

THE BOURGEOIS TRIANGLE

A FIRST-PRINCIPLES LINEAR MEDIUM FOR PRIME WHEELS, RAMANUJAN LANES, AND THE VON MANGOLDT DOORWAY

BLAIZE ROUYEA AND COREY BOURGEOIS

ABSTRACT. The Bourgeois triangle is a Pascal-type array with a unit left border, a signed right border, and ordinary addition in the interior. It started as scratch-work: put the first prime, the unit, and alternating prime direction into one picture and see what holds it together. This paper rebuilds that object from the ground up and keeps a hard line between what is proved and what is not.

Four mechanisms are exact. The row polynomials obey a one-line recursion and a Green-cone formula. The alternating row value $P_n(-1)$ is a perfect selector: it sees only the newest right-edge impulse. Under the von Mangoldt doorway $T(n, n) = \psi(n)$, the near-edge difference is exactly $\Lambda(n)$, so the Euler identity $-\zeta'/\zeta = \sum \Lambda(n)n^{-s}$ enters the medium with no approximation. And the survivor border is a Ramanujan–Fourier wheel signal: on $[p_k, p_{k+1}^2)$ it is exactly the Q_k -wheel, with lane amplitudes equal to Ramanujan sums of relative strength $1/(p-1)$. The foundational lane $p=2$ is the medium’s spectral zero at $\omega = \pi$.

Then we draw the boundary. A structural acceleration operator $\mathcal{D} = (I - S)^2(I + S)$ collapses the Li–Laguerre prime trace to a standard Laguerre sampling sequence, and boundedness of that filtered trace would exclude interior poles under the map $z = 1 - 1/s$. We prove every reduction leading there. We state the remaining boundedness assertion as a conjecture, not a theorem. That is the honest endpoint: a rigorous account of what the triangle is, what it proves, and the exact place where its link to zeta-zero confinement becomes a sharp open estimate.

CONTENTS

1. The genesis and the first-principles problem	2
2. Definition of the triangle	3
3. The affine clock and the seed	4
4. Green cones and the exact selector	4
5. The von Mangoldt doorway	5
6. The wheel-lane bridge	5
6.1. One eliminator at a time	5
6.2. The wheel-onset law	6
6.3. Lane amplitudes are Ramanujan sums	6
6.4. The same data, two axes	7
6.5. Controls and the honesty line	7
7. The residual-bilinear identity	8
8. The completed Li–Laguerre trace	8
9. Residual form of the prime trace	9
10. Structural acceleration and the filtered trace	10
11. The zero-mode detector	10
12. Why the obvious bound fails	11
13. The prime-Laguerre frame conjecture	11
14. Numerical evidence	12
15. Conclusion	12

Date: Draft: May 29, 2026.

1. THE GENESIS AND THE FIRST-PRINCIPLES PROBLEM

The scratch-work was simple on purpose. Put a 1 down the left, a 2 at the apex, alternate the sign on the right border, and ask one thing: can the primes be made to show up inside an ordinary Pascal interior? Written out, it looked like this:

$$2, 1, 1, 1, \dots, \quad -3, 5, -7, 11, \dots$$

where the signs only mark direction. Drawing it was never the hard part. The hard part was the seed. Could we explain those first numbers without just declaring them to be primes?




FIGURE 1. The triangle's first rows under the prime border $B_n = p_{n+1}$. The right border (dark) carries the signed survivor impulses, the left border (soft) is the unit affine clock, and the interior is plain Pascal addition.

This paper answers that question as far as the elementary structure allows, and it answers it by separating three things:

- (i) the left border is 1, because that is the only constant that makes the affine clock tick at unit speed;
- (ii) the apex s_0 is gauge — it touches no forward row;
- (iii) the first signed right value is -3 once the first eliminator is taken to be the least non-unit integer and the border is read at the alternating frequency $\omega = \pi$.

So the seed is not magic. The value 2 is the least non-unit integer after the unit normalization. Once that first eliminator is chosen, its first self-hit is $2^2 = 4$, so 3 is forced to survive before any second eliminator can act. And the minus sign is $(-1)^1$, because the medium's one distinguished frequency is the alternating one. That is the whole story of the seed: not an independent axiom, but a normalization and its first consequence.

We keep the language plain the rest of the way. The triangle is a linear medium. The right border feeds it impulses. The Pascal interior carries those impulses outward. The row-Fourier axis reads them as lanes, and the column axis stores prime-power memory. Everything below is one of those four things looked at from a different side.

Remark 1.1 (What is derived, and what is not). Be exact about the claim. The *seed* $(2, 1, -3)$ is genuinely derived: unit clocking forces the left border (Corollary 3.2), the apex is gauge (Theorem 3.3), and “least non-unit eliminator, first self-hit at 4, alternating sample” fixes the first right value (Remark 3.4). The wheel-onset law (Theorem 6.1) is different. It takes the primes $p_1 < p_2 < \dots$ as *given* and characterizes the survivor stream on each window $[p_k, p_{k+1}^2)$ by coprimality. It does not grow that stream out of the operator alone. Deriving the whole survivor sequence from pure medium dynamics, with no sieve in hand, stays open. We do not claim it. What the triangle does give is the fusion (Section 6.5): once the border is handed to us, its gaps, its onset times p^2 , and its lane amplitudes are all intrinsic to this one medium.

2. DEFINITION OF THE TRIANGLE

Definition 2.1 (Bourgeois triangle). Let $(B_n)_{n \geq 1}$ be a sequence in an abelian group and let $s_0 \in \mathbb{C}$. Set $B_0 = 0$. The Bourgeois triangle associated to (B_n, s_0) is the array $T(n, k)$, $0 \leq k \leq n$, fixed by

- (1) $T(0, 0) = s_0$,
- (2) $T(n, 0) = 1$ $(n \geq 1)$,
- (3) $T(n, n) = (-1)^n B_n$ $(n \geq 1)$,
- (4) $T(n, k) = T(n-1, k-1) + T(n-1, k)$ $(1 \leq k \leq n-1)$.

Outside $0 \leq k \leq n$, put $T(n, k) = 0$. We call B_n the *unsigned boundary data* and $T(n, n) = (-1)^n B_n$ the *visible right edge*. If s_0 and all B_n are integers, then every entry $T(n, k)$ is an integer, since the construction uses only the integer border values and integer addition.

Proposition 2.2 (Well-posedness). *For every right border (B_n) and apex s_0 , exactly one array satisfies the defining relations.*

Proof. Induct on rows. Row 0 is the apex. Given row $n-1$, the borders fix the two edge entries of row n and the Pascal rule fixes every interior entry. Existence and uniqueness together. \square

Definition 2.3 (Row polynomial). For $n \geq 0$, set

$$P_n(z) = \sum_{k=0}^n T(n, k) z^k.$$

Lemma 2.4 (Row recursion). *For $n \geq 2$,*

$$P_n(z) = (1+z)P_{n-1}(z) + u_n z^n, \quad u_n = (-1)^n (B_n + B_{n-1}),$$

and $P_1(z) = 1 - B_1 z$.

Proof. The interior rule contributes $(1+z)P_{n-1}(z)$. The coefficient of z^n there is $(-1)^{n-1}B_{n-1}$, so the correction needed to land the top coefficient on $(-1)^n B_n$ is

$$(-1)^n B_n - (-1)^{n-1} B_{n-1} = (-1)^n (B_n + B_{n-1}) = u_n. \quad \square$$

3. THE AFFINE CLOCK AND THE SEED

Theorem 3.1 (Affine clock). *For every $n \geq 1$,*

$$T(n, 1) = (n - 1) - B_1.$$

If the left border were a constant c instead of 1, then

$$T(n, 1) = c(n - 1) - B_1.$$

Proof. $T(1, 1) = -B_1$, and for $n \geq 2$,

$$T(n, 1) = T(n - 1, 0) + T(n - 1, 1) = 1 + T(n - 1, 1).$$

Iterating gives $T(n, 1) = (n - 1) - B_1$. Replacing 1 by c gives the general form. \square

Corollary 3.2 (Unit inflow). *The value 1 is the unique constant left border for which the first interior column is a unit-speed affine clock.*

Theorem 3.3 (Seed irrelevance). *Fix (B_n) . Two triangles with the same borders but different apex values agree on every row $n \geq 1$.*

Proof. Row 1 is fixed by the borders alone: $T(1, 0) = 1$, $T(1, 1) = -B_1$. If row $n - 1$ is apex-independent, so is row n , since borders and the Pascal rule build it. Induct. \square

Remark 3.4 (The seed, plainly). The apex $s_0 = 2$ is a convenient historical gauge, nothing more — it changes no forward row. The first non-unit eliminator is 2 because 2 is the least integer above the unit. Its first self-hit is 4, so 3 survives before a second eliminator exists. Reading the border at the alternating frequency makes that first value -3 . So 2, 1, and -3 are not arbitrary marks. They are the least non-unit, the unit inflow, and the alternating sample of the next survivor.

4. GREEN CONES AND THE EXACT SELECTOR

The triangle is linear, so each right-edge impulse sends out its own Pascal cone.

Definition 4.1 (Boundary impulses). For $n \geq 1$, set $u_n = (-1)^n (B_n + B_{n-1})$, with $B_0 = 0$.

Theorem 4.2 (Green representation). *For $n \geq 1$ and $0 \leq k \leq n$,*

$$T(n, k) = \binom{n-1}{k} + \sum_{m=1}^n u_m \binom{n-m}{k-m},$$

with binomials outside their range read as 0.

Proof. The right side obeys the Pascal rule because binomials do. At $k = 0$ only the first term survives, giving 1. At $k = n$ the first term vanishes and the sum telescopes:

$$\sum_{m=1}^n u_m = \sum_{m=1}^n (-1)^m (B_m + B_{m-1}) = (-1)^n B_n.$$

By well-posedness this is T . \square

Theorem 4.3 (Alternating selector). *For $n \geq 2$,*

$$P_n(-1) = \sum_{k=0}^n (-1)^k T(n, k) = B_n + B_{n-1}.$$

Proof. Evaluate Lemma 2.4 at $z = -1$. The factor $1 + z$ vanishes, leaving

$$P_n(-1) = u_n(-1)^n = B_n + B_{n-1}. \quad \square$$

Remark 4.4 (The spectral zero). The column step has symbol

$$\lambda(\omega) = 1 + e^{i\omega},$$

whose only zero on the unit circle is $\omega = \pi$. The selector is not a trick, then. It is the row-Fourier sample taken at the one frequency the medium annihilates.

5. THE VON MANGOLDT DOORWAY

Let $\Lambda(n)$ be the von Mangoldt function and $\psi(x) = \sum_{m \leq x} \Lambda(m)$ the Chebyshev function.

Definition 5.1 (Doorway). Put ψ on the diagonal:

$$T(n, n) = \psi(n), \quad \text{equivalently} \quad B_n = (-1)^n \psi(n).$$

Proposition 5.2 (Near-edge selector). *Under the doorway,*

$$T(n, n) - T(n-1, n-1) = \psi(n) - \psi(n-1) = \Lambda(n),$$

and therefore, for $\text{Re } s > 1$,

$$\sum_{n \geq 1} \frac{T(n, n) - T(n-1, n-1)}{n^s} = \sum_{n \geq 1} \frac{\Lambda(n)}{n^s} = -\frac{\zeta'(s)}{\zeta(s)}.$$

Proof. The first line is the definition of ψ . The Dirichlet identity is the Euler product for ζ on $\text{Re } s > 1$ (Apostol, *Intro. to Analytic Number Theory*). \square

Theorem 5.3 (Column-sieve law). *Under the doorway, for all $n \geq k \geq 0$,*

$$T(n, k) = \binom{n-1}{k} + \sum_{m=2}^k \Lambda(m) \binom{n-m}{k-m}.$$

Column k remembers exactly the prime powers $m \leq k$.

Proof. Write $D_m = \psi(m) - \psi(m-1) = \Lambda(m)$. By Theorem 4.2 each increment D_m seeds a cone with column- k trace $\binom{n-m}{k-m}$; the left border gives the baseline $\binom{n-1}{k}$. Since $\Lambda(1) = 0$, the sum starts at $m = 2$. \square

6. THE WHEEL-LANE BRIDGE

This is the part that explains why the border breaks where it breaks, and why its spectrum carries exact arithmetic amplitudes.

6.1. One eliminator at a time. The first eliminator is 2, and its first deletion is 4. So every integer below 4, past the unit, survives it — which forces 3 before any second eliminator exists. That is the first signed value -3 , read at the alternating frequency. The early gaps 1, 2, 2 all live in this one-wheel world. The first real break comes exactly when the next eliminator reaches its square.

6.2. The wheel-onset law.

Theorem 6.1 (Wheel-onset law). *Let $p_1 < p_2 < \dots$ be the primes and $Q_k = \prod_{i=1}^k p_i$. On the interval*

$$[p_k, p_{k+1}^2),$$

an integer n survives the first k eliminators p_1, \dots, p_k if and only if

$$\gcd(n, Q_k) = 1.$$

So on this window the survivor stream coincides with the residue wheel modulo Q_k , matching the ordinary prime stream up to p_{k+1}^2 , where it first fails.

Proof. If $n < p_{k+1}^2$ is composite, it has a factor at most $\sqrt{n} < p_{k+1}$, hence one of p_1, \dots, p_k . Conversely any n divisible by one of those is not a survivor. So survival on this window is coprimality to Q_k , and the first composite missed by the first k primes is p_{k+1}^2 . \square




FIGURE 2. The wheel-onset law on $[2, 30]$. Dark discs survive, light discs are eliminated. Each eliminator p first bites at p^2 . The first four-gap $7 \rightarrow 11$ is created when $9 = 3^2$ is deleted, fusing the run $7, 9, 11$.

Corollary 6.2 (The first four-gap). *On $[3, 25)$ the survivors are the residues 1 and 5 modulo 6. The jump $7 \rightarrow 11$ is the in-wheel spacing*

$$5 - 1 = 4,$$

created the moment $9 = 3^2$ is deleted.

6.3. Lane amplitudes are Ramanujan sums. Let $w_Q(a) = \mathbf{1}_{\gcd(a, Q)=1}$ be the coprime indicator on $\mathbb{Z}/Q\mathbb{Z}$. With the unnormalized discrete Fourier transform

$$\widehat{w}_Q(j) = \sum_{a=1}^Q w_Q(a) e^{2\pi i a j / Q} = \sum_{\substack{a=1 \\ (a, Q)=1}}^Q e^{2\pi i a j / Q},$$

this is exactly the Ramanujan sum $c_Q(j)$ (Ramanujan, 1918); the closed forms below are stated for this convention, and the relative amplitudes are independent of any overall normalization.

Theorem 6.3 (Ramanujan lane law). *For each frequency $\omega_j = 2\pi j / Q$,*

$$|\widehat{w}_Q(j)| = |c_Q(j)|.$$

With $d = \gcd(Q, j)$, Hölder's identity gives

$$c_Q(j) = \mu\left(\frac{Q}{d}\right) \frac{\varphi(Q)}{\varphi(Q/d)}.$$

If the lane's exact denominator is a prime p , meaning $Q/d = p$, then

$$|c_Q(j)| = \frac{\varphi(Q)}{p-1}.$$

Proof. The first line is the definition of the Ramanujan sum; the closed form is Hölder's identity. If $Q/d = p$ then $\mu(p) = -1$ and $\varphi(p) = p - 1$, giving $\varphi(Q)/(p - 1)$. \square

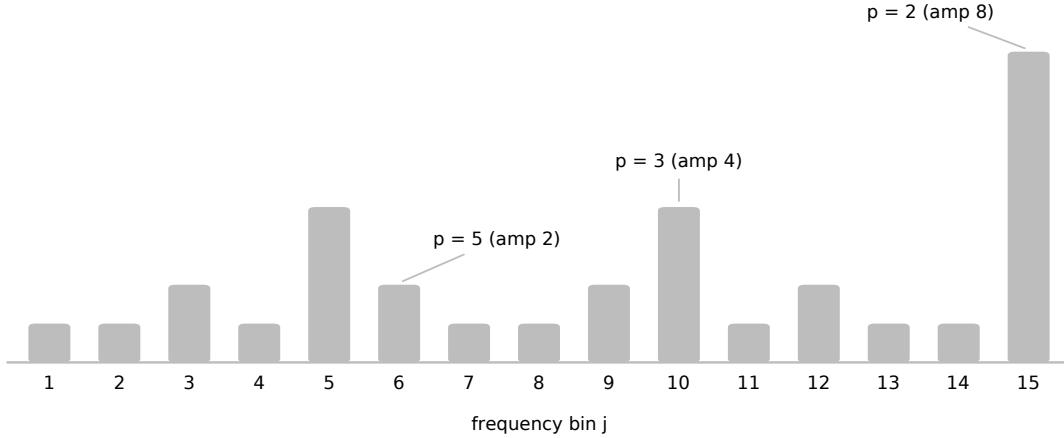


FIGURE 3. Lane spectrum of the mod-30 wheel, bins $j = 1, \dots, 15$, heights $|c_{30}(j)|$. The single-prime lanes follow the law: $p = 2$ at $j = 15$ ($\omega = \pi$) is tallest at 8, $p = 3$ at $j = 10$ rings at 4, and $p = 5$ at $j = 6$ at 2. Bins of composite exact denominator (such as $j = 5$, denominator 6) carry the combined-wheel amplitude and are not single-prime lanes. The single-prime heights are the $\varphi(Q)/(p - 1)$ hierarchy — relative strength $1/(p - 1)$.

Corollary 6.4 (The volume law). *Each prime denominator $p \mid Q$ marks a lane class — the bins j with $Q/\gcd(Q, j) = p$ — of relative amplitude*

$$\frac{1}{p - 1}.$$

The foundational denominator $p = 2$ has relative amplitude 1 and sits at $\omega = \pi$, the spectral zero of the Pascal step.

6.4. **The same data, two axes.** The lanes are not bolted on from outside; the triangle already carries them on both axes. On the row axis,

$$P_n(-1) = B_n + B_{n-1} \quad (\text{the alternating selector}),$$

and on the column axis, Theorem 5.3 stores $\Lambda(m)$ for $m \leq k$ in column k . The border is a Ramanujan–Fourier wheel signal; the interior is its Pascal–Green propagation. Same prime data, read two ways.

6.5. **Controls and the honesty line.** None of the Ramanujan machinery is ours, and we should say so plainly. Ramanujan sums, Hölder’s identity, the primorial wheel, and Ramanujan–Fourier expansions are classical; the Ramanujan–Fourier expansion of modified von Mangoldt data is a developed subject, with the sums going back to Ramanujan (1918) and a worked prime-pair application in Gadiyar–Padma (1999). We claim none of it. The contribution here is the fusion: this one linear medium binds that wheel spectrum to a distinguished spectral zero, a row selector, and a columnwise von Mangoldt memory. The standard literature works with $\Lambda(n)$ directly — no triangle, no annihilating frequency.

The lane law on its own says nothing about zeta zeros. Getting from amplitudes to zero location is the filtered boundedness problem of Section 13, and we keep that wall where it belongs.

As a check, the Ramanujan signature was correlated against matched null models at $Q = 30$. The match statistic is the Pearson correlation between the absolute amplitude vectors $(|\hat{w}(j)|)_{1 \leq j < Q}$ and $(|c_Q(j)|)_{1 \leq j < Q}$, computed over 200 trials for the randomized models. A density-matched random stream and the Cramér model both scored about zero, a merely periodic but non-coprime pattern

about 0.08, and the true coprime wheel scored 1. The $1/(p-1)$ law detects coprimality, not density and not mere periodicity.

7. THE RESIDUAL-BILINEAR IDENTITY

The doorway hands us a second exact object: an identity for the Chebyshev error. We use the right-continuous convention for ψ , so the jump at an integer n is included in $\psi(n)$. Let $R(x) = \psi(x) - x$, and for $X \geq 1$ set

$$S(X) = \sum_{n \leq X} \frac{\Lambda(n)R(n)}{n}, \quad C(X) = \sum_{n \leq X} \frac{\Lambda(n)^2}{n}.$$

Theorem 7.1 (Residual-bilinear identity). *For every $X \geq 1$,*

$$S(X) = \int_1^X \frac{R(t)}{t} dt + \frac{R(X)^2}{2X} - \frac{1}{2} + \frac{C(X)}{2} + \frac{1}{2} \int_1^X \frac{R(t)^2}{t^2} dt.$$

Proof. Start from the Stieltjes form

$$S(X) = \int_1^X \frac{R(t)}{t} d\psi(t).$$

Since $d\psi = dt + dR$,

$$S(X) = \int_1^X \frac{R(t)}{t} dt + \int_1^X \frac{R(t)}{t} dR(t).$$

Integrate the second piece by parts for right-continuous functions of bounded variation. The shared jumps of R/t and R at prime powers contribute $\sum_{n \leq X} \Lambda(n)^2/n = C(X)$, and solving gives

$$\int_1^X \frac{R(t)}{t} dR(t) = \frac{R(X)^2}{2X} - \frac{1}{2} + \frac{C(X)}{2} + \frac{1}{2} \int_1^X \frac{R(t)^2}{t^2} dt. \quad \square$$

Remark 7.2. The term $C(X)/2$ is the jump bracket, and it is not optional. Treat R as smooth and you lose it.

8. THE COMPLETED LI-LAGUERRE TRACE

Let ζ be the completed zeta function,

$$\zeta(s) = \frac{1}{2} s(s-1) \pi^{-s/2} \Gamma(s/2) \zeta(s),$$

and let λ_n be the Li coefficients (Li, 1997),

$$\lambda_n = \frac{1}{(n-1)!} \frac{d^n}{ds^n} \left[s^{n-1} \log \zeta(s) \right]_{s=1}, \quad \sum_{n \geq 1} \lambda_n z^n = \frac{z}{(1-z)^2} \frac{\zeta'(s)}{\zeta(s)}, \quad s = \frac{1}{1-z}.$$

Splitting the logarithmic derivative (the generating form above follows by substituting $s = 1/(1-z)$ into Li's derivative definition and summing the resulting Taylor series; see Li (1997) and Lagarias (2007) for this standard form),

$$\frac{\zeta'(s)}{\zeta(s)} = \frac{1}{s} + \frac{1}{s-1} - \frac{1}{2} \log \pi + \frac{1}{2} \psi_0\left(\frac{s}{2}\right) + \frac{\zeta'(s)}{\zeta(s)},$$

with $\psi_0 = \Gamma'/\Gamma$ the digamma function, gives the completed split

$$\sum_{n \geq 1} \lambda_n z^n = A(z) - \mathcal{P}(z),$$

where

$$A(z) = \frac{1+z}{1-z} - \frac{\log \pi}{2} \frac{z}{(1-z)^2} + \frac{z}{2(1-z)^2} \psi_0\left(\frac{1}{2(1-z)}\right), \quad \mathcal{P}(z) = \frac{z}{(1-z)^2} \sum_{m \geq 1} \frac{\Lambda(m)}{m^{1/(1-z)}}.$$

Writing $A(z) = \sum A_n z^n$ and $\mathcal{P}(z) = \sum \mathcal{P}_n z^n$,

$$\lambda_n = A_n - \mathcal{P}_n.$$

Proposition 8.1 (Archimedean coefficients). For $n \geq 1$,

$$A_n = 2 - \frac{n}{2} \log \pi + \frac{1}{2} \sum_{r=0}^{n-1} \binom{n}{r+1} \frac{\psi_0^{(r)}(1/2)}{r! 2^r}.$$

Using $\psi_0(1/2) = -\gamma - 2 \log 2$ and, for $r \geq 1$,

$$\psi_0^{(r)}(1/2) = (-1)^{r+1} r! (2^{r+1} - 1) \zeta(r+1),$$

A_n is an explicit finite expression in γ , $\log 2$, $\log \pi$, and $\zeta(2), \dots, \zeta(n)$.

Proof. Expand $\psi_0\left(\frac{1}{2(1-z)}\right)$ at $1/2$ using $\frac{1}{2(1-z)} - \frac{1}{2} = \frac{z}{2(1-z)}$; coefficient extraction gives the binomial sum. \square

Definition 8.2 (Abel regularization). For the prime-power series we use the Abel limit

$$\mathcal{P}_{n,\varepsilon} = \sum_{m \geq 1} \frac{\Lambda(m)}{m^{1+\varepsilon}} L_{n-1}^1(\log m), \quad \mathcal{P}_n = \lim_{\varepsilon \downarrow 0} \mathcal{P}_{n,\varepsilon},$$

the limit taken through the generating function $\mathcal{P}(z)$, where it agrees with coefficient extraction. We write Abel \sum for this regularized sum.

Proposition 8.3 (Abel–Laguerre prime trace). For $n \geq 1$,

$$\mathcal{P}_n = \text{Abel} \sum_{m \geq 1} \frac{\Lambda(m)}{m} L_{n-1}^1(\log m),$$

with L_k^α the generalized Laguerre polynomial and the sum read as the Abel limit of Definition 8.2.

Proof. Use the generating function

$$\frac{1}{(1-z)^2} \exp\left(-\frac{xz}{1-z}\right) = \sum_{k \geq 0} L_k^1(x) z^k, \quad x = \log m,$$

together with $m^{-1/(1-z)} = m^{-1} \exp\left(-\frac{z \log m}{1-z}\right)$, and read off the coefficient of z^n . \square

9. RESIDUAL FORM OF THE PRIME TRACE

Theorem 9.1 (Laguerre residual identity). For $n \geq 1$, in the Abel-regularized sense,

$$\mathcal{P}_n = n + 1 + \int_0^\infty R(e^u) e^{-u} L_{n-1}^2(u) du = n + 1 + \int_1^\infty R(t) \frac{L_{n-1}^2(\log t)}{t^2} dt.$$

Proof. Write $\mathcal{P}_n = \text{Abel} \int_{1^+}^\infty \frac{L_{n-1}^1(\log t)}{t} d\psi(t)$, with $d\psi = dt + dR$. For the flat part, the Abel parameter ε of Definition 8.2 gives, after the substitution $u = \log t$,

$$\int_1^\infty t^{-1-\varepsilon} L_{n-1}^1(\log t) dt = \int_0^\infty e^{-\varepsilon u} L_{n-1}^1(u) du.$$

The Laplace transform of the Laguerre polynomial is $\int_0^\infty e^{-pu} L_{n-1}^1(u) du = p^{-1} (1 - 1/p)^{n-1}$ for $\text{Re } p > 0$, here with $p = \varepsilon$, and the regularized value extracted through $\mathcal{P}(z)$ is 1 as $\varepsilon \downarrow 0$; this is the constant flat contribution. For the dR part, integrate by parts: the lower limit gives n since $R(1^+) = -1$ and $L_{n-1}^1(0) = n$, and

$$\frac{d}{dt} \left(\frac{1}{t} L_{n-1}^1(\log t) \right) = -\frac{1}{t^2} L_{n-1}^2(\log t)$$

follows from $dL_{n-1}^1/du = -L_{n-2}^2$ and $L_{n-1}^1 + L_{n-2}^2 = L_{n-1}^2$. Combining the flat constant, the boundary term, and the dR integral gives the stated identity. \square

10. STRUCTURAL ACCELERATION AND THE FILTERED TRACE

Let S be the backward shift, $(Sy)_n = y_{n-1}$, and define

$$\mathcal{D} = (I - S)^2(I + S), \quad \mathcal{D}y_n = y_n - y_{n-1} - y_{n-2} + y_{n-3}.$$

This is the structural acceleration stencil of the affine clock triangle.

Theorem 10.1 (Laguerre collapse). *For $n \geq 3$,*

$$\mathcal{D}L_{n-1}^1(x) = L_{n-1}(x) - L_{n-3}(x),$$

and for the residual kernel,

$$\mathcal{D}L_{n-1}^2(x) = L_{n-1}(x) + L_{n-2}(x).$$

Proof. Use the index-lowering identity $L_k^\alpha - L_{k-1}^\alpha = L_k^{\alpha-1}$. For the type-1 kernel, $(I - S)L_{n-1}^1 = L_{n-1}$, and the remaining factor $(I - S)(I + S) = I - S^2$ gives

$$\mathcal{D}L_{n-1}^1 = (I - S^2)L_{n-1} = L_{n-1} - L_{n-3}.$$

For type 2, two backward differences give L_{n-1} , and $I + B$ then gives $L_{n-1} + L_{n-2}$. \square

Corollary 10.2 (Filtered prime trace). *For $n \geq 3$,*

$$\mathcal{D}\mathcal{P}_n = \text{Abel} \sum_{m \geq 1} \frac{\Lambda(m)}{m} [L_{n-1}(\log m) - L_{n-3}(\log m)].$$

11. THE ZERO-MODE DETECTOR

The same filtered trace has a meaning on the zero side. Put

$$z = 1 - \frac{1}{s}, \quad s = \frac{1}{1-z}, \quad |z|^2 = 1 - \frac{2\sigma - 1}{\sigma^2 + t^2} \quad (s = \sigma + it),$$

so that

$$\sigma = \frac{1}{2} \iff |z| = 1, \quad \sigma > \frac{1}{2} \iff |z| < 1.$$

The generating function of $\mathcal{D}\mathcal{P}_n$, up to harmless initial terms, is

$$H(z) = -z(1+z) \frac{\zeta'(1/(1-z))}{\zeta(1/(1-z))},$$

and a zero ρ of ζ becomes a pole at $z_\rho = 1 - 1/\rho$.

Theorem 11.1 (Filtered boundedness implies confinement). *If $\mathcal{D}\mathcal{P}_n = O(1)$, then ζ has no nontrivial zero with $\text{Re } \rho > \frac{1}{2}$. By the functional equation, all nontrivial zeros lie on $\text{Re } s = \frac{1}{2}$.*

Proof. Under $z = 1 - 1/s$ the critical line maps to $|z| = 1$, the half-plane $\text{Re } s > \frac{1}{2}$ to the open disk $|z| < 1$, and $\text{Re } s < \frac{1}{2}$ to $|z| > 1$.

We use only the poles of H arising from nontrivial zeros. The pole of ζ at $s = 1$ maps to $z = 0$ and is removed by the filtered construction, up to the displayed initial terms; the trivial zeros $s = -2, -4, \dots$ map to points with $|z| > 1$, outside the disk relevant to $\text{Re } s > \frac{1}{2}$. The prefactor $-z(1+z)$ vanishes only at $z = 0, -1$, and neither is a z_ρ for a nontrivial zero. So any nontrivial zero ρ with $\text{Re } \rho > \frac{1}{2}$ creates an uncancelled pole of H at $z_\rho = 1 - 1/\rho$ in the open disk $|z| < 1$.

Now suppose such a zero exists. Fix one, ρ_0 with $\text{Re } \rho_0 > \frac{1}{2}$, and write $z_0 = 1 - 1/\rho_0$, so $|z_0| < 1$. Choose a radius r with

$$|z_0| < r < 1$$

such that H has no pole on the circle $|z| = r$ (possible since poles are isolated). Because H is meromorphic, only finitely many poles lie in the closed disk $|z| \leq r$. Decompose H there as

$$H(z) = \sum_{\text{poles } z_\rho, |z_\rho| < r} (\text{principal part at } z_\rho) + H_{\text{hol}}(z),$$

where H_{hol} is holomorphic on $|z| < r$ and hence has Taylor coefficients that are $O(\rho^{-n})$ for any $\rho < r$, in particular $O(r^{-n})$. Each interior principal part contributes a finite sum of terms $C_{\rho,j} n^j z_\rho^{-n}$ with $|z_\rho| < r < 1$, and the pole at z_0 has a nonzero principal part. A finite nonzero combination of such exponential-growth terms cannot have bounded coefficients: the term of largest growth rate dominates along a subsequence and is not cancelled by the others. Therefore the Taylor coefficients of H , which equal \mathcal{DP}_n up to finitely many initial terms, are unbounded, contradicting $\mathcal{DP}_n = O(1)$.

Hence no nontrivial zero has $\text{Re } \rho > \frac{1}{2}$. The functional equation $\zeta(s) = \zeta(1-s)$ reflects nontrivial zeros by $\rho \mapsto 1-\rho$, so the absence of right-of-line zeros forces the absence of left-of-line zeros, and every nontrivial zero in the critical strip lies on $\text{Re } s = \frac{1}{2}$. \square

Remark 11.2 (Where the rigor rests). The argument is sound *conditional on* $\mathcal{DP}_n = O(1)$. It needs no global minimal pole: a single off-line zero already forces a nonzero interior principal part inside some circle $|z| = r < 1$, and one nonzero interior exponential mode is enough to break boundedness. The one nonelementary input is standard — an isolated interior pole dictates exponential growth of the Taylor coefficients, and the on-circle poles accumulating at $z = 1$ do not interfere because they all sit at modulus $1 > r$. The single real obstruction is Conjecture 13.1.

12. WHY THE OBVIOUS BOUND FAILS

It is tempting to close $\mathcal{DP}_n = O(1)$ for free. The trace is a Laguerre projection of finite-energy prime-power data, so why not write

$$\sum_{n \geq 3} |\mathcal{DP}_n|^2 \leq C \sum_{m \geq 1} \frac{\Lambda(m)^2}{m^2} < \infty$$

and call it bounded? Because that step is wrong, and it is worth saying exactly why. The filtered trace is the pairing of the Laguerre functions with the discrete weighted distribution

$$\nu = \sum_{m \geq 1} \frac{\Lambda(m)}{m} \delta_{\log m}.$$

The standard Laguerre polynomials are orthonormal for the continuous measure $e^{-x} dx$ on $[0, \infty)$, and orthonormality there says nothing about orthonormality against ν . The finite-energy quantity $\sum_m \Lambda(m)^2 / m^2$ belongs to the *coefficient weights* of this distribution — it says the coefficients $\Lambda(m)/m$ are ℓ^2 — not to any orthogonality measure for the Laguerre basis. It does not say the Laguerre family is a Bessel sequence for the sampling distribution ν , and that is precisely the gap.

The failure is concrete. Taking $G_{ij} = \sum_{m \leq X} (\Lambda(m)/m) L_i(\log m) L_j(\log m)$ with $m \leq 2 \times 10^5$ and prime powers included, the Gram matrix is nowhere near the identity: its diagonal climbs into the thousands with off-diagonals to match (this computation tests the obstruction; it is not used in any proof). Bessel does not apply. Boundedness of \mathcal{DP}_n would be a real arithmetic cancellation theorem, not a free Hilbert-space lemma. That distinction is the whole game.

13. THE PRIME-LAGUERRE FRAME CONJECTURE

Conjecture 13.1 (Filtered prime-Laguerre estimate). *The filtered prime trace is bounded:*

$$\mathcal{DP}_n = O(1),$$

equivalently

$$\text{Abel} \sum_{m \geq 1} \frac{\Lambda(m)}{m} [L_{n-1}(\log m) - L_{n-3}(\log m)] = O(1).$$

By Theorem 11.1, this implies critical-line confinement of the nontrivial zeros of ζ . That is the exact boundary of the paper. The triangle, the wheel lanes, the residual-bilinear identity, the Li–Laguerre trace, and the filtered detector are all proved. What remains is the boundedness of one explicit filtered prime-power Laguerre trace — nothing more, and nothing less.

14. NUMERICAL EVIDENCE

We computed \mathcal{P}_n from the completed split $\mathcal{P}_n = A_n - \lambda_n$, taking A_n from the explicit archimedean formula of Proposition 8.1 and the Li coefficients λ_n from high-precision tabulated values, then formed

$$\mathcal{DP}_n = \mathcal{P}_n - \mathcal{P}_{n-1} - \mathcal{P}_{n-2} + \mathcal{P}_{n-3}, \quad 4 \leq n \leq 500.$$

The largest value seen was

$$\max_{4 \leq n \leq 500} |\mathcal{DP}_n| = 0.234778 \dots, \quad \text{at } n = 4.$$

n	\mathcal{DP}_n
4	0.234778
5	0.171131
10	0.031458
20	-0.015552
50	0.009423
100	-0.008606
200	0.003894
300	0.011993
400	0.013256
500	0.014799

The partial square mass over $4 \leq n \leq 500$ was about 0.207866. This is not a proof of Conjecture 13.1. It is a falsifiable numerical signature, and it behaves the way the conjecture would predict. These computations test the predicted signatures; they are not used in any proof.

15. CONCLUSION

The Bourgeois triangle is a minimal linear medium, and almost everything about it is exact. The unit left border gives an affine clock. The signed right border gives a row-Fourier selector at $\omega = \pi$. The Green cones carry the boundary impulses inward. Under the von Mangoldt doorway the diagonal difference is exactly $\Lambda(n)$, and the columns remember prime powers. The survivor border is a Ramanujan–Fourier wheel signal whose prime-denominator lane classes have relative amplitude $1/(p-1)$, and the foundational lane $p = 2$ is the medium’s own spectral zero.

The same object reaches into the completed zeta world. The Li coefficients split as $\lambda_n = A_n - \mathcal{P}_n$, with A_n an explicit archimedean coefficient and \mathcal{P}_n an Abel–Laguerre prime trace, and the structural acceleration operator collapses that trace to a standard Laguerre sampling sequence. If that filtered trace is bounded, interior poles under $z = 1 - 1/s$ are impossible and the nontrivial zeros are confined to the critical line.

So one challenge is left standing:

$$\mathcal{DP}_n = O(1).$$

Everything else is structure. The seed is explained, the lanes are measured, the doorway is exact, and the last obstruction is one clean prime-Laguerre estimate. That is exactly where we stop, and exactly why we stop there.

REFERENCES

- [1] T. M. Apostol, *Introduction to Analytic Number Theory*, Springer, 1976.
- [2] J. B. Conrey, *The Riemann hypothesis*, Notices Amer. Math. Soc. **50** (2003), 341–353.
- [3] P. Flajolet and R. Sedgewick, *Analytic Combinatorics*, Cambridge University Press, 2009.
- [4] H. G. Gadiyar and R. Padma, *Ramanujan–Fourier series, the Wiener–Khinchine formula and the distribution of prime pairs*, Physica A **269** (1999), 503–510.
- [5] J. C. Lagarias, *Li coefficients for automorphic L-functions*, Ann. Inst. Fourier **57** (2007), 1689–1740.
- [6] X.-J. Li, *The positivity of a sequence of numbers and the Riemann hypothesis*, J. Number Theory **65** (1997), 325–333.
- [7] H. L. Montgomery and R. C. Vaughan, *Multiplicative Number Theory I. Classical Theory*, Cambridge University Press, 2007.
- [8] S. Ramanujan, *On certain trigonometrical sums and their applications in the theory of numbers*, Trans. Cambridge Philos. Soc. **22** (1918), 259–276.
- [9] R. P. Stanley, *Enumerative Combinatorics, Vol. 1*, Cambridge University Press, 2nd ed., 2011.
- [10] E. C. Titchmarsh, *The Theory of the Riemann Zeta-Function*, 2nd ed., rev. D. R. Heath-Brown, Oxford University Press, 1986.