansion contain, once each, the reciprocal of every positive integer built only from those primes. In particular, it contains $1 + 1/2 + \dots + 1/y$. Those harmonic partial sums grow without bound.

If the sum of reciprocal primes converged, the logarithms of the product factors would have bounded total and these products would remain bounded. They do not. The primes cannot be too thin for their reciprocal sum to converge. Their density still tends to zero: vanishing proportion and divergent reciprocal sum are compatible.

To find primes rather than prove their existence, use the sieve of Eratosthenes. Write the integers from 2 to a chosen limit. Keep 2 and cross out its larger multiples. Keep the next uncrossed number, 3, and cross out its larger multiples. Continue with 5, then 7. Once the current prime exceeds the square root of the limit, every remaining composite has already been crossed out, because a composite number must have a factor no larger than its square root.

For numbers up to 100, it is enough to sieve with 2, 3, 5 and 7. The survivors are the twenty-five primes up to 100. The method finds a whole interval efficiently because it shares work: crossing out multiples of one prime removes many composites at once. Trial division treats each candidate almost from scratch.

The same idea scales through segmented sieves, which process manageable blocks rather than store every integer up to a vast limit. Modern prime searches use more specialised tests when the candidate has special form, but sieving remains the natural way to enumerate ordinary ranges.

Euclid and Eratosthenes answer different questions. One shows that any claimed final list must fail. The other shows how a bounded list can be produced without testing every divisor of every number. Together they

establish a pattern that recurs throughout number theory: proof can settle an infinite statement while computation still has to organise finite work carefully.

The Crowd Thins by a Law

Count the primes up to a number x and call the result $\pi(x)$. There are 4 primes up to 10, 25 up to 100, 168 up to 1,000, and 78,498 up to one million. Across these scales, the proportion falls. Four out of ten numbers are prime at the first scale; fewer than eight in a hundred are prime by one million.

The prime number theorem describes the fall:

$$\pi(x) \sim x/\log x.$$

The symbol $\sim$ means that the ratio of the two quantities approaches one as x grows. Log denotes the natural logarithm here. It does not say they are equal. At one million, $x/\log x$ is about 72,382, below the exact count by more than six thousand. The relative error is already modest and tends to zero. The logarithmic integral $\text{Li}(x)$, which accumulates $1/\log t$ rather than freezing the density at the endpoint, usually gives a closer estimate.

Its leading term suggests a local density model. Around a large number x , the rough density of primes is about $1/\log x$. Near one million, $\log x$ is about 13.8, so a rough model says one integer in fourteen is prime. Since every prime above two is odd, the rough density among odd candidates is about $2/\log x$, or roughly one in seven there.

This is a statistical statement about a deterministic sequence. No coin is tossed to decide whether 1,000,003 is prime. Its status was fixed before anyone asked. Probability language enters because the average density gives a useful model when a candidate is sampled from a large region. Treating that model as independence would be a mistake. Divisibility creates strong local structure: apart from three itself, no prime is divisible by three; residues modulo small primes rule out patterned subsets; nearby primality events influence one another.

The thinning can be understood heuristically. A large integer avoids

divisibility by 2 with frequency about $1/2$, by 3 with frequency about $2/3$, by 5 with frequency about $4/5$, and so on. Multiplying such survival proportions suggests a logarithmic scale, although dependencies and normalising constants make the naive product insufficient as a proof. Sieving intuition sees why the density falls. The prime number theorem establishes the exact leading law.

The same model suggests average spacing. If the density near x is about $1/\log x$, the average gap between neighbouring primes there is about $\log x$. Near one million, that is around fourteen. Near a 2048-bit number, whose natural logarithm is close to 1,420, the average numerical gap is around 1,420, equivalent to about 710 odd candidates. This is why software can generate a large probable prime by choosing odd candidates and testing them: the search does not usually need millions of attempts.

Average gap is not maximum gap. Primes can appear close together, as 1,000,000,007 and 1,000,000,009 do, then leave a longer desert. There are arbitrarily long runs of composite numbers. For any chosen n , the numbers

$$(n + 1)! + 2, (n + 1)! + 3, \dots, (n + 1)! + (n + 1)$$

are all composite, because each term is divisible by the added integer. The prime number theorem does not rule out such deserts. It controls the total population, not local spacing.

That distinction is the heart of distribution theory. A smooth curve can describe the crowd while individual arrivals remain erratic. The theorem tells you how many primes to expect below a vast threshold. It does not offer an exact schedule. Mathematics has global order here without local regularity.

The Error Has a Hidden Spectrum

The prime number theorem gives the main trend. Number theory then asks how far the true count can wander from its smooth approximation. That error leads to the Riemann hypothesis.

Euler supplied the bridge. For a complex number s with real part greater than one, the zeta function is defined by

$$\zeta(s) = 1 + 1/2^s + 1/3^s + 1/4^s + \dots$$

Unique factorisation lets the same function be written as a product over primes:

$$\zeta(s) = \prod_p (1 - p^{-s})^{-1}.$$

Expand each factor as a geometric series. Choosing $1/p^k$ from the factor for p records that p appears k times. Multiplying choices across all primes creates $1/n^s$ once for every positive integer n , because n has one prime factorisation. The Euler product therefore compresses the full multiplicative structure of the integers into one analytic object.

Riemann extended the function beyond the region where the original series converges and studied its zeros, the values of s for which $\zeta(s) = 0$. A complex number has a horizontal real coordinate and a vertical imaginary coordinate, so these zeros can be plotted as points on a plane. Some occur at negative even integers and are called trivial. The non-trivial zeros lie in a vertical strip. The Riemann hypothesis says that every one of them has real part exactly $1/2$.

Why should zeros of a complex function govern primes on the ordinary number line? Formulas for weighted prime counts have a smooth main term plus oscillations contributed by the zeros. A zero with real part β contributes on a scale tied to x^β , though the exact formula carries weights and logarithmic factors. Its imaginary part supplies an oscillatory frequency. Zeros farther to the right allow larger deviations. If every non-trivial zero lies on the line $1/2$, one standard consequence is that, for sufficiently large x , the error in approximating $\pi(x)$ by $\text{Li}(x)$ is bounded by a constant times $\sqrt{x} \log x$. That is much tighter control of the count, not an address for the next prime.

This resembles a spectrum more than a codebook. The zeros do not label primes one by one. They organise oscillations in the gap between actual weighted counts and average growth. The analogy has limits, but it captures why imaginary parts act like frequencies and real parts govern scale. Knowing the hypothesis would sharpen many bounds because those estimates depend on how strongly prime-related sums can fluctuate.

A common explicit formula uses a weighted count that includes prime powers as well as primes, because logarithmic differentiation of the Euler product brings down powers naturally. The pole at one supplies the main term x . Each non-trivial zero contributes an oscillating correction, while the trivial zeros and other analytic features supply smaller terms. Recovering an unweighted prime count then requires further conversion. This machinery explains why statements about the zeros translate into bounds for primes without pretending that $\zeta(s)$ is a hidden list in disguise.

The hypothesis remains unproved. Enormous numbers of zeros have been checked computationally on the critical line, but no finite verification can settle an infinite statement. A proof must explain why every non-trivial zero, including those beyond any computation, lies there. A counterexample needs only one zero off the line, though finding one could lie beyond present searches.

The hypothesis is also neither the whole of prime theory nor a magic generator. Many important results are known without it. Some theorems have conditional versions that become stronger if it is true. Even a proof would not make the next prime easy to read from a closed expression. It would tighten control over aggregate error and strengthen consequences whose assumptions currently include the hypothesis.

The deeper lesson is that prime positions are studied by transforming the problem. Directly staring at 2, 3, 5, 7, 11 and their successors reveals little. Encoding all integers in $\zeta(s)$, extending that function into the complex plane and examining where it vanishes turns irregular counting into analysis. The primes look local. Their fluctuations are controlled by a global object.

Patterns Survive the Irregularity

Irregular does not mean patternless. It means that each proposed pattern must survive arithmetic obstructions before anyone worries about proof.

Consider an arithmetic progression $a, a + q, a + 2q$, and so on. If a and q share a factor greater than one, every term shares it. The progression $6n + 3$ contains only one prime, namely 3, because every later term is divisible by

three. If a and q are coprime, Dirichlet's theorem says the progression contains infinitely many primes. There are infinitely many primes congruent to one modulo four and infinitely many congruent to three modulo four. Every allowable residue class receives an infinite supply.

Green and Tao proved a different statement: the primes themselves contain arithmetic progressions of every finite length. Somewhere in the sequence are primes $p, p + d, p + 2d, \dots, p + (k - 1)d$ for any chosen k . This does not contradict thinning. A set can have density tending to zero and still contain long internal patterns, provided enough structure survives.

Gaps expose the opposite behaviour. Twin primes differ by two: 11 and 13, 29 and 31, 101 and 103. The twin-prime conjecture says there are infinitely many such pairs. It remains open. Work beginning with Yitang Zhang and sharpened through Maynard-Tao methods and the Polymath8 collaboration established in published work that infinitely many pairs of consecutive primes differ by at most 246. A first-version arXiv preprint posted on 31 August 2026 claims 240. It had not appeared in a peer-reviewed publication when this book was checked, so 246 is the published benchmark and 240 a live claim. Neither is the twin-prime conjecture. Each says the sequence cannot spread out without repeatedly returning to a fixed-width window.

Large gaps are easier to force. The factorial construction creates any desired finite run of composites. Stronger research asks how unusually large natural prime gaps can be relative to the average logarithmic spacing. Small and large gaps coexist: the sequence thins globally, permits deserts of arbitrary finite length and still produces bounded clusters infinitely often.

Additive questions are stranger because primes are defined multiplicatively. The strong Goldbach conjecture says every even integer greater than two is the sum of two primes. Large finite ranges have been checked computationally, and heuristic models support the claim, but neither can replace proof for every even integer. The weak or ternary version, that every odd integer greater than five is the sum of three primes, was proved by Harald Helfgott. One extra prime changes the problem enough to cross the boundary from conjecture to theorem.

Local checks explain why some patterns are plausible. An odd prime above three must be congruent to one or five modulo six. A proposed constellation of linear expressions must avoid covering every residue modulo some prime, otherwise one expression is always divisible by that prime. Such an obstruction can kill an infinite pattern before any deep analysis begins. Passing every local test does not prove the pattern occurs infinitely often, but failing one proves it cannot.

This is why primes can feel random while obeying more structure than independent random points would. Congruences impose exclusions. Analytic estimates govern density. Combinatorial theorems force patterns. Many central conjectures ask whether locally admissible prime patterns occur with the frequencies predicted by refined heuristics.

The right mental model is constrained irregularity. The primes are neither a periodic design nor a random spray. Multiplication forbids many arrangements, global laws predict totals, and the remaining freedom is difficult enough to keep elementary questions open.

Easy to Verify, Hard to Reverse

Suppose someone hands you a thousand-digit integer and asks two questions. Is it prime? If it is composite, what are its prime factors? These may sound like the same task. Computationally, they are sharply different.

Trial division answers both by testing possible factors up to the square root. It is hopeless at cryptographic sizes. Primality testing has much better routes because a proof of primality need not exhibit every failed divisor. Modern tests exploit congruences, group structure, polynomial identities and certificates.

Miller-Rabin is a common example. It tests whether a candidate behaves as a prime should under repeated squaring modulo that candidate. Every odd composite number has many bases that expose it: at most one quarter of possible bases can be strong liars. Independent random choices therefore drive the chance of a composite surviving k rounds below $4^{-(k)}$, under the standard sampling assumptions. In practice, probable-prime tests are

combined with small-prime sieving and enough rounds or deterministic base choices for the relevant range.

A probable prime is not the same thing as a proof, but deterministic polynomial-time primality testing exists. The Agrawal-Kayal-Saxena result, announced in 2002, established that primality belongs to the class P. More practical proving systems can produce certificates that another program checks efficiently. The conceptual result is decisive: no unproved hardness assumption is needed to say that primality itself has an efficient algorithmic solution.

Factorisation has no comparable general result. The leading general classical methods for arbitrary large integers are far faster than trial division, with subexponential heuristic running times rather than polynomial bounds in the number of input bits. No classical polynomial-time factoring algorithm is known, and no proof shows that none exists. A 200-digit semiprime, a product of two primes, can therefore be easy to recognise as composite while its factors remain expensive to recover.

Special forms change the calculation. A Mersenne number has the form $2^p - 1$. If it is prime, p must be prime, though many prime exponents still produce composite Mersenne numbers. The Lucas-Lehmer test decides primality for this form with a compact recurrence, which is why distributed searches can certify record primes with tens of millions of digits. The largest known prime when checked on 3 September 2026 was $2^{136,279,841} - 1$, found in 2024 and independently verified. Its record reflects a searchable family and a specialised test, not a belief that record-sized primes are generally Mersenne.

RSA turns factor recovery into a security assumption. Its public modulus is the product of two selected large primes. Multiplication is cheap; recovering the factors from the product is not known to be cheap on classical computers. The private operation depends on information derived from those factors. Real RSA security also depends on padding, key generation, implementation and protocol design, so the slogan that primes secure the internet is far too broad.

Shor's algorithm would factor integers in polynomial time on a fault-tolerant quantum computer with enough reliable logical qubits and operations, breaking the relevant RSA assumption. The algorithm is established; the required machine is an engineering condition, not part of the proof. This makes the asymmetry contingent rather than eternal. Elliptic-curve systems face the related discrete-logarithm version of the algorithm.

The causal loop is now complete. Unique factorisation made every composite a unique product of prime atoms. That same uniqueness gives reversal a definite answer: there is one factorisation to recover. Yet the forward operation and the reverse search have different known computational costs. Arithmetic's cleanest order creates one of computation's most useful asymmetries. The atoms are easy to assemble, easy to verify under the right tests, and sometimes hard to retrieve from the object they made.

How It Actually Works

Euclid turns divisibility into proof

Around 300 BCE, the *Elements* gathered Greek mathematics into definitions, propositions and demonstrations. Its most famous modern reputation is geometric, but Books VII to IX concern whole numbers. They contain algorithms for greatest common divisors, propositions about prime divisibility, results on ratios and a proof that primes cannot be exhausted.

The Euclidean algorithm begins with a practical question. To find the greatest common divisor of 252 and 105, divide and keep the remainder:

$$252 = 2 \times 105 + 42$$

$$105 = 2 \times 42 + 21$$

$$42 = 2 \times 21.$$

The final non-zero remainder is 21. The method works because replacing a pair by the smaller number and the remainder does not change their

common divisors. It reaches the answer without factorising either number first. More than two millennia later, the same algorithm sits inside computer algebra and cryptographic arithmetic.

Euclid's work did not present the modern fundamental theorem of arithmetic in one familiar sentence. It supplied much of the machinery from which later mathematicians could state and prove uniqueness cleanly. This distinction matters because histories often project a finished theorem backwards onto an older text whose concepts and notation differ from ours.

The infinitude argument has survived the change of language because its mechanism is independent of notation. Any finite list defeats itself when its product is increased by one. It is an early example of a proof that creates the object needed to escape a proposed boundary.

Lists become sieves

A proof of infinite supply does not tell you which numbers in a given range are prime. For that, the ancient method associated with Eratosthenes removes composites rather than confirming candidates one at a time.

Imagine a table from 2 to 50. Circle 2 and cross out 4, 6, 8 and the other multiples. Circle 3 and cross out 6, 9, 12 and its remaining multiples. Four has already gone. Circle 5, then 7. Once the next uncrossed number would exceed $\sqrt{50}$, the work is complete. Any composite at most 50 must have one factor at most $\sqrt{50}$; otherwise both factors would be larger and their product would exceed 50.

This square-root boundary appears everywhere in elementary primality work. If 221 is composite, one factor is at most $\sqrt{221}$, which is under 15. Test the primes 2, 3, 5, 7, 11 and 13. Thirteen divides it, giving 17. The method is conclusive, but its cost grows badly when the candidate has hundreds of digits.

Sieve thinking evolved instead. Rather than store an impossible table up to a high bound, a segmented sieve marks one interval at a time using primes up to its square root. More advanced sieves attach weights, estimate

survivors and search for patterns such as almost-primes. The action remains ancient: use small prime divisibility to remove whole families together.

Fermat makes remainders speak

Pierre de Fermat worked in the seventeenth century without publishing a systematic number-theory treatise. He sent claims in letters, challenged correspondents and left notes in margins. In 1640 he wrote that if p is prime and a is not divisible by p , then p divides $a^p - a$. This became Fermat's little theorem, to distinguish it from his last theorem about powers.

The result offered a new style of prime evidence. Instead of looking for divisors, calculate a power modulo the candidate and see whether prime-like behaviour appears. If $a^{n-1} \not\equiv 1 \pmod{n}$ is not congruent to one modulo an odd n , then n is certainly composite. Passing does not prove primality, because composites can mimic the result.

Fermat also proposed that numbers of the form

$$F_n = 2^{2^n} + 1$$

might all be prime. The first five are. Leonhard Euler ended the pattern by showing in 1732 that

$$F_5 = 4,294,967,297 = 641 \times 6,700,417.$$

The failure is instructive. A formula can produce several primes and still conceal composite behaviour beyond the visible sample. Prime research has repeatedly punished induction from small data.

Euler went further than correction. He generalised Fermat's congruence using the totient function, which counts residues coprime to a modulus. If $\gcd(a, n) = 1$, then $a^{\phi(n)} \equiv 1 \pmod{n}$. He also turned unique factorisation into analysis through the product now attached to the zeta function and proved that the reciprocal primes have a divergent sum.

Euler also completed an older link between primes and perfect numbers. A perfect number equals the sum of its positive proper divisors: $6 = 1 + 2 + 3$,

and $28 = 1 + 2 + 4 + 7 + 14$. Euclid had shown that if $2^p - 1$ is prime, then $2^{p-1}(2^p - 1)$ is perfect. Euler proved the converse for even perfect numbers. Finding every even perfect number is therefore equivalent to finding every Mersenne prime. Whether an odd perfect number exists remains unknown.

Gauss gives the atoms a notation

The nineteenth-century story also depends on a change in language. Gauss's *Disquisitiones Arithmeticae*, published in 1801, made congruence notation systematic. Writing $a \equiv b \pmod{m}$ turns a verbal statement about equal remainders into an object that can be manipulated. It separates the residue pattern from the size of the integers carrying it.

Gauss also gave a prominent modern proof of unique factorisation and organised number theory around congruences. The shift was not cosmetic. A problem about a large integer could be projected into several modular systems, solved or obstructed there, and then related back to the original question. Quadratic residues, which ask whether $x^2 \equiv a \pmod{p}$ can be solved, became a central study. Their pattern led Gauss to quadratic reciprocity, a theorem relating whether one odd prime is a square modulo another.

This is where the prime as modulus becomes as important as the prime as factor. Factorisation decomposes a number into local components. Congruences ask how a number behaves when viewed through each prime. Much of later number theory moves between those directions.

Gauss counts the thinning crowd

By the late eighteenth century, tables contained enough primes for a different question: how many are there below x ? As a teenager, Carl Friedrich Gauss examined the data and conjectured a logarithmic law. Adrien-Marie Legendre proposed a related approximation. Neither possessed the later tools required for proof, and their formulas differed in detail, but the central observation was right: prime density declines like the reciprocal of a logarithm.

Gauss's insight changed the scale of the subject. An exact list is jagged. A counting function can reveal a smooth trend. This is the same move that turns individual collisions into gas laws or individual births into demography, except that no physical randomness is present.

In 1837, Peter Gustav Lejeune Dirichlet proved that every arithmetic progression $a \bmod q$ with $\gcd(a, q) = 1$ contains infinitely many primes. His proof introduced functions now called Dirichlet L-functions. The advance joined congruence classes to analysis and showed that Euclid's infinite supply is distributed across every arithmetically permissible route, not dumped into a few favoured remainders.

Dirichlet's theorem also exposes why conditions matter. The progression $2 \bmod 4$ contains one prime and then only even composites. The progression $1 \bmod 4$ has infinitely many primes. A local common factor completely changes the global conclusion.

Estimates arrive before asymptotics

Between Gauss's conjecture and the 1896 proof, Pafnuty Chebyshev established firm upper and lower bounds of the correct order for $\pi(x)$. He also proved Bertrand's postulate: for every integer n greater than one, at least one prime lies between n and $2n$. The result gives a local guarantee far stronger than infinitude. However far you travel, doubling the scale always catches another prime.

Such intermediate theorems matter because mathematical progress is rarely a jump from guess to final proof. A conjectured asymptotic can be surrounded by bounds, special cases and equivalent formulations before its exact form yields. Chebyshev could show that the prime count behaved like a constant multiple of $x/\log x$ within fixed limits without proving that the constant settles to one.

Riemann moves off the number line

Bernhard Riemann published one short paper on prime numbers in 1859. It changed the field. He began from Euler's product, continued the zeta function into the complex plane and linked its zeros to the distribution of primes.

Complex numbers add a second direction to the ordinary line. Write $s = \sigma + it$, with real part σ and imaginary part t . The zeta function's non-trivial zeros lie in the strip $0 < \sigma < 1$, and symmetry reflects them around the line $\sigma = 1/2$. Riemann suggested that every non-trivial zero lies on that line.

The paper did more than state a conjecture. It supplied an explicit connection between prime-counting functions and the zeros. The smooth logarithmic estimate is the main signal. The zeros contribute corrections. Their positions determine how large the discrepancy can become.

Riemann's insight explains why the hypothesis has consequences far beyond improving a graph of $\pi(x)$. Prime numbers enter factorisation, congruences, character sums and estimates throughout number theory. Better control of prime fluctuations strengthens many dependent bounds. The hypothesis is a shared load-bearing assumption in conditional results, although each consequence must still be stated at its proper scope.

The average law is proved

In 1896, Jacques Hadamard and Charles-Jean de la Vallée Poussin independently proved the prime number theorem. A decisive step was showing that the zeta function has no zeros on the line with real part one. That zero-free boundary is enough to control the analytic expression for prime counts and establish $\pi(x) \sim x/\log x$.

The proof confirmed Gauss's century-old numerical vision without settling Riemann's stronger claim. The average law needs less information about the zeros than the sharpest error estimates do. This is a recurring hierarchy: one can know the dominant behaviour while the fluctuations remain mysterious.

In 1948, work by Atle Selberg and Paul Erdős led to what became known as an elementary proof, meaning one that avoided complex analysis. Separate papers appeared in 1949 after a priority dispute. Elementary does not mean easy. Their arguments used intricate identities and estimates. The episode showed that a theorem can outgrow the route by which it was first reached. Complex analysis remains central to understanding finer distribution, but it is not logically indispensable for the leading asymptotic.

Tables continued to improve alongside theory. Mechanical calculators and computers extended exact counts, tested conjectures and made error terms visible at larger scales. Computation did not replace proof. It changed which questions could be inspected before proof and which false patterns could be killed quickly.

Sieves learn their limits

Sieve methods underwent a related change. Viggo Brun developed a sieve capable of studying pairs such as n and $n + 2$. He proved that the sum of reciprocals of twin primes converges, unlike the divergent sum over all primes. This result did not settle whether infinitely many twins exist. A finite collection would also have a convergent reciprocal sum. It showed that twins, if infinite, are sparse enough to leave a finite total.

Modern sieves can detect numbers with few prime factors and prove bounded-gap results when combined with other ideas. They also face a parity problem: many sieve weights struggle to distinguish a prime from a number with an even or odd count of prime factors in the way a target theorem demands. This is not a declaration that sieve methods can never prove the twin-prime conjecture. It is a precise obstruction for broad families of existing sieve arguments, and it helps explain why removing composites is easier than certifying the exact prime pattern left behind.

Computers split testing from factoring

In the twentieth century, primality became an algorithmic subject. Édouard Lucas developed tests for special forms in the nineteenth century, and Derrick Henry Lehmer refined the Mersenne test in the 1930s. For an odd prime exponent p , the Lucas-Lehmer recurrence can decide whether $2^p - 1$ is prime using repeated squaring and reduction modulo the candidate. Its fit to binary computation made Mersenne numbers natural record targets.

General candidates required other tools. Gary Miller gave a deterministic polynomial-time primality test conditional on the extended Riemann hypothesis for Dirichlet L-functions. Michael Rabin turned the underlying idea into an unconditional randomised test with a rigorous error bound. A composite candidate may imitate prime behaviour for some bases, but most bases expose it. Repetition makes false acceptance negligible without searching for a factor.

In 2002, Manindra Agrawal, Neeraj Kayal and Nitin Saxena announced a deterministic polynomial-time primality algorithm. Its famous summary, “PRIMES is in P”, settled the complexity classification of primality testing. The algorithm is more important as theory than as the fastest routine for large practical inputs. Modern software usually combines sieving, fast probable-prime tests and certificate-producing methods suited to the job.

Factorisation followed a different curve. Pollard methods, elliptic-curve factorisation and the number field sieve improved performance dramatically, but the general problem has no known classical polynomial-time solution. This gap supports RSA, while never proving that RSA must remain secure. Algorithmic history is full of problems that looked hard until a new representation changed the cost.

Shor's quantum algorithm supplies such a representation in principle. A large fault-tolerant quantum computer could factor efficiently by using quantum period finding. The obstacle is engineering a machine of adequate scale and reliability, not a missing mathematical algorithm.

A record is a chain of checks

A record-prime announcement begins with candidate selection. For Mersenne searches, software assigns odd prime exponents p and tests $M_p = 2^p - 1$. Trial factoring and other screening can remove many candidates cheaply. Modern GIMPS searches then use a probable-prime test for the surviving candidate; a probable-prime result receives definitive Lucas-Lehmer confirmations. The Lucas-Lehmer recurrence begins with 4 and repeatedly squares, subtracts 2 and reduces modulo M_p . After $p - 2$ steps, M_p is prime exactly when the final residue is zero.

The computation is long, and hardware can make mistakes. A credible record therefore needs replication. The 2024 discovery used one probable-prime computation to identify the candidate, followed by Lucas-Lehmer confirmations using independent programs and hardware. The mathematical theorem says what residue establishes primality. Engineering establishes that the claimed residue was computed faithfully.

The search is also incomplete in a controlled way. The current record exponent lies far above the previous one, and not every exponent between them has been eliminated. A smaller untested exponent could therefore yield another Mersenne prime. It would not displace the size record, but it would change the record prime's ordinal position among Mersenne primes. “Largest known” reports the search record, not proof that every smaller candidate has been settled.

This chain shows how modern number theory combines proof and experiment. The test converts primality into a finite certificate condition. Distributed computation performs the work. Independent verification protects against machine or software error. Publication records the claim. None of those stages can be removed without changing what has been established.

The frontier breaks into patterns

Prime research no longer has one frontier. It has a family of elementary questions connected by different methods.

The conjecture emerged from Goldbach's 1742 correspondence with Euler, in formulations shaped by the period's treatment of one. Its modern strong form says every even integer greater than two is a sum of two primes. The ternary version for odd integers yielded to circle-method refinements and was completed by Harald Helfgott in the 2010s. The binary statement remains open.

Twin primes resisted every attempt to prove a fixed gap until 2013, when Yitang Zhang showed that some gap below seventy million occurs infinitely often. The number was huge compared with two, but finiteness was the breakthrough. James Maynard and Terence Tao introduced stronger methods for finding several primes in bounded intervals, while the Polymath collaboration reduced the unconditional bound for infinitely recurring consecutive gaps to 246.

In another direction, Ben Green and Tao proved that primes contain arithmetic progressions of any finite length. Their work imported tools from additive combinatorics and built on a theorem about dense subsets of the integers. Since primes have density tending to zero, they needed a way to transfer enough pseudorandom behaviour into a sparse setting.

These results do not merge into one master technique. Sieve methods, harmonic analysis, algebraic ideas, probability-inspired models, computation and combinatorics each see different parts. The subject stays alive because the definition is small while its consequences distribute themselves across mathematics.

How we know

Claims about ancient practice survive through later manuscripts rather than original working documents. The sieve is traditionally attributed to Eratosthenes, but no complete account by him survives. Euclid's *Elements* is read through a long transmission history, and modern statements such as unique factorisation use later notation and conceptual packaging.

From the seventeenth century onward, letters and publications make attribution firmer, though priority disputes and changing terminology still

require care. Fermat's little theorem appears in his 1640 correspondence; Euler's factorisation of the fifth Fermat number and Riemann's 1859 paper are directly documented.

Modern theorems rest on published proofs open to checking. Computational records require another layer: software, hardware and independent verification. The 2024 Mersenne record was confirmed by separate implementations, while its discovery reflects a search of a special family rather than a survey of every integer. Open-problem status was checked against current mathematical institutions and the research literature on 3 September 2026. Finite computation supplies evidence and catches errors; it cannot prove an infinite conjecture by exhaustion.

What People Get Wrong

“One is prime”

One is divisible only by one, so it looks as though it should qualify. The modern definition requires exactly two positive divisors, one and the number itself. For one, those are the same divisor.

The exclusion is structural rather than bureaucratic. Prime factorisation is meant to be unique apart from order. If one were prime, 30 could be written as $2 \times 3 \times 5$, or $1 \times 2 \times 3 \times 5$, or with any chosen number of extra factors of one. Every theorem using the number of prime factors would need an exception that removed those meaningless insertions.

One is instead a unit: it has a multiplicative inverse among the integers, namely itself. In more general algebra, units and primes remain different roles. Units can be inserted or removed without changing divisibility in the meaningful sense; primes are the irreducible factors that carry information.

Definitions are judged by the structure they reveal. Moving one out of the primes makes factorisation, exponents and generalisation cleaner. It was not always classified consistently in older writing, but the modern choice is not arbitrary housekeeping.

“Primes are random”

No randomness decides whether a number is prime. A positive integer either has a non-trivial divisor or it does not. The answer follows from multiplication with complete determinism.

The random impression comes from local unpredictability. Prime gaps vary, clusters appear and no short periodic rule lists exactly the primes. Probability models can still be useful. Near x , the density is about $1/\log x$, so treating a sampled candidate as prime with that rough probability predicts search effort and average counts.

The model has constraints that independent coin tosses lack. Every prime above two is odd. Every prime above three lies in one of two residue classes modulo six. Candidate patterns can be forbidden because one member is always divisible by a small prime. These correlations matter in precise conjectures and sieve calculations.

Calling primes random can therefore help if it means statistically modelled after arithmetic restrictions. It misleads if it means causeless, independent or unknowable. The useful surprise is sharper: a deterministic sequence can support probability-like laws because global density is easier to describe than individual position.

“Every Euclid number is prime”

Euclid's proof multiplies a finite list of primes and adds one. School summaries often turn this into a machine for making a new prime. It is not.

Using 2, 3, 5, 7, 11 and 13 gives

$$2 \times 3 \times 5 \times 7 \times 11 \times 13 + 1 = 30,031 = 59 \times 509.$$

The constructed number is composite. What matters is that none of the primes in the original list divides it, since division always leaves remainder one. Its prime factors must therefore be absent from the list.

The mistaken version is attractive because it makes the proof feel more constructive than it is. The real argument is stronger in a different way. It

defeats every finite list without needing the new number itself to be prime. Unique factorisation guarantees that some new prime appears inside it.

This correction teaches a general habit. Do not remember a proof by its props. Remember the logical burden. Euclid needs a number not divisible by any listed prime, then the fact that every number above one has a prime divisor. Anything beyond that is decoration and may be false.

“There is no formula for primes”

There are many formulas that output primes, encode prime positions or test primality. Some use floor functions, factorials, trigonometric expressions or constants defined in terms of the primes themselves. The bare claim that no formula exists has no precise mathematical meaning until “formula” and “useful” are specified.

Wilson's theorem gives a clean example: n greater than one is prime exactly when $(n - 1)! \equiv -1 \pmod n$. That is a perfect criterion. It is usually a poor way to test huge candidates because computing the factorial modulo n costs far more than modern alternatives.

Other expressions hide the answer in their constants or require work comparable to enumerating primes. A closed-looking line can therefore be mathematically valid while offering no predictive advantage.

The genuine absence is more modest. No known elementary expression gives the next prime with low computation and transparent structure in the way a linear formula gives the next term of an arithmetic progression. Prime theory has exact characterisations, efficient tests, asymptotic counts and specialised generators. It lacks one universal shortcut that collapses all of those jobs into effortless prediction.

“Testing primality means finding the factors”

Finding a factor proves compositeness, but failure to find one does not prove primality unless the search is exhaustive or attached to a valid certificate. The tasks can be separated.

Miller-Rabin asks whether modular powers behave in ways required of primes. Most bases expose any given odd composite. Repeated tests can make false acceptance negligible. Certificate-producing algorithms establish primality through verifiable mathematical conditions without listing all failed divisors.

AKS goes further at the level of complexity theory: primality can be decided deterministically in polynomial time. No general classical polynomial-time factoring algorithm is known. A large number may therefore be recognised as composite long before its factors are recovered, and a prime may be certified without testing every possible divisor.

A compositeness witness may be far cheaper to find than a factor. One failed modular identity can prove that n is not prime while revealing nothing useful about how n splits. Conversely, a factorisation is a certificate of compositeness that multiplication checks at once. Discovery cost and verification cost need not match, and the cheapest certificate depends on which claim is being made.

The confusion survives because small classroom numbers encourage trial division. For 91, spotting 7×13 both answers “composite?” and supplies the factorisation. At hundreds of digits, the computational paths diverge. Recognising a property can be easier than reconstructing the hidden components responsible for it.

“The Riemann hypothesis will reveal the next prime”

The Riemann hypothesis concerns the zeros of the zeta function and the size of fluctuations in prime-related counting functions. It would give far tighter control over error terms. It is not a lookup rule for the next prime.

Even under the hypothesis, exact prime location would still require computation or more detailed information. Knowing that $\pi(x)$ stays close to a smooth approximation narrows the possible population in large intervals; it does not label each integer prime or composite.

The myth grows from language about the hypothesis “explaining” prime distribution. Explanation here is spectral and aggregate. The zeros contribute oscillations to explicit formulas for weighted prime counts. Their real parts control the scale of possible deviations from the main trend.

A proof would have major consequences because many theorems currently say “assuming the Riemann hypothesis” before giving sharper bounds. Those consequences differ by problem and would not turn all number theory into routine calculation. The hypothesis would discipline the error. The local sequence would keep its irregularity.

“Prime numbers keep the internet secure”

Some major public-key systems use arithmetic built from primes. RSA publishes a product of two selected large primes and keeps factor-derived information private. Finite-field Diffie-Hellman uses a carefully chosen prime modulus. Elliptic-curve systems often work over finite fields and select groups with a large prime-order component.

That does not make primality a security spell. A large prime can sit inside a weak protocol, a reused key, broken random generation or vulnerable software. RSA needs safe padding and adequate parameters.

Diffie-Hellman needs authentication and valid groups. Many symmetric ciphers and hash functions do not depend on factoring at all.

The hard problem must also be named correctly. RSA relies on the absence of an efficient known way to recover the private relation from the public modulus under its design assumptions. Primality testing itself is efficient. Keeping the primes secret is part of key construction, not evidence that primes are hard to recognise.

Quantum computing narrows the slogan further. Shor's algorithm would threaten factoring and discrete logarithms on a sufficiently capable fault-tolerant machine, while other cryptographic assumptions respond differently. Security comes from a complete system whose mathematical and operational assumptions remain defensible. Primes supply structure. They do not supply safety alone.

Use It

Choose the operation before the metaphor

Calling primes atoms is useful only after asking what counts as assembly. Positive integers are assembled from primes by multiplication. Under addition, every number can be assembled from ones and prime status adds little. In Gaussian integers, ordinary primes may split. In polynomial rings over a field, irreducible polynomials take the prime-like role.

This question prevents metaphors from becoming false ontology. What is the operation? Which objects count as units that can be inserted without changing the structure? What does irreducible mean there? Does every object factor, and is the factorisation unique?

Prime numbers teach a disciplined version of reduction. Break an object into components because a theorem says the decomposition exists and is informative, not because small pieces always explain a whole. The atom claim earns its force from unique factorisation. Without uniqueness, a decomposition may be one description among several rather than the description.

Reduce a question to remainders

A large integer can be hard to inspect in full and easy to reject modulo a small number. If its digit sum is divisible by three, so is the integer. If its final digit is even, it cannot be an odd prime. These school rules are modular arithmetic disguised as tricks.

The stronger habit is to choose a modulus fitted to the claim. To show that no square is congruent to three modulo four, calculate the residues of 0^2 , 1^2 , 2^2 and 3^2 . They are only zero or one. Any equation forcing $x^2 \equiv 3 \pmod{4}$ is impossible, no matter how large x might be.

Remainders can also organise search. A prime greater than three must be one or five modulo six, so other candidates can be skipped. A progression $a + nq$ can contain infinitely many primes only if a and q are coprime. Before

attacking the full integers, inspect the small residue systems where divisibility patterns become finite.

This does not mean that every global problem is settled locally. Goldbach and twin primes pass obvious modular tests and remain difficult. Modular reduction is a filter. It can prove impossibility, expose necessary conditions and compress calculation. Passing the filter means the deeper work begins.

Separate local surprise from global law

The prime number theorem estimates a population, not an itinerary. It can tell you that roughly $x/\log x$ primes lie below x and that average gaps near x have logarithmic scale. It cannot say exactly whether the next candidate is prime.

Keep this distinction whenever an average law is applied to an individual case. Aggregate control and event prediction are different achievements, even when the aggregate estimate is exceptionally accurate.

Prime gaps make the lesson sharp because there is no hidden physical noise. The sequence is deterministic, yet local variation persists beside an increasingly accurate global trend. Uncertainty about the next event can therefore come from computational and structural difficulty rather than from chance in the generating process.

Ask two questions separately. What can be said about the total or average over a large range? What can be said about this next object? A strong answer to the first should not be inflated into an answer to the second. Nor should local irregularity be used to deny a global law that has been proved.

Ask which computational job is hard

“Working with a thousand-digit number” names a scale, not a task. Multiplying two such numbers, testing a candidate for probable primality, proving primality and factorising a composite have different algorithms and costs.

Input size should be measured in digits or bits, not by the numerical value

itself. Trial division to $\sqrt{n}$ is exponential in the bit length of n because $\sqrt{n}$ grows as $2^{(k/2)}$ for a k -bit input. An algorithm polynomial in k can still handle numbers whose values are beyond physical counting.

Then distinguish decision from search. “Is n prime?” asks for a yes or no. “What are n ’s factors?” asks for hidden structure. A certificate may let somebody verify the decision faster than discovering the certificate. Multiplication may be cheap while reversal remains expensive.

Do not ask whether a problem sounds difficult. Ask what information is given, what output is required, how cost grows with input length, and whether the claim concerns the best known method or a proved lower bound. Prime factorisation is useful in cryptography because no efficient classical method is known, not because every efficient method has been proved impossible.

Check the obstruction before hunting the pattern

Suppose you want infinitely many primes among values of a formula. Begin by asking whether a small prime divides one of those values for every possible input.

The formula $n^2 + n + 2$ always gives an even number. It equals two at $n = 0$ and is composite for every positive integer n . No amount of computation will uncover an infinite prime family there. The obstruction is built into parity.

For a set of linear forms, inspect residues modulo each prime. If every residue choice makes at least one form divisible by that prime, the proposed prime pattern is inadmissible. Twin primes survive this test: for each prime q , some residue avoids divisibility of both n and $n + 2$. Survival is necessary and still far from proof.

This ordering saves effort. First remove logical impossibility. Then use computation to inspect plausible cases. Then decide which theorem or heuristic applies to the remaining question. Evidence from many examples can raise confidence in an admissible pattern, but examples cannot overrule a missed obstruction or prove an infinite claim.

Prime research rewards disciplined pessimism at the start. Look for the small divisor that kills the grand pattern. Only after none appears is optimism mathematically licensed.

The limits

Primes do not explain every property of the integers. They govern multiplication, divisibility and structures built from them. Additive questions such as Goldbach require methods that cannot be read directly from the factorisation of each summand. Order, approximation, geometry and combinatorics bring their own ideas.

The subject also tempts overconfidence in heuristics. Treating primality near x as an event with probability $1/\log x$ predicts many averages, but dependencies must be corrected before precise conjectures are trusted. A persuasive random model is evidence about scale, not a proof.

Computers have a defined role. They can verify enormous finite ranges, find counterexamples, search structured families and check certificates. They cannot turn a trillion confirmed cases into a theorem about every integer. Conversely, a proof of existence may offer no practical route to locating a small example. Mathematical knowledge includes both demonstration and computation, and neither replaces the other.

The atom metaphor has its final limit in security. Prime arithmetic can create useful hard problems, but a deployed system also contains code, devices, users, standards and adversaries. Factorisation may remain hard while a key leaks through weak randomness or an error message. Pure structure enters the world through machinery, and machinery adds failure modes.

The one thing to keep

Keep the split between factorisation and location.

Every positive integer greater than one has one prime factorisation, apart from order. That is a severe form of order. The factors do not negotiate, and there is no second valid receipt. Yet the primes themselves do not occupy the number line according to a comparably transparent local schedule. Their

total count follows a logarithmic law, their residue classes obey exact restrictions, and their gaps still support questions that resist proof.

Those facts belong together. The mystery is not that mathematics has found chaos inside arithmetic. It is that exact rules can create a sequence whose global behaviour is more accessible than its next event, whose components are easier to assemble than recover, and whose simplest questions open into several branches of mathematics.

When you meet a system that looks irregular, ask what kind of order has already been proved. Is there a unique decomposition? A conservation rule? A distribution law? A local obstruction? A cheap verification method? The absence of a simple next-step prediction may coexist with strong structure at another level.

Prime numbers permanently change the image of certainty. Mathematical certainty is not the same as effortless foresight. We can prove that primes never end, estimate how many lie below a vast boundary, certify a forty-one-million-digit example and still fail to prove that pairs two apart continue for ever. Knowledge arrives in layers, each with a different question and standard.

The subtitle is now exact. Primes are atoms only for multiplication in the positive integers, where unique factorisation gives every number one receipt. That theorem promises neither a simple timetable for where the atoms appear nor a cheap route from a product back to its factors. Once those distinctions are clear, primes stop looking like a bag of awkward numbers. They become a model of exact structure whose consequences can remain difficult.

Terms

Prime

An integer above one whose only positive divisors are one and the number itself. Integers above one divide into prime and composite classes. Two is the sole even prime.

Composite

A positive integer greater than one that is not prime. It can be written as a product of smaller positive integers and ultimately decomposed into prime factors. Its prime factors may repeat.

Factor

An integer that divides another with no remainder. Factor and divisor are interchangeable here. Prime factorisation records each prime factor together with the number of times it occurs.

Unit

An element with a multiplicative inverse in the number system being used. In the integers the units are 1 and -1, and neither is treated as prime. Units can be inserted without changing meaningful factorisation data.

Fundamental theorem of arithmetic

The theorem that every integer greater than one has a prime factorisation unique apart from factor order. It is the classification rule behind ordinary multiplicative arithmetic. Many arguments reduce integers to prime powers because of it.

Prime factorisation

The expression of an integer as a product of prime powers, such as $360 = 2^3 \times 3^2 \times 5$. Its exponents compactly record divisibility.

Coprime

Two integers whose greatest common divisor is one. They need not be prime themselves: 8 and 15 are composite, yet no prime divides them both. Coprimality belongs to the pair, not either number alone.

Greatest common divisor

The largest positive integer dividing each of two or more integers. It can be found from prime exponents or without factorisation through the Euclidean algorithm. Bézout's identity also expresses it as an integer linear combination.

Euclidean algorithm

A procedure repeatedly replacing a pair of integers by the smaller one and the remainder after division. The final non-zero remainder is their greatest common divisor. The extended form finds the corresponding Bézout coefficients.

Sieve of Eratosthenes

A method for listing primes up to a chosen limit by crossing out multiples of each surviving prime. Sieving need continue only through the square root. Segmented versions process large intervals in manageable blocks.

Congruence

A relation stating that two integers leave the same remainder modulo a chosen number. The notation $a \equiv b \pmod{m}$ means that m divides $a - b$. Congruence respects both addition and multiplication.

Modulus

The number defining a system of remainders. Arithmetic modulo m identifies integers differing by a multiple of m and leaves exactly m residue classes. Prime moduli avoid non-zero zero divisors.

Residue class

All integers congruent to one another under a modulus. Modulo five, the class of two contains ..., -3, 2, 7, 12, ... as different representatives.

Multiplicative inverse

For a modulo m , a residue b satisfying $ab \equiv 1 \pmod{m}$. Such an inverse exists exactly when a and m are coprime.

Finite field

A finite system in which addition, subtraction, multiplication and division by non-zero elements obey field rules. Integers modulo a prime give the most immediate example. Every finite field has prime-power size.

Fermat's little theorem

For prime p , $a^p \equiv a \pmod{p}$ for every integer a . When p does not divide a , this becomes $a^{p-1} \equiv 1$. The converse fails for pseudoprimes.

Euler's totient function

Written $\phi(n)$, the number of residue classes modulo n that are coprime to n . It supplies the exponent in Euler's generalisation of Fermat's theorem.

Prime-counting function

Written $\pi(x)$, the number of primes less than or equal to x . It is a step function whose long-run growth follows the prime number theorem.

Prime number theorem

The result $\pi(x) \sim x/\log x$. It states that the ratio of the true prime count to $x/\log x$ approaches one as x grows.

Logarithmic integral

A function commonly written $\text{Li}(x)$, obtained by accumulating $1/\log t$. It often approximates $\pi(x)$ more closely than using the endpoint expression $x/\log x$.

Arithmetic progression

A sequence with a constant difference, such as 5, 11, 17, 23. Prime behaviour in such sequences depends first on the starting value and step being coprime.

Dirichlet's theorem

The theorem that every arithmetic progression $a + nq$ with $\gcd(a, q) = 1$ contains infinitely many primes. It guarantees supply in every allowable residue class.

Twin primes

Pairs of primes differing by two, such as 29 and 31. Two is the smallest possible gap between distinct odd primes; infinitude remains conjectural.

Prime gap

The difference between consecutive primes. Average gaps near x have logarithmic scale, while particular gaps may be far smaller or larger than that local average.

Mersenne prime

A prime of the form $2^p - 1$. The exponent p must be prime, but many prime exponents still produce composite Mersenne numbers.

Perfect number

A positive integer equal to the sum of its proper positive divisors. Every even perfect number comes from a Mersenne prime through the Euclid-Euler formula. No odd perfect number has been found or ruled out.

Zeta function

The function $\zeta(s)$, first defined by the series $\sum n^{-s}$ where it converges and then continued more widely. Its analytic behaviour carries information about primes.

Euler product

The identity $\zeta(s) = \prod_p (1 - p^{-s})^{-1}$ for real part greater than one. It encodes every positive integer once through unique prime factorisation.

Riemann hypothesis

The conjecture that every non-trivial zero of the zeta function has real part one half. It would sharply constrain errors in several statements about prime distribution.

Primality test

An algorithm deciding whether an integer is prime. Tests may be deterministic, probabilistic or certificate-producing, and a compositeness result need not reveal any factor. Testing and full factorisation are separate computational problems.

Go Deeper

Vicky Neale, Closing the Gap: The Quest to Understand Prime Numbers

Begin here for the modern story. Neale explains why bounded gaps mattered, how Zhang's breakthrough changed the question and how Maynard, Tao and the Polymath collaboration improved the result. She writes for interested non-specialists without pretending that the hard parts are elementary. The book appeared in 2017, so later numerical refinements and current records are outside it. Its deeper value is methodological: you see a research community turn one unexpected proof into several new routes through an old problem. Watch the distinction between proving that some bounded gap recurs and proving that a named gap, such as two, does.

Barry Mazur and William Stein, Prime Numbers and the Riemann Hypothesis

Read this to understand how counting primes leads to the zeta function and why its zeros matter. The authors build the argument visually and computationally before increasing the analytic depth. A reader comfortable with graphs, logarithms and complex numbers will gain most, though the early chapters remain accessible without specialist training. It is focused rather than comprehensive. Goldbach, Mersenne searches and cryptographic applications appear only where they support the central distribution problem. Work through the plots rather than skipping to the hypothesis: the book's strength is showing how the smooth count and oscillating corrections fit together.

Euclid, The Thirteen Books of the Elements, translated by T. L. Heath

Use Books VII to IX as primary evidence for ancient number theory. You will find the Euclidean algorithm, prime divisibility machinery, the infinitude proof and the link between Mersenne-type primes and perfect numbers in their older conceptual language. Heath's notes are extensive and historically valuable, but the translation and commentary reflect scholarship of their period. Do not expect modern notation or the fundamental theorem of arithmetic packaged in its present textbook form. Read the propositions beside Heath's commentary, while remembering that the commentary itself belongs to early twentieth-century classical scholarship.

G. H. Hardy and E. M. Wright, *An Introduction to the Theory of Numbers*, sixth edition

This is the bridge from an inviting survey to serious number theory. It covers primes, congruences, arithmetic functions, Diophantine equations, continued fractions and analytic methods, with later editors updating a classic text. The prose is compact and the exercises assume mathematical maturity. Read selected chapters rather than forcing a straight march. It is the book for checking how the elementary statements in this one sit inside the wider discipline and where an attractive heuristic must give way to proof. The sixth edition was revised by D. R. Heath-Brown and J. H. Silverman, with a foreword by Andrew Wiles. Keep paper nearby: many of its shortest paragraphs contain an argument that this book has expanded across several pages.

Notes and Sources

Scope, definitions and the atom metaphor

This book uses prime to mean a positive integer greater than one with exactly two positive divisors. Negative primes are sometimes discussed as associates of positive primes, but signs add nothing needed here. The distinction among units, irreducibles and primes, and the warning that factorisation changes in larger rings, follow standard algebraic number theory. The example $5 = (2 + i)(2 - i)$ is in the Gaussian integers. Hardy and Wright, Ireland and Rosen, and Stillwell supply the main general treatments.

The fundamental theorem of arithmetic is stated for ordinary positive integers and uniqueness is understood apart from factor order. Euclid's *Elements* contains the prime-divisibility proposition now associated with Euclid's lemma at Book VII, Proposition 30, and the infinitude proof at Book IX, Proposition 20. Euclid did not state the modern theorem in one proposition using present notation. The body therefore distinguishes ancient machinery from later packaging.

The current record prime

The Great Internet Mersenne Prime Search reported that $2^{136,279,841} - 1$ was found prime on 12 October 2024. It has 41,024,320 decimal digits. At verification it was one of 52 known Mersenne primes; GIMPS marks its ordinal position in the exponent sequence as provisional because not every intervening exponent has been eliminated. A probable-prime result was followed by a Lucas-Lehmer confirmation and further independent checks with different implementations and hardware. GIMPS still listed it as the largest known prime on 3 September 2026.

The reading-time comparison is arithmetic: 41,024,320 seconds is about 474.8 continuous days. The printing comparison is illustrative rather than a publishing estimate. At roughly 3,000 digits per page, the decimal expansion would occupy more than 13,600 pages before front matter or spacing.

Prime factorisation and divisors

The divisor-count calculation for 360 uses the standard rule: if $n = p_1^{a_1} \dots p_k^{a_k}$, then n has $(a_1 + 1) \dots (a_k + 1)$ positive divisors. Greatest common divisors take minimum prime exponents and least common multiples take maximum exponents. These consequences and the termination argument for factorisation follow Hardy and Wright and Ireland and Rosen.

One is excluded because it is a unit and because arbitrary insertions of one would destroy literal uniqueness. Historical writers did not always use the modern classification, so the body presents this as the settled structural convention rather than an eternal definition.

Modular arithmetic and finite fields

For prime p , the residue classes modulo p form a field. The proof that non-zero classes have inverses may be given through Euclid's lemma, Bézout's identity or finite cancellation. Composite moduli can have zero divisors, as $2 \times 4 \equiv 0 \pmod{8}$ shows.

Fermat stated the theorem now bearing his name in a letter to Frénicle de Bessy in October 1640. Euler supplied published proofs and the totient

generalisation. Carmichael numbers show why the congruence $a^{(n-1)} \equiv 1 \pmod n$ for every a coprime to n does not characterise primes. Korselt gave the standard criterion for Carmichael numbers; the body needs only the counterexample class.

Infinitude, sieving and reciprocal primes

The product-plus-one argument proves that no finite list contains every prime. It does not prove that the constructed number is prime. The factorisation $30,031 = 59 \times 509$ was checked directly.

The sieve is traditionally attributed to Eratosthenes, but no complete description written by him survives. The account here states the classical method without claiming direct textual provenance. The square-root stopping rule follows because a composite $n = ab$ cannot have both a and b greater than $\sqrt{n}$.

Euler's proof that the sum of reciprocal primes diverges can be formulated through finite Euler products, logarithms and comparison with the harmonic series. The body uses finite products and suppresses the elementary inequalities comparing $-\log(1 - 1/p)$ with $1/p$. Unique factorisation establishes which reciprocal integers occur in each expansion; the limiting argument turns that structure into divergence.

Counting primes

The exact values used are $\pi(10) = 4$, $\pi(100) = 25$, $\pi(1,000) = 168$ and $\pi(1,000,000) = 78,498$. The approximation $1,000,000/\log(1,000,000)$ is about 72,382, using the natural logarithm.

The prime number theorem states $\pi(x) \sim x/\log x$. Interpreting $1/\log x$ as local prime density is a standard heuristic consequence, not an assertion that neighbouring primality events are independent. Among odd candidates the corresponding rough density is $2/\log x$. A 2048-bit magnitude has natural logarithm near $2048 \log 2$, about 1,419.6, so the rough average number of odd candidates examined per prime is about 710.

The factorial construction gives a composite run of length n : for k from 2

through $n + 1$, $(n + 1)! + k$ is divisible by k . It proves arbitrarily long prime-free intervals without estimating the largest gaps that occur naturally near a chosen scale.

Gauss's recollection of his early prime-counting observations appears in an 1849 letter to Encke. Legendre published related approximations.

Chebyshev proved bounds of the correct order and Bertrand's postulate before Hadamard and de la Vallée Poussin independently proved the prime number theorem in 1896.

Zeta, zeros and the Riemann hypothesis

The series for $\zeta(s)$ and its Euler product converge absolutely when the real part of s exceeds one. Analytic continuation defines the function more widely, apart from a simple pole at one. The non-trivial zeros lie in the critical strip, and the Riemann hypothesis places them on the line with real part one half.

The body's spectral language compresses explicit formulas involving weighted counts such as the Chebyshev functions. Prime powers appear naturally after logarithmic differentiation of the Euler product. A zero with real part β contributes on a scale related to x^β , with qualifications concerning weights, summation and logarithmic factors. Under the Riemann hypothesis, one standard equivalent consequence is $\pi(x) = \text{Li}(x) + O(\sqrt{x} \log x)$. The body states this as a constant-times bound for sufficiently large x rather than implying an exact formula. Edwards, Davenport, Montgomery and Vaughan, and Mazur and Stein support this account.

The Clay Mathematics Institute continued to list the Riemann hypothesis as unsolved on 3 September 2026. Numerical checks of zeros are evidence and error detection, not a proof of the infinite statement. The book does not use a finite verification count because that record depends on definition, method and subsequent computation.

Progressions, gaps and additive problems

Dirichlet's theorem gives infinitely many primes in $a + nq$ when $\gcd(a, q) = 1$. Green and Tao proved that the primes contain arithmetic progressions of every finite length. Their theorem concerns finite progressions inside a set of zero density, not an infinite arithmetic progression of primes.

Zhang's published theorem gave a finite unconditional bound of 70,000,000 for gaps occurring infinitely often. Maynard's independent refinement gave 600, and the D. H. J. Polymath project reduced the published unconditional bound for consecutive primes to 246. These statements concern $H_1 = \liminf(p_{n+1} - p_n)$, so they guarantee that some even gap at most the bound recurs infinitely often. They do not identify which gap.

Julia Stadlmann posted arXiv version 1 of a preprint on 31 August 2026 claiming $H_1 \leq 240$ by combining the Bombieri-Vinogradov theorem with newer equidistribution estimates for smooth moduli. At verification on 3 September 2026, no peer-reviewed publication was identified. The body therefore labels 240 as a preprint claim and retains 246 as the published benchmark. This is the newest material in the manuscript and the strongest current-source risk.

The twin-prime conjecture and modern strong binary Goldbach conjecture remained open at verification. The latter emerged from a 1742 exchange whose formulations reflected the period's treatment of one, so the body distinguishes the historical correspondence from the modern statement. Helfgott's work proves that every odd integer greater than five is a sum of three primes. The formulation permits repeated primes. The body avoids treating finite computation for binary Goldbach as proof of every case.

The sieve parity problem is described narrowly. Broad classes of combinatorial sieves have difficulty separating numbers according to the parity of their count of prime factors, which obstructs direct extraction of some prime patterns. It is not a theorem that no future sieve-based argument can contribute to twin primes.

Primality, factorisation and certificates

For an odd composite n , the number of strong Miller-Rabin liar bases is at most one quarter of the relevant residue choices. Uniform independent rounds therefore give the familiar upper bound $4^{(-k)}$ on false acceptance, assuming correct implementation and sampling. Miller's earlier polynomial-time bound was conditional on the extended Riemann hypothesis for Dirichlet L-functions; Rabin's randomised version removed that hypothesis. Actual software may use deterministic base sets over bounded input ranges or combine tests.

Agrawal, Kayal and Saxena announced their deterministic polynomial-time primality result in 2002; the journal paper appeared in 2004. "PRIMES is in P" concerns polynomial growth in the input bit length. It does not say AKS is the fastest practical primality test.

No polynomial-time classical algorithm is known for general integer factorisation, and no theorem proves that one cannot exist. The general number field sieve is the leading classical method for large unstructured integers in the relevant size range, with a subexponential heuristic running-time model. Special forms can admit different algorithms.

For $M_p = 2^p - 1$ to be prime, p must be prime, but the converse fails. The Lucas-Lehmer test gives a necessary and sufficient condition for Mersenne numbers with prime exponent. The description of candidate screening, probable-prime testing, definitive Lucas-Lehmer testing and independent confirmation follows official GIMPS documentation and Crandall and Pomerance.

The Miller-Rabin discussion distinguishes compositeness witnesses from factors. A failed congruence may certify compositeness without yielding a useful divisor. A complete factorisation is another certificate whose product is cheap to verify.

RSA and quantum factoring

RSA uses a public modulus formed from selected large primes and private information derived from the factors. The security claim is narrower than “factoring is hard”: secure use also requires correct key generation, parameter selection, encoding, implementation and protocol context. The body does not claim that every conceivable RSA break is formally equivalent to factoring.

Shor's 1994 algorithm gives polynomial-time quantum algorithms for factoring and discrete logarithms on a suitable quantum computer. The conditional phrase “sufficiently capable fault-tolerant” matters. The mathematical algorithm is established; the machine resources and engineering timeline remain uncertain. Current cryptographic migration and protocol consequences belong primarily to *Cryptography in a Hurry* and *Quantum Computing in a Hurry*.

Historical sequence and attribution

The Euclidean algorithm appears in *Elements* Book VII. The ancient text survives through later manuscript traditions. Eratosthenes' sieve attribution is later. Weil and Dickson were used to check the movement from Greek divisibility through Fermat, Euler, Legendre and Gauss without forcing modern terminology onto earlier authors.

Euler's factorisation $2^{32} + 1 = 641 \times 6,700,417$ refuted Fermat's conjecture that every number $2^{(2^n)} + 1$ is prime. Euclid proved that a Mersenne prime produces an even perfect number; Euler proved that every even perfect number has that form. The existence of an odd perfect number remains open.

Riemann's 1859 paper is the primary source for the complex-analytic programme and hypothesis. Hadamard and de la Vallée Poussin proved the prime number theorem in 1896. Work by Selberg and Erdős produced elementary proofs in 1948, with separate papers appearing in 1949 amid a priority dispute. The body credits both and does not assign sole authorship.

The modern bounded-gap, Green-Tao, Goldbach, primality and quantum claims are tied to the research papers listed below. Current open-problem and record status was rechecked on 3 September 2026.

Bibliography

Primary texts and research papers

Agrawal, Manindra, Neeraj Kayal, and Nitin Saxena. "PRIMES Is in P." *Annals of Mathematics* 160, no. 2 (2004): 781-793.

D. H. J. Polymath. "Variants of the Selberg Sieve, and Bounded Intervals Containing Many Primes." *Research in the Mathematical Sciences* 1 (2014), article 12.

Erdős, Paul. "On a New Method in Elementary Number Theory Which Leads to an Elementary Proof of the Prime Number Theorem." *Proceedings of the National Academy of Sciences of the United States of America* 35, no. 7 (1949): 374-384.

Euclid. *The Thirteen Books of the Elements*. Translated with introduction and commentary by Thomas L. Heath. 2nd ed. Cambridge University Press, 1926. Reprinted by Dover, 1956.

Green, Ben, and Terence Tao. "The Primes Contain Arbitrarily Long Arithmetic Progressions." *Annals of Mathematics* 167, no. 2 (2008): 481-547.

Helfgott, Harald Andrés. "The Ternary Goldbach Conjecture Is True." arXiv:1312.7748, 2013, revised subsequently.

Maynard, James. "Small Gaps Between Primes." *Annals of Mathematics* 181, no. 1 (2015): 383-413.

Miller, Gary L. "Riemann's Hypothesis and Tests for Primality." *Journal of Computer and System Sciences* 13, no. 3 (1976): 300-317.

Rabin, Michael O. "Probabilistic Algorithm for Testing Primality." *Journal of Number Theory* 12, no. 1 (1980): 128-138.

Riemann, Bernhard. "Ueber die Anzahl der Primzahlen unter einer gegebenen Grösse." *Monatsberichte der Königlich Preussischen Akademie der Wissenschaften zu Berlin* (1859): 671-680.

Rivest, Ronald L., Adi Shamir, and Leonard Adleman. "A Method for Obtaining Digital Signatures and Public-Key Cryptosystems." *Communications of the ACM* 21, no. 2 (1978): 120-126.

Selberg, Atle. "An Elementary Proof of the Prime-Number Theorem." *Annals of Mathematics* 50, no. 2 (1949): 305-313.

Shor, Peter W. "Algorithms for Quantum Computation: Discrete Logarithms and Factoring." In *Proceedings of the 35th Annual Symposium on Foundations of Computer Science*, 124-134. IEEE, 1994.

Stadlmann, Julia. "Bounded Gaps Between Primes." arXiv:2608.31126, posted 31 August 2026. Preprint.

Zhang, Yitang. "Bounded Gaps Between Primes." *Annals of Mathematics* 179, no. 3 (2014): 1121-1174.

Books and scholarly syntheses

Crandall, Richard, and Carl Pomerance. *Prime Numbers: A Computational Perspective*. 2nd ed. Springer, 2005.

Davenport, Harold. *Multiplicative Number Theory*. 3rd ed. Revised by Hugh L. Montgomery. Springer, 2000.

Dickson, Leonard Eugene. *History of the Theory of Numbers*. Vol. 1, *Divisibility and Primality*. Carnegie Institution of Washington, 1919.

Edwards, H. M. *Riemann's Zeta Function*. Academic Press, 1974. Reprinted by Dover, 2001.

Hardy, G. H., and E. M. Wright. *An Introduction to the Theory of Numbers*. 6th ed. Revised by D. R. Heath-Brown and J. H. Silverman, with a foreword by Andrew Wiles. Oxford University Press, 2008.

Ireland, Kenneth, and Michael Rosen. *A Classical Introduction to Modern*

Number Theory. 2nd ed. Springer, 1990.

Mazur, Barry, and William Stein. *Prime Numbers and the Riemann Hypothesis*. Cambridge University Press, 2016.

Montgomery, Hugh L., and Robert C. Vaughan. *Multiplicative Number Theory I: Classical Theory*. Cambridge University Press, 2007.

Neale, Vicky. *Closing the Gap: The Quest to Understand Prime Numbers*. Oxford University Press, 2017.

Ribenboim, Paulo. *The Little Book of Bigger Primes*. 2nd ed. Springer, 2004.

Stillwell, John. *Elements of Number Theory*. Springer, 2003.

Weil, André. *Number Theory: An Approach through History from Hammurapi to Legendre*. Birkhäuser, 1984.

Current institutional sources

Clay Mathematics Institute. “Riemann Hypothesis.” Millennium Problems materials. Status checked 3 September 2026.

Great Internet Mersenne Prime Search. “GIMPS Discovers Largest Known Prime Number: $2^{136,279,841} - 1$ ” and current known Mersenne prime list. Checked 3 September 2026.