Improved Bounded Gaps Between Primes

Shiva Kintali*

September 1, 2026

Abstract

We give a computer-assisted extension of Stadlmann's recent GPY/Maynard–Tao framework for bounded gaps between primes. A one-band support with $k = 48$, smoothness threshold 0.0182, and uniform cap 0.1899 satisfies the cited Type I, IIa, IIb, and III hypotheses; the Type-IIc parameter interval is empty. An exact rational evaluation of a public 1,295-term symmetric-polynomial witness gives $48J/I = 1.000774701187468444763475111497864\ldots > 1$. Together with an explicit admissible 48-tuple of diameter 236, this proves $H_1 \leq 236$. Every new finite step is checked using exact integer or rational arithmetic and a fail-closed reproducibility manifest. We state precisely the analytic results imported from the literature and the local editorial conventions required to apply the source preprint.

Use of AI: We prepared a summary of proof strategies and techniques used to prove bounded gaps between primes. We fed this to a *meta harness* (developed by the author) to control multiple AI tools (Claude, ChatGPT and two open-weight models with open-source harnesses) and iteratively prompt these AI tools to decrease the bound. When an AI tool was hallucinating, stuck, or going in the wrong direction, we stopped that AI tool, corrected the intermediate `.md` files, added more specific new ideas and switched to a different AI tool with more detailed fresh prompts. Equations and computations written by one AI tool were critiqued, verified and expanded by a completely different AI tool in an iterative manner, till we were satisfied with their correctness and details. The code was written by Codex and verified by the author. The author is solely responsible for the correctness of the proofs and the code.

1. Introduction

A finite set $\mathcal{H} \subset \mathbb{Z}$ is **admissible** if, for every prime p, the reduction of $\mathcal{H}$ modulo p omits at least one residue class. Write

$$H_1 := \liminf_{n\to\infty}(p_{n+1} - p_n)$$

and let $H(k)$ be the least possible diameter $\max \mathcal{H} - \min \mathcal{H}$ of an admissible k-element set.

Theorem 1 (Main theorem)**.** $H_1 \leq 236$.

The proof is computer assisted only in finite, exactly reproducible calculations. It certifies the following three premises:

1. a legal one-band D182 support for $k = 48$, with Stadlmann's Type-IIc range empty;
2. an exact rational variational certificate

$$\frac{48J(F)}{I(F)} = 1.000774701187468444763475111497864\ldots > 1;$$

3. an explicit admissible 48-tuple of diameter 236.

*Email: `shiva.kintali@gmail.com` Homepage: shivakintali.github.io.

Stadlmann's GPY proposition then gives $H_1 \le H(48)$, while the explicit tuple gives $H(48) \le 236$. Section 7 supplies the formal proof of Theorem 1. A separate exhaustive calculation also proves the stronger ancillary statement $H(48) = 236$, but minimality is not needed for the prime-gap bound.

All displayed decimals arising from the certificate are explanatory only: every sign and comparison is made with integers or reduced rational numbers.

The word "unconditional" is used in the ordinary theorem-citation sense. No conjectural distribution hypothesis is assumed; the precise imported analytic results and five source-level corrections are stated below.

Stadlmann's preprint [8] proves $H_1 \le 240$, improving the earlier Polymath8b bound 246 [6]. It begins with Maynard–Tao/GPY square weights [3, 4]. Proposition 1 uses the integrals in the paper's Definition 5. In the prime-indicator specialization $c_1 = c_2 = 0$, and for the one-band support used here, the positivity criterion is exactly

$$\frac{kJ(F)}{I(F)} > 1, \qquad I(F) = \int_{\mathbb{R}^k_{\ge 0}} F(t)^2 \, dt,$$

and

$$J(F) = \int_{\substack{x \in \mathbb{R}^{k-1}_{\ge 0} \\ |x|_1 \le A_1 - \varepsilon}} \left(\int_0^\infty F(x,t)\, dt \right)^2 dx.$$

The integrands are extended by zero outside the support. Thus I is the squared L^2-norm and, in this one-band case, J is the squared norm of the one-coordinate marginal restricted to the epsilon-enlargement base. For a general multi-band support, Definition 5 is instead a sum over ordered band pairs with band-dependent base cutoffs. Its additional K-term is multiplied by c_2 and vanishes in the prime-indicator specialization used here.

The paper's main analytic innovation is a support described by total-sum bands and caps on the sum of coordinates larger than a cutoff δ. The cap shape alone is not sufficient. Once the global inequalities and tuple-partition clauses of Proposition 3 are verified, its proof and Lemmas 11–13 place every relevant modulus from Definition 2 in an applicable Type I, Type IIa, Type IIb, Type IIc, or Type III factorization range, with the classical sub-half range supplied by Bombieri–Vinogradov. For D182 the required numerical hypotheses are checked exactly and the Type-IIc parameter interval is empty. At the displayed Harman parameters the minorant specializes to the prime indicator $\mathbf{1}_{\mathbb{P}}$, with no density or negativity loss.

The paper's displayed numerical choice is

$$k = 49, \qquad \varepsilon = 0.0075, \qquad A = 0.253, \qquad \delta = 0.028,$$

with caps 0.15 for one or two large coordinates and 0.17 thereafter, and $(\xi_1, \xi_2, \xi_3) = (0.38, 0.4, 0.4)$. It states that a matrix computation gives $49J/I > 1$; a diameter-240 admissible tuple then completes its argument. Version 1 does not supply the final matrices, rational vector, quotient, or code, so that numerical claim is not reproducible from the arXiv bundle alone.

The source also has a literal Type-IIc defect: its displayed condition permits $\omega_0 < 0$ while requiring a nonnegative subset sum to be at most $8\omega_0 < 0$. The construction below neither repairs nor uses that case. It makes the entire Type-IIc exponent interval empty.

2. The D182 support

Set

$$\varepsilon = \frac{3}{400}, \quad A = \frac{253}{1000}, \quad U = A + \varepsilon = \frac{521}{2000}, \quad L = A - \varepsilon = \frac{491}{2000},$$

and take

$$\delta = \frac{91}{5000} = 0.0182, \qquad B_0 := 0, \qquad B_m = \frac{1899}{10000} = 0.1899 \quad (1 \le m \le 54).$$

Here $54 = \lfloor 1/\delta \rfloor$. In Stadlmann's band notation this is $n = 1$, $A_0 = -\varepsilon$, and $A_1 = A$. The denominator band is $[0, U)$, while the fixed-coordinate marginal is integrated over bases of total mass at most L. The support is

$$T = \left\{ t \in \mathbb{R}^{48}_{\ge 0} : |t| * 1 < U, \quad R * \delta(t) := \sum_{i : t_i > \delta} t_i \le B_{\#\{i : t_i > \delta\}} \right\}.$$

Use the analytic parameters

$$(\xi_1, \xi_2, \xi_3) = \left(\frac{19}{50}, \frac{2}{5}, \frac{2}{5} \right), \qquad \eta = 10^{-10}.$$

There is one ordered pair of bands, and its excess exponent is

$$\omega = \frac{A + A}{2} - \frac{1}{4} = \frac{3}{1000}.$$

Proposition 2 (D182 analytic support). *With the zero-group convention $B_0 = 0$ explained below, the support T satisfies every global and tuple-partition hypothesis needed for the Type I, Type IIa, Type IIb, and Type III branches of Stadlmann's Proposition 3. The Type-IIc range is empty. Consequently the equidistribution conclusion of that proposition holds for the moduli attached to T.*

Moreover, Proposition 2 of the same source specializes at $\xi_2 = \xi_3 = 2/5$ to $\rho = \mathbf{1}_{\mathbb{P}}$ and $c_1 = c_2 = 0$; its roughness exponent is $\beta = 1 - 2\xi_2 = 1/5 > B_1$.

Proof. We verify the required alternatives and inequalities in five steps.

The Type-IIc range is empty

After the symmetry reduction $\gamma \le 1/2$, the actual Type-II exponent range begins at

$$\gamma_{\mathrm{lo}} = \xi_2 - \eta = \frac{3999999999}{10^{10}} = 0.3999999999.$$

The printed upper endpoint of the Type-IIc branch is

$$\gamma_{\mathrm{IIc,hi}} = \frac{1}{3} + 8\omega + \frac{7\delta}{3} + 3\eta = \frac{3998000003}{10^{10}} = 0.3998000003.$$

Hence its closed parameter interval is empty, with exact gap

$$\gamma_{\mathrm{lo}} - \gamma_{\mathrm{IIc,hi}} = \frac{499999}{2500000000} > 0.$$

The Type-IIb branch covers every operative exponent below

$$\tau = \frac{2}{5} + \frac{24\omega}{5} + \frac{7\delta}{5} + 2\eta = \frac{2199400001}{5000000000} = 0.4398800002,$$

and Type IIa covers $\gamma \ge \tau$. The strict upper endpoint in IIb meets the weak lower endpoint in IIa. Thus the two branches cover the whole operative Type-II range, and neither the Type-IIc estimate nor its defective negative-ω_0 clause is used.

Global inequalities

For reference, the following table gives the left side minus the required right side in each global inequality of Stadlmann's Proposition 3:

branch	margin expression	exact value
I_1	$\xi_1 - 4A + \frac{2}{3} - 2\eta - \delta$	$\frac{246999997}{15000000000}$
I_2	$\frac{9}{7} - \frac{34A}{7} - 2\eta - \delta$	$\frac{1352999993}{35000000000}$
II_0	$\frac{19}{2} - 36A - 13\delta + 100\eta$	$\frac{15540001}{100000000}$
II_1	$\frac{\xi_2}{10} - \frac{32A}{10} + \frac{8}{10} - 2\eta - \delta$	$\frac{60999999}{5000000000}$
II_2	$\frac{\xi_2}{4} + \frac{11}{16} - 3A - 2\eta - \delta$	$\frac{51499999}{5000000000}$
III	$\frac{11}{8} - \frac{7A}{2} - \frac{9\xi_3}{8} - 2\eta - \delta$	$\frac{106499999}{5000000000}$

Every entry in the last column is positive; no floating-point comparison is used.

Tuple-partition inequalities

For $B_1, B_2 \geq 0$ and positive integers m, m', write

$$\Xi(B_1, B_2, m, m', \delta) = \left\{ y \in [\delta, 1]^{m+m'} : \sum_{i=1}^{m} y_i \leq B_1, \ \sum_{i=m+1}^{m+m'} y_i \leq B_2 \right\}.$$

This is the tuple set appearing in Stadlmann's Proposition 3. It records the logarithmic prime-factor exponents contributed by the two rough groups. For every positive pair m, m', a tuple in $\Xi(B, B, m, m', \delta)$ has total mass at most $2B = 0.3798$. Place every coordinate in the first partition bin. The resulting first-bin slacks and the capacities of the unused bins are

branch	first-bin slack	unused-bin capacities
I	$\frac{999999}{5000000000}$	$\frac{2319999997}{15000000000}$
IIa	$\frac{300399999}{5000000000}$	$\frac{2139999993}{35000000000}$
IIb	$\frac{49999999}{2500000000}$	$\frac{135299999}{2500000000}, \frac{724099993}{17500000000}$
III	$\frac{10999999}{5000000000}$	$\frac{787499999}{5000000000}$

All entries are positive. The separate $\gamma > 1/2$ Type-I passage in the source proof has first-bin slack 570999999/5000000000 and unused-bin capacity 1989999993/35000000000, so it is also covered.

The uniform row obeys $\delta < B_m$, and $B_m \leq B_{m+1} \leq B_m + \delta$ whenever both columns are defined. The first inequality has margin 1717/10000; the two adjacent-column margins are 0 and 91/5000.

Zero rough groups and the Harman specialization

Definition 1 of the source prints only $B_{j,m}$ for $m \geq 1$, whereas its relevant-moduli union permits $m = 0$. We use the unique natural convention $B_{j,0} = 0$, with empty products equal to 1 and empty sums equal to 0. The partition lemmas manipulate only the rough exponents that are present, so their proofs remain valid after deleting an absent group; the smooth factor is unchanged. A one-zero tuple has total mass at most B, rather than $2B$, and hence strictly more slack. The both-zero case is the empty-tuple case treated as trivial in the source proof.

Finally, Stadlmann's five Harman conditions, with their exact slacks, are

condition	slack
$2\xi_1 + 3\xi_2 < 2$	$\frac{1}{25}$
$\xi_2 \le \xi_3$	0
$\xi_1 + 9\xi_2 < 4$	$\frac{1}{50}$
$2\xi_1 + \xi_2 > 1$	$\frac{4}{25}$
$17\xi_2 < 7$	$\frac{1}{5}$

where the zero belongs to the permitted weak inequality $\xi_2 \le \xi_3$. The source explicitly states that $\xi_2 \le 2/5$ gives $\rho = \mathbf{1}_{\mathbb{P}}$; thus $c_1 = c_2 = 0$. The rough-prime condition in the GPY proposition holds because

$$\beta - B_1 = \frac{1}{5} - \frac{1899}{10000} = \frac{101}{10000} > 0.$$

This proves Proposition 2. The exact substitutions are recomputed by the D182 checker, whose rational output records every margin. Replacing strict support boundaries by weak ones in the later integrals changes only null sets. $\square$

3. Explicit variational witness

For a partition $\alpha = (\alpha_1, \ldots, \alpha_\ell)$ of positive integers, write $|\alpha| = \sum_i \alpha_i$. The monomial symmetric polynomial $m_\alpha(z_1, \ldots, z_k)$ is the sum of the distinct monomials obtained by permuting $(\alpha_1, \ldots, \alpha_\ell, 0, \ldots, 0)$ among the k variables. This convention fixes the normalization of every coefficient below.

The coefficient file contains 1,295 exact triples $(a, \alpha, c_{a,\alpha})$, with $a \ge 0$ and $a + |\alpha| \le 25$.

It is byte-identical to the public AxiomMath PrimeGapsLib witness at commit `1faa7b14...b106477`. Write

$$P(z) = \sum_{a,\alpha} c_{a,\alpha} \left(\frac{26}{25} - |z| * 1\right)^a m * \alpha(z).$$

Put $z = (26/(25U))t$, let $P_U(t) = P(z)$, and define

$$F(t) = P_U(t)\mathbf{1}_T(t).$$

The rescaling is exact. If $d = a + |\alpha|$, homogeneity gives

$$\left(\frac{26}{25} - |z| * 1\right)^a m * \alpha(z) = \left(\frac{26}{25U}\right)^d (U - |t| * 1)^a m * \alpha(t).$$

Multiplication of the whole polynomial by the common factor $(25U)^{25}$ does not change J/I. Equivalently, after passing to integer coordinates $x = 10000t$ with $U_{\text{int}} = 2605$, multiplication by $(25U_{\text{int}})^{25}$ differs only by another common amplitude factor and produces the integer coefficient

$$c_{a,\alpha} 26^{a+|\alpha|} (25U_{\text{int}})^{25-a-|\alpha|}.$$

The provenance and scaling audit reproduces the public $k = 50$ certificate. Two independently assembled exact calculations agree on the $k = 48$, $U = 0.2605$, $L = 0.2455$ cap-free values. Specifically, put

$$T_0 = \{t \in \mathbb{R}^{48}_{\ge 0} : |t| * 1 < U\}, \qquad F_0 = P_U \mathbf{1} * T_0,$$

$$I_0 = \int_{T_0} F_0(t)^2 \, dt, \qquad J_0 = \int_{\substack{x \in \mathbb{R}^{47}_{\ge 0} \\ |x|_1 \le L}} \left(\int_0^\infty F_0(x, t) \, dt\right)^2 dx.$$

Then

$$q_0 := \frac{48J_0}{I_0} = 1.0009087041759875057075373324887 25\ldots > 1.$$

The exact reduced rational and its positive 409-digit sign integer are in the canonical cap-free artifact.

Lemma 3 (Approximation)**.** *The discontinuities of $F = P_U\mathbf{1}_T$ along cap faces do not obstruct the GPY variational argument.*

Proof. The function F is bounded, symmetric, and supported in the bounded simplex T_0, hence belongs to L^2. Contracting every defining inequality by a factor $1-\theta$ and smoothing produces symmetric compactly supported functions F_θ with $\|F_\theta - F\|_2 \to 0$ as $\theta \downarrow 0$. For the marginal operator

$$(\mathcal{M}G)(x) = \int_0^\infty G(x,t)\,dt,$$

Cauchy–Schwarz on each fiber of length at most U gives $\|\mathcal{M}G\|_2 \le \sqrt{U}\,\|G\|*2$. Consequently both I and J are continuous under this approximation. Since the certified inequality is strict, it persists for all sufficiently small $\theta > 0$. The compact support of $F*\theta$ lies a positive distance from every support face. After taking the repeated antiderivative used in the GPY weights, Stone–Weierstrass approximation on that compact set gives finite sums of products of smooth one-variable functions; symmetrizing those sums preserves the quotient and the support. Differentiating the approximants recovers admissible smooth tensor-product sieve weights. This is precisely the translation–mollification–Stone–Weierstrass argument in the proof of Lemma 1 of [8], which is stated for symmetric square-integrable F. □

4. Exact capped denominator

Retain the preceding notation and put $I = \|F\|_2^2$.

The elementary moment identity underlying the calculation is the following. If ν is a partition, let $\ell(\nu)$ be its length, let $m_q(\nu)$ be the multiplicity of the part q, and put

$$\mathrm{emb}*k(\nu) = \frac{k!}{(k-\ell(\nu))!\prod*q \ge 1 m_q(\nu)!}.$$

For every nonnegative integer b,

$$\int_{\substack{t_i \ge 0\\ |t|_1 \le U}} (U - |t|*1)^b m*\nu(t)\,dt \frac{U^{k+b+|\nu|}\,b!\,\mathrm{emb}*k(\nu)\prod*v \in \nu v!}{(k+b+|\nu|)!}. \tag{4.1}$$

This is the ordinary Dirichlet integral, summed over the distinct embeddings of ν.

Proposition 4 (Exact denominator certificate)**.** *The function $F = P_U\mathbf{1}_T$ satisfies*

$$\frac{I}{I_0} = 0.99989030538386807968349797790 4532\ldots.$$

The fraction represented by this decimal is evaluated exactly in the supplementary certificate.

Proof. Expand P_U^2 into monomial-symmetric signatures. For a fixed number r of coordinates exceeding δ, choose those coordinates, write $t_i = \delta + u_i$ on them, and write the remaining coordinates as $v_j \in [0,\delta]$. The two support inequalities become

$$\sum_i u_i \le B_r - r\delta, \qquad \sum_i u_i + \sum_j v_j \le U - r\delta.$$

Finite inclusion–exclusion removes the upper bounds $v_j \leq \delta$. The remaining iterated simplex moments are instances of (4.1), or its one-dimensional beta-integral consequence. All parameters, exponents, and factorials are integral after the common coordinate and factorial scaling; there is no numerical quadrature.

The expansion contains 7,338 distinct output signatures. Eight disjoint shards evaluate each signature once, and the combiner checks their geometry, coverage, scales, and hashes before summing the integers. Independent rational-polygon integrations in dimensions one and two verify the radial, cap, inclusion–exclusion, and embedding factors. This produces the stated exact fraction. □

Thus the cap removes

$0.00010969461613192031650202209546 7\ldots$ of the denominator energy.

The fail-closed verifier independently checks the eight committed shard hashes and recomputes the combined total.

5. A diagnostic bound not used in the proof

This section records a useful consistency check and explains why the direct marginal calculation in Section 6 is necessary. Put $G = F_0 - F$. Since F is the restriction of F_0, their supports give $\|G\|_2^2 = I_0 - I$.

For a fixed base point x, the deleted part of the final-coordinate fiber has length at most

$$C = U - B = 0.0706.$$

Indeed, if the rough mass of x already exceeds B, then $U - |x| * 1 < U - B$. Otherwise deletion begins no earlier than $t = B - R * \delta(x)$, and its length is at most $U - |x| * 1 - B + R * \delta(x) \leq U - B$. Fiberwise Cauchy–Schwarz therefore gives

$$\|\mathcal{M}G\|_2^2 \leq C(I_0 - I).$$

The reverse triangle inequality applied to $\mathcal{M}F = \mathcal{M}F_0 - \mathcal{M}G$ yields

$$\frac{48J(F)}{I(F)} \geq \frac{\left(\max\{0, \sqrt{q_0} - \sqrt{s}\}\right)^2}{r}, \qquad r = \frac{I}{I_0}, \quad s = 48C(1 - r). \tag{5.1}$$

Set $d = q_0 - r - s$. The right side of (5.1) exceeds 1 precisely when $d > 0$ and $d^2 - 4rs > 0$. Exact arithmetic gives

$$d = 0.00064666567697157445547730210 7\ldots > 0, \qquad d^2 - 4rs = -0.00148635117560815450352222283 1\ldots < 0.$$

Consequently this method supplies only

$$\frac{48J(F)}{I(F)} \geq 0.96280777270898265428473165 0\ldots,$$

which is insufficient. This is only a failed lower-bound method, not an upper bound on the true quotient. No later inference uses (5.1).

6. Exact capped variational certificate

We first justify the geometric decomposition used by the exact program. Fix the first 47 coordinates x and write

$$S = |x| * 1, \quad r = \#\{i : x_i > \delta\}, \quad R = \sum *i : x_i > \delta x_i, \quad V = S - R, \quad H = U - S.$$

If $R \leq B_r$, the allowed endpoint of the last-coordinate fiber is

$$h(x) = \min\{H, \max\{\delta, B_{r+1} - R\}\}. \tag{6.1}$$

If $R > B_r$, the base is illegal and the capped marginal is zero. On a legal base one has $r\delta < R \leq B = 0.1899$ when $r > 0$, hence $r \leq 10$; consequently the index $r + 1$ in (6.1) is always among the defined cap columns. (For $r = 0$, the same conclusion is immediate.) Formula (6.1) follows directly by separating $t \leq \delta$, which adds no rough coordinate, from $t > \delta$, for which the cap is $R + t \leq B_{r+1}$.

Put $C_x = B_{r+1} - R$. The difference between the capped and cap-free marginals is supported on exactly three disjoint regions:

1. **Fixed-cap face:** $R \leq B_r$, $C_x > \delta$, and $H > C_x$.
2. **Fixed-δ shell:** $R \leq B_r$, $C_x \leq \delta$, and $H > \delta$.
3. **Illegal base:** $R > B_r$, where the entire cap-free marginal is removed.

The first two regions satisfy $S < U - \delta = 0.2423 < L = 0.2455$, so they lie automatically inside the marginal-base domain. The illegal-base integration retains the explicit condition $S \leq L$. Thus the three regions are disjoint and exhaustive.

For a typical last-coordinate monomial $t^q(H - t)^a$, the complete primitive is

$$g_{a,q}(H) = \int_0^H t^q(H - t)^a \, dt = \frac{q!\, a!}{(a + q + 1)!} H^{a+q+1}. \tag{6.2}$$

On the fixed-cap face, set $G_x = H - C_x$; the deleted primitive is

$$b_C = \sum_{m=0}^{q} \binom{q}{m} C_x^{q-m} G_x^{a+m+1} \frac{a!\, m!}{(a + m + 1)!}. \tag{6.3}$$

On the fixed-δ shell, set $W = H - \delta$; the deleted primitive is

$$b_\delta = \sum_{m=0}^{q} \binom{q}{m} \delta^{q-m} W^{a+m+1} \frac{a!\, m!}{(a + m + 1)!}. \tag{6.4}$$

Equations (6.3) and (6.4) follow by the substitutions $t = C_x + z$ and $t = \delta + z$, respectively. If g denotes the complete marginal and b the deleted part, the correction on a legal base is

$$(g - b)^2 - g^2 = b^2 - 2gb; \tag{6.5}$$

on an illegal base it is $-g^2$. These identities account for every sign and every diagonal/off-diagonal factor in the contraction.

Proposition 5 (Exact variational certificate)**.** *For the explicit function $F = P_U \mathbf{1}_T$,*

$$\frac{J}{J_0} = 0.999756438739933773404567680738685\ldots$$

and the reduced rational quotient satisfies

$$q_{\mathrm{cap}} := \frac{48J(F)}{I(F)} = \frac{N}{D} = 1.0007747011874684447634751114978 64\ldots > 1,$$

with

$$\frac{N-D}{D} = 0.00077470118746844476347511149786 4\ldots.$$

Proof. The marked-coordinate identity

$$m_\alpha(x,t) = m_\alpha(x) + \sum_{q\in\operatorname{supp}\alpha} t^q\, m_{\alpha-q}(x) \tag{6.6}$$

uses the convention that $\alpha - q$ is the partition obtained by deleting one part equal to q. It reduces each witness term to finitely many primitives of the forms (6.2)–(6.4). Products of monomial-symmetric polynomials are expanded using their exact integer structure constants. The 272 marginal signatures give $272 \cdot 273/2 = 37{,}128$ unordered products and 7,338 distinct output signatures.

After multiplying marginal coefficients by 26!, every coefficient in (6.2)–(6.4) is integral. The 47-dimensional base moments use the common factor $(99!)^2$. Since the computation uses coordinates $x = 10^4 t$, the physical quotient is exactly

$$\frac{48J_{\mathrm{physical}}}{I_{\mathrm{physical}}} = \frac{48J_{\mathrm{stored}}}{10^4\, 99^2 (26!)^2 I_{\mathrm{stored}}}. \tag{6.7}$$

Eight disjoint shards cover the 7,338 output signatures once each. The combiner checks their deterministic blocks, geometry, scales, and external hashes, sums the three correction regions as integers, and reduces (6.7). It obtains the fraction N/D in the statement, with $N - D$ a positive 479-digit integer. The final verifier recomputes $N - D$ from N and D, rather than trusting a stored sign.

Independent checks verify (6.2)–(6.6), all 37,128 symmetric products at three exact specializations, 210 rational-polygon moments at the D182 geometry, a constant-function end-to-end calculation, and the complete factorial and coordinate normalization. This proves the proposition. ☐

The exact numerator N, denominator D, and reduced fraction are the fields `capped_quotient_numerator` and `capped_quotient_denominator` in the completed capped-J artifact and the direct verification manifest. The exact sign field `capped_quotient_minus_one_sign_numerator` stores $N - D$, a positive 479-digit integer; the verifier recomputes it from the reduced fraction rather than trusting the stored sign.

The combined capped-J artifact has SHA-256

`07267dbc5c6c26aa95444d70aed77f62132fe4b7eb516e11b33496b4c8d5459d`

The hostile full-J audit checks the primitive identities, all signature products at three exact specializations, 210 rational-polygon moments at the D182 geometry, a constant-function end-to-end reconstruction, and the complete factorial and coordinate normalization. The fail-closed direct verifier records status `PASSED`.

7. The diameter-236 tuple and theorem

Take

$$\begin{aligned}\mathcal{H} = \{&0, 6, 8, 14, 18, 24, 26, 48, 50, 54, 56, 60, 66, 68, 74, 78, 80, 84,\\ &90, 96, 98, 104, 110, 116, 120, 126, 134, 138, 144, 150, 158, 164, 168, 176,\\ &180, 186, 188, 194, 200, 204, 206, 210, 216, 224, 228, 230, 234, 236\}.\end{aligned}$$

Lemma 6. *The set $\mathcal{H}$ is admissible and has diameter* 236*.*

Proof. The following table gives one omitted residue a_p modulo every prime $p \le 47$:

p	2	3	5	7	11	13	17	19
a_p	1	1	2	2	9	10	4	13

p	23	29	31	37	41	43	47
a_p	13	15	1	7	10	2	5

Direct reduction of the displayed 48 integers verifies that $h \not\equiv a_p \pmod p$ for every $h \in \mathcal{H}$. For every prime $p \ge 53$, cardinality alone prevents a 48-element set from covering all p residue classes. Hence $\mathcal{H}$ is admissible. Its smallest and largest elements are 0 and 236. □

Proposition 7 (Ancillary minimality certificate)**.** *The exact finite search proves $H(48) = 236$.*

Proof. Only the lower bound needs explanation. Translate a hypothetical 48-element admissible set of diameter at most 235 into $[0, 235]$ so that it contains 0. For each prime $p \le 19$, choose a residue omitted by the set. That residue is nonzero because 0 belongs to the set. There are exactly

$$(2-1)(3-1)(5-1)(7-1)(11-1)(13-1)(17-1)(19-1) = 1{,}658{,}880$$

possible omitted-residue vectors. For each vector, delete from $[0, 235]$ every integer lying in one of its omitted classes. Any admissible candidate must be contained in the resulting survivor set.

Exact bit-mask enumeration finds survivor size at most 48. Exactly four vectors attain 48, and in each case the 48 survivors occupy every residue modulo 23. Such a survivor set is not admissible, while every other survivor set has fewer than 48 elements. Therefore no admissible 48-element set has diameter at most 235. Lemma 6 supplies the matching upper bound. The verifier stores the four extremal survivor sets and their SHA-256 digest, so the enumeration is independently recheckable. □

Proof of Theorem 1. Proposition 2 supplies the equidistribution hypothesis for the D182 support. At $\xi_2 = \xi_3 = 2/5$, Stadlmann's minorant proposition gives $\rho = \mathbf{1}_{\mathbb{P}}$, $c_1 = c_2 = 0$, and roughness exponent $\beta = 1/5 > B_1$. Proposition 5 and Lemma 3 supply a symmetric admissible test function with $48J(F)/I(F) > 1$.

Stadlmann's GPY proposition therefore gives $H_1 \le H(48)$. Lemma 6 gives $H(48) \le 236$. Hence

$$H_1 \le 236 < 240.$$

Equivalently, the GPY conclusion produces infinitely many translates of $\mathcal{H}$ containing two primes; between any two such primes lies a pair of consecutive primes separated by at most $\max \mathcal{H} - \min \mathcal{H} = 236$. □

10. Data and code availability

The required files are file are in the directory `comutation/` in the public github repository
github.com/shivakintali/Prime-Gaps

References

[1] R. C. Baker and A. J. Irving, *Bounded intervals containing many primes*, Math. Z. **286** (2017), 821–841.

[2] E. Bombieri, *On the large sieve*, Mathematika **12** (1965), 201–225.

[3] D. A. Goldston, J. Pintz, and C. Y. Yıldırım, *Primes in tuples I*, Ann. of Math. (2) **170** (2009), 819–862.

[4] J. Maynard, *Small gaps between primes*, Ann. of Math. (2) **181** (2015), 383–413.

[5] D. H. J. Polymath, *New equidistribution estimates of Zhang type*, Algebra Number Theory **8** (2014), no. 9, 2067–2199.

[6] D. H. J. Polymath, *Variants of the Selberg sieve, and bounded intervals containing many primes*, Res. Math. Sci. **1** (2014), Art. 12, 83 pp.

[7] J. Stadlmann, *On primes in arithmetic progressions and bounded gaps between many primes*, Adv. Math. **468** (2025), Paper No. 110190.

[8] J. Stadlmann, *Bounded gaps between primes*, arXiv:2608.31126v1 (2026).

[9] A. I. Vinogradov, *The density hypothesis for Dirichlet L-series*, Izv. Akad. Nauk SSSR Ser. Mat. **29** (1965), 903–934.

[10] AxiomMath, *PrimeGapsLib*, commit `1faa7b14e82ddebc2772dfb9153922f01b106477`, https://github.com/AxiomMath/PrimeGapsLib.