ABSOLUTE FERMAT PSEUDOPRIMES, THE DISTRIBUTION THEOREMS OF LARSEN, AND THEIR CONSEQUENCES FOR CRYPTOGRAPHIC PRIME GENERATION

ISHAAN MISHRA

Abstract. The Fermat primality test rejects a candidate integer n whenever it exhibits a base a with $a^{n-1} \not\equiv 1 \pmod{n}$. Its Achilles' heel is the existence of composite integers that satisfy Fermat's congruence for every coprime base: the *Carmichael numbers*, or absolute Fermat pseudoprimes, which are invisible to the test across its entire useful range of bases. This paper is a self-contained and rigorous exposition of these numbers, developed from the algebraic foundations up to their most recent distributional theory and its engineering significance. We begin with complete proofs of Fermat's little theorem and Euler's theorem, formalize the Fermat test through a taxonomy of witnesses and liars, and establish the equivalence between the coprime-base definition of a Carmichael number and the universal congruence $a^n \equiv a \pmod{n}$. We then prove the classical structural constraints—oddness, at least three distinct prime factors, squarefreeness, and $\gcd(n, \varphi(n)) = 1$—and give a two-directional proof of Korselt's criterion. Turning to distribution, we state the Alford–Granville–Pomerance infinitude theorem with its lower bound $C(x) > x^{2/7}$, and describe Daniel Larsen's two breakthroughs: a Bertrand-type theorem placing Carmichael numbers in short intervals, obtained via Maynard's multidimensional sieve and Fourier analysis on finite abelian groups, and a complete classification of the arithmetic progressions containing infinitely many Carmichael numbers. Finally, we translate this theory into cryptographic terms: we prove the group-theoretic Fermat witness bound, explain the Carmichael blind spot and how Miller–Rabin removes it with a uniform $(1/4)^k$ error bound, and give a rigorous analysis—including a no-go obstruction—of the extent to which Larsen's progression dichotomy can inform an optimized prime-generation pipeline.

1. Introduction

Modern public-key cryptography rests on the routine ability to produce large prime numbers. The RSA cryptosystem [18], for instance, requires two secret primes of several hundred decimal digits each; these are generated on demand, typically by drawing a random integer of the desired size and subjecting it to a primality test. The efficiency and correctness of this test therefore sit at the foundation of the entire construction, and the mathematics that governs which composite integers can masquerade as primes is not a curiosity but a security-relevant concern.

The oldest primality test, trial division, is deterministic and unconditionally correct but hopelessly slow: to certify a 600-digit number one would in the worst case divide by primes up to its square root, a quantity vastly larger than the number of atoms in the observable universe. What makes prime generation feasible in practice is the shift from *proving* primality by exhaustion to *testing* it probabilistically, using an algebraic identity that all primes satisfy and that most composites violate. The identity is Fermat's little theorem: if p is prime and $p \nmid a$, then $a^{p-1} \equiv 1 \pmod{p}$. Its contrapositive yields a fast test for compositeness—exhibit

Date: July 2026.

a base a that breaks the congruence and n is proved composite—and this test can be executed in time polynomial in the number of digits of n by fast modular exponentiation.

The difficulty is that the converse of Fermat's little theorem is false. There exist composite integers n for which $a^{n-1} \equiv 1 \pmod{n}$ holds for *every* base a coprime to n. Such an integer defeats the Fermat test completely: no coprime base can ever reveal its compositeness, and the only bases that betray it are those sharing a factor with n, which for a number with only a few enormous prime factors are exceedingly rare. These are the *Carmichael numbers*, the absolute Fermat pseudoprimes, and they are the central object of this paper. The smallest is $561 = 3 \cdot 11 \cdot 17$; it was first exhibited by Šimerka in 1885 [19], rediscovered by Carmichael in the early twentieth century [3, 4], and named after the latter by Beeger in 1950 [6].

For a long time even the most basic global question about these numbers was open: are there infinitely many? This was settled affirmatively in 1994 by Alford, Granville, and Pomerance [1], who proved not merely infinitude but the quantitative lower bound $C(x) > x^{2/7}$ for the count $C(x)$ of Carmichael numbers up to x. Their method, however, is intrinsically non-local: it produces Carmichael numbers of some size in a large range but cannot control where they land. The two questions that this leaves—whether Carmichael numbers appear in every sufficiently long short interval, and exactly which arithmetic progressions contain them—were resolved only recently by Daniel Larsen [12, 13], whose work forms the modern heart of this exposition.

The purpose of the paper is to develop this material rigorously and in order, so that each result is a consequence of what precedes it, and then to trace its consequences for the engineering of primality tests and prime generators. Section 2 builds the algebraic foundations: Fermat's little theorem and Euler's theorem with complete proofs, the formalization of the Fermat test, and the definitions of pseudoprimes and absolute pseudoprimes. Section 3 establishes the structural theory of Carmichael numbers, culminating in a two-directional proof of Korselt's criterion. Section 4 treats their global distribution: the Alford–Granville–Pomerance theorem and Larsen's two theorems. Section 5 returns to computer science, developing the witness taxonomy, the group-theoretic witness bound, the Carmichael vulnerability, and the Miller–Rabin remedy. Section 6 analyzes optimized prime generation and the precise sense in which Larsen's progression dichotomy does—and does not—bear on it. Sections 7 and 8 synthesize and conclude.

2. Primality, Pseudoprimality, and Fermat's Little Theorem

2.1. **Notation and standing conventions.** Throughout, n denotes an integer greater than 1, and p and q denote primes. We write $a \mid b$ for "a divides b", $\gcd(a, b)$ or simply (a, b) for the greatest common divisor, and lcm for the least common multiple. The ring of integers modulo n is $\mathbb{Z}/n\mathbb{Z}$, and its group of units—the residues coprime to n under multiplication—is $(\mathbb{Z}/n\mathbb{Z})^\times$. Euler's totient function $\varphi(n) = \#\{1 \le a \le n : \gcd(a, n) = 1\} = |(\mathbb{Z}/n\mathbb{Z})^\times|$ counts these units. The Carmichael function $\lambda(n)$ denotes the exponent of $(\mathbb{Z}/n\mathbb{Z})^\times$, that is, the largest order of an element of $(\mathbb{Z}/n\mathbb{Z})^\times$. We say n is *squarefree* if no prime square divides it, and $\omega(n)$ denotes the number of distinct prime factors of n. All congruences are taken modulo the indicated integer.

2.2. **Fermat's little theorem and Euler's theorem.** The engine of every test we discuss is the following classical theorem.

Theorem 2.1 (Fermat's little theorem)**.** *Let p be prime and let a be an integer with $p \nmid a$. Then*

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

Equivalently, for every integer a one has $a^p \equiv a \pmod{p}$.

Proof. We prove the multiplicative form by the classical argument of permuting a reduced residue system. Consider the $p-1$ integers

$$a,\ 2a,\ 3a,\ \ldots,\ (p-1)a.$$

We claim their residues modulo p are exactly the residues $1, 2, \ldots, p-1$ in some order. First, none is divisible by p: if $p \mid ja$ for some $1 \le j \le p-1$, then since p is prime and $p \nmid a$ we would have $p \mid j$, impossible for $0 < j < p$. Second, the residues are pairwise distinct: if $ia \equiv ja \pmod{p}$ with $1 \le i, j \le p-1$, then because $p \nmid a$ we may cancel a (as a is a unit modulo p) to obtain $i \equiv j \pmod{p}$, hence $i = j$. Thus the map $j \mapsto ja \bmod p$ is an injection from $\{1, \ldots, p-1\}$ to itself, therefore a bijection, and

$$\prod_{j=1}^{p-1} (ja) \equiv \prod_{j=1}^{p-1} j \pmod{p}.$$

The left-hand side equals $a^{p-1}\,(p-1)!$ and the right-hand side equals $(p-1)!$, so

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

Since p is prime, each of $1, 2, \ldots, p-1$ is coprime to p, so $(p-1)!$ is coprime to p and may be cancelled. This yields $a^{p-1} \equiv 1 \pmod{p}$, as claimed. Multiplying by a gives $a^p \equiv a \pmod{p}$ when $p \nmid a$; and when $p \mid a$ both sides are $\equiv 0$, so $a^p \equiv a \pmod{p}$ holds for every integer a. □

Fermat's theorem is the special case of a statement valid for every modulus, in which the exponent $p-1$ is replaced by the order of the unit group.

Theorem 2.2 (Euler)**.** *Let $n > 1$ and let a be an integer with $\gcd(a, n) = 1$. Then*

$$a^{\varphi(n)} \equiv 1 \pmod{n}.$$

Proof. Let $u_1, u_2, \ldots, u_{\varphi(n)}$ be the distinct residues in $\{1, \ldots, n-1\}$ that are coprime to n; these represent the elements of the group $(\mathbb{Z}/n\mathbb{Z})^\times$. Because $\gcd(a, n) = 1$, multiplication by a maps units to units: for each i, $\gcd(au_i, n) = 1$, so $au_i \bmod n$ is again one of the u_j. This map is injective on $(\mathbb{Z}/n\mathbb{Z})^\times$, since $au_i \equiv au_j \pmod{n}$ permits cancellation of the unit a and forces $u_i \equiv u_j$, hence $i = j$. An injection of a finite set to itself is a bijection, so the multiset $\{au_1 \bmod n, \ldots, au_{\varphi(n)} \bmod n\}$ is a permutation of $\{u_1, \ldots, u_{\varphi(n)}\}$. Taking the product of all elements in each description,

$$\prod_{i=1}^{\varphi(n)} (au_i) \equiv \prod_{i=1}^{\varphi(n)} u_i \pmod{n}, \qquad \text{i.e.} \qquad a^{\varphi(n)} \prod_{i=1}^{\varphi(n)} u_i \equiv \prod_{i=1}^{\varphi(n)} u_i \pmod{n}.$$

The product $U = \prod_i u_i$ is a product of units, hence itself a unit modulo n, and may be cancelled. We conclude $a^{\varphi(n)} \equiv 1 \pmod{n}$. □

Remark 2.3. In the language of group theory, Theorem 2.2 is Lagrange's theorem applied to the finite abelian group $(\mathbb{Z}/n\mathbb{Z})^\times$ of order $\varphi(n)$: the order of any element divides the group order, so $a^{\varphi(n)} = (a^{\operatorname{ord}(a)})^{\varphi(n)/\operatorname{ord}(a)} = 1$. Fermat's little theorem is the case $n = p$, where $\varphi(p) = p - 1$. This group-theoretic viewpoint is the one we exploit repeatedly below.

2.3. **The Fermat primality test, witnesses, and liars.** Fermat's little theorem is a statement that all primes obey. Reading its contrapositive gives a test for compositeness.

Definition 2.4 (Fermat test)**.** Let $n > 1$ and let a be an integer with $1 \le a \le n-1$. If

$$a^{n-1} \not\equiv 1 \pmod{n},$$

then n is composite, and a is said to *witness* this. The *Fermat primality test* consists of choosing bases a (deterministically or at random) and evaluating $a^{n-1} \bmod n$; if any base yields a value other than 1, the test outputs COMPOSITE, and otherwise it outputs PROBABLY PRIME.

The soundness of the test is immediate: by Theorem 2.1, a prime n satisfies $a^{n-1} \equiv 1 \pmod{n}$ for all $1 \le a \le n-1$, so the test never declares a prime composite. The subtlety is entirely on the other side—the test may fail to detect a composite.

Definition 2.5 (Fermat witness and Fermat liar)**.** Let n be composite. A base a with $1 \le a \le n-1$ is a *Fermat witness* for n if $a^{n-1} \not\equiv 1 \pmod{n}$; such a base proves n composite. A base a with $\gcd(a, n) = 1$ and $a^{n-1} \equiv 1 \pmod{n}$ is a *Fermat liar* for n; such a base fails to detect the compositeness of n.

We follow the convention of [7] in which any a with $a^{n-1} \not\equiv 1 \pmod{n}$ counts as a Fermat witness regardless of $\gcd(a, n)$; a base sharing a factor with n is automatically a witness, since $\gcd(a, n) > 1$ forces $a^{n-1} \not\equiv 1 \pmod{n}$ (a unit power can never be a non-unit). The bases that matter for the strength of the test are the coprime ones, which split into liars and non-obvious witnesses.

2.4. **Fermat pseudoprimes and absolute Fermat pseudoprimes.** If a composite n has a Fermat liar a, then relative to that particular base n imitates a prime.

Definition 2.6 (Fermat pseudoprime)**.** A composite integer n is a *Fermat pseudoprime to base* a if $\gcd(a, n) = 1$ and $a^{n-1} \equiv 1 \pmod{n}$.

Example. The number $341 = 11 \cdot 31$ is composite yet $2^{340} \equiv 1 \pmod{341}$, so 341 is a Fermat pseudoprime to base 2. It is not, however, a pseudoprime to base 3: one computes $3^{340} \equiv 56 \not\equiv 1 \pmod{341}$, so the base $a = 3$ is a Fermat witness that unmasks 341 after a single additional trial. The pseudoprimality of 341 is thus *base-specific*: a fixed small set of alternative bases exposes it.

Example 2.4 shows the ordinary failure mode of the Fermat test—a composite that fools one base but not another. The dangerous failure mode is a composite that fools *all* coprime bases simultaneously.

Definition 2.7 (Absolute Fermat pseudoprime / Carmichael number)**.** A composite integer n is an *absolute Fermat pseudoprime*, or a *Carmichael number*, if

$$a^{n-1} \equiv 1 \pmod{n} \qquad \text{for every integer } a \text{ with } \gcd(a, n) = 1.$$

Equivalently, n is composite and has no Fermat liar that fails; every coprime base is a liar.

The two definitions in play—the coprime-base congruence of Definition 2.7 and the universal congruence $a^n \equiv a \pmod{n}$ valid for *all* integers—are equivalent, a fact we now prove. The universal form is often taken as the definition because it dispenses with the coprimality hypothesis.

Proposition 2.8 (Equivalence of the two formulations)**.** *Let n be composite. Then n is a Carmichael number in the sense of Definition 2.7 if and only if*

$$a^n \equiv a \pmod{n} \qquad \text{for every integer } a \in \mathbb{Z}.$$

Proof. Suppose first that $a^n \equiv a \pmod{n}$ for all $a \in \mathbb{Z}$. If $\gcd(a, n) = 1$, then a is a unit modulo n and may be cancelled from $a^n \equiv a \pmod{n}$ to give $a^{n-1} \equiv 1 \pmod{n}$. Since n is composite, it is a Carmichael number.

For the converse, assume n is a Carmichael number. We will use the fact—established in Corollary 3.6 below and depending only on Definition 2.7, not on the present proposition—that n is squarefree, say $n = p_1 \cdots p_r$ with the p_i distinct primes. To prove $a^n \equiv a \pmod{n}$ for a given integer a, it suffices by the Chinese Remainder Theorem to prove $a^n \equiv a \pmod{p_i}$ for each i. Fix a prime $p = p_i$. If $p \mid a$, then $a^n \equiv 0 \equiv a \pmod{p}$ trivially. If $p \nmid a$, then Fermat's little theorem gives $a^{p-1} \equiv 1 \pmod{p}$; and since $(p-1) \mid (n-1)$ for every prime factor of a Carmichael number (Theorem 3.5 below, again independent of this proposition), we may write $n - 1 = (p-1)t$ and compute

$$a^{n-1} = (a^{p-1})^t \equiv 1^t = 1 \pmod{p},$$

whence $a^n \equiv a \pmod{p}$. As this holds for each prime factor p_i, the Chinese Remainder Theorem yields $a^n \equiv a \pmod{n}$. □

Remark 2.9 (The nature of the vulnerability). The contrast between Definitions 2.6 and 2.7 is the crux of this paper. A base-a Fermat pseudoprime is a local deception: it survives base a but is caught by some other base, as 341 is caught by 3. A Carmichael number is a global deception: it survives *every* coprime base at once. From the standpoint of a tester who samples coprime bases at random, a Carmichael number is a systematic blind spot, indistinguishable from a prime no matter how many bases are tried—unless one happens to sample a base sharing a factor with n, which for a Carmichael number with a few large prime factors is vanishingly unlikely. Making this quantitative is the subject of Section 5.

3. The Structure of Carmichael Numbers and Korselt's Criterion

Definition 2.7 is a statement about infinitely many congruences and is not directly checkable. The theory of this section replaces it by a finite condition on the prime factorization—Korselt's criterion—and along the way extracts the rigid structural constraints that every Carmichael number must obey. These constraints are what later make the numbers both rare and, paradoxically, tractable to construct.

3.1. **Elementary structural constraints.**

Proposition 3.1 (Oddness)**.** *Every Carmichael number n is odd.*

Proof. Since $\gcd(n-1, n) = 1$, the base $a = n - 1$ is coprime to n, so Definition 2.7 gives $(n-1)^{n-1} \equiv 1 \pmod{n}$. But $n - 1 \equiv -1 \pmod{n}$, so this says $(-1)^{n-1} \equiv 1 \pmod{n}$. Now $(-1)^{n-1}$ equals $+1$ or -1 as an integer. If it were -1, we would have $-1 \equiv 1 \pmod{n}$, i.e. $n \mid 2$, forcing $n \leq 2$; but a Carmichael number is composite and hence at least 4. Therefore $(-1)^{n-1} = +1$, which means $n - 1$ is even and n is odd. □

Proposition 3.2 (At least three distinct prime factors)**.** *Every Carmichael number n has at least three distinct prime factors.*

Proof. By Corollary 3.6 below, n is squarefree, so it is a product of distinct primes; being composite, it has at least two. Suppose, for contradiction, that it had exactly two, say $n = pq$ with primes $p > q$. Korselt's criterion (Theorem 3.5) gives $(p-1) \mid (n-1)$. Compute, using $n - 1 = pq - 1$,

$$\frac{n-1}{p-1} = \frac{pq-1}{p-1} = \frac{(p-1)q + (q-1)}{p-1} = q + \frac{q-1}{p-1}.$$

For $(p-1) \mid (n-1)$ the fractional part $\frac{q-1}{p-1}$ must be an integer, i.e. $(p-1) \mid (q-1)$. But $0 < q - 1 < p - 1$, and a positive integer smaller than $p-1$ cannot be divisible by $p-1$. This contradiction shows n cannot have exactly two prime factors, so it has at least three, all distinct by squarefreeness. □

Proposition 3.3 (Coprimality with the totient)**.** *Every Carmichael number n satisfies* $\gcd(n, \varphi(n)) = 1$.

Proof. Write $n = p_1 \cdots p_r$ with distinct primes p_i (squarefree by Corollary 3.6, with $r \geq 3$ by Proposition 3.2). Since the p_i are distinct, $\varphi(n) = \prod_{i=1}^{r}(p_i - 1)$. Suppose $\gcd(n, \varphi(n)) > 1$. Then some prime factor p_i of n divides $\varphi(n) = \prod_j (p_j - 1)$, and since p_i is prime it must divide one of the factors, say $p_i \mid (p_j - 1)$ for some index j. By Korselt's criterion, $(p_j - 1) \mid (n-1)$, and by transitivity of divisibility $p_i \mid (n-1)$. But $p_i \mid n$ as well, so p_i divides both n and $n-1$, hence divides their difference 1—impossible for a prime. Therefore $\gcd(n, \varphi(n)) = 1$. □

Remark 3.4. The converse implication is worth recording: $\gcd(n, \varphi(n)) = 1$ forces n to be squarefree, because if $p^2 \mid n$ then $p \mid \varphi(n)$ (the prime p divides $\varphi(p^2) = p(p-1)$ and hence $\varphi(n)$), giving $p \mid \gcd(n, \varphi(n))$. Thus squarefreeness is subsumed by the totient-coprimality condition [6].

3.2. **Korselt's criterion.** We now state and prove the theorem, due to Korselt in 1899 [11], that converts the infinite congruence condition of Definition 2.7 into two finite arithmetic conditions on the factorization of n. Korselt found the criterion a decade before Carmichael exhibited examples, but produced none himself, which is why the numbers bear Carmichael's name [6].

Theorem 3.5 (Korselt's criterion)**.** *A composite integer $n > 1$ is a Carmichael number if and only if*

(i) *n is squarefree, and*
(ii) *$(p-1) \mid (n-1)$ for every prime p dividing n.*

Proof. **(⇒) Necessity.** Assume n is a Carmichael number.

Squarefreeness. Suppose toward a contradiction that n is not squarefree, so $p^2 \mid n$ for some prime p. Write $n = p^k n_0$ with $k \geq 2$ and $p \nmid n_0$. By the Chinese Remainder Theorem, choose an integer a with

$$a \equiv 1 + p \pmod{p^k} \qquad \text{and} \qquad a \equiv 1 \pmod{n_0}.$$

Then $\gcd(a, n) = 1$: modulo p we have $a \equiv 1 \not\equiv 0$, and modulo any prime dividing n_0 we have $a \equiv 1$. By Definition 2.7, $a^{n-1} \equiv 1 \pmod{n}$, and in particular $a^{n-1} \equiv 1 \pmod{p^2}$ since $p^2 \mid n$. Now reduce $a \equiv 1 + p \pmod{p^2}$ (which follows from $a \equiv 1 + p \pmod{p^k}$ and $k \geq 2$) and expand by the binomial theorem modulo p^2:

$$a^{\,n-1} \equiv (1+p)^{\,n-1} = \sum_{j=0}^{n-1} \binom{n-1}{j} p^j \equiv 1 + (n-1)p \pmod{p^2},$$

because every term with $j \geq 2$ carries a factor $p^j \equiv 0 \pmod{p^2}$. Since $p \mid n$, we have $n \equiv 0 \pmod{p}$, so $(n-1)p \equiv -p \pmod{p^2}$, giving

$$a^{n-1} \equiv 1 - p \pmod{p^2}.$$

Combining with $a^{n-1} \equiv 1 \pmod{p^2}$ yields $1 - p \equiv 1 \pmod{p^2}$, i.e. $p \equiv 0 \pmod{p^2}$, which is false because $0 < p < p^2$. This contradiction proves n is squarefree.

The divisibility condition. Let p be any prime dividing n. Because n is squarefree, $\gcd(p, n/p) = 1$. A primitive root exists modulo any prime, so choose an integer b whose residue modulo p has multiplicative order exactly $p-1$. By the Chinese Remainder Theorem choose a with

$$a \equiv b \pmod{p} \qquad \text{and} \qquad a \equiv 1 \pmod{n/p}.$$

Then $\gcd(a, n) = 1$, so Definition 2.7 gives $a^{n-1} \equiv 1 \pmod{n}$, and reducing modulo p gives $b^{n-1} \equiv 1 \pmod{p}$. Since b has order $p-1$ modulo p, the exponent $n-1$ must be a multiple of that order: $(p-1) \mid (n-1)$. As p was an arbitrary prime factor of n, condition (ii) holds.

($\Leftarrow$) Sufficiency. Assume n is composite, squarefree, and $(p-1) \mid (n-1)$ for every prime $p \mid n$. Let a be any integer with $\gcd(a, n) = 1$. Fix a prime $p \mid n$. Then $\gcd(a, p) = 1$, so Fermat's little theorem gives $a^{p-1} \equiv 1 \pmod{p}$. Writing $n - 1 = (p-1)t$ using condition (ii),

$$a^{n-1} = \left(a^{p-1}\right)^t \equiv 1^t = 1 \pmod{p}.$$

This congruence holds for every prime p dividing n. Because n is squarefree, it is the product of these distinct primes, so by the Chinese Remainder Theorem the simultaneous congruences $a^{n-1} \equiv 1 \pmod{p}$ combine to $a^{n-1} \equiv 1 \pmod{n}$. Since n is composite by hypothesis, n is a Carmichael number. □

Corollary 3.6 (Squarefreeness, extracted)**.** *Every Carmichael number is squarefree.*

Proof. This is the first half of the necessity direction of Theorem 3.5, whose proof uses only Definition 2.7 and no result that depends on it; hence the earlier forward references are not circular. □

Example (Chernick's construction)*.* Korselt's criterion makes it possible to manufacture Carmichael numbers. Chernick [5] observed that if k is a positive integer for which $6k+1$, $12k+1$, and $18k+1$ are all prime, then

$$n = (6k+1)(12k+1)(18k+1)$$

is a Carmichael number. Squarefreeness is clear. For the divisibility condition it suffices, since the three factors are distinct primes, to verify $(6k) \mid (n-1)$, $(12k) \mid (n-1)$, and $(18k) \mid (n-1)$. Reducing n modulo $12k$ gives $n \equiv (6k+1)(6k+1)(0+1) \equiv 1 \pmod{12k}$ after noting $18k+1 \equiv 6k+1 \pmod{12k}$ and $(6k+1)^2 = 36k^2 + 12k + 1 \equiv 1 \pmod{12k}$; similarly $n \equiv (6k+1)(12k+1) \equiv 1 \pmod{18k}$; and $n \equiv 1 \pmod{6k}$ follows because $6k \mid 12k$. Hence each of $6k$, $12k$, $18k$ divides $n-1$, and n is Carmichael. Taking $k = 1$ gives the primes $7, 13, 19$ and the Carmichael number $7 \cdot 13 \cdot 19 = 1729$.

Example (The smallest Carmichael numbers)*.* The first five Carmichael numbers are

$$561 = 3 \cdot 11 \cdot 17, \quad 1105 = 5 \cdot 13 \cdot 17, \quad 1729 = 7 \cdot 13 \cdot 19, \quad 2465 = 5 \cdot 17 \cdot 29, \quad 2821 = 7 \cdot 13 \cdot 31.$$

Each has exactly three prime factors, consistent with Proposition 3.2, and each is odd and squarefree. For 561, Korselt's condition reads $(3-1) \mid 560$, $(11-1) \mid 560$, $(17-1) \mid 560$, i.e. $2, 10, 16$ all divide $560 = 2^4 \cdot 5 \cdot 7$, which they do.

4. The Global Distribution of Carmichael Numbers

Having characterized Carmichael numbers structurally, we ask how they are distributed among the integers. Let

$$C(x) = \#\{\, n \le x : n \text{ is a Carmichael number} \,\}.$$

The behavior of $C(x)$ was, for most of the twentieth century, understood only through upper bounds and heuristics; the decisive lower-bound results are recent.

4.1. Infinitude: the Alford–Granville–Pomerance theorem.

Theorem 4.1 (Alford–Granville–Pomerance, 1994 [1])**.** *There are infinitely many Carmichael numbers. More precisely, there is an x_0 such that for all $x \ge x_0$,*

$$C(x) > x^{2/7}.$$

We do not reproduce the proof, which is long and analytic; we describe its architecture because Larsen's later work is best understood as overcoming a specific limitation of it. The strategy [1,12] is to choose a highly composite modulus L such that the unit group $(\mathbb{Z}/L\mathbb{Z})^\times$ has small exponent—concretely, L is built as a product of primes p for which $p-1$ has only small prime factors, so that $\lambda(L) = \operatorname{lcm}_{p\mid L}(p-1)$ is small. One then seeks an integer k for which many of the linear forms $dk+1$, as d ranges over divisors of L, are simultaneously prime. Character theory on the finite abelian group $(\mathbb{Z}/L\mathbb{Z})^\times$ produces a subset of these primes whose product Π is $\equiv 1 \pmod{L}$; if k is chosen coprime to L, one arranges $\Pi \equiv 1 \pmod{kL}$. Since each prime factor $p = dk+1$ of Π satisfies $(p-1) = dk \mid kL$, and $\Pi \equiv 1 \pmod{kL}$ gives $(p-1) \mid (\Pi - 1)$, Korselt's criterion certifies Π as a Carmichael number. A pigeonhole argument over the pairs (k, d) guarantees enough such Π to force the bound $x^{2/7}$.

Remark 4.2. The exponent $2/7$ is not optimal; it has since been improved to $C(x) > x^{1/3}$ by Harman [10]. The structural point for our purposes is that the pigeonhole step relinquishes control over the *sizes* of the primes $dk+1$, and hence over the size of the resulting Π. The construction can guarantee Carmichael numbers somewhere in a wide range but cannot confine them to a prescribed short interval. This is precisely the obstacle that Larsen's first theorem removes.

4.2. Larsen's first theorem: a Bertrand postulate for Carmichael numbers.

Bertrand's postulate asserts that for every $x > 1$ there is a prime between x and $2x$; equivalently, primes never leave a multiplicative gap. Alford, Granville, and Pomerance explicitly asked whether the analogue holds for Carmichael numbers—whether there is always a Carmichael number between x and $2x$ for large x [12]. Larsen answered this in the affirmative, and in a far stronger quantitative form: Carmichael numbers appear not merely in every interval $[x, 2x]$ but in intervals that shrink toward x.

Theorem 4.3 (Larsen, 2022 [12])**.** *For every $\delta > 0$ and every x sufficiently large in terms of δ, the number of Carmichael numbers in the interval*

$$\left(x,\ x + \frac{x}{(\log x)^{\frac{1}{2}+\delta}}\right)$$

is at least

$$\exp\left(\frac{\log x}{(\log\log x)^{2+\delta}}\right).$$

In particular, the ratio of consecutive Carmichael numbers tends to 1*, and a Bertrand-type postulate holds: for all sufficiently large x there is a Carmichael number between x and $2x$.*

The mechanism by which Larsen defeats the size-control problem of Theorem 4.1 is a substitution of modern sieve technology for the pigeonhole principle. In the Alford–Granville–Pomerance argument, one has no lower bound on how many admissible primes $dk+1$ fall into any given size band, so the primes may all cluster at one scale and their products may miss a target interval. Larsen partitions the divisors of L into M-tuples and, invoking Maynard's 2016 sieve for dense clusters of primes in subsets [15]—itself a descendant of the bounded-gaps breakthroughs of Zhang [21] and Maynard [14]—obtains quantitative lower bounds on the number of triples (k, d, d') with d, d' in a common tuple and both $dk+1$ and $d'k+1$ prime. Crucially, the sieve can be arranged so that the resulting primes do not cluster near powers of any fixed $\alpha > 1+\epsilon$; that is, the primes are guaranteed to be spread across scales. A standard argument in Fourier analysis on the finite abelian group $(\mathbb{Z}/kL\mathbb{Z})^\times$ then selects, from among the products of these primes, ones whose magnitude lands in a controllably narrow interval, each such product being Carmichael by the Korselt mechanism of Theorem 4.1. Sliding the starting parameter to infinity yields the stated density in short intervals. The interplay is worth emphasizing: the classical algebra (characters, Korselt) supplies the *Carmichael-ness*, while the analytic sieve supplies the *location*.

4.3. **Larsen's second theorem: all admissible arithmetic progressions.** The finest distributional question is which residue classes contain Carmichael numbers. For primes, Dirichlet's theorem gives a clean answer: the progression $r \pmod m$ contains infinitely many primes exactly when $\gcd(r, m) = 1$. For Carmichael numbers the answer is more delicate, because unlike a prime a Carmichael number can be divisible by a fixed m while remaining composite. Larsen's second paper resolves the question completely.

Definition 4.4 (Carmichael compatibility [13])**.** Let $r \pmod m$ be an arithmetic progression. Set $g = \gcd(r, m)$, let $\lambda(g)$ be the maximal order of an element of $(\mathbb{Z}/g\mathbb{Z})^\times$, and set $h = \gcd(\lambda(g), m)$. The progression $r \pmod m$ is *Carmichael incompatible* if at least one of the following holds:

(a) $\gcd\big(g,\, 2\varphi(g)\big) > 1$;
(b) $h \nmid (r-1)$;
(c) $36 \mid m$, $\quad r \equiv 3 \pmod{12}$, $\quad$ and $\quad r/g \equiv 5 \text{ or } 7 \pmod{12}$.

Otherwise the progression is *Carmichael compatible*.

Theorem 4.5 (Larsen, 2025 [13])**.** *An arithmetic progression contains infinitely many Carmichael numbers if it is Carmichael compatible, and none at all if it is Carmichael incompatible. Quantitatively, if $r \pmod m$ is Carmichael compatible, then for every $\epsilon > 0$ and every x sufficiently large in terms of ϵ, the number of Carmichael numbers $n \le x$ with $n \equiv r \pmod m$ exceeds $x^{1/168-\epsilon}$. In particular, if $\gcd(m, 2\varphi(m)) = 1$ then there are infinitely many Carmichael numbers divisible by m.*

The necessity half—that an incompatible progression contains no Carmichael numbers—is elementary and we prove it in Section 6, since it is exactly the ingredient the cryptographic analysis requires. The sufficiency half, that every compatible progression contains infinitely many, is the deep direction and rests on the sieve and Fourier machinery developed for Theorem 4.3, adapted to prescribe residues.

Corollary 4.6 (Resolution of an Alford–Granville–Pomerance question [13])**.** $\liminf_{\substack{n\to\infty \\ n \text{ Carmichael}}} \dfrac{\varphi(n)}{n} = 0$.

Corollary 4.6 follows because Theorem 4.5 permits Carmichael numbers divisible by any m with $\gcd(m, 2\varphi(m)) = 1$; choosing m to be a product of many such primes drives $\varphi(n)/n \le \varphi(m)/m$ arbitrarily close to 0. The corollary answers a question left open by Alford, Granville, and Pomerance about how "prime-poor" a Carmichael number can be.

Example. Since $\gcd(m, \varphi(m)) = 1$ holds for $m = 3, 5, 7, 15, 35$, there are infinitely many Carmichael numbers divisible by each of these. By contrast, no Carmichael number is divisible by 21, because $\gcd(21, \varphi(21)) = \gcd(21, 12) = 3 > 1$; this is the incompatibility condition (a) applied to the progression $0 \pmod{21}$, and it also follows directly from Proposition 3.3, since $21 \mid n$ would force $\gcd(n, \varphi(n)) \ge 3$.

5. Cryptographic Primality Testing

We now return to the computational setting and quantify the phenomena of Section 2. The guiding question is how many bases a probabilistic test must sample to be confident of a candidate's primality, and how Carmichael numbers sabotage this analysis for the Fermat test.

5.1. A taxonomy of witnesses.

Compositeness of n can be certified by bases of increasing subtlety. We collect the standard notions [7, 9].

Definition 5.1 (Witness taxonomy)**.** Let n be composite and $1 \le a \le n-1$.

- a is a *divisor (trial-division) witness* if $a > 1$ and $a \mid n$.
- a is a *gcd witness* if $\gcd(a, n) > 1$.
- a is a *Fermat witness* if $a^{n-1} \not\equiv 1 \pmod{n}$.
- a is an *Euler witness* (in the Solovay–Strassen sense) if $\gcd(a, n) = 1$ and $a^{(n-1)/2} \not\equiv \left(\frac{a}{n}\right) \pmod{n}$, where $\left(\frac{a}{n}\right)$ is the Jacobi symbol.
- a is a *root (strong / Miller–Rabin) witness* if it certifies compositeness in the Miller–Rabin test of Definition 5.6 below—that is, a exposes a nontrivial square root of unity modulo n.

These classes are nested. For $a \ge 2$ the implications

$$a^{n-1} \equiv 1 \pmod{n} \implies \gcd(a, n) = 1 \implies a \nmid n$$

give, on passing to contrapositives,

$$\underbrace{a \mid n}_{\text{divisor}} \implies \underbrace{\gcd(a, n) > 1}_{\text{gcd}} \implies \underbrace{a^{n-1} \not\equiv 1}_{\text{Fermat}},$$

so every divisor witness is a gcd witness, and every gcd witness is a Fermat witness. Moreover every Fermat witness is an Euler witness, and every Euler witness is a strong witness [9]. The chain of inclusions among the *witness sets* is therefore

$$\{\text{divisor}\} \subseteq \{\text{gcd}\} \subseteq \{\text{Fermat}\} \subseteq \{\text{Euler}\} \subseteq \{\text{strong}\},$$

and it is exactly the widening of these sets that makes successive tests more powerful. Figure 1 depicts the hierarchy.

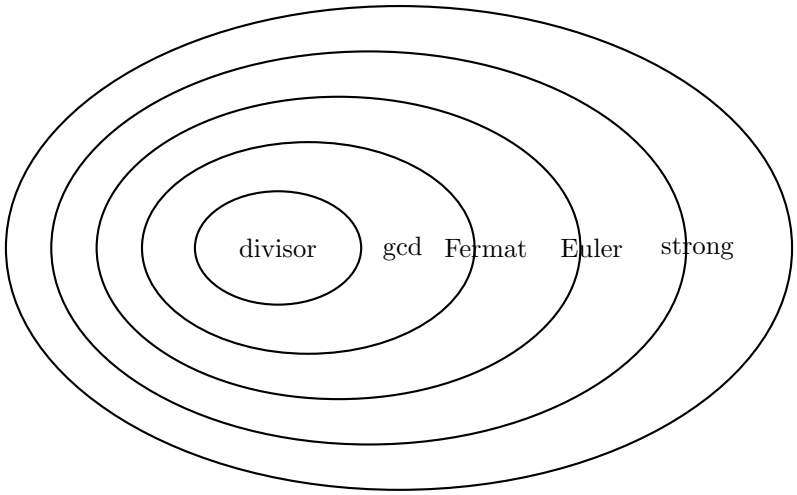

Figure 1. The nested hierarchy of witnesses to compositeness. Each larger class contains the previous one; a base that certifies compositeness by a weaker test certifies it by all stronger tests. The gaps between the sets are where the tests differ in power, and the outermost gap is what Carmichael numbers exploit against the Fermat test.

5.2. **The Fermat witness bound.** The value of the Fermat test rests on the following dichotomy: for a composite that is not Carmichael, Fermat witnesses are abundant.

Theorem 5.2 (Fermat witness bound)**.** *Let n be composite. If there exists a base b with* $\gcd(b, n) = 1$ *and* $b^{n-1} \not\equiv 1 \pmod{n}$*—that is, if n is not a Carmichael number—then more than half of the residues in* $\{1, 2, \ldots, n-1\}$ *are Fermat witnesses for n.*

Proof. Work inside the group $G = (\mathbb{Z}/n\mathbb{Z})^\times$ of units modulo n, which has order $\varphi(n)$. Consider the subset

$$A = \{\, a \in G : a^{n-1} \equiv 1 \pmod{n} \,\}.$$

We claim A is a subgroup of G. It contains 1; it is closed under multiplication, since $a^{n-1} \equiv 1$ and $c^{n-1} \equiv 1$ give $(ac)^{n-1} = a^{n-1}c^{n-1} \equiv 1$; and it is closed under inverses, since $a^{n-1} \equiv 1$ gives $(a^{-1})^{n-1} = (a^{n-1})^{-1} \equiv 1$. Thus $A \leq G$.

By hypothesis there is a unit $b \in G$ with $b^{n-1} \not\equiv 1$, so $b \in G \setminus A$; hence A is a *proper* subgroup of G. By Lagrange's theorem the index $[G : A]$ is an integer at least 2, so

$$|A| = \frac{\varphi(n)}{[G : A]} \leq \frac{\varphi(n)}{2}.$$

Now count Fermat witnesses. Every residue $a \in \{1, \ldots, n-1\}$ with $\gcd(a, n) > 1$ is a Fermat witness (a non-unit cannot satisfy $a^{n-1} \equiv 1$); there are $(n-1) - \varphi(n)$ of these, counting that $n - 1 \geq$ the units and the residue $n - 1$ appropriately, but we may argue more cleanly as follows. The residues in $\{1, \ldots, n-1\}$ that are *not* Fermat witnesses are exactly the units a with $a^{n-1} \equiv 1$, namely the elements of A (note $n - 1 \notin \{\text{non-witnesses}\}$ contributes nothing extra since we only bound from above). Hence the number of non-witnesses is $|A| \leq \varphi(n)/2 \leq (n-1)/2$, and the number of Fermat witnesses is at least

$$(n-1) - |A| \;\geq\; (n-1) - \frac{n-1}{2} = \frac{n-1}{2}.$$

Since $b \in G \setminus A$ is itself a witness not counted among the non-units, the inequality is strict: more than half of $\{1, \dots, n-1\}$ are Fermat witnesses. □

Corollary 5.3. *If n is composite and not a Carmichael number, then a single base chosen uniformly at random from $\{2, \dots, n-1\}$ is a Fermat witness with probability exceeding $1/2$; consequently k independent random bases fail to detect the compositeness of n with probability less than 2^{-k}.*

Corollary 5.3 is what would make the Fermat test a genuine probabilistic primality test—*if* the Carmichael hypothesis could be discharged.

5.3. **The Carmichael vulnerability.** For a Carmichael number the subgroup A of the preceding proof is the whole of G: every unit satisfies $a^{n-1} \equiv 1$, so $A = G$ is not proper and the index bound gives nothing. The only Fermat witnesses are the gcd witnesses, the non-units.

Proposition 5.4 (Density of witnesses for a Carmichael number)**.** *If n is a Carmichael number, then the Fermat witnesses in $\{1, \dots, n-1\}$ are exactly the residues a with $\gcd(a, n) > 1$, and their proportion is*

$$\frac{(n-1) - \varphi(n)}{n-1} = 1 - \frac{\varphi(n)}{n-1}.$$

For a Carmichael number $n = p_1 p_2 p_3$ with three large prime factors this proportion is close to $\frac{1}{p_1} + \frac{1}{p_2} + \frac{1}{p_3}$, which can be made arbitrarily small.

Proof. By Definition 2.7 every unit is a Fermat liar, so no unit is a witness; and every non-unit is a witness as noted above. Hence the witnesses are precisely the $(n-1) - \varphi(n)$ non-units, giving the stated proportion. For $n = p_1 p_2 p_3$ squarefree, $\varphi(n) = \prod_i (p_i - 1)$ and by inclusion–exclusion the density of non-units is $1 - \prod_i (1 - 1/p_i) = \sum_i 1/p_i - \sum_{i<j} 1/(p_i p_j) + \cdots$, whose leading term is $\sum_i 1/p_i$; when the p_i are large this is tiny. □

Example. For the Carmichael number $n = 13079177569$ one finds that a Fermat witness coincides with a gcd witness, and the proportion of witnesses is about 0.14% [7, 8]. A random search therefore requires on the order of hundreds of trials to stumble upon a witness; running the Fermat test even ten times will, with overwhelming probability, wrongly report the number as prime. This is the concrete danger: Corollary 5.3 silently fails, and no number of Fermat trials repairs it.

By Theorem 4.1 there are infinitely many Carmichael numbers, so this blind spot is not an artifact of small cases that one could tabulate and exclude; it persists at every scale, including the cryptographic scale of hundreds of digits. The Fermat test is therefore not a sound probabilistic primality test.

5.4. **The Miller–Rabin remedy.** The repair, due to Miller [16] and Rabin [17], is to test not merely $a^{n-1} \equiv 1$ but the finer structure of how the value 1 is reached along the chain of repeated squarings. The idea rests on a property that distinguishes primes from all composites uniformly: modulo a prime, the only square roots of 1 are ± 1, whereas a composite with at least two distinct prime factors always has nontrivial square roots of unity.

Lemma 5.5 (Roots of unity)**.** *If p is prime and $x^2 \equiv 1 \pmod{p}$, then $x \equiv \pm 1 \pmod{p}$. If n has at least two distinct odd prime factors, then $x^2 \equiv 1 \pmod{n}$ has at least four solutions, hence a solution $x \not\equiv \pm 1 \pmod{n}$.*

Proof. Modulo a prime p, $x^2 \equiv 1$ means $p \mid (x-1)(x+1)$, so $p \mid x-1$ or $p \mid x+1$, giving $x \equiv \pm 1$. If $n = n_1 n_2$ with $n_1, n_2 > 1$ coprime (possible when n has two distinct prime factors), the Chinese Remainder Theorem identifies solutions of $x^2 \equiv 1 \pmod{n}$ with pairs of solutions modulo n_1 and modulo n_2; taking $x \equiv 1 \pmod{n_1}$ and $x \equiv -1 \pmod{n_2}$ produces an x with $x \not\equiv \pm 1 \pmod{n}$, a nontrivial square root of unity. □

Definition 5.6 (Miller–Rabin test)**.** Let $n > 2$ be odd and write $n - 1 = 2^s d$ with d odd and $s \geq 1$. For a base a with $1 < a < n-1$, compute the sequence

$$a^d,\ a^{2d},\ a^{4d},\ \ldots,\ a^{2^{s-1}d} \pmod{n}.$$

The base a is a *strong liar* for n (fails to detect compositeness) if either $a^d \equiv 1 \pmod{n}$, or $a^{2^i d} \equiv -1 \pmod{n}$ for some $0 \leq i < s$. Otherwise a is a *root witness*, and n is composite. The test samples bases a and declares COMPOSITE on encountering any root witness.

The soundness of the test is the observation that if n is prime, then $a^{n-1} \equiv 1$ and the sequence of repeated square roots of 1 must, by the first part of Lemma 5.5, pass through -1 before reaching 1 (or start at 1); a prime therefore always presents as a strong liar, so the test never misclassifies a prime. The decisive advantage over the Fermat test is the following uniform bound, with no exceptional inputs.

Theorem 5.7 (Monier–Rabin bound [17])**.** *Let $n > 9$ be odd and composite. Then the number of strong liars a in $\{1, \ldots, n-1\}$ is at most $\varphi(n)/4$. Consequently at least $3/4$ of the bases are root witnesses, and k independent random bases fail to detect the compositeness of n with probability at most $(1/4)^k$.*

We indicate the mechanism rather than reproducing the full count. The strong liars are contained in a set defined by congruence conditions coming from the factorization $n = \prod p_i^{e_i}$; a careful group-theoretic enumeration, using that a nontrivial square root of unity exists whenever $\omega(n) \geq 2$ (Lemma 5.5), shows this set has size at most $\varphi(n)/4$ for every odd composite $n > 9$. The essential contrast with the Fermat test is that *this bound has no exceptions*: whereas the Fermat witness bound of Theorem 5.2 is voided precisely by Carmichael numbers, the Miller–Rabin bound holds for Carmichael numbers too. Indeed, a Carmichael number is squarefree with $\omega(n) \geq 3$, so Lemma 5.5 furnishes nontrivial square roots of unity in abundance; the very structure that makes a Carmichael number an absolute Fermat pseudoprime (many prime factors) is what guarantees many root witnesses. The Miller–Rabin test thus converts the Carmichael numbers from a fatal blind spot into ordinary composites.

Corollary 5.8. *For any odd composite $n > 9$, including every Carmichael number, running Miller–Rabin with k random bases has error probability at most $(1/4)^k$. Choosing, e.g., $k = 64$ makes the error below 2^{-128}, negligible for cryptographic purposes.*

6. Optimized Algorithmic Prime Generation

We now assemble the preceding theory into the practical task of generating a large prime and analyze rigorously the extent to which Larsen's progression dichotomy (Theorem 4.5) can accelerate it.

6.1. **The standard generator.** The textbook RSA key-generation routine [18] produces a prime of a prescribed bit-length b by rejection sampling.

```
Algorithm STANDARD-GENERATE(b, k):
  input:  bit-length b, number of Miller-Rabin rounds k
  output: a probable prime of b bits
  1  repeat
  2      n <- random odd integer with exactly b bits
  3      if n is divisible by any prime in a small trial-division list:
  4          continue                          // cheap pre-screen
  5      if MILLER-RABIN(n, k) reports COMPOSITE:
  6          continue
  7      return n                              // probable prime, error <= (1/4)^k
```

By the prime number theorem a random b-bit odd integer is prime with probability $\Theta(1/b)$, so the loop runs $O(b)$ times in expectation; the trial-division pre-screen (line 3) cheaply discards the many candidates with small factors, and Miller–Rabin (line 5) adjudicates the rest with the uniform guarantee of Corollary 5.8. Correctness does not depend on any distributional input beyond Theorem 5.7; in particular Carmichael numbers pose no threat here, because Miller–Rabin handles them like any other composite.

6.2. **The incompatibility certificate.** The temptation raised by Theorem 4.5 is to restrict candidates to a Carmichael-*incompatible* progression, so that no candidate can possibly be a Carmichael number. The certificate underlying this idea is the necessity half of Theorem 4.5, which we now prove; it is elementary and rests only on the structural results of Section 3.

Proposition 6.1 (Incompatible progressions are Carmichael-free [13])**.** *Let $r \pmod m$ be a Carmichael incompatible progression in the sense of Definition 4.4. Then no Carmichael number is congruent to $r \pmod m$.*

Proof. Suppose n is a Carmichael number with $n \equiv r \pmod m$; we show none of conditions (a),(b),(c) can hold, so the progression is in fact compatible. Set $g = \gcd(r, m) = \gcd(n, m)$, $\lambda(g)$ the exponent of $(\mathbb{Z}/g\mathbb{Z})^\times$, and $h = \gcd(\lambda(g), m)$.

For (a): by Proposition 3.3 and Proposition 3.1, n is odd and $\gcd(n, \varphi(n)) = 1$, and moreover, if primes $p, q \mid n$ satisfied $p \mid q - 1$ then $p \mid q - 1 \mid n - 1$ (Korselt) while $p \mid n$, contradicting $\gcd(n, n - 1) = 1$; hence $\gcd(n, 2\varphi(n)) = 1$. Since $g \mid n$ and $\varphi(g) \mid \varphi(n)$, it follows that $\gcd(g, 2\varphi(g)) \mid \gcd(n, 2\varphi(n)) = 1$, so (a) fails.

For (b): because $g \mid n$ and n is Carmichael, the exponent $\lambda(g)$ divides $\lambda(n)$, and by Korselt's criterion $\lambda(n) = \operatorname{lcm}_{p|n}(p - 1)$ divides $n - 1$. Hence $\lambda(g) \mid n - 1$. Also $h = \gcd(\lambda(g), m) \mid \lambda(g)$, so $h \mid n - 1$; and since $n \equiv r \pmod m$ with $h \mid m$, we have $n \equiv r \pmod h$, whence $h \mid r - 1$ iff $h \mid n - 1$, which holds. Thus (b) fails.

For (c): suppose $36 \mid m$ and $r \equiv 3 \pmod{12}$. Then $n \equiv 3 \pmod{12}$, so $n - 1 \equiv 2 \pmod{12}$ is divisible by neither 3 nor 4; by Korselt every prime $p \mid n$ has $(p - 1) \mid n - 1$, so no $p - 1$ is divisible by 3 or 4, i.e. every prime factor of n/g is $\equiv 11 \pmod{12}$. As $12 \mid m/g$ in this configuration, reducing gives $r/g \equiv n/g \pmod{12}$, and a product of primes each $\equiv 11 \pmod{12}$ is $\equiv 1$ or $11 \pmod{12}$; hence $r/g \equiv 1$ or 11, not 5 or 7, so (c) fails.

All three incompatibility conditions fail, contradicting the assumption that $r \pmod m$ is incompatible. Therefore no Carmichael number lies in an incompatible progression. □

6.3. **A structural obstruction, and the correct engineering conclusion.** Proposition 6.1 is a genuine, unconditional certificate: a candidate known to lie in an incompatible progression is provably not a Carmichael number. The question is whether this certificate can be harnessed to accelerate prime generation—and here rigor demands care, because the incompatible progressions are, with a single degenerate exception, exactly the progressions that contain no large primes either.

Proposition 6.2 (Obstruction)**.** *Let $r \pmod m$ be an arithmetic progression that contains at least one prime exceeding m. Then $r \pmod m$ is Carmichael compatible. Equivalently: no Carmichael-incompatible progression contains a prime larger than its modulus.*

Proof. Let $p > m$ be a prime with $p \equiv r \pmod m$. Then $\gcd(r, m) = \gcd(p, m)$, and since p is a prime exceeding m, no prime factor of m equals p, so $\gcd(p, m) = 1$; thus $g = \gcd(r, m) = 1$. We check that none of the incompatibility conditions of Definition 4.4 can hold. With $g = 1$ we have $\varphi(g) = \varphi(1) = 1$, so $\gcd(g, 2\varphi(g)) = \gcd(1, 2) = 1$, and (a) fails. The exponent $\lambda(1) = 1$, so $h = \gcd(1, m) = 1$, and $1 \mid r - 1$ always, so (b) fails. Finally (c) requires $r \equiv 3 \pmod{12}$; but $\gcd(r, m) = 1$ together with $36 \mid m$ forces $\gcd(r, 12) = 1$, so $r \equiv 1, 5, 7$, or $11 \pmod{12}$, never 3, and (c) fails. Hence the progression is compatible. □

The consequence is a precise no-go statement. To generate large primes one must draw from a progression $r \pmod m$ with $\gcd(r, m) = 1$—otherwise every candidate is divisible by the common factor $g > 1$ and is composite (indeed no element exceeding m can be prime). But by Proposition 6.2 every such prime-admitting progression is Carmichael compatible, so the incompatibility certificate of Proposition 6.1 is never available for it. Restricting the random number generator to Carmichael-incompatible progressions does indeed guarantee a 0% chance of drawing a Carmichael number—vacuously, because it simultaneously guarantees a 0% chance of drawing a large prime. The two guarantees are the same guarantee.

This is not a defect of Larsen's theorem but a clarification of its cryptographic meaning. Theorem 4.5 draws the exact boundary of Carmichael contamination, and Proposition 6.2 shows that boundary lies entirely inside the prime-free region. The correct engineering reading is therefore a *negative* one with a constructive corollary:

Corollary 6.3 (Why a Carmichael-robust test is indispensable)**.** *Every arithmetic progression from which large primes can be sampled contains, by Theorem 4.5, infinitely many Carmichael numbers. No choice of sampling modulus can exclude them. Hence a prime generator cannot outsource Carmichael-avoidance to its candidate distribution and must rely on a primality test whose error bound is uniform over all composites—precisely the property that distinguishes Miller–Rabin (Theorem 5.7) from the Fermat test.*

We record the honest residual value of the arithmetic viewpoint. The structural facts $\gcd(n, 2\varphi(n)) = 1$ and squarefreeness that every Carmichael number satisfies (Section 3) do support cheap pre-screens—trial division by small primes already removes every candidate divisible by a small prime, which is where the densest strong-liar composites concentrate—and these pre-screens, not any progression restriction, are what reduce the expected number of Miller–Rabin rounds. The following generator makes the certificate explicit as an assertion, emphasizing that for a coprime sampling class it always returns "compatible," i.e. provides no filtering, so its role is documentation of correctness rather than acceleration.

```
Algorithm LARSEN-AWARE-GENERATE(b, k, m, r):
  input:  bit-length b, MR rounds k, modulus m, residue r with gcd(r,m)=1
```

```
output: a probable prime of b bits, n = r (mod m)
precondition asserted:
    // By Prop 6.3, gcd(r,m)=1  =>  r (mod m) is Carmichael COMPATIBLE.
    // The class therefore CANNOT be used to exclude Carmichael numbers;
    // Carmichael-robustness is delegated entirely to MILLER-RABIN.
1  repeat
2      n <- random b-bit integer with n = r (mod m)      // wheel / stride
3      if n divisible by any prime in small trial-division list:
4          continue
5      if MILLER-RABIN(n, k) reports COMPOSITE:
6          continue
7      return n
```

6.4. **Complexity of the filtering step.** Whichever generator is used, the dominant cost per candidate is the modular exponentiation inside Miller–Rabin. We bound it.

Proposition 6.4 (Cost of a Miller–Rabin round)**.** *Let n have $\ell = \Theta(\log n)$ bits. A single Miller–Rabin round—computing $a^d \bmod n$ and up to $s-1$ subsequent squarings, where $n-1 = 2^s d$—uses $O(\log n)$ modular multiplications, hence $O(\log^3 n)$ bit operations with schoolbook arithmetic. Testing a candidate with k rounds costs $O(k \log^3 n)$ bit operations.*

Proof. The exponent $d \le n-1$ has $O(\log n)$ bits, so binary (square-and-multiply) exponentiation computes $a^d \bmod n$ using at most $2\log_2 d = O(\log n)$ modular multiplications: one squaring per bit and one conditional multiply per one-bit. The subsequent chain $a^{2d}, a^{4d}, \ldots, a^{2^{s-1}d}$ adds $s-1 < \log_2 n = O(\log n)$ further squarings. Thus a round performs $O(\log n)$ modular multiplications. Each multiplication of two residues modulo n—numbers of $O(\log n)$ bits—costs $O(\log^2 n)$ bit operations by the schoolbook algorithm (including the reduction modulo n, a division of a 2ℓ-bit number by an ℓ-bit number, itself $O(\ell^2) = O(\log^2 n)$). Multiplying the count of multiplications by the cost of each gives $O(\log n) \cdot O(\log^2 n) = O(\log^3 n)$ bit operations per round, and $O(k \log^3 n)$ for k rounds. □

Remark 6.5. With fast integer multiplication the per-multiplication cost drops to $\tilde{O}(\log n)$, improving the round cost to $\tilde{O}(\log^2 n)$; the schoolbook bound $O(k \log^3 n)$ is the standard conservative figure and shows the filtering is comfortably polynomial. For context, deterministic polynomial-time primality proving is possible in principle by the AKS algorithm [2], but its exponents make it uncompetitive with Miller–Rabin for key generation, where a probabilistic guarantee of error $\le (1/4)^k$ is entirely satisfactory.

6.5. **Engineering trade-offs.** We close the section by weighing the two levers a designer can pull: the number k of Miller–Rabin rounds and the choice of candidate distribution.

The round count k trades running time against soundness linearly in time and geometrically in error: by Proposition 6.4 each round costs $O(\log^3 n)$, while Corollary 5.8 shows k rounds bound the error by $(1/4)^k$. Because this bound is a worst-case bound over adversarially chosen composites, it is pessimistic for the *random* candidates that arise in key generation: sharper average-case analyses show that for a random b-bit candidate the error after even a few rounds is far below $(1/4)^k$, so implementations safely use small k (a few dozen rounds suffice for cryptographic margins). Crucially, none of this reduction in k relies on excluding Carmichael numbers by arithmetic means—by Corollary 6.3 that is

impossible—but rather on the uniform Miller–Rabin bound together with the abundance of primes.

The candidate distribution offers the second lever, and here the analysis of this section sets a firm boundary. Restricting candidates to a fixed residue $r \pmod m$ with $\gcd(r, m) = 1$ (a "wheel") is standard and beneficial: it skips candidates with small factors and slightly raises the prime density among those tested, reducing the expected number of loop iterations. It costs a small amount of key-space entropy, since candidates are confined to one of the $\varphi(m)$ reduced classes rather than all odd residues; concretely the entropy loss is $\log_2\big((m/2)/\varphi(m)\big)$ bits per candidate, negligible for the small moduli used in practice. What such a restriction *cannot* do, by Propositions 6.1 and 6.2, is exclude Carmichael numbers: any wheel fine enough to admit primes admits Carmichael numbers as well. A designer contemplating a Larsen-informed restriction therefore trades entropy for prime-density and small-factor avoidance—worthwhile in themselves—but obtains *no* reduction in the Carmichael risk, which remains the responsibility of the Miller–Rabin test. The clean lesson is that distribution shaping and Carmichael-robustness are orthogonal concerns: the former tunes efficiency, the latter is delegated wholesale to a test with a uniform error bound.

7. Discussion

The arc of this paper runs from a single algebraic identity to a modern distribution theorem and back to an engineering guarantee, and the pivot at every stage is the group $(\mathbb{Z}/n\mathbb{Z})^\times$. Fermat's little theorem and Euler's theorem are statements about the order of this group; the Fermat test reads compositeness off a failure of the group identity; the witness bound of Theorem 5.2 is a coset count; and the Carmichael numbers are exactly the composites for which the relevant subgroup $A = \{a : a^{n-1} \equiv 1\}$ fills the whole group. Korselt's criterion translates this group-theoretic degeneracy into the arithmetic conditions of squarefreeness and $(p-1) \mid (n-1)$, and those conditions in turn drive both the constructions of Carmichael numbers (Chernick, Alford–Granville–Pomerance) and the classification of the progressions that contain them (Larsen).

Larsen's two theorems occupy complementary positions. The Bertrand-type result (Theorem 4.3) is a statement about *vertical* distribution—how Carmichael numbers are spaced along the number line—and its novelty is methodological: it is the first transfer of the bounded-gaps sieve technology of Zhang and Maynard into the study of Carmichael numbers, converting a pigeonhole argument that loses size control into a sieve argument that retains it. The arithmetic-progressions result (Theorem 4.5) is a statement about *horizontal* distribution—which residue classes are populated—and its novelty is the sharp compatibility dichotomy, together with the resolution of the Alford–Granville–Pomerance question on $\liminf \varphi(n)/n$. Together they bring the distributional theory of Carmichael numbers to a level of resolution approaching that of the primes themselves.

The cryptographic analysis of Section 6 is, we believe, the most instructive part of the story precisely because the naive hope fails. It is tempting to read Theorem 4.5 as licensing a prime generator that sidesteps Carmichael numbers by construction. Propositions 6.1 and 6.2 show why this cannot work: the incompatible progressions and the prime-free progressions coincide except for degenerate cases, so the certificate that would exclude Carmichael numbers also excludes primes. Far from diminishing the theorem's relevance, this pins down its role: it certifies that Carmichael contamination is unavoidable in any usable candidate stream, which is exactly the fact that makes a uniform-error test like Miller–Rabin not a convenience but

a necessity. The pure-mathematical distribution law and the engineering requirement meet at this point.

Several questions remain open or lie beyond our scope. The optimal exponent in the lower bound for $C(x)$ is unknown; the passage from $2/7$ to $1/3$ [10] and the constant $1/168$ in Larsen's progression theorem [13] are both understood to be improvable. Whether Carmichael numbers with a bounded number of prime factors are infinite remains conjectural, tied to prime-tuple conjectures [12]. And the average-case error analysis of Miller–Rabin for structured (non-uniform) candidate distributions, of direct relevance to standardized key-generation procedures, continues to receive attention.

8. Conclusion

We have given a complete and self-contained development of the theory of Carmichael numbers, from the algebraic foundations to the frontier of their distribution theory and its cryptographic consequences. Starting from full proofs of Fermat's little theorem and Euler's theorem, we formalized the Fermat primality test and its taxonomy of witnesses and liars, and we established the equivalence of the coprime-base and universal-congruence definitions of an absolute Fermat pseudoprime. We proved the rigid structural constraints on Carmichael numbers—oddness, at least three distinct prime factors, squarefreeness, and coprimality with the totient—and gave a two-directional proof of Korselt's criterion, the finite test that replaces the infinite congruence condition.

On the distributional side we stated the Alford–Granville–Pomerance infinitude theorem with its bound $C(x) > x^{2/7}$, and we described Daniel Larsen's two decisive advances: a Bertrand-type theorem placing dense clusters of Carmichael numbers in short intervals through Maynard's multidimensional sieve and Fourier analysis, and a complete classification of the arithmetic progressions containing infinitely many Carmichael numbers, resolving an open question on $\liminf \varphi(n)/n$. On the computational side we proved the group-theoretic Fermat witness bound and its failure for Carmichael numbers, and we showed how the Miller–Rabin test restores a uniform $(1/4)^k$ error bound with no exceptional inputs, precisely by exploiting the multi-factor structure that makes Carmichael numbers absolute Fermat pseudoprimes in the first place.

Finally, we analyzed rigorously the interaction between Larsen's progression dichotomy and prime generation, and found a clean obstruction: every progression from which large primes can be drawn is Carmichael compatible, so no candidate distribution can exclude Carmichael numbers without excluding primes. The mathematical significance of the subject is thus twofold. Internally, the Carmichael numbers are the exact composites on which the most natural primality identity degenerates, and their distribution has now been mapped with tools imported from the analytic theory of primes. Externally, they stand as the reason that a primality test used for cryptographic key generation must guarantee its error bound uniformly over all composites—a requirement met by Miller–Rabin and quantified by the theorems developed here.

References

[1] W. R. Alford, A. Granville, and C. Pomerance, *There are infinitely many Carmichael numbers*, Annals of Mathematics **140** (1994), 703–722.

[2] M. Agrawal, N. Kayal, and N. Saxena, *PRIMES is in P*, Annals of Mathematics **160** (2004), no. 2, 781–793.

[3] R. D. Carmichael, *Note on a new number theory function*, Bulletin of the American Mathematical Society **16** (1910), 232–238.
[4] R. D. Carmichael, *On composite numbers P which satisfy the Fermat congruence $a^{P-1} \equiv 1 \pmod{P}$*, American Mathematical Monthly **19** (1912), 22–27.
[5] J. Chernick, *On Fermat's simple theorem*, Bulletin of the American Mathematical Society **45** (1939), 269–274.
[6] K. Conrad, *Carmichael numbers and Korselt's criterion*, expository notes, University of Connecticut.
[7] K. Conrad, *Fermat's test*, expository notes, University of Connecticut.
[8] K. Conrad, *Fermat's little theorem*, expository notes, University of Connecticut.
[9] K. Conrad, *The Solovay–Strassen test*, expository notes, University of Connecticut.
[10] G. Harman, *Watt's mean value theorem and Carmichael numbers*, International Journal of Number Theory **4** (2008), no. 2, 241–248.
[11] A. R. Korselt, *Problème chinois*, L'Intermédiaire des Mathématiciens **6** (1899), 142–143.
[12] D. Larsen, *Bertrand's postulate for Carmichael numbers*, International Mathematics Research Notices (2023); arXiv:2111.06963.
[13] D. Larsen, *Carmichael numbers in all possible arithmetic progressions*, preprint, 2025; arXiv:2504.09056.
[14] J. Maynard, *Small gaps between primes*, Annals of Mathematics **181** (2015), 383–413.
[15] J. Maynard, *Dense clusters of primes in subsets*, Compositio Mathematica **152** (2016), 1517–1554.
[16] G. L. Miller, *Riemann's hypothesis and tests for primality*, Journal of Computer and System Sciences **13** (1976), 300–317.
[17] M. O. Rabin, *Probabilistic algorithm for testing primality*, Journal of Number Theory **12** (1980), 128–138.
[18] R. L. Rivest, A. Shamir, and L. Adleman, *A method for obtaining digital signatures and public-key cryptosystems*, Communications of the ACM **21** (1978), 120–126.
[19] V. Šimerka, *Zbytky z arithmetické posloupnosti* (On the remainders of an arithmetic progression), Časopis pro pěstování matematiky a fysiky **14** (1885), 221–225.
[20] R. Solovay and V. Strassen, *A fast Monte-Carlo test for primality*, SIAM Journal on Computing **6** (1977), 84–85.
[21] Y. Zhang, *Bounded gaps between primes*, Annals of Mathematics **179** (2014), 1121–1174.